

# COMPARISON OF DIFFERENT SOLVERS AND GEOMETRY REPRESENTATION STRATEGIES FOR DNS OF ROUGH WALL CHANNEL FLOW

Florian Theobald<sup>1</sup>, Kay Schäfer<sup>2</sup>, Jiasheng Yang<sup>2</sup>, Bettina Frohnäpfel<sup>2</sup>, Matthias Stripf<sup>1</sup>, Pourya Forooghi<sup>3</sup> and Alexander Stroh<sup>2</sup>

<sup>1</sup> Thermo Fluid Dynamics and Efficient Energy Conversion, Karlsruhe University of Applied Sciences, Moltkestr. 30, 76133 Karlsruhe, Germany

<sup>2</sup> Institute of Fluid Mechanics, Karlsruhe Institute of Technology, Kaiserstr. 10, 76131 Karlsruhe, Germany

<sup>3</sup> Department of Mechanical and Production Engineering, Aarhus University, Inge Lehmanns Gade 10, 8000 Aarhus, Denmark

corresponding author: [florian.theobald@hs-karlsruhe.de](mailto:florian.theobald@hs-karlsruhe.de)

**Key words:** DNS, Nek5000, OpenFOAM, SIMSON, Xcompact3d, immersed boundary method

**Abstract.** In the present study we investigate an incompressible turbulent channel flow with heat transfer at  $Re_\tau = 180$  with a deterministic surface topography consisting of truncated cones. Two solvers for each of the two boundary handling strategies are considered. With Nek5000 and OpenFOAM the influence of the roughness elements is directly accounted for by an unstructured body fitted mesh, whereas Xcompact3d and SIMSON utilize the immersed boundary method (IBM) to deal with the 3D geometry. The main focus of this work is on an evaluation of the usability of the IBM and a comparison of the parallel performance of the different solvers. Since usability is an ambiguous definition, various quantities are compared: global statistics like Nusselt number and friction coefficient, one-dimensional wall-normal profiles for first and second order statistics, as well as three-dimensional averages over roughness sections. In addition, the computational effort for each method is documented.

## 1 INTRODUCTION

The choice of numerical method and simulation code has a strong impact on accuracy and computational costs of direct numerical simulations. With simple geometric configurations, high performance codes working on structured grids can be used. However, in most cases geometry is complex and cannot be meshed with structured grids. However, body conforming unstructured meshes are feasible despite considerable effort. In some cases, body conforming meshes with sufficient quality are even impossible. In order to reduce pre and post processing costs, the immersed boundary method (IBM), introduced by Peskin [1], can be applied. With the IBM the geometry is not treated explicitly but implicitly by modifying the source terms of the Navier-Stokes equations. Therefore, structured meshes with advantages of easy pre- and postprocessing can be used. However, due to modified source terms using the IBM, uncertainties may be introduced. In this work uncertainties are addressed by comparing the IBM to the results obtained with solvers working on body conforming meshes.

There is a wide variety of solvers suitable for scale resolving simulations. In the present work, four representative codes frequently used by the authors have been chosen. Nek5000 [2] is based on the spectral element method introduced by Patera [3]. Due to the variable order of the polynomials, the behaviour can be adjusted to the problem at hand. The multipurpose simulation toolbox OpenFOAM (Weller et al. [4]) has a large userbase and is used for many different applications. It is based on the finite volume method implemented for unstructured grids allowing arbitrary meshes but features only low order numerical schemes. SIMSON (Chevalier et al. [5]) is based on the pseudo-spectral method. The code is therefore known to be highly accurate and requires relatively low computational resources for large simulations. However, due to the method it is restricted to simple geometries and periodic boundary condition in the flow-parallel plane. Therefore, it was extended by Forooghi et al. [6] with the IBM based on the approach introduced by Goldstein [7] to allow for more complex geometries. Xcompact3d (Bartholomew et al. [8]) is a DNS code based on high-order compact finite differences scheme, which approach the accuracy of spectral methods. The code features an IBM based on the polynomial reconstruction approach introduced by Gautier et al. [9].

