

# Impact of the floating-point precision and interpolation scheme on the results of DNS of turbulence by pseudo-spectral codes

Holger Homann, Jürgen Dreher and Rainer Grauer

*Institut für Theoretische Physik I, Ruhr-Universität-Bochum, 44780 Bochum, Germany*

---

## Abstract

In this paper we investigate the impact of the floating-point precision and interpolation scheme on the results of direct numerical simulations (DNS) of turbulence by pseudo-spectral codes. Three different types of floating-point precision configurations show no differences in the statistical results. This implies that single precision computations allow for increased Reynolds numbers due to the reduced amount of memory needed. The interpolation scheme for obtaining velocity values at particle positions has a noticeable impact on the Lagrangian acceleration statistics. A tri-cubic scheme results in a slightly broader acceleration probability density function than a tri-linear scheme. Furthermore the scaling behavior obtained by the cubic interpolation scheme exhibits a tendency towards a slightly increased degree of intermittency compared to the linear one.

*Key words:* Direct numerical simulations, Spectral methods, Interpolation methods, Turbulent flows, High-Reynolds-number turbulence, Isotropic turbulence; homogeneous turbulence

*PACS:* 47.27.ek, 47.27.er, 02.60.Ed, 47.27 -i, 47.27.Jv, 47.27.Gs

---

## 1 Introduction

Due to the increasing speed and memory of computers, numerical simulations of the Navier-Stokes equations have become a valuable tool for studying turbulence. To investigate the intrinsic properties of turbulence such as the energy cascade and intermittency one often uses periodic boundary conditions in order to keep the influence of the geometry on the flow structure as small as possible. For this kind of turbulence simulations pseudo-spectral codes are widely used [1–10]. These treat the Navier-Stokes equations in Fourier space

	fields	convolutions	interpolation
RUN1	double precision	double precision	tri-cubic
RUN2	single precision	double precision	tri-cubic
RUN3	single precision	single precision	tri-cubic
RUN4	double precision	double precision	tri-linear

Table 1

List of floating-point precision configurations and interpolation schemes for RUN1 – RUN4.

and compute the convolutions in real space. A fast Fourier transformation is used to switch between the two spaces.

A fundamental feature of turbulence is the inertial range of scales. Within this range energy solely cascades by the inertial interaction of eddies, which causes the turbulence to be scale-independent. Attached to this range at small scales is the dissipation range where dissipative effects are important and damp the turbulent motion. The size of the inertial range is directly connected to the Reynolds number  $Re = VL/\nu$ , where  $V$  and  $L$  denote the velocity and size of the large scale motion and  $\nu$  the kinematic viscosity of the fluid. Many theories of turbulence deal with an infinite Reynolds number [11] which means that they focus on the properties of the inertial range. Unfortunately the Reynolds number  $Re$  is connected to the degrees of freedom  $N$  of the turbulent flow according to  $Re \sim N^{9/4}$ . Therefore the memory and speed of computer limits the achievable Reynolds number and size of the inertial range in direct numerical simulations.

Nearly all numerical simulations are performed with double floating-point precision. Recently, the largest numerical simulation worldwide [6] performed on the Earth Simulator used single precision data for the velocity field and double precision data for the calculation of the convolution sums in order to reduce the amount of memory needed. For this reason it was possible to set up a simulation of  $4096^3$  grid points.

In this paper we examine the impact of the floating-point precision on the numerical results. Therefore we performed simulations using three different configurations of floating-point precision (see Table 1). The first configuration, which is the most common approach, computes all fields with double precision (RUN1). The second corresponds to the configuration on the Earth Simulator which uses single precision for the velocity fields and double precision for the convolutions (RUN2). The third uses single precision for all fields (RUN3), which halves the needed amount of memory compared to RUN1 and therefore allows for an increased Reynolds number.

Interesting quantities in turbulence are obtained from statistical averages in

space such as the energy spectrum and Eulerian structure functions or in time such as the acceleration statistics of passive tracers and Lagrangian structure functions. Especially for the Lagrangian statistics a lot of effort has been invested into numerical simulations [7–10] and experiments [12,13] of tracers and heavy particles and in turbulent flows. The integration of the tracers (and particles) requires the interpolation of the velocity field at the tracer positions. In order to obtain reliable Lagrangian statistics it is necessary to achieve accurate trajectories from the simulations. The crucial point is the interpolation scheme used. There are three different schemes applied in the literature. The first is based on cubic splines [7,8], the second is a tri-cubic scheme [14] and the third a tri-linear scheme [9]. In this paper we investigate the impact of the interpolation scheme on the Lagrangian statistical results. The scheme using splines has the drawback that it is difficult to parallelize for distributed memory computer due to its lack of locality. Most of the current massive parallel computers are of this type of memory so we will not consider interpolation schemes using splines. For comparison we performed a simulation with a tri-cubic (RUN1 – RUN3) and a tri-linear (RUN4) interpolation scheme.

