Comparison of computational codes for direct numerical simulations of turbulent Rayleigh–Bénard convection

Gijs L. Kooij^a, Mikhail A. Botchev^{d,e}, Edo M.A. Frederix^a, Bernard J. Geurts^{a,h,*},
Susanne Horn^{b,i}, Detlef Lohse^{c,f}, Erwin P. van der Poel^c, Olga Shishkina^f,
Richard J.A.M. Stevens^c, Roberto Verzicco^{g,c}

^a Multiscale Modeling and Simulation, Faculty EEMCS, University of Twente, Enschede, Netherlands

^b Applied Mathematics and Mathematical Physics, Imperial College London, UK

^c Physics of Fluids, Faculty of Science and Technology, University of Twente, Enschede, Netherlands

^d Mathematics of Computational Science, Faculty EEMCS, University of Twente, Enschede, Netherlands

^e Keldysh Institute of Applied Mathematics, Russian Academy of Sciences, Moscow, Russia

^f Max Planck Institute for Dynamics and Self-Organization, Göttingen, Germany

^g Dipartimento di Ingegneria Meccanica, Università di Roma “Tor Vergata”, Rome, Italy

^h Faculty of Applied Physics, Fluid Dynamics Laboratory, Eindhoven University of Technology, Eindhoven, Netherlands

ⁱ Earth, Planetary, and Space Sciences, University of California, Los Angeles, USA

ARTICLE INFO

Article history:

Received 2 August 2017

Revised 9 January 2018

Accepted 13 January 2018

Available online 31 January 2018

Keywords:

Direct numerical simulations

Rayleigh–Bénard convection

Heat transfer

ABSTRACT

Computational codes for direct numerical simulations of Rayleigh–Bénard (RB) convection are compared in terms of computational cost and quality of the solution. As a benchmark case, RB convection at $Ra = 10^8$ and $Pr = 1$ in a periodic domain, in cubic and cylindrical containers is considered. A dedicated second-order finite-difference code (AFID/RBFlow) and a specialized fourth-order finite-volume code (GOLDFISH) are compared with a general purpose finite-volume approach (OPENFOAM) and a general purpose spectral-element code (Nek5000). Reassuringly, all codes provide predictions of the average heat transfer that converge to the same values. The computational costs, however, are found to differ considerably. The specialized codes AFID/RBFlow and GOLDFISH are found to excel in efficiency, outperforming the general purpose flow solvers Nek5000 and OPENFOAM by an order of magnitude with an error on the Nusselt number Nu below 5%. However, we find that Nu alone is not sufficient to assess the quality of the numerical results: in fact, instantaneous snapshots of the temperature field from a near wall region obtained for deliberately under-resolved simulations using Nek5000 clearly indicate inadequate flow resolution even when Nu is converged. Overall, dedicated special purpose codes for RB convection are found to be more efficient than general purpose codes.

© 2018 Published by Elsevier Ltd.

1. Introduction

Rayleigh–Bénard (RB) convection is the flow driven by buoyancy forces when a fluid layer is heated from below and cooled from above [1–4]. The main governing parameter for RB convection is the Rayleigh number Ra which is the ratio between the destabilizing buoyancy and the stabilizing viscous and diffusive effects. For sufficiently large Ra , RB flow becomes turbulent. To understand the flow physics, direct numerical simulation (DNS) is in principle a straightforward approach that can be used to study tur-

bulence dynamics and heat transfer. For low Ra this is now routinely done. For higher Ra however, it is much more challenging, though there are many scientific questions. For example, how does the heat transfer scale with the Ra number in the regime of very high Ra numbers ($Ra \gtrsim 10^{14}$)? This regime, referred to as ‘ultimate’, is characterized by an enhanced heat transfer and associated with a transition to fully turbulent boundary layers [5–7]. DNS, provided it achieves proper accuracy, could be very useful in providing a deeper insight into the nature of this transition and the properties of this ultimate state. Unfortunately, DNS becomes exceedingly demanding as Ra increases, since turbulence produces smaller flow scales that need finer spatial resolution and proportionally small time steps to track their dynamics. For Ra where the ultimate regime is expected the needed computational power is,

* Corresponding author.

E-mail addresses: g.l.kooij@utwente.nl (G.L. Kooij), b.j.geurts@utwente.nl (B.J. Geurts).

at present, prohibitive and understanding which numerical code is most cost efficient is a key issue to establish a roadmap for the computer simulations of turbulent RB convection.

Over the years, several codes suitable for DNS of turbulent RB convection have been developed. We compare four of these codes. The first is based on the work by Verzicco and Orlandi [8] and Verzicco and Camussi [9] in which a second-order energy conserving finite-difference method is applied. For the periodic domain simulation we use the AFID code, which was developed by van der Poel et al. [10]. For the cylindrical simulations we use the latest version of RBFlow, which is an optimized version of the code used by Stevens et al. [11,12]. The second code is GOLDFISH by Shishkina et al. [13], Shishkina and Wagner [14] and Shishkina and Horn [15], which is based on a finite-volume approach and uses discretization schemes of fourth-order in space. GOLDFISH can be used to study turbulent thermal convection in cylindrical and parallelepiped domains. The third code is a general purpose open-source code Nek5000, based on the spectral-element method described by Fischer [16]. This code is designed to handle a large variety of flow problems, and was also used in the context of RB convection [17,18]. The fourth code is an open-source software package OPENFOAM [19]. More precisely, its widely used second-order finite-volume scheme was selected for the comparison.