Ohlsson et al. [10] carried out a comparison study between Nek5000 and SIMSON. As part of their comparison a smooth wall channel flow was studied. Due to the smooth walls, no IBM was necessary. In their study SIMSON was considerably faster than Nek5000 given similar accuracy. They also noted that for more complex cases the performance of the spectral element method, i.e. Nek5000 may be favourable due to higher geometrical flexibility. Kooij et al. [11] compared the multipurpose codes Nek5000 and OpenFOAM with other more specialized codes designed for Rayleigh-Bénard convection applications in simple geometries without roughness. In their study the specialized codes clearly outperformed Nek5000 and OpenFOAM while Nek5000 was more efficient than OpenFOAM for a given accuracy. They also highlighted that for proper comparison, it may not be sufficient to just look at global statistics. The authors also propose examining local fields for unphysical oscillations. Schäfer et al. [12] compared codes featuring an IBM, namely SIMSON and the former version of Xcompact3d, namely Incompact3d, against a reference simulation performed using Nek5000. Channel flow featuring a transverse bar was used as a test case. The IBM was applied for codes SIMSON and Xcompact3d to account for introduction of the transverse bar. They achieved comparable results in terms of accuracy and computational performance for the codes using the IBM. While the study of a single two-dimensional rectangular obstacle is a good starting point, it remains unclear how the IBM performs for a more complex geometry in 3D.

Taking into account previous research and its findings the test case in this study is chosen to be more complex, i.e. multiple three dimensional non-rectangular roughness elements. In addition to global statistics and one dimensional profiles, we also compare the local flow fields. In total, friction coefficient, Nusselt number, profiles for velocity, temperature, Reynolds stresses, dispersive stresses, turbulent heat fluxes, dispersive heat fluxes, temperature variance, dispersive temperature variance and two dimensional slices for velocity and Reynolds shear stress are compared.

## 2 METHODOLOGY

In the present study we consider an incompressible turbulent flow in a rough wall channel featuring deterministic roughness elements as used by Stripf et al. [13] (Fig. 1a). The friction Reynoldsnumber  $Re_\tau$  is 180. The dimensionless roughness height is eighteen wall units. The flow is driven by a constant pressure gradient  $P_x$  which would correspond to  $Re_\tau = 180$  in a smooth channel. In addition to the flow

field, the temperature field is solved as passive scalar with a Prandtl number  $Pr = 0.71$  and Dirichlet boundary conditions at the walls, i. e. constant temperature difference  $\Delta T_w$  between the upper and the lower wall. Since deterministic roughness is studied it is possible to average the results inside the whole channel onto a smaller roughness unit (Fig. 1b), which allows for a more compact data representation and reduces the necessary simulation time due to a faster temporal convergence. Application of this averaging operation is indicated by a hat symbol in the variable names.




Figure 1: Simulation domain sketch with domain and roughness dimensions. Slope angle of roughness elements is  $55^\circ$ , melt-down height is  $h_{\text{eff}} = 0.0066\delta$ .

In order to obtain physical understanding of the observed changes in the global flow quantities we utilize further data analysis based on the triple decomposition proposed by Reynolds & Hussain [14]. The triple decomposition enables an extraction of the temporal and spatial deviation from the in-plane mean value:

$$\phi(x, y, z, t) = \langle \hat{\phi} \rangle(z) + \hat{\phi}''(x, y, z) + \phi'(x, y, z, t) \quad (1)$$

In order to reach minimum simulation time in terms of CPU hours, it is necessary to determine the coarsest grid resolution from which results are accurately predicted. Therefore, three different mesh sizes are considered for each code. All the results are compared to one reference simulation computed with Nek5000. For this simulation a very fine grid is used. An overview of mesh and simulation parameters is given in Tab. 1. The integration time for each simulation is chosen, so that one dimensional profiles are converged. For Nek5000 case n3 temporal convergence was investigated showing the error for the peak value of the streamwise Reynolds stress to be less than one percent. Since integration time used with the other codes is higher than the one used with Nek5000, it is assumed that the temporal averaging error is low in all considered cases as well.

In Nek5000 polynomial order seven is chosen for all solution variables using the so-called  $\mathbb{P}_N - \mathbb{P}_N$  method. For time integration a second order scheme is used, see section 3.2 for further details. In OpenFOAM the second order *linear* scheme in space and the second order *backward* scheme in time is used. For Nek5000 and OpenFOAM pure hexahedral meshes are used. In SIMSON trigonometric polynomials are used in periodic directions and Chebychev polynomials in wall normal direction. For time integration explicit third order Runge-Kutta and second order implicit Crank-Nicolson is used.

