The 8th European Congress on Computational Methods in Applied Sciences and Engineering
ECCOMAS Congress 2022
5–-9 June 2022, Oslo, Norway

AUTOMATIC ESTIMATION OF INITIAL TRANSIENT IN A TURBULENT FLOW TIME SERIES

DONNATELLA XAVIER, SALEH REZAEIRAVESH, RICARDO VINUESA AND PHILIPP SCHLATTER

SimEX/FLOW, KTH Engineering Mechanics
Royal Institute of Technology
SE-100 44 Stockholm, Sweden
e-mail: dgxavier@mech.kth.se, www.flow.kth.se

Key words: Variance, initial transient, time series, stationarity, turbulent flow

Abstract. An automatic method is proposed for the removal of the initialization bias that is intrinsic to the output of any statistically stationary simulation. The general techniques based on optimization approaches such as Beyhaghi *et al.* [1] following the Marginal Standard Error Rules (MSER) method of White *et al.* [16] were observed to be highly sensitive to the fluctuations in a time series and resulted in frequent overprediction of the length of the initial truncation. As fluctuations are an innate part of turbulence data, these techniques performed poorly on turbulence quantities, meaning that the local minima was often wrongly interpreted as the minimum variance in the time series and resulted in different transient point predictions for any increments to the sample size. This limitation was overcome by considering the finite difference of the slope of the variance computed in the optimization algorithm. The start of the zero slope region was considered as the initial transient truncation point. This modification to the standard approach eliminated the sensitivity of the scheme, and led to consistent estimates of the transient truncation point, provided that the finite difference time interval was chosen large enough to cover the fluctuations in the time series. Therefore, the step size for the finite difference slope was computed using both visual inspection of the time series and trial and error. We propose the Augmented Dickey–Fuller test as an automatic and reliable method to determine the truncation point, from which the time series is considered stationary and without an initialization bias.

1 INTRODUCTION AND OVERVIEW

In most turbulence simulations, one is interested in the *converged* or *fully-developed* state that the flow will reach after a certain time. This so-called *initial transient* is a collection of states in an unsteady simulation that arises solely due to the imposition of a specific initial condition. Since it represents an unwanted albeit physical state of the solution, it is very important to ensure that the entire initial transient has been discarded from the solution, before any meaningful statistical postprocessing can be done on the solution. For fluid dynamics simulations, the initial transient period is also a very sensitive regime for the solution, when the flow evolves from the user defined initial condition to a flow that is representative of the physics and has been observed to lead to growth of spurious numerical perturbations [12]. Time averages of statistically stationary processes like turbulence computed without discarding the transient can be severely biased. Even a small amount of the initial transient can lead to a large bias in the estimation of the sample means. The problem of determining this warm-up length, or the truncation point of the transient has been well researched in econometrics and systems engineering [7, 8, 15, 16, 17], but is still not thoroughly studied, or even quantified in the turbulence community.

The traditional approach to initial transient estimation for turbulence simulations has been typically through visual and graphical monitoring of the variables. For wall-bounded flows, the integral quantities such as the wall shear stress and the shape factor are graphically inspected during the simulation and the truncation point is usually considered through human judgment, after the values have reached a certain reference value or oscillate about it. This approach although extremely simplistic, relies a lot on the experience of the analyst and is not appropriate for all the variables. For example, it may work for the shape factor, in which there is a smooth decay of the transient, but might be very difficult to determine for a time series that has high intrinsic variance such as a turbulent signal or even if the time series is noisy.

Another approach that was widely used especially in the Direct Numerical Simulation (DNS) of turbulent channel flows to identify the initial transient and convergence to steady state was the assessment of the linear profile of the total shear stress and a quasi-periodic total kinetic energy (*e.g.* Refs. [10, 11]). The linear total shear stress is a direct result of the conservation of momentum and hence indicates that flow has reached the developed turbulent state. The total shear stress is computed as $\tau^+_{xy} = Re^{-1}\partial U^+/\partial y^+ - \overline{uv}^+$ where $U+$ is the mean streamwise velocity normalized in wall units, x and y are the stream- and spanwise coordinates, while $\overline{uv}^+$ is the Reynolds shear stress component in wall units for the channel. Vinuesa *et al.* [14] further modified this criterion to include the momentum balance for ducts and explained the clear distinction between the initial transient time and averaging times in a simulation. This study was one of the first to emphasize on the need for developing a systematic way of finding the transient times by pointing out the different residual values obtained for the same starting conditions, in a wide range of literature (see Tables 3 and 4 in Ref. [14]). Although these methods of identifying the initial transient are more systematic than visual inspections, they are focused on the mean momentum balance and do not provide a comprehensive assessment of the convergence. The linear profile of total shear stress is a good overall indicator of the flow having reached turbulent state for simple flows, but these methods do not assert the stationarity of the process, from a statistical point of view. The mean and higher-order moments may still contain an initialization bias. Another limitation for linear total shear stress is that it can only be applied to turbulent channels, pipes and ducts with periodic boundary conditions in the streamwise direction.