In this study, we compare the four codes in terms of computational efficiency and quality of the results, with a special focus on the heat transport by the turbulent flow, measured by the Nusselt number (Nu). The efficiency of the codes is assessed in relation to their computational costs and the capability to achieve grid converged results. We simulate RB convection in three different geometries, i.e., a periodic domain, a cubic container, and a cylindrical container. Experiments for RB convection are typically conducted in a cylindrical tank. A major challenge to the DNS is to handle sharp gradients in the boundary layers near the walls as well as to capture thermal structures (plumes) that protrude far into the bulk of the flow.

In this paper we perform a convergence test at $Ra = 10^8$ and $Pr = 1$ in which we compare several levels of mesh refinement. We show that all four codes produce the same Nu number when appropriate spatial resolution is used. At high resolutions, the results become practically identical, taking into account a small uncertainty due to the finite averaging time. When we increase the Ra number for a fixed spatial resolution, we observe, not surprisingly, that for all codes eventually the resolution becomes insufficient for accurately resolving the turbulent flow. However, the Nu number calculation seems to be more robust against deliberate under-resolution in the higher order codes like Nek5000 than in the lower order codes. In this context, Nek5000 follows the theoretical scaling of Nu versus Ra better than the others. This might suggest that, for a given number of grid points, the Nek5000 code is capable of correctly capturing the flow physics even when the other codes fail. A direct inspection of some instantaneous snapshots of temperature in the near wall region, however, clearly shows that this is not the case since the temperature distribution displays the footprint of the underlying discretization. The conclusion is that the evaluation of the Nu alone is not a sufficient criterion to assess the quality of the results that, instead, should be assessed by evaluating more than one quantity. In this paper we also discuss some other advantages and drawbacks of the compared codes.

The remainder of this paper is organized as follows. In Section 2, we describe the governing equations of RB convection and the geometries of the domains included in this study. The codes are described in more detail in Section 3 and the results are compared in Section 4. Last, a summary and conclusions are given in Section 5.

2. Governing equations and evaluation of the Nusselt number

In this section we present the mathematical model and introduce the methods adopted to evaluate the Nu number. We consider RB convection in three different geometries: a periodic domain, a cube, and a cylinder. Every considered RB cell is characterized by a width D and a height H . A flow in any RB cell is determined by the dimensionless parameters, which are the Rayleigh number $Ra = g\beta\Delta H^3/(\nu\kappa)$, the Prandtl number $Pr = \nu/\kappa$, and the aspect-ratio $\Gamma = D/H$. Here g is gravitational acceleration, β the thermal expansion coefficient, Δ the temperature difference between the upper and lower plate, ν the kinematic viscosity, and κ the thermal diffusivity of the fluid. In this study, we consider $\Gamma = 1$ for all geometries, and $Pr = 1$, which means that the inner length scales of the velocity and the temperature fields are of similar order.

In the numerical simulations, we solve the incompressible Navier–Stokes equation with the Boussinesq approximation to account for buoyancy effects. The governing equations read

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = \sqrt{\frac{Pr}{Ra}} \nabla^2 \mathbf{u} - \nabla p + \theta \mathbf{e}_z, \quad (1)$$

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

$$\frac{\partial \theta}{\partial t} + \mathbf{u} \cdot \nabla \theta = \frac{1}{\sqrt{PrRa}} \nabla^2 \theta, \quad (3)$$

where $\mathbf{u}$ is the velocity, p the pressure, θ the temperature, and $\mathbf{e}_z$ the unit vector in the vertical direction anti-parallel to the gravitational acceleration. Here lengths are expressed in terms of H , velocities in terms of free fall velocity $U = \sqrt{\beta g \Delta H}$, and temperatures in terms of Δ . No-slip and constant temperature conditions are imposed at the plates, and no-slip and adiabatic conditions at the sidewall. We do not simulate all geometrical configurations with all codes, since not every geometry is feasible in every code. For the periodic domain, we compare AFID and Nek5000, and for the cubic container GOLDFISH and Nek5000. The cylindrical domain is simulated with all four codes included in this study.

One of the main aspects of RB convection is the heat transported by the turbulent flow from the lower to the upper plate. The heat transfer is quantified by the dimensionless heat flux, i.e. the Nu number which is the ratio of the actual specific heat flux to the purely conductive counterpart. Following [11] we consider several ways to compute Nu . First we consider those of them, which are related directly to the gradient of the temperature and to the convective heat transport. As a function of the vertical coordinate z , Nu is defined as the average heat flux through a horizontal cross section of the domain [20], $Nu(z) = -\langle \partial_z \theta \rangle_A + \sqrt{PrRa} \langle u_z \theta \rangle_A$ where $\langle \cdot \rangle_A$ denotes the average over a horizontal cross section A and in time. From the no-slip boundary conditions, it follows that the Nu numbers at the lower and upper plate, denoted by Nu_{lo} and Nu_{up} respectively, can be calculated from the average temperature gradient at the plates only. We also define the average of the two as $Nu_{pl} := (Nu_{lo} + Nu_{up})/2$. The third definition is obtained using the volume average $Nu_{vol} := 1 + \sqrt{RaPr} \langle u_z \theta \rangle_V$, where $\langle \cdot \rangle_V$ denotes the average over the complete volume of the domain. Note that these definitions of Nu are averaged over time as well.

Two more definitions can be obtained from the global balance of energy. We can derive a relation between the Nu number and the kinetic and thermal dissipation rates [21]. The kinetic and thermal dissipation rates are $\varepsilon := \sqrt{Pr/Ra} (\nabla \mathbf{u})^2$, $\varepsilon_\theta := 1/(\sqrt{PrRa}) (\nabla \theta)^2$ and the Nu number can be calculated from the kinetic and thermal dissipation rate respectively as follows: $Nu_{kin} := 1 + \sqrt{PrRa} \langle \varepsilon \rangle_V$, and $Nu_{th} := \sqrt{PrRa} \langle \varepsilon_\theta \rangle_V$. These relations from the global balance of energy are sometimes used to assess

the quality of DNS of RB convection. If the simulation is well resolved, the global balance of energy is respected accurately. When averaged over time Nu calculated from the dissipation rates agrees with the other definitions of the Nu number. Note that the converse is not necessarily true as will be illustrated in Section 4.1. In particular, if some definitions of Nu agree with each other then this does not automatically imply that the resolution is adequate.