## 2 The Numerics

The simulations are performed by numerically solving the incompressible Navier-Stokes equations,

$$\partial_t \mathbf{u} + (\mathbf{u} \cdot \nabla) \mathbf{u} = -\nabla p + \nu \Delta \mathbf{u}, \quad (1)$$

$$\nabla \cdot \mathbf{u} = 0, \quad (2)$$

in a periodic cube with  $256^3$  collocation points with a pseudo-spectral method using spherical mode truncation to reduce aliasing effects [3]. The non-linear term in (1) would be a convolution in Fourier space and therefore is treated in real space. The dissipation term can be computed exactly in Fourier space. The equation of continuity (2) determines the pressure  $p$  and is satisfied by first neglecting the pressure term and afterwards projecting the velocity on its solenoidal part.

We parallelize the computations via slab geometry. This means every process holds the data of a sub-slice of the whole cube. The parameters of the simulations RUN1 – RUN4 are listed in Table 2. For the inter-process communication we use the Message Passing Interface (MPI). The FFTs are performed by the MPI-parallel C-library FFTW (version 2.1.5) [15]. The time integration of the velocity field is done by a strongly stable Runge-Kutta third order scheme [16].

In order to advance the tracers the velocity field has to be interpolated at

$Re$	$u_0$	$\epsilon$	$\nu$	$dx$	$\eta$	$\tau_\eta$	$L$	$T_L$	$N^3$	$N_p$
1575	0.15	0.0015	0.0002	0.0245	0.0086	0.37	2.1	14	$256^3$	$1 \cdot 10^5$

Table 2

Parameters of the numerical simulations.  $Re$ : Reynolds number  $VL/\nu$ ,  $u_0 = \sqrt{2/3E_k}$ ,  $E_k$ : kinetic energy,  $\epsilon$ : mean energy dissipation rate,  $\nu$ : kinematic viscosity,  $\eta$ : dissipation length scale  $(\nu^3/\epsilon)^{1/4}$ ,  $\tau_\eta$ : Kolmogorov time scale  $(\nu/\epsilon)^{1/2}$ ,  $L = (2/3E)^{3/2}/\epsilon$ : integral scale,  $T_L = L/u_0$ : large-eddy turnover time,  $N^3$ : number of collocation points,  $N_p$ : number of particles

the tracer positions. The velocity field is given on a equally spaced grid and evolves according to the Navier-Stokes equations. The tracers have to be integrated at run-time due to the limited amount of hard-disk space. This is done using the same Runge-Kutta third order scheme as for the integration of the velocity field. To interpolate the velocity field at the positions of the individual particles, we implemented a tri-cubic and a tri-linear scheme. The tri-cubic interpolation relies on prescribing the values of the function, its first, mixed second and mixed third derivative at the eight corners of the particle surrounding cubic grid cell (see [17] for the 2D version). For calculating the derivatives we use a first order centered scheme. The tri-linear interpolation just needs the values of the function at the corners. Both schemes parallelize very efficiently.

The turbulence is driven by keeping constant the Fourier modes with  $1 \leq |\mathbf{k}| \leq 2$ . The starting point of the different runs is the same snapshot of a statistically stationary flow. Subsequently we injected  $1 \cdot 10^5$  particles into the flow and integrated the system for approximately 4 large eddy turnover times. The statistics are computed from the last 3 large eddy turnover times for all runs.

### 3 Eulerian statistics

A central point of many theories of turbulence [11] are the scaling laws of the energy spectrum within the inertial range. Figure 1 shows the computed energy spectra from RUN1 – RUN3. Because of the limited Reynolds number, no clear scaling range is visible. More important for the purpose of this paper is the fact that the spectra of the different runs are hardly distinguishable. Even higher order statistics such as the Eulerian longitudinal structure functions

$$S_p(l) = \left\langle |(\mathbf{u}(\mathbf{x} + \mathbf{l}) - \mathbf{u}(\mathbf{x})) \cdot \hat{\mathbf{l}}|^p \right\rangle, \quad (3)$$

angular brackets denoting spatial averaging, look the same for all types of floating-point precisions (see Figure 2). A more subtle comparison yields the according Eulerian probability density function (PDF) of velocity increments.