As mentioned earlier, a large contribution to the literature of estimating the initial transient (also referred to as warm-up period or starting length) comes from domains outside the fluid dynamics community in the context of discrete event simulation [7, 8, 13, 16, 17]. Among these, the truncation rule known as Marginal Standard Error Rules (MSER) developed by White (1997) [16] and its batch means version (MSER-m) by Spratt (1998) [13] are renowned techniques used for correlated time series data [15]. If $\{u_i\}_{i=1}^n$ are samples of u obtained from a numerical simulation of a time-dependent process (where $i = 1$ is the initial condition), then the truncation point $i = \hat{k}$ for the removal of initial transient is found by solving the following optimization problem [17]:

$$\hat{k} = \underset{n \gg k \geq 1}{\arg\min} \left[\frac{1}{(n-k-1)^2} \sum_{i=k+1}^{n} (u_i - \bar{u}_{k,n})^2 \right] , \tag{1}$$

where

$$\bar{u}_{k,n} = \frac{1}{n-k} \sum_{i=k+1}^{n} u_i \ .$$

Beyhaghi *et al.* (2018) [1] were the first to test this truncation rule on a real turbulent time series. Eq. 1 was however modified to start from the series of $n/2$ samples rather than an arbitrary value of $k \ll n$. The method entails discarding k initial samples sequentially in order to minimize an estimate of $\sigma^2/(n-k)$, which is (within a multiplicative constant) an estimate of the squared uncertainty of the averaged value of the remaining dataset [1]. When tested on an autoregressive process of order 6, the optimization rule found only one sample to belong to the initial transient, out of 1000 samples of the process. Although the authors have used this method to determine the transient time in the evolution of turbulent kinetic energy of a channel flow at friction Reynolds number of 180, certain details regarding the implementation, the number of iterations and the number of discarded samples out of the total number of observations are missing. In our own implementation of the optimization procedure of Ref. [1] for the turbulent channel velocity time series, we observed the method to be very sensitive to fluctuations in the time series and, more often than not, resulted in overpredicting the transient time. Our observation was also consistent with those of Refs. [7, 16], who claim that this estimation technique is very sensitive to *outliers* in the steady state data. Turbulent flow is characterized by fluctuations and a method that is sensitive to fluctuations in the data cannot be directly applicable.

One immediate way to overcome the sensitivity of the scheme, which is not mentioned in literature, is to compute the change of the variance, *e.g.* by simple finite differences: $\Delta\sigma^2/\Delta t$ of the σ^2 in the iterations of Eq. 1. However, the finite difference time step size Δt is not easy to compute. The chosen Δt has to be large enough to cover the fluctuations in the variance while not too large to remove significant changes in the trend of the time series samples. Therefore, we chose the step through some trial and error. The start of the zero slope region was then taken to be the initial transient truncation point. Once a correct Δt was found, the estimates of this method were consistent with the same prediction of the initial transient value even upon additions to the samples or outliers in the steady state region. Thus the sensitivity could be completed eliminated using this approach, and a cost-effective implementation was possible, even though the predictions of $\hat{k}$ were still somewhat conservative.

A method that is common in econometrics and statistics, but has, to our knowledge, not been used in the fluid dynamics literature, is the **Augmented Dickey–Fuller test (ADF)**. ADF is a unit root test to check if a given time series is stationary. The way we used this test to determine the initial transient was by feeding samples of the time series sequentially starting from the either the stationary or the non stationary, until the initial transient was detected. The ADF test computes the p-values for the null-hypothesis that a unit root is present in the series. A p-value that was less than 0.01 was taken as the indicator for rejection of the null-hypothesis.

With ADF, it was possible to obtain a highly accurate estimate of the stationarity in the series. This method gave a value of truncation point earlier in the series than the method of Ref. [1] and closer to the finite-difference method's predictions, but with a 99% certainty of stationarity in the truncated series. Aside from the provisions for an automatic procedure, the ADF tests are also suited for turbulence time-series analysis, because their underlying detection relies on modelling the time series as an autoregressive model [5] and the finite correlation time within turbulence is clearly aligned with this.