Being an integral quantity, Nu is one of the main characteristics in RB convection and is, therefore, a natural quantity to investigate. Of course, besides Nu , there are other quantities describing the turbulent RB flow that one could include in a comparison. A good prediction of Nu does not automatically guarantee that other quantities are approximated accurately, in particular higher order moments will converge less easily. However, the converse holds, i.e., Nu predictions will correspond closely if the solution is accurately captured. In the present comparison study, we focus mostly on Nu , because it is one of the most important quantities and it gives a first indication of how well different codes perform.

3. Numerical methods

In this section, we provide a brief description of the four codes that are compared. Detailed information can be found in the mentioned references.

3.1. AFID/RBFlow

The second-order finite-difference scheme has initially been developed by Verzicco and Orlandi [8] and Verzicco and Camussi [9] for cylindrical containers. Time integration is performed with a third-order Runge–Kutta method, in combination with a second-order Crank–Nicolson scheme for the viscous terms. RBFlow, which is used for the simulations in the cylindrical domain, computes all viscous terms implicitly. The open-source code AFID, specialized for domains with two periodic horizontal directions, uses an explicit scheme in the non-bounded directions to improve scalability of the code [10]. In AFID, the pressure is solved using a fast Fourier Transform (FFT) in the horizontal directions by means of a 2D pencil decomposition [10].

3.2. Nek5000

The open-source package Nek5000 is based on the spectral element method, which is an essential extension of the standard finite element method to the case of higher-order basis functions. In this case, the basis functions for the velocity and the pressure are tensor product Lagrange polynomials of order $\mathcal{N}$. Details of the code are found in Ref. [16]. The spectral element method has been used successfully for DNS of RB convection [17,18]. We use the so-called $P_N - P_N$ formulation, based on the splitting scheme in Ref. [22]. This means that the spectral elements for the velocity components and the pressure are both order $\mathcal{N}$. In the more traditional $P_N - P_{N-2}$ formulation the pressure is treated with order $\mathcal{N} - 2$. Slightly more accurate results can be obtained with the higher order approximation of the pressure in the $P_N - P_N$ formulation for moderately resolved turbulent flows. In our simulations, we use $N = 8$, which is similar to what was used in DNS of other turbulent flows in [17,23,24]. The viscous term is treated implicitly with the second-order backward differentiation formula, in combination with an explicit second-order extrapolation scheme for the convective and other terms. The linear system for the velocity is solved with the conjugate gradient method using Jacobi preconditioning. The linear system for the pressure is solved with the generalized minimal residual method, preconditioned with an additive Schwarz method.

3.3. GOLDFISH

The computational code GOLDFISH is based on a finite-volume approach. To calculate the velocity and temperature at the surfaces of each finite volume, it uses higher-order discretization schemes in space, up to the fourth order in the case of equidistant meshes. GOLDFISH has been used to study thermal convective flows in different configurations [13–15], in cylindrical and parallelepiped domains. For the time integration, the leapfrog scheme is used for the convective term and the explicit Euler scheme for the viscous term. Although formally first order in time, the accuracy is close to second-order in convection dominated flows [25, Section 5.8]. Note that due to the von Neumann numerical stability of the chosen scheme, the fourth-order spatial discretization requires asymptotically at least 4/3 times finer time stepping than the second-order scheme [26]. Due to the regularity of the used computational meshes, direct solvers are applied to compute the pressure in cylindrical and Cartesian coordinate systems. Thus, when the RB container is a cylinder, FFT is used in two directions. In the case of a parallelepiped RB container, the grid regularity also allows separation of variables. In this case, the corresponding eigenvalues and eigenvectors for the pressure solver are calculated and stored at the beginning of the simulations. The code is quite flexible in parallelization, including parallel I/O, and is characterized by high modularity and is applicable to different configurations of turbulent thermal convective flows.

3.4. OPENFOAM

OPENFOAM is a widely used open-source second-order finite-volume software package [19]. Although OPENFOAM offers many different options, we use the implementation of a standard solver in OPENFOAM, which would be representative for typical engineering applications. A linear interpolation scheme is used for the convective term. The equations are solved with the PISO algorithm. The default implementation of the second-order Crank–Nicolson scheme is used for time integration. Furthermore, we do not use the “non-orthogonal” correction for the non-orthogonality of the mesh.

4. Performance comparison

In this section, we present the results of the simulations using two specialized RB convection codes (AFID/RBFlow and GOLDFISH) and two general purpose codes (Nek5000 and OPENFOAM). Our findings shed some light on the issue of the relevance of general purpose codes for moderate Ra number turbulence. We first discuss a convergence test for a moderate $Ra = 10^8$, for which a fully resolved DNS is easily affordable. A comparison in terms of quality of results and cost is made in Section 4.1. Finally, some results at higher Ra numbers and fixed spatial resolution are shown in Section 4.2, illustrating the inevitable loss of accuracy with significant rise of Ra .

The simulations with AFID/RBFlow, Nek5000, and OPENFOAM, are all performed on Cartesius (SURFsara). The simulations with GOLDFISH are performed on SuperMUC of the Leibniz-Rechenzentrum (LRZ). All the simulations run on Cartesius are performed on the same type of ‘thin’ nodes with Intel Haswell processors clocked at 2.6 Ghz. Similar nodes were also used on SuperMUC. Hence, we expect similar performance on these computing platforms.