Fig. 1. Energy spectra from RUN1 – RUN3




Fig. 2. Eulerian longitudinal structure functions  $S_2$ ,  $S_4$  and  $S_6$  (see (3)) from top to bottom from RUN1 – RUN3

Figure 3 shows the PDFs for the variable  $u_x(x + 2dx) - u_x(x)$  normalized to unit variance. The PDFs differ just within the statistical errors due to the finite statistical ensemble, but the overall shape is identical. Therefore the floating-point precision has no impact on the considered Eulerian statistical properties of turbulent flows. One can suspect that also other Eulerian statistical quantities are unaffected by the underlying floating-point precision.

Now the question will be addressed up to which grid resolution these results remain true. A rough estimate can be given by determining the threshold at




Fig. 3. Eulerian probability density function for  $u_x(x + 2dx) - u_x(x)$  from RUN1 – RUN3, normalized to unit variance




Fig. 4. Energy spectra for simulations with single floating-point precision reduced by several bits

which the statistics are affected by the floating-point precision. Starting from single floating-point precision for all fields we artificially reduced the precision by cutting off several least significant bits of the mantissa. Figure 4 shows energy spectra from simulations with single and reduced floating-point precision. Reducing the floating-point precision by nine bits yields unaffected statistical results, while a reduction of 11 bits slightly spoils the spectrum. Cutting off 13 bits yields a clearly modified energy spectrum. These findings are similar for high-order structure functions. The threshold is therefore approximately nine




Fig. 5. Fluctuating total energy for simulations with single floating-point precision reduced by several bits

bits which can also be seen from Figure 5. Here the temporal evolution of the total energy is depicted. Besides the fluctuations, the dynamics are affected at a reduction of 11 bits.

In order to give an estimate up to which grid size high Reynolds number simulations can be performed in single precision, where the ratio of the grid spacing  $dx$  to the dissipation length scale  $\eta$  is kept constant, we assume a Kolmogorov scaling of the energy spectrum. This means that the energy range will be enlarged by a factor of  $n^{5/3}$  when increasing the grid resolution by a factor of  $n$  in each direction. Enlarging a number by nine bits yields a 512 times larger number. We therefore expect that it is still possible to use single floating-point precision data with  $4096^3$  grid points.

#### 4 Lagrangian statistics

Apart from the Eulerian formulation where the evolution of the velocity field is considered at fixed points in space, the Lagrangian coordinates follow the fluid elements in time. Figure 6 shows trajectories of a single particle starting at  $x$  for RUN1 – RUN4. The particles are seeded into the flow at the same time and same positions for all runs and are followed for approximately 4 large-eddy turnover times. The trajectories differ substantially. RUN2 stays close to RUN1 for the longest time as one would expect. The sudden aberration of RUN2, RUN3 and RUN4 from RUN1 is due to the chaotic character of the turbulent flow. The trajectory of RUN4 deviates first from the others. This




Fig. 6. Trajectories of a single particle from RUN1 – RUN4

already implies that the interpolation scheme has a greater impact on the trajectories than the floating-point precision.

To analyze the impact of differing single particle trajectories on the statistics of an ensemble of fluid elements we computed the PDFs of the Lagrangian acceleration (acceleration of fluid elements), which equals the PDF of the Lagrangian velocity increments

$$|(\mathbf{u}(t + \tau) - \mathbf{u}(t)) \cdot \hat{\mathbf{e}}| \quad (4)$$

for small time increments  $\tau$ . The velocity increment after a time-lag  $\tau$  is projected onto a coordinate axis  $\hat{\mathbf{e}}$ . These PDFs for RUN1 – RUN4 are shown in Figure 7. The differences between the runs with different floating-point precisions differ just within the statistical fluctuations. The PDF computed with the tri-linear interpolation in RUN4 is slightly more narrow than the PDF with the tri-cubic interpolation in RUN1. This is because the tri-cubic interpolation scheme is more capable to follow the trajectories of the nearly singular structures (vortex tubes) which are presumably responsible for the stretched tails of the PDFs [12]. As the broadness reflects the degree of intermittency and the Reynolds number of the turbulent flow [3], the tri-linear interpolation scheme might underestimate the degree of intermittency. Figure 8 shows the corresponding Lagrangian structure functions

$$S_p(\tau) = \langle |(\mathbf{u}(t + \tau) - \mathbf{u}(t)) \cdot \hat{\mathbf{e}}|^p \rangle, \quad (5)$$

angular brackets denoting temporal averaging. These functions clearly exhibit no differences for RUN1 – RUN3, i.e. depending on the floating-point precision. Concerning the interpolation scheme, the Lagrangian structure functions slightly differ for RUN1 and RUN4. As for the PDFs, the interpolation scheme has a small impact on the shape of the measured Lagrangian structure func-