The structure of the paper is as follows: Section 2 reviews in more detail the optimization method of Ref. [1], and our finite difference modification for fast computation of the initial transient. Section 3 discusses the theory behind the Augmented Dickey–Fuller unit root test for stationarity and the algorithm describing our implementation for the computation of the p-values. A brief description of the generation of the time series used in this study which was obtained through the spectral-element solver Nek5000 follows in Section 4. Finally, the initial transient computation using the ADF test for various turbulence flow quantities has been shown in Section 5. The predictions have been compared with those from the method of Beyhaghi *et al.* [1] and from the finite-difference scheme. Note that although this method of initial transient identification has been illustrated only on a turbulent channel flow, our suggested framework can used on any unsteady simulation and, unlike the original MSER or MSER-5, can handle noisy data well. Unlike the method by Vinuesa *et al.* [14], no analytical relations are needed, thus any complicated flow case may be studied.

2 OPTIMIZATION-BASED APPROACHES FOR INITIAL TRANSIENT ESTIMATION

In the optimization-based approaches used to discard the initial transient, the time instant k^* is sought for truncating the time series, which best balances the trade-off between improved accuracy and decreased precision. This implies that the optimal truncation point k^* is one which minimizes the variance of the sample mean of the truncated series, without too much reduction in the sample size. Following Beyhaghi et al. [1], given a sequence of realizations $\{u_1, u_2, u_3, \cdots, u_n\}$ sampled from a simulation, the optimal truncation point k^* for reduction in the initialization bias is obtained by minimizing the variance as follows:

$$k^* = \arg\min_{1\leq k\leq \frac{n}{2}} \frac{1}{(n-k-1)^2} \sum_{i=k+1}^{n} (u_i - \bar{u}_{k,n})^2 \ , \tag{2}$$

where

$$\bar{u}_{k,n} = \frac{1}{n-k} \sum_{i=k+1}^{n} u_i \ . \tag{3}$$

Observe that the truncation point is only sought in the first half of the time series, *i.e.* $k \leq n/2$. Fig. 1 shows the variance for a turbulence time series of streamwise velocity, containing the initial transient region for two sample sizes. The variance has been computed using Eq. 2 by discarding samples one by one until $n/2$. The optimization method seeks for the minimum variance and from Fig. 1, it is evident that the method is very sensitive to the oscillations in the

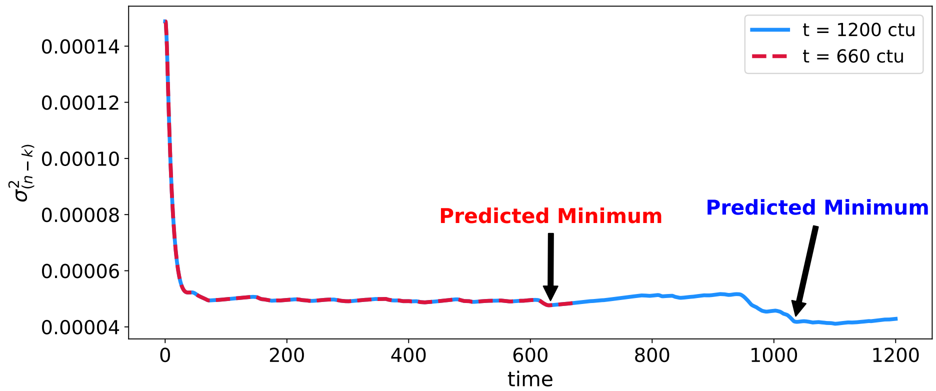

Figure 1: Sensitivity of the predicted minimum using Eq. 2 on the number of samples.

time series, hence resulting in a different value of minimum with a different sample size. The steady state region in the plot, through visual inspection itself, can be seen to start at around 100 convective time units, when the variance plateaus to a zero slope. But the direct application of Eq. 2 ignores these macro changes and the method instead focuses on the various local minima in the curve. It is no surprise that upon varying the length of the samples, the prediction of the minimum variance changes according to the new minimum value that it found based on the local changes in the curve. The method thus, is found to give inconsistent estimates of the initial transient point, k^*. Fig. 1 shows this problem for two sample sizes. It is also obvious from Fig. 1 that the truncation point k^* was often predicted near the end of the time series. Such a restrictive estimate of the transient implies that long simulation run times would be needed for obtaining statistically significant samples. These observations also explain why the minimum is only sought in the first half of the signal as discussed above.