4.1. Convergence test at $Ra = 10^8$

Here we present a classical convergence test in which a number of grid refinements is undertaken to assess the sensitivity of

Fig. 1. Nu against the number of grid points for different geometries: (a) Periodic domain (b) Cubic domain (c) Cylindrical domain. Markers: $\times$ AFID, $\diamond$ GOLDFISH, $\square$ NEK5000, $\circ$ OPENFOAM. Lines: — Nu_{vol} , --- Nu_{pl} , ... Nu_{kin} , - · - Nu_{th} .

Fig. 2. Computational cost against the number of grid points for different geometries: (a) Periodic domain (b) Cubic domain (c) Cylindrical domain. $\times$ AFID/RBFlow, $\diamond$ GOLDFISH, $\square$ NEK5000, $\circ$ OPENFOAM, --- $O(N^{4/3})$ and ... $O(N^{5/3})$.

the results to spatial resolution. Because of the fundamental differences between the spatial discretization techniques, the simulations on the different levels of grid refinement need to be performed with slightly different meshes. For a fair comparison, at any grid refinement level, the total number of grid points or degrees of freedom (which for simplicity we also call grid points) is kept similar for all codes. The mesh refinement is undertaken near the plates. The boundary layer thicknesses are estimated a priori using the scaling theory by Grossmann and Lohse (GL theory) [5,11,27,28]. In the studied case of $Pr = 1$, the thicknesses of the kinetic and thermal boundary layers are similar. At any considered level of grid refinement, the total number of grid points, N , and the number of grid points inside the boundary layers, N_{BL} , of the chosen meshes are similar for all codes. The spatial resolutions and the corresponding values of N and N_{BL} are listed in the Appendix.

The Nu numbers are averaged over 300 dimensionless time units after the solution approaches a statistically stationary state, which takes about 200 time units, depending on the initial conditions. Nu , obtained at different levels of grid refinement, versus the number of grid points, is presented in Fig. 1. We observe that at high grid resolutions all codes converge to the same result within a small time averaging error of about 0.5%. Different ways to calculate Nu and different codes lead to different convergence of the obtained Nu with increasing grid resolution. For example, the results for NEK5000 at very coarse resolutions are quite inaccurate, but they converge quickly to the final value when the resolution is increased. We can interpret the results in Fig. 1 as a good indication of convergence to nearly grid independent results, achieved by all codes, albeit at different spatial resolutions. Ultimate convergence assessment is hampered by the degree of time averaging uncertainty that remains, we come back to this momentarily.

For a fixed statistical averaging interval, the costs are proportional to the mesh size, both in space and time. For this cost-estimate to hold, it is required that the iterative solvers converge in a number of steps that is approximately constant during the time-interval over which the statistical averaging is performed. With suitable preconditioners such can be realized, as was observed for the corresponding codes. The Courant–Friedrichs–Lewy (CFL) condition and the numerical stability of the simulations were the two restrictions on the time stepping that were taken into account in all conducted simulations. As one can see in Fig. 2, the computational costs scale from $O(N^{4/3})$ to $O(N^{5/3})$ in all cases, in a full accordance to the von Neumann stability of the schemes used in the corresponding codes. For the schemes, which are optimal with respect to the von Neumann stability, the time step size τ is taken proportional to the mesh width h , which in turn is proportional to $N^{-1/3}$. This leads to the scaling of the computational costs with the mesh size as $O(N^{4/3})$. Apart from the von Neumann stability, there exists also another restriction on the time stepping in accurate DNS, which is the resolution of the Kolmogorov time microscales; we will come back to this issue in Section 5.

Further, in Fig. 2, we observe that the general purpose codes designed for unstructured grids in complex geometries are more expensive than those for structured grids, exploiting the periodic directions of the geometry or, at least the possibility to separate variables due to the regularity of the grids. For example, AFID and OPENFOAM are both second-order accurate, but OPENFOAM is much closer to NEK5000 in terms of computational cost. Apart from the grid organization, the order of the schemes, used in the codes, influences the computational load. Obviously, the higher-order schemes need more operations per time step than the lower-order schemes. Also the way to solve the Poisson equation for the pressure-like function determines the efficiency. In general purpose

Fig. 3. Average error of Nu against the computational cost for different geometries: (a) Periodic domain (b) Cubic domain (c) Cylindrical domain. $\times$ AFID/RBFlow, $\diamond$ GOLDFISH, $\square$ NEK5000, $\circ$ OPENFOAM. The dashed line indicates the 95% confidence level of the reference value given in Tables 1–3.

Table 1

Nu obtained with the highest spatial resolutions for the periodic domain, where $N_x \times N_y \times N_z = 384 \times 384 \times 384$ for AFID and $N_x \times N_y \times N_z = 379 \times 379 \times 379$ for NEK5000.

	Nu_{lo}	Nu_{up}	Nu_{vol}
AFID	32.24	32.27	32.18
NEK5000	32.29	32.41	32.54
Average	32.32 ± 0.24 (0.73%)		

Table 2

Nu obtained with the highest spatial resolutions for the cubic container, where $N_x \times N_y \times N_z = 384 \times 384 \times 384$ for GOLDFISH and $N_x \times N_y \times N_z = 379 \times 379 \times 379$ for NEK5000.