Fig. 7. PDFs of Lagrangian velocity increments (4) from RUN1 – RUN4, normalized to unit variance

tions. The inset of 8 shows the logarithmic derivative of the second and fourth order structure function. Due to the limited Reynolds number no clear scaling region is observable and no absolute scaling exponents can be extracted. We computed the relative scaling exponents using the assumption of extended self-similarity (ESS) (see [18]). Figure 9 shows the relative scaling exponents. Despite the large error-bars the linear interpolation scheme systematically yields less intermittent statistics than the cubic interpolation scheme. This is in agreement with the observation concerning the Lagrangian PDFs. The tricubic interpolation scheme results in more stretched tails than the tri-linear scheme. An increased degree of intermittency is reflected in a more intermittent scaling behavior of the Lagrangian structure functions. The higher the order of the scaling exponent the larger is the difference between the interpolation schemes. High-order structure functions are determined by the most intense events resulting from the most singular structures in the flow. These are spiral motions close to strong vortex filaments. Our findings show that the interpolation scheme has an influence on the computation of the Lagrangian scaling behavior. It has to be stressed that these findings also remain true for simulations using less or more collocation points if the ratio of the dissipation length scale and the grid spacing is similar to that of the simulations presented. This ratio determines the smoothness of discretized velocity field. For higher ratios the differences due to the interpolation scheme might become smaller. However, high Reynolds number simulations commonly use ratios similar to the one used in this work. Thus, the accuracy of the interpolation scheme will have an influence on the statistical results independently of the number of collocation points.




Fig. 8. Lagrangian structure functions (5) from RUN1 – RUN4, inset: logarithmic derivative of  $S_2$  (bottom) and  $S_4$  (top)

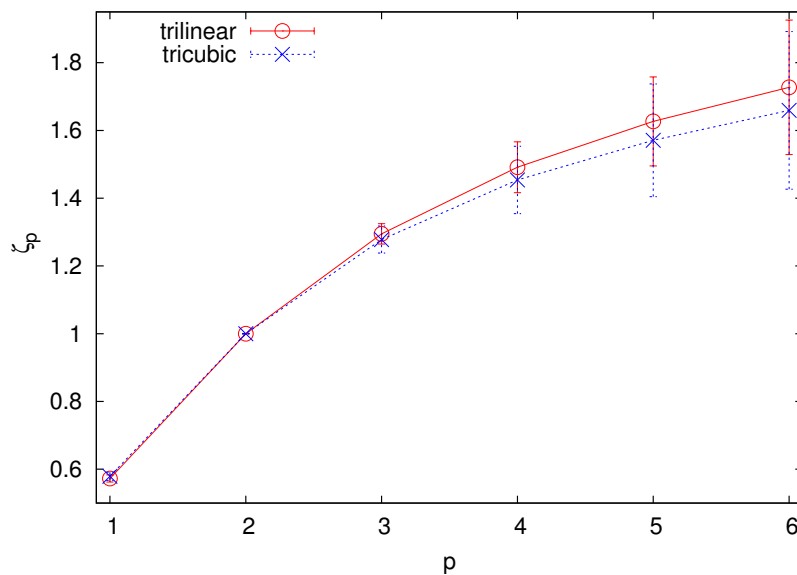


Fig. 9. Lagrangian scaling exponents using ESS for the tri-linear (RUN4) and tri-cubic (RUN1) interpolation scheme

## 5 Conclusions

We investigated the impact of the floating-point precision and the interpolation scheme on the results of DNS of turbulence by pseudo-spectral codes. To analyze the influence of the floating-point precision, we performed three runs. First we used double precision for all fields (RUN1). Second, we used single precision for the velocity fields and double precision for the convolu-

tion fields (RUN2), like the currently largest simulation of turbulence on the Earth Simulator did. The third simulation uses single precision for all fields (RUN3), which halves the amount of memory needed compared to RUN1 and therefore allows for an increased Reynolds number. Although the trajectories of tracer particles differ depending on the floating-point precision which is due to the chaotic nature of turbulence, the statistical quantities such as the energy spectrum, the Eulerian and Lagrangian structure functions and the corresponding PDFs show only minor, negligible differences. Therefore it is possible to achieve the same statistical results by halving the amount of memory needed.