To avoid the high sensitivity of the optimization approach to the fluctuations in the time series, the change of the variance using a finite-difference slope was computed, iteratively in the optimization routine. Note that Δt in the finite difference step size, *i.e.* $\dfrac{\Delta\sigma^2_{(n-k)}}{\Delta t}$ is a decisive parameter in eliminating the sensitivity to the oscillations in the steady state region. This Δt when chosen correctly, can result in a quick prediction of the transient point. The prediction of the truncation point for the transient should be at the start of the zero slope in the variance curve (*i.e.* in the beginning of the plateau), approximated by a threshold value. This is shown in Fig. 2

The finite-difference approximation of the variance removes the fluctuations in the curve and the start of the zero-slope region is predicted accurately. Hence, the time sample $t_k = \hat{k}dt$ at which $\left|\dfrac{\Delta\sigma^2_{(n-k)}}{\Delta t_k}\right| < \epsilon$ (where ϵ is a small value) is predicted by the finite-difference algorithm as the initial transient. This modification to the method of Ref. [1] was found to be more stable because it provided a consistent estimate for the initial transient point, independent of the sample size. However, two parameters remain: i) the step size for the variance calculation and ii) the threshold when the signal can be assumed zero. The problem of finding the step size that sufficiently smoothens the fluctuations in σ^2 is based on visual inspection and trial and

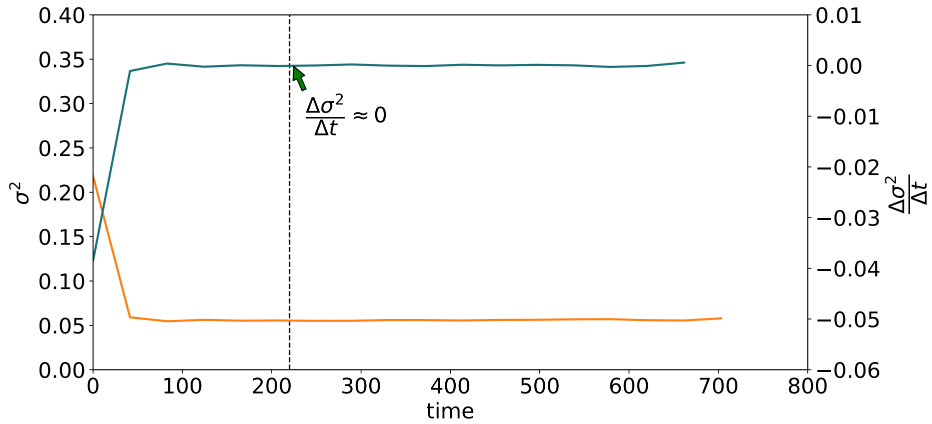

Figure 2: —— Variance (σ^2) and —— finite-difference slope of variance ($\Delta\sigma^2/\Delta t$) for turbulent velocity time series data.

error. Similarly, the threshold value is problem and quantity-dependent. Therefore, this method is not automatic and requires plotting the time series graphically, however it may give some a-posteriori justification of the chosen length of the initial transient, together with a quantification of the variance fluctuations allowed in the statistically stationary part of the time series.

Because of the heuristic parameters in the scheme, the possibility of the unit root test as a viable option for determining the initial transient was investigated. The following section discusses the Augmented Dickey–Fuller test, and how it works to verify the stationarity in the time series.

3 AUGMENTED DICKEY–FULLER TEST

The Augmented Dickey–Fuller (ADF) test is a robust mathematical technique to ascertain the stationarity of a given time series. It is more advanced than the original Dickey–Fuller test [2] in which the time series is approximated by a first order autoregressive process. Since turbulence quantities contain finite temporal correlations, a higher-order autoregressive process is required for modeling the underlying time series and is considered by the ADF test. The ADF technique tests the null hypothesis that a unit root is present in the time series. The presence of a unit root indicates that the time series is not stationary.

The given time series is first modelled by an autoregressive process of order p:

$$u_t = \alpha + \sum_{i=1}^{p} \phi_i u_{t-i} + \epsilon_t \tag{4}$$

where α is a constant, ϕ_i are the coefficients of the autoregressive process until lag p and ϵ_t is a white-noise error term (see Eq.(3.4.31) in Ref. [6]). A regression model can be obtained from Eq. 4 by subtracting the first lag from the equation and reformulating the autoregressive process in differences.

$$\Delta u_t = \alpha + \delta u_{t-1} + \sum_{i=1}^{p} \beta_i \Delta u_{t-i} + \epsilon_t \tag{5}$$

where Δ is the first difference operator and $\delta \equiv \phi_1 - 1$. The unit root test is carried out under the null hypothesis that there exists a unit root and $\delta = 0$. The alternative hypothesis is that $\delta < 0$. It is well known in time series analysis that the existence of a unit root implies a high probability of an explosive time series [5] and non-stationarity. The ADF test starts with this assumption and validates the null hypothesis through the computation of the t-statistic. The t-statistic for δ is computed using the null hypothesis, and its value compared to the Dickey–Fuller distribution,

$$t_{(\hat{\delta})} = \frac{\hat{\delta}}{SE(\hat{\delta})} \ , \tag{6}$$

where $SE(\hat{\delta})$ is the standard error in the estimate. The t-statistic is compared to the relevant critical value for the Dickey–Fuller distribution [2, 5]. If the computed t-statistic is less than the critical value then the null hypothesis is rejected and the time series is proved to be stationary.