	Nu_{lo}	Nu_{up}	Nu_{vol}
GOLDFISH	31.56	31.49	31.47
NEK5000	31.53	31.58	31.53
Average	31.53 ± 0.08 (0.24%)		

codes only iterative computationally intensive solvers can be employed, while specialized codes can use direct solvers, which are much more efficient. Also the parallel scalability of the codes influences the total computational costs. OPENFOAM, for example, is characterized by quite modest scalability, compared to the other considered codes [29,30]. In our simulations with OPENFOAM, we used a sufficiently low number of cores, such that we operate only in the range of good parallel efficiency. In that case, the measurements of computational time are not affected significantly by possible effects of non-ideal scalability. The Nu values obtained in the highest resolution simulations are listed in Tables 1–3, for the periodic domain, cube, and cylinder, respectively. These data can be used to further quantify the error in Nu versus computational costs. To arrive at this we undertake a number of steps. First, we calculate the average Nu as reference points. We calculate the standard deviation from the data given in the tables and use twice the standard deviation as a 95% confidence interval for the average values. These average Nu are subsequently used as a reference value to calculate the error in Nu_{lo} , Nu_{up} , and Nu_{vol} separately for the different codes and grids. After that, we take the average of those individual errors, and show the average error against the computational cost in Fig. 3. The error decreases with increasing cost, until it becomes comparable to the confidence bounds. At that point, the error is dominated by the time averaging error, and no longer due to the spatial discretization. Since the computational costs for simulations of 300 dimensionless time units are already considerably large we did not pursue a further reduction in the time averaging error. In fact, such time averaging error will tend to zero at a rate inversely proportional to the square root of the simulation

Table 3

Nu obtained with the highest spatial resolutions for the cylindrical container, where $N_r \times N_\theta \times N_z = 192 \times 512 \times 384$ for RBFlow/GOLDFISH, $N_{xy} \times N_z = 85,009 \times 384$ for NEK5000, and $N_{xy} \times N_z = 110,592 \times 384$ for OPENFOAM.

	Nu_{lo}	Nu_{up}	Nu_{vol}
RBFlow	32.08	32.15	32.24
GOLDFISH	32.19	32.31	32.33
NEK5000	32.26	32.23	32.16
OPENFOAM	32.16	32.10	32.13
Average	32.20 ± 0.14 (0.44%)		

time. Hence, only at extreme costs one could perceive a significant reduction of the time averaging error. Such resources are not available for this study and are also not required to establish the main conclusions.

In the periodic domain, AFID was found to be much faster than NEK5000. At a given computational cost, a much higher number of grid points can be afforded with AFID. On the other hand, this significant difference in computational cost between NEK5000 and AFID when counting the number of grid points only, is considerably reduced when counting the actually achieved level of precision of the Nu prediction. Clearly, the higher-order method used in NEK5000 is beneficial at reducing the gap with AFID in error versus cost considerations. This is illustrated concisely in Fig. 3. For the cubic container we observe a similar situation: the specialized code GOLDFISH is much faster than NEK5000, see Fig. 2(b). And again, the efficiency of NEK5000 becomes closer to that by GOLDFISH when the convergence of Nu is taken into account. When the error of the Nu calculation is above the confidence bound, GOLDFISH calculates Nu up to tenfold more accurately than NEK5000, for given computational costs. In the cylindrical container, RBFlow and GOLDFISH are up to a factor ten faster than OPENFOAM for a given level of accuracy, while RBFlow and OPENFOAM are both of second order. This illustrates the penalty that comes with the use of a general purpose code compared to a dedicated specialized code. NEK5000 falls roughly in between RBFlow and OPENFOAM. Overall, Fig. 3 shows that the large differences in speed that appear in Fig. 2 (with AFID/RBFlow the most efficient when it comes to costs of simulation with a certain number of grid points) decrease when counting the error in Nu versus computational costs due to the usage of higher-order schemes in the other codes.

In Fig. 3 one can also see that the accuracy of the general purpose codes behave non-monotonically with increasing computational time. The strong oscillations in the behavior of the Nu -error versus the costs are explained not only by the restricted time of statistical averaging, but mainly by the usage of the iterative solvers within these codes. While the dedicated codes AFID/RBFlow and GOLDFISH use direct solvers and get the corre-

Fig. 4. Nu versus Ra in cylindrical container of $\Gamma = 1$. These results of deliberately under-resolved DNS are obtained at a fixed computational mesh, which is too coarse for large Ra . $+ Nu_{vol}$ (RBFlow), $\circ Nu_{pl}$ (RBFlow), $* Nu_{vol}$ (GOLDFISH), $\square Nu_{pl}$ (GOLDFISH), $\times Nu_{vol}$ (NEK5000), $\diamond Nu_{pl}$ (NEK5000), $---$ Grossmann–Lohse theory [28]. Note that a correct Nu does not imply a well-resolved flow, see Fig. 5.

sponding solutions at the machine accuracy, the iterative solvers of the general purpose codes stop iterations with a certain residual error. Furthermore, the used iterative solvers generally do not guarantee a monotonic reduction of the errors with increasing number of conducted iterations. This holds in particular for the generalized minimal residual method used in NEK5000. This makes prediction of the accuracy versus computation costs for general purpose codes less trivial and somewhat uncertain.

4.2. Robustness against under-resolution

At $Ra = 10^8$, we are able to compute an accurate reference solution that is converged with respect to the spatial resolution independent of which code was adopted, (see Fig. 1). At higher Ra numbers, the computation of such a reference solution for all

codes becomes too expensive, however. As an alternative, we can compare the codes in a different, somewhat more qualitative, way by increasing the Ra number while keeping the spatial resolution fixed. In this case, we use the meshes with the highest resolution adopted for $Ra = 10^8$ in the cylindrical container. As Ra increases, the effect of insufficient resolution will unavoidably show up sooner or later, which indicates the robustness of the codes, against under-resolution. Nu , compensated with $Ra^{-1/3}$, is plotted against Ra in Fig. 4. Initially, the results of the three codes at $Ra = 10^8$ are all very close to each other with the differences between the results less than 0.5%. The values of Nu are also very close to the prediction by the GL theory with the deviations between the simulation results and the GL predictions less than 1%. As Ra increases the different robustnesses of the various codes against deliberate under-resolution become apparent. NEK5000 shows the smallest deviation from the theoretical scaling of Nu , and RBFlow the largest. However, we emphasize that this robustness of Nu for NEK5000 against deliberate underresolution does not imply that other flow features would still be well represented. E.g., in Fig. 5 we show temperature snapshots for well resolved and deliberately underresolved simulations with NEK5000. The latter clearly show a pronounced imprint of the computational grid at $Ra = 10^{10}$ even though the Nu number predictions are not affected much. NEK5000 does not show the imprint of the mesh at $Ra = 10^9$, but the effect of underresolution can be seen in very subtle ripples near high gradients. Similar plots obtained from RBFlow also show an inadequacy of the grid resolution at $Ra = 10^9$. Underresolution appears to affect NEK5000 predictions somewhat less than is seen for RBFlow.