Table 1: Table of simulation parameters for all codes ( $n_{t_b} = T_{\text{sim}}/t_b$  with  $t_b = \delta/U_b$ )

Code	Case	$N_x$	$N_y$	$N_z$	$N_{\text{total}}$ in $10^6$	$\Delta x^+$	$\Delta y^+$	$\Delta z_{\text{Wall}}^+$	$\text{CFL}_{\text{max}}$	$\Delta t^+$ in $10^{-2}$	$n_{t_b}$
Nek5000	Ref	-	-	128	16.22	-	-	0.20	4	3.34	729
	n1	-	-	32	2.65	-	-	1.15	4	3.34	435
	n2	-	-	64	5.31	-	-	0.58	4	3.34	
	n3	-	-	80	7.08	-	-	0.58	4	3.34	
OpenFOAM	o1	-	-	-	6.07	-	-	1.00	0.30	1.40	1200
	o2	-	-	-	11.36	-	-	0.70	0.30	1.14	1000
	o3	-	-	-	27.45	-	-	0.46	0.30	0.65	934
SIMSON	s1	288	256	129	9.51	5.04	2.52	0.05	2.26	12.58	1813
	s2	360	320	201	23.16	4.03	2.01	0.02	2.26	9.65	2888
	s3	576	512	257	75.79	2.51	1.26	0.01	2.26	9.18	702
Xcompact3d	x1	256	128	193	6.32	5.67	5.06	1.00	0.45	7.67	700
	x2	384	192	193	14.23	3.77	3.36	1.00	0.45	7.67	699
	x3	512	256	193	25.30	2.83	2.52	1.00	0.45	7.67	696

Details regarding numerical schemes in Xcompact3d are given in [8].

Using body conforming meshes global statistics like friction coefficient and Nusselt number could be computed by evaluating surface integrals on fluid-solid interfaces. However, since evaluation of surface information for IBM is difficult due to no explicit fluid-solid interface, the quantities are consistently computed via volume integrals for each code. The global skin friction coefficient is then given by:

$$C_f = \frac{2\tau_{\text{eff}}}{\rho U_b^2} \quad \text{with} \quad \tau_{\text{eff}} = \delta_{\text{eff}} P_x. \quad (2)$$

Here,  $U_b$  is the bulk velocity given by  $U_b = 1/2\delta \int_0^{2\delta} \langle \bar{u} \rangle dz$ . In this section the total spatial average is used:  $\langle \phi \rangle = 1/A_{xy} \int_{A_{xy}} \phi dx dy$ , with  $A_{xy} = L_x L_y$ .  $\delta_{\text{eff}} = (\delta - h_{\text{eff}})$  represents the effective channel half height, where  $h_{\text{eff}}$  denotes the melt-down height of the surface structure. Hence,  $\tau_{\text{eff}}$  measures the average shear stress on the entire solid surface. It corresponds to the well-known definition of the wall shear stress  $\tau_w$  in the case of a smooth channel when  $h_{\text{eff}} = 0$ . The corresponding friction velocity and viscous length scale are calculated as  $u_\tau = \sqrt{\tau_{\text{eff}}/\rho}$  and  $\delta_v = \nu/u_\tau$  respectively. Quantities non-dimensionalized by these values are denoted with a superscript + sign.

Following the definition by Nagano et al. [15] the Nusselt number on the lower and upper wall is

$$\text{Nu}_{l,u} = \frac{4\delta_{\text{eff}} \cdot q_{\text{tot}}}{\Delta\theta_{l,u} \lambda} \quad \text{with} \quad \Delta\theta_l = \frac{1}{U_b \delta} \int_0^\delta \langle \bar{u} \rangle \langle \bar{\theta} \rangle dz \quad \text{and} \quad \Delta\theta_u = \frac{1}{U_b \delta} \int_\delta^{2\delta} \langle \bar{u} \rangle (1 - \langle \bar{\theta} \rangle) dz, \quad (3)$$