In addition to this, the t-statistic can also be computed for the other coefficients $\{\beta_i\}_{i=1}^{p}$ and verified to be less than the critical values of the Dickey–Fuller distribution,

$$t_{\hat{\beta}_i} = \frac{\hat{\beta}_i}{SE(\hat{\beta}_i)} \ . \tag{7}$$

In practice, the ADF is straightforward to implement. The library `statsmodels.tsa.stattools` contains a built-in function named `adfuller` which can be used with the relevant choice of arguments. The function computes the p-value for the null hypothesis. A p-value lower than 0.01 (99% confidence interval) was chosen as the indicator for stationarity in the series. When the initial transient is encountered the p-value switches to a high value. The algorithm used to compute the initial transient is given in **Algorithm 1**.

Algorithm 1: ADF test for calculating initial transient

```
Initialize: time_series, num_samples
m ← 1000 ;                                   // granularity for iterations
initial_transient ← 0
num_iter ← ceiling(num_samples/m)
for i ← 0 to num_iter do
    p_value ← adfuller(time_series[: (i + 1) * m])
    if p_value < tolerance then
        initial_transient ← (i + 1) * m
        break
    end
end
// Truncate the first initial_transient time steps in the series.
modified_series ← time_series[initial_transient :]
```

For a given time series, the ADF test first models it as an autoregressive process. The number of lags that are needed for this autoregressive model depends on the maximum autocorrelation within the time series. For our ADF tests, the minimum number of lags needed for the consistent estimate was found to be 200, and any value larger than this made no difference in the prediction. Note that this of course depends on the exact nature of the time series and its sampling characteristics. For an inhomogeneous turbulent flow, different locations in the flow may have different maximum lags.

4 GENERATION OF TURBULENT TIME SERIES USING Nek5000

The techniques for identification of the initial transient will now be illustrated on time series of turbulent channel flow quantities. In particular, we consider the centreline velocity, the wall-friction velocity and the shape factor, computed on a instantaneous spatial average in the domain. To obtain the time history of these quantities, a well-resolved numerical simulation of the channel was performed using the high-order spectral-element solver Nek5000 [4]. The solver is based on the spectral element method (SEM) that is well-known for its (spectral) accuracy, favourable dispersion properties, and efficient parallelization [3]. In Nek5000, the incompressible Navier—Stokes equations are approximated in space using a Legendre polynomial based SEM, wherein the equations are cast into weak form and discretised by the Galerkin approximation. The basis chosen for the velocity space are typically Nth-order Lagrange polynomial interpolants on Gauss—Lobatto-–Legendre (GLL) points while for Lagrangian interpolants of order $N-2$ are used on Gauss-–Legendre quadrature points are used for pressure space (also known as the P_N-P_{N-2} formulation). For our simulation of the channel flow, a polynomial basis of order 7 was used in each spatial dimension within each element, while the time marching was done with an implicit third-order backward-difference scheme (BDF3) for the viscous terms and third-order extrapolation EXT3 scheme for the nonlinear terms of the Navier–Stokes equations. The turbulent channel was simulated at $Re_\tau = \dfrac{u_\tau h}{\nu} = 300$ where u_τ is the wall friction velocity, h the channel half-height and ν, the viscosity. Since the Reynolds number for our simulation was similar to that of Iwamoto *et al.* [9], the three-dimensional domain was also constructed to match that of Ref. [9]. The simulation was run for about 2000 convective time units (ctu) or equivalently 120 large eddy turnover times (ETT), where ETT is related to the convective time t as $u_\tau t/h$ (here 1 ETT $\approx$ 16 ctu). The initial condition chosen for the simulation was a turbulent flow profile developed by Reichardt disturbed by oblique harmonic modes of 4% amplitude. This initial condition leads to rapid transition to a turbulent state.