This comparison clearly indicates that the agreement of Nu with the theoretical prediction (and among the values obtained from the various definitions) is not enough to assess the adequacy of the spatial resolution of the numerical simulation. Additional quantities have to be analyzed, such as the instantaneous temperature snapshots or rms profiles, in order to clarify this issue.

Fig. 5. Temperature field at $z/H = 0.0152$ from the NEK5000 code (top row) and RBFlow (bottom row) at (from left to right) $Ra = 10^8$, $Ra = 10^9$ and $Ra = 10^{10}$. For the latter two Ra values the chosen grid resolution is insufficient. Note that in the NEK5000 snapshots the imprint of the computational grid is clearly visible in the higher Ra number cases, even though the Nu number from the simulations looks reasonable, see Fig. 4. The ripples in the RBFlow snapshots are observed near sharp gradients when the resolution is insufficient.

5. Conclusions and outlook

In this paper, we have compared several codes for the simulation of turbulent RB convection in a number of typical geometries. Particular attention has been given to the heat transport in the turbulent flow, which is quantified by Nu . The computational efficiency of the codes is determined with reference to fully converged simulations at relatively high spatial resolutions. We observed significant differences between the codes in terms of computational costs, i.e. the specialized AFID/RBFlow and GOLDFISH code clearly outperform NEK5000 and OPENFOAM. Thus, we note that a considerable saving in computational costs can be achieved by employing an optimized code for a simple geometry compared to general purpose codes designed for complex geometries. The benefit of general purpose codes like NEK5000 and OPENFOAM is of course that they are much wider applicable than specialized codes, which need to be specifically tuned per case (Tables 4–11).

The usage of unstructured grids in the general purpose codes requires iterative solutions of the governing equations on each time step. This leads to higher computational costs and makes these codes less predictable with respect to the accuracy of the calculation of Nu with growing computational costs. Also the scalability of the OPENFOAM codes on supercomputers leaves much to be desired. All this leads to the fact that AFID/RBFlow, being also the second-order as OPENFOAM in the considered configuration, is at least ten times faster than OPENFOAM, while providing the same level of accuracy. Therefore we conclude that OPENFOAM, at least in the analyzed configuration, which is the most popular in engineering, is not optimal for scientific investigations of high Ra number thermal convection.

Also, among the other codes, AFID/RBFlow is clearly the fastest one. It is up to tenfold faster than GOLDFISH and up to hundredfold faster than NEK5000. However, when the accuracy of the Nu calculation is taken into account, the efficiency of GOLDFISH and AFID/RBFlow is similar, while NEK5000 and OPENFOAM are up to 10 times slower. When in a certain numerical study the point of interest is an integral quantity (zero moment), like Nu or Reynolds number, or when the profiles of the mean temperature or velocity are aimed to be studied (first moments), the advantages of the usage of the second-order code AFID/RBFlow are clearly pronounced. It is extremely fast and calculates these quantities precisely on sufficiently fine meshes. This is partly thanks to the implementation of AFID/RBFlow, which is highly optimized, and scales excellently on large number of cores [31].

Finally, we give a general estimate of the complexity of the DNS of turbulent RB convection in the classical regime and in the ultimate regime, which is to be studied in the future. As we already mentioned in Section 4.1, apart from the CFL-condition and the von Neumann stability, there exists also another restriction on the time stepping in accurate DNS, which is the resolution of the Kolmogorov time microscales. Note that the Kolmogorov microscale in space, $\eta \equiv (\mu^3 / \langle \varepsilon \rangle_V)^{1/4}$, and the microscale in time, $\eta_\tau \equiv (\mu / \langle \varepsilon \rangle_V)^{1/2}$, are related as $\mu \eta_\tau \sim \eta^2$ with $\mu \equiv \sqrt{Pr/Ra}$. Thus, the optimal (not over-resolved but accurate) DNS, which resolve both, the Kolmogorov time microscale η_τ and the Kolmogorov spatial microscales η , will lead to the scaling of the computational costs with the grid size N at least as $\mathcal{O}(\mu N^{5/3})$. Since $\langle \varepsilon \rangle_V = (Nu - 1) / \sqrt{PrRa}$, for a fixed Pr , the computational costs in accurate DNS must grow at least as $\mathcal{O}(Nu^{5/4} Ra^{3/4})$. Therefore, in the classical regime, where $Nu \sim Ra^{1/3}$, the cost will increase with Ra at least as $\mathcal{O}(Ra^{7/6})$, while for the ultimate regime, where the scaling $Nu \sim Ra^{1/2}$ is expected, the anticipated computational costs in accurate DNS are at least $\mathcal{O}(Ra^{11/8})$.

Before concluding this paper we wish to point once more out that comparing Nu obtained by the numerical simulation with the expected value is not a reliable criterion to assess its validity. In fact, we showed that deliberately under-resolved simulations performed with higher order codes show a small error in Nu while producing temperature fields with strong unphysical oscillations. Instantaneous snapshots of temperature and profiles of higher order moments have to be evaluated, together with Nu , in order to establish the quality of a numerical simulation.