The differences according to the interpolation scheme are more pronounced. Again, the single particle trajectories differ substantially between the tri-cubic and the tri-linear scheme. Furthermore the Lagrangian PDFs have slightly differing shapes. The tri-cubic interpolation scheme reflects a higher degree of intermittency than the tri-linear one. This is because the first is more capable to resolve the nearly singular structures (vortex tubes). These are responsible for the intermittency of the flow. The Lagrangian structure functions and their logarithmic derivative show slightly different shapes and scaling behavior. The tri-cubic interpolation scheme yields a slightly more intermittent statistic than the tri-linear scheme.

## Acknowledgements

Access to the JUMP multiprocessor computer at the FZ Jülich was made available through project HB022. Part of the computations were performed on an Linux-Opteron cluster supported by HBFG-108-291. This work benefitted from support through SFB 591 of the Deutsche Forschungsgesellschaft.

## References

- [1] D. Gottlieb, S. A. Orszag, Numerical Analysis of Spectral Methods: Theory and Applications., SIAM-CBMS (1977), Philadelphia
- [2] C. Canuto, M. Y. Hussaini, A. Quarterioni, T. A. Zang, Spectral Methods in Fluid Dynamics, Springer-Verlag (1988), New York
- [3] A. Vincent, M. Meneguzzi, The spatial structure and statistical properties of homogeneous turbulence, *J. Fluid Mech.* 225 (1991) 1
- [4] Z.-S. She, S. Chen, G. Doolen, R. H. Kraichnan, S. A. Orszag, Reynolds Number Dependence of Isotropic Navier–Stokes Turbulence, *Phys. Rev. Lett.* 70 (1993) 3251.

- [5] W.-C. Müller, D. Biskamp, Scaling properties of three-dimensional magnetohydrodynamic turbulence, *Phys. Rev. Lett.* 84 (1999) 475.
- [6] M. Yokokawa, K. Itakura, A. Uno, T. Ishihara, Y. Kaneda, 16.4-tflops direct numerical simulation of turbulence by a fourier spectral method on the earth simulator., in: *SC*, 2002, pp. 1–17.
- [7] L. Chevillard, S. G. Roux, E. Leveque, N. Mordant, J.-F. Pinton, A. Arneodo, Intermittency of velocity time increments in turbulence, *Phys. Rev. Lett.* 95 (2005) 064501.
- [8] P. K. Yeung, M. S. Borgas, Relative dispersion in isotropic turbulence. part 1. direct numerical simulations and reynolds-number dependence, *Journal of Fluid Mechanics* 503 (2004) 93.
- [9] L. Biferale, G. Boffetta, A. Celani, B. J. Devenish, A. Lanotte, F. Toschi, Multifractal statistics of lagrangian velocity and acceleration in turbulence, *Phys. Rev. Lett.* 93 (2004) 4502.
- [10] J. Bec, L. Biferale, G. Boffetta, A. Celane, M. Cencini, A. Lanotte, S. Musaccio, F. Toschi, Acceleration statistics of heavy particles in turbulence, *J. Fluid Mech.* 550 (2006) 349.
- [11] A. N. Kolmogorov, The local structure of turbulence in incompressible viscous fluid for very large reynolds numbers, *Royal Society of London Proceedings Series A* 434 (1991) 9.
- [12] N. Mordant, E. E. Lévêque, J. F. Pinton, Experimental and numerical study of the lagrangian dynamics of high reynolds turbulences, *New Journ. Phys.* 6 (2004) 116.
- [13] A. L. Porta, G. A. Voth, A. M. Crawford, J. Alexander, E. Bodenschatz, Fluid particle accelerations in fully developed turbulence, *Nature* 409 (2001) 1017.
- [14] H. Homann, R. Grauer, A. Busse, W. Müller, Lagrangian Statistics of Navier-Stokes- and MHD-Turbulence, submitted to *Phys. Rev. E*.
- [15] <http://www.fftw.org>
- [16] C. Shu, S. Osher, Efficient implementation of essentially non-oscillatory shock-capturing schemes, *J. Comput. Phys.* 77 (1988) 439.
- [17] W. H. Press, S. A. Teucholsky, W. T. Vetterling, B. P. Flannery, *Numerical Recipes in C*, 2nd Edition, Cambridge University Press, 1992.
- [18] R. Benzi, S. Ciliberto, R. Tripiccone, C. Baudet, F. Massaioli, S. Succi, Extended self-similarity in turbulent flows, *Phys. Rev. E* 48 (1993) R29