5 AUTOMATIC ESTIMATION OF THE INITIAL TRANSIENT FOR TURBULENT CHANNEL TIME SERIES

The initial transient truncation point was computed for three crucial quantities typically reported for turbulent channel flow, namely the centreline velocity u_{CL}, the shape factor and wall-friction velocity. The centreline velocity is representative of the bluntness of the profile and known to be sensitive when it comes to convergence. The wall-friction velocity is related to the pressure drop necessary to drive the flow, and is also used for scaling of various turbulence statistics, so a biased value of u_τ could lead to errors in the computation of other derived quantities. The shape factor is defined as $H = \delta^*/\theta$ where $\delta^* = \int_{-h}^{h} \left(1 - u/u_{CL}\right) dy$ is the displacement thickness and $\theta = \int_{-h}^{h} u/u_{CL}\left(1 - u/u_{CL}\right) dy$ is the momentum thickness. Since

H is an integral quantity, it typically shows less fluctuations. H is representative of the mean profile shape, and known to be sensitive to a proper developed state. Hence, for the flow through a channel, the statistical stationarity of these three quantities are good indicators for the fully-developed state of turbulence.

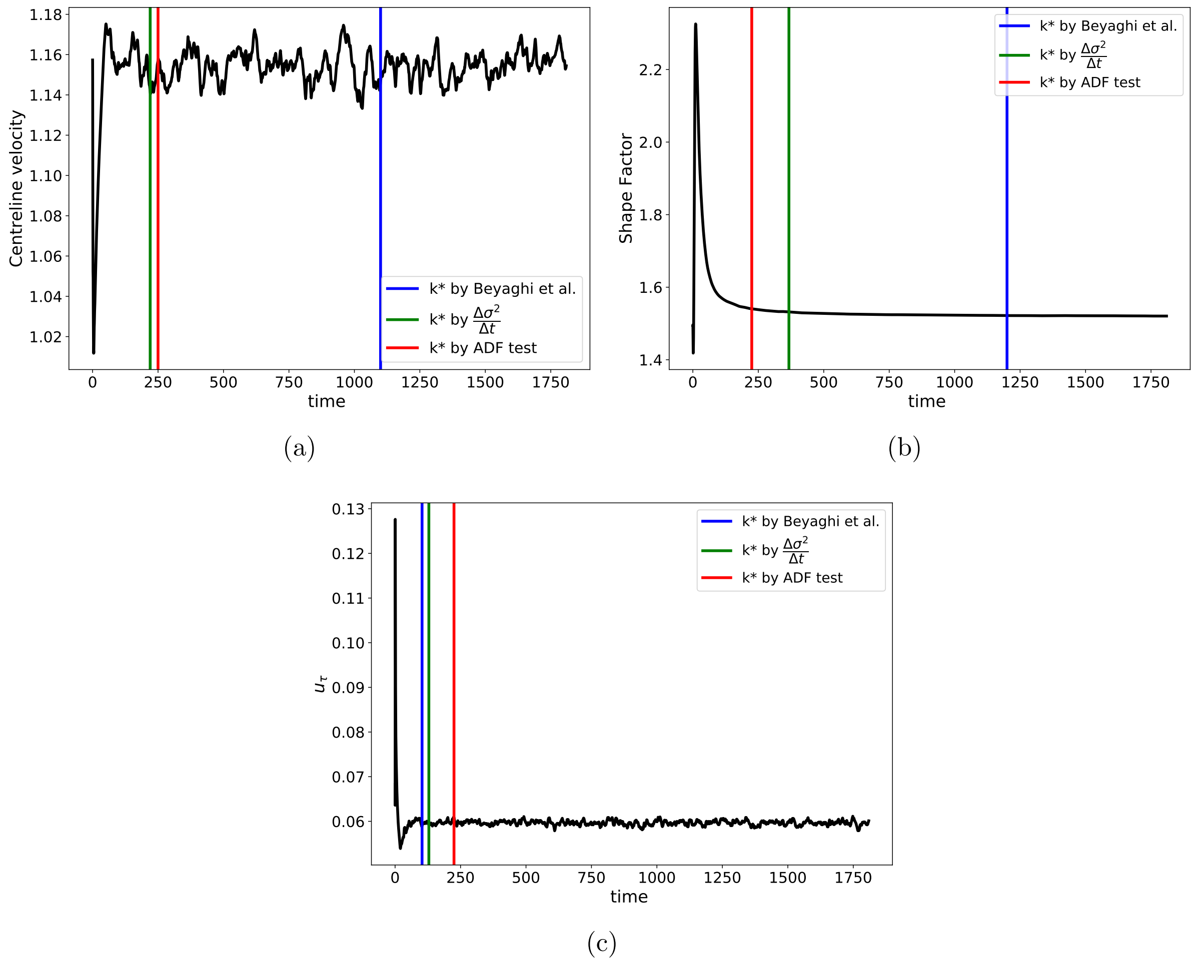

Figure 3: Initial transient prediction for the channel flow at $Re_\tau = 300$ in the time series of (a) centreline velocity, (b) shape factor and (c) wall friction velocity.