Acknowledgements

GLK is supported financially as part of program No. 142 “Towards ultimate turbulence” by Foundation for Fundamental Research on Matter (FOM), part of the Netherlands Organization for Scientific Research (NWO). DL and RV are funded by the Netherlands Center for Multiscale Catalytic Energy Conversion (MCEC), an NWO Gravitation program funded by the Ministry of Education, Culture and Science of the government of the Netherlands. OS and SH are funded by the Deutsche Forschungsgemeinschaft (DFG) under the grants Sh 405/4-2 and Ho 5890/1-1, respectively. DL and OS thank DFG Priority Programme SPP 1881 “Turbulent Superstructures”. We thank NWO for granting us computational time on Cartesius cluster from the Dutch Supercomputing Consortium SURFsara under grants SH-061 and SH-015 and also are grateful to the Leibniz Rechenzentrum (LRZ) for providing us computational resources under the grant pr84pu.

Appendix A. Spatial resolution and Nusselt number at $Ra = 10^8$ and $\Gamma = 1$.

Table 4

Spatial resolution and Nu numbers for AFID in the periodic domain.

N	N_{BL}	Nu_{lo}	Nu_{up}	Nu_{vol}	Nu_{kin}	Nu_{th}
64^3	3	37.02	37.00	37.06	36.94	37.24
96^3	4	33.73	33.81	33.74	33.32	33.80
128^3	5	33.11	33.12	33.27	32.76	33.12
192^3	8	32.57	32.62	32.76	32.24	32.60
256^3	11	32.57	32.62	32.75	32.69	32.59
384^3	16	32.24	32.27	32.18	32.10	32.25

Table 5

Spatial resolution and Nu numbers for NEK5000 in the periodic domain.

E	N	N_{BL}	Nu_{lo}	Nu_{up}	Nu_{vol}	Nu_{kin}	Nu_{th}
5^3	36^3	2	43.78	43.63	33.47	28.01	32.25
7^3	50^3	2	36.47	36.48	32.86	29.72	31.29
9^3	64^3	3	32.01	32.04	32.27	30.70	31.33
14^3	99^3	4	31.57	31.53	32.22	31.83	31.95
18^3	127^3	5	32.30	32.22	32.39	32.19	32.23
27^3	190^3	8	32.41	32.46	32.41	32.38	32.41
36^3	253^3	11	32.31	32.43	32.45	32.35	32.38
54^3	379^3	16	32.29	32.41	32.54	32.35	32.37

Table 6

Spatial resolution and Nu numbers for GOLDFISH in the cubic container.

N	N_{BL}	Nu_{lo}	Nu_{up}	Nu_{vol}	Nu_{kin}	Nu_{th}
64^3	3	34.93	34.99	34.91	30.13	31.47
96^3	4	32.38	32.36	32.34	29.88	30.46
128^3	5	31.56	31.70	31.66	30.16	30.44
192^3	8	31.67	31.51	31.58	30.89	31.01
256^3	11	31.65	31.63	31.64	31.24	31.28
384^3	16	31.56	31.49	31.47	31.31	31.36

Table 7
Spatial resolution and Nu numbers for NEK5000 in the cubic container.

E	N	N_{BL}	Nu_{lo}	Nu_{up}	Nu_{vol}	Nu_{kin}	Nu_{th}
5^3	36^3	2	43.44	43.38	33.29	29.58	32.86
7^3	50^3	2	36.13	36.28	32.47	30.55	31.48
9^3	64^3	3	31.76	31.88	31.83	30.77	31.12
14^3	99^3	4	30.93	30.93	31.55	31.36	31.38
18^3	127^3	5	31.53	31.64	31.68	31.56	31.62
27^3	190^3	8	31.40	31.41	31.37	31.36	31.38
36^3	253^3	11	31.56	31.71	31.60	31.62	31.63
54^3	379^3	16	31.53	31.58	31.53	31.52	31.55

Table 8
Spatial resolution and Nu numbers for RbFlow in the cylindrical container.

N_r	N_ϕ	N_z	N_{BL}	Nu_{lo}	Nu_{up}	Nu_{vol}
48	128	96	4	34.95	34.86	34.85
96	256	192	8	32.59	32.76	32.58
192	512	384	16	32.08	32.15	32.24

Table 9
Spatial resolution and Nu numbers for GOLDFISH in the cylindrical container.

N_r	N_ϕ	N_z	N_{BL}	Nu_{lo}	Nu_{up}	Nu_{vol}
48	128	96	4	33.20	33.02	33.18
96	256	192	8	32.40	32.19	32.26
192	512	384	16	32.19	32.31	32.33

Table 10
Spatial resolution and Nu numbers for NEK5000 in the cylindrical container. E_{xy} denotes the number of spectral elements in a horizontal cross section, and E_z in the vertical direction.

E_{xy}	E_z	N_{xy}	N_z	N_{BL}	Nu_{lo}	Nu_{up}	Nu_{vol}
48	5	2409	36	2	41.78	41.85	33.10
48	7	2409	50	2	36.91	36.62	32.87
48	9	2409	64	3	32.68	33.06	32.52
108	14	5377	99	4	31.47	31.41	32.29
192	18	9521	127	5	32.07	32.12	32.24
432	27	36,100	190	8	32.14	32.18	32.16
768	36	64,009	253	11	32.11	32.08	32.03
1728	54	85,009	379	16	32.26	32.23	32.16

Table 11
Spatial resolution and Nu numbers for OPENFOAM in the cylindrical container.

N_{xy}	N_z	N_{BL}	Nu_{lo}	Nu_{up}	Nu_{vol}
3072	64	3	35.43	35.30	37.10
6912	96	4	33.33	33.27	34.00
12,288	128	5	32.62	32.52	32.78
27,648	192	8	32.17	32.09	32.21
49,152	256	11	31.90	32.10	32.00
110,592	384	16	32.16	32.10	32.13

References

[1] Kadanoff LP. Turbulent heat flow: structures and scaling. *Phys Today* 2001;54:34.
[2] Ahlers G, Grossmann S, Lohse D. Heat transfer and large scale dynamics in turbulent Rayleigh-Bénard convection. *Rev Mod Phys* 2009;81:503.

[3] Lohse D, Xia KQ. Small-scale properties of turbulent Rayleigh-Bénard convection. *Ann Rev Fluid Mech* 2010;42:335.
[4] Chilla F, Schumacher J. New perspectives in turbulent Rayleigh-Bénard convection. *Eur Phys J E* 2012;35:58.
[5] Grossmann S, Lohse D. Scaling in thermal convection: a unifying view. *J Fluid Mech* 2000;407:27.
[6] Grossmann S, Lohse D. Multiple scaling in the ultimate regime of thermal convection. *Phys Fluids* 2011;23:045108.
[7] He X, Funfschilling D, Nobach H, Bodenschatz E, Ahlers G. Transition to the ultimate state of turbulent Rayleigh-Bénard convection. *Phys Rev Lett* 2012;108:024502.
[8] Verzicco R, Orlandi P. A finite-difference scheme for three-dimensional incompressible flow in cylindrical coordinates. *J Comput Phys* 1996;123:402.
[9] Verzicco R, Camussi R. Numerical experiments on strongly turbulent thermal convection in a slender cylindrical cell. *J Fluid Mech* 2003;477:19.
[10] van der Poel EP, Ostilla-Mónico R, Donners J, Verzicco R. A pencil distributed finite difference code for strongly turbulent wall-bounded flows. *Comput Fluids* 2015;116:10.
[11] Stevens RJAM, Verzicco R, Lohse D. Radial boundary layer structure and Nusselt number in Rayleigh-Bénard convection. *J Fluid Mech* 2010;643:495.
[12] Stevens RJAM, Lohse D, Verzicco R, Prandtl and Rayleigh number dependence of heat transport in high rayleigh number thermal convection. *J Fluid Mech* 2011;688:31.
[13] Shishkina O, Horn S, Wagner S, Ching ESC. Thermal boundary layer equation for turbulent Rayleigh-Bénard convection. *Phys Rev Lett* 2015;114:114302.
[14] Shishkina O, Wagner S. Prandtl-number dependence of heat transport in laminar horizontal convection. *Phys Rev Lett* 2016;116:024302.
[15] Shishkina O, Horn S. Thermal convection in inclined cylindrical containers. *J Fluid Mech* 2016;790:R3.
[16] Fischer PF. An overlapping Schwarz method for spectral element solution of the incompressible Navier-Stokes equations. *J Comput Phys* 1997;133:84.
[17] Scheel JD, Emran MS, Schumacher J. Resolving the fine-scale structure in turbulent Rayleigh-Bénard convection. *New J Phys* 2013;15:113063.
[18] Kooij GL, Botchev MA, Geurts BJ. Direct numerical simulation of Nusselt number scaling in rotating Rayleigh-Bénard convection. *Int J Heat Fluid Flow* 2015;55:26.
[19] Weller HG, Tabor G, Jasak H, Fureby C. A tensorial approach to computational continuum mechanics using object-oriented techniques. *Comput Phys* 1998;12:620.
[20] Verzicco R, Camussi R. Prandtl number effects in convective turbulence. *J Fluid Mech* 1999;383:55.
[21] Shraiman BI, Siggia ED. Heat transport in high-Rayleigh number convection. *Phys Rev A* 1990;42:3650.
[22] Tomboulides AG, Lee JCY, Orszag SA. Numerical simulation of low Mach number reactive flows. *J Sci Comput* 1997;12:139.
[23] El Khoury GK, Schlatter P, Noorani A, Fischer PF, Brethouwer G, Johansson AV. Direct numerical simulation of turbulent pipe flow at moderately high Reynolds numbers. *Flow, Turbul Combust* 2013;91:475.
[24] Ohlsson J, Schlatter P, Fischer PF, Henningson DS. Direct numerical simulation of separated flow in a three-dimensional diffuser. *J Fluid Mech* 2010;650:307.
[25] Wesseling P. Principles of computational fluid dynamics. Springer, Berlin; 2001.
[26] Shishkina O. The Neumann stability of high-order symmetric schemes for convection-diffusion problems. *Sib Math J* 2007;48:1141.
[27] Shishkina O, Stevens RJAM, Grossmann S, Lohse D. Boundary layer structure in turbulent thermal convection and its consequences for the required numerical resolution. *New J Phys* 2010;12:075022.
[28] Stevens RJAM, van der Poel EP, Grossmann S, Lohse D. The unifying theory of scaling in thermal convection: the updated prefactors. *J Fluid Mech* 2013;730:295.
[29] Axtmann G, Rist U. Scalability of OpenFOAM with large eddy simulations and DNS on high-performance systems. *High Performance Computing in Science and Engineering*. Springer; 2016.
[30] Rivera O, Furlinger K, Kranzlmüller D. Investigating the scalability of OpenFOAM for the solution of transport equations and large eddy simulations. In: International Conference on Algorithms and Architectures for Parallel Processing, Springer; 2011.
[31] Zhu X, Phillips E, Spandan V, Donners J, Ruetsch G, Romero J, et al. AFiD-GPU: a versatile Navier-Stokes solver for wall-bounded turbulent flows on GPU clusters. *Comput Phys Commun* 2017. Submitted.