The predictions by the three approaches to determine the initial transient in the turbulent time series are shown in Fig. 3(a), (b) and (c). The plots show that the optimization method of Beyhaghi *et al.* from Eq. 2 results in a large overprediction of the initial transient truncation point k^* for the centreline velocity and shape factor accordant with the discussion in Section 2. The estimates of initial transient by the finite difference of the variance and the ADF test do not differ much from each other.

An important observation that can be made regarding the shape of the curves in Fig. 3 is that the time series for the centreline velocity and u_τ are very different from the smooth curve of the shape factor. Since the shape factor is based on a volume integral, its curve tends to be smoother and even *non-stationary* with a weak trend in the steady state part. On the other hand, for the instantaneous centreline velocity and u_τ, past the initial transient, the quantities are found to fluctuate about a stationary mean in the steady state part, with no obvious trend, as compared to the shape factor. The treatment of these two kinds of series, namely one with a slight trend (*i.e.* for H) and one without the trend (*i.e.* for u_τ and centreline velocity u_{CL}) in the ADF test is not the same. For the standard stochastic process, like centreline velocity and u_τ, the time series is checked for stationary by sequentially feeding samples from the beginning of the time series and incrementing the samples. The corresponding p-values then start out as non-stationary and then suddenly switch to stationary, once the initial transient region has passed. For the time series pertaining to the shape factor, the samples are fed starting from the stationary part of the series and moved towards the region containing the initial transient, working backward through the samples in time. These arguments are also reflected in Fig. 4 where the p-values have been plotted for the three different time series. The p-values corresponding to the stationary regime of the flow is clearly visible in this plot from 250 time units as a nearly horizontal line.

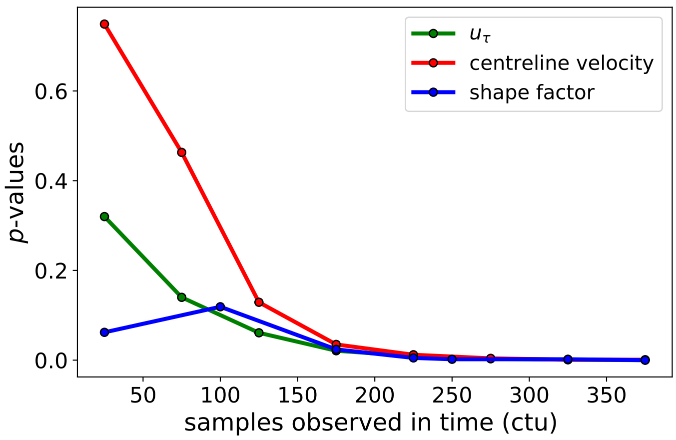

Figure 4: p-values from the ADF test for the turbulent time series.

Thus, it is clear from Fig. 3 and the corresponding p-values of Fig. 4, that the predictions of the initial transient by the ADF tests for the three considered time series lies around 250 convective time units. The finite difference method predicts transient region similar to the ADF for the centreline velocity but a more conservative value for the shape factor (around 370 ctu) and an earlier value (around 125 ctu) for the u_τ.

6 CONCLUSIONS

The initial transient period in a simulation aiming at reproducing a statistically stationary state is not representative of the developed flow state and hence must be discarded completely before starting any meaningful statistical postprocessing on the solution. Failing to do so would result in an initialisation bias in the time averages of the solution. For turbulent time series, the optimization-based approaches for identifying the initial transient truncation point were found to be highly sensitive to fluctuations in the computed variance and resulted in inconsistent estimates of the transient point. By computing the slope of the variance of the samples using finite differences, it was possible to eliminate the sensitivity to outliers in the series. The time series is truncated from the start of an approximately-zero slope in the variance. Although this method was insensitive to outliers and provided consistent estimates of the transient truncation point, its main drawback was the heuristic approach used to choose the time step for the finite difference slope calculation, and the threshold for zero variance slope. The method could involve many iterations of trial and error before converging on the time step that gave a reasonable estimate of the initial transient.

The Augmented Dickey–Fuller (ADF) test overcomes these limitations of the optimization approaches, by using a purely statistical time series analysis approach to determining the transient part of the series. The ADF test works by estimating the t-statistics for the series under the null hypothesis that the time series contains a unit root and is non-stationary. The practical implementation of ADF involves computing the p-values which indicates with 99% confidence whether the given time series was stationary or not. By feeding parts of the time series from the beginning or the end, it was possible to get a more accurate, less conservative estimate of the initial transient point, past which the remaining series had reached statistically steady state. For the turbulent channel flow at $Re_\tau = 300$, the ADF tests predicted the initial transient to exist until at least 225 convective time units, for the three turbulent quantities that were monitored, namely the centreline velocity, shape factor and wall friction velocity. These values were also closer to the predictions made by the finite difference method, confirming the reliability of the estimates.

The ADF tests are straightforward to implement, their only drawback being that a certain minimum number of samples is needed to obtain the correct estimates, because the underlying the autoregressive model has to be able to model the time series accurately. In turbulence time series, the autocorrelations extend until integral timescales, and hence autoregressive models need sufficient number of lags to represent the time series accurately. Future work will involve understanding and tuning the hyperparameters governing the ADF tests, for better control of the algorithm and extensions to more complex quantities and applications.

Acknowledgments: Financial support by the Knut and Alice Wallenberg foundation, the Foundation for Strategic Research and the EXCELLERAT project (grant agreement No 823691 from the European Union's Horizon 2020 research and innovation programme) are gratefully acknowledged. The computations were/was enabled by resources provided by the Swedish National Infrastructure for Computing (SNIC), partially funded by the Swedish Research Council through grant agreement no. 2018-05973.

REFERENCES

[1] Beyhaghi P., Alimo S. R. and Bewley T. R., *A multiscale, asymptotically unbiased approach to uncertainty quantification in the numerical approximation of infinite time-averaged statistics*, arXiv:1802.01056v1 (2018)

[2] Dickey, D. A. and Fuller, W. A., *Distribution of the Estimators for Autoregressive Time Series with a Unit Root*, Journal of the American Statistical Association, 74 (366), pp:427–431 (1979).

[3] El Khoury G. K., Schlatter P., Noorani A., Fischer P. F., Brethouwer G., Johansson A.V. *Direct Numerical Simulation of Turbulent Pipe Flow at Moderately High Reynolds Numbers*, Flow Turbul. Combust. 91:475 (2013)

[4] Fischer P. F., Lottes J. W., Kerkemeier S. G., *Nek5000: Open source spectral element CFD solver* `http://nek5000.mcs.anl.gov` (2008)

[5] Fuller, W. A., *Introduction to Statistical Time Series*, New York: John Wiley and Sons. ISBN 0-471-28715-6 (1976)

[6] Hamilton J. D., *Time Series Analysis*, First Edition, Princeton University Press (1994).

[7] Hoad, K., Robinson, S., and Davies, R., *Automating warm-up length estimation*, Proceedings of the 2008 Winter Simulation Conference.

[8] Hoad, K. and Robinson, S., *Implementing MSER-5 in commercial simulation software and its wider implications*, Proceedings of the 2011 Winter Simulation Conference (WSC), pp.:495-503 (2011).

[9] Iwamoto, K., Suzuki, Y., and Kasagi, N., *Reynolds number effect on wall turbulence: toward effective feedback control*, Int. J. Heat Fluid Flow 23, 678 (2002).

[10] Kim J., Moin P., and Moser R. D., *Turbulence statistics in fully developed channel flow at low Reynolds number*. J. Fluid Mech. 177:133 (1987).

[11] Lee M. K., Moser R. D., *Direct numerical simulation of turbulent channel flow up to $Re_\tau = 5200$*. J Fluid Mech. 774:395 (2015)

[12] Sánchez Abad, N., Vinuesa, R., Schlatter, P., Andersson, M., and Karlsson, M., *Simulation strategies for the Food and Drug Administration nozzle using Nek5000*, AIP Adv. 10, 025033 (2020); https://doi.org/10.1063/1.5142703

[13] Spratt, S. C., *Heuristics for the startup problem*, M.S. Thesis, Department OS Systems Engineering, University of Virginia (1998)

[14] Vinuesa R., Prus C., Schlatter P. Nagib H. M., *Convergence of numerical simulations of turbulent wall-bounded flows and mean cross-flow structure of rectangular ducts*, Meccanica 51:3025–3042 (2016)

[15] White, K. P. and Spratt, S. C., *A comparison of five steady-state truncation heuristics for simulation*, Proceedings of the 2000 Winter Simulation Conference, pp:755-760 (2000)

[16] White, K. P., *An effective truncation heuristic for bias reduction in simulation output*, Simulation, 69(6):323-334 (1997)

[17] White, K. P. and Robinson, S.,*The problem of the initial transient (again) or why MSER works: Proceedings of the 2009 INFORMS Simulation Society Research Workshop*, J. Simulation (2011)