where  $4\delta_{\text{eff}}$  represents the hydraulic diameter and  $\lambda$  is the thermal conductivity.  $\theta$  is the reduced temperature  $\theta = (T - T_1)/\Delta T_w$ . In the present configuration definitions for  $\text{Nu}_l$  and  $\text{Nu}_u$  yield the same value due to the constant wall-normal heat flux and the symmetry of the configuration, which is referred to as  $\text{Nu}$  throughout the manuscript. The same applies to the bulk mean temperature  $\Delta\theta_u$  and  $\Delta\theta_l$ , which is referred to as  $\Delta\theta$ . The total heat flux  $q_{\text{tot}}$  is estimated as the sum of the viscous and total fluctuation

Table 2: Global statistics; Numbers in parantheses are relative changes to the reference case "Ref"

Code	Case	Friction coefficient $c_f$	Nusselt number Nu
Nek5000	Ref	$1.648 \times 10^{-2}$	36.46
	n1	$1.812 \times 10^{-2}$ (+10.0%)	39.37 (+8.0%)
	n2	$1.688 \times 10^{-2}$ (+2.4%)	37.18 (+2.0%)
	n3	$1.658 \times 10^{-2}$ (+0.63%)	36.73 (+0.74%)
OpenFOAM	o1	$1.508 \times 10^{-2}$ (-8.5%)	34.91 (-4.2%)
	o2	$1.576 \times 10^{-2}$ (-4.3%)	35.82 (-1.7%)
	o3	$1.644 \times 10^{-2}$ (-0.22%)	36.49 (+0.091%)
SIMSON	s1	$1.663 \times 10^{-2}$ (+0.94%)	36.35 (-0.30%)
	s2	$1.650 \times 10^{-2}$ (+0.17%)	36.69 (+0.63%)
	s3	$1.635 \times 10^{-2}$ (-0.77%)	36.98 (+1.4%)
Xcompact3d	x1	$1.649 \times 10^{-2}$ (+0.079%)	35.19 (-3.5%)
	x2	$1.661 \times 10^{-2}$ (+0.83%)	35.73 (-2.0%)
	x3	$1.669 \times 10^{-2}$ (+1.3%)	36.05 (-1.1%)

contributions in the center of the channel:

$$q_{\text{tot}} = \lambda \left. \frac{d\langle \bar{\theta} \rangle}{dz} \right|_{\delta} - \rho c_p \langle w' \theta' \rangle \Big|_{\delta}. \quad (4)$$

Here  $c_p$  denotes the specific heat capacity. As mentioned before, the heat flux is constant in wall-normal direction due to the choice of thermal boundary conditions. Hence, the heat flux  $q_{\text{tot}}$  evaluated in the channel center is equal to the one on the structured wall surface. The friction temperature for inner scaling is defined as  $\theta_{\tau} = q_{\text{tot}} / (\rho c_p u_{\tau})$ .

### 3 RESULTS

In the following, we will first conduct a refinement study and choose a suitable mesh for each code, by analysing global statistics as well as one dimensional profiles for velocity and streamwise Reynolds stress. Then, using these meshes, differences between the four codes for quantities which were mentioned in the introduction are presented. Unlike in the previous section intrinsic spatial averages are used instead of total spatial averages. The spatial averages are then given by  $\langle \phi \rangle = 1/A_f \int_{A_f} \phi dx dy$ , with  $A_f$  the area occupied by fluid.

#### 3.1 Refinement Study

Due to the rather small roughness elements, the minimum number of spectral elements necessary to resolve the geometry in streamwise and spanwise direction is higher than necessary in order to predict the quantites discussed in this study. Therefore, for Nek5000 cases n1, n2 and n3 the resolution in wall-parallel direction is identical and only the wall normal resolution is varied (Tab. 1). The reference case ("Ref"), features a finer mesh also in wall parallel direction, i.e. additional elements forming a thin boundary layer surrounding the roughness circumference are added. The grid of n1 features only one spectral element from wall to roughness crest and another element from roughness crest to channel centerline. This leads to an extremely underresolved bulk flow. Nevertheless, no numerical instabilities




Figure 2: Refinement study; gray shaded area highlights one percent positive and negative change to the reference case. (a) streamwise velocity, (b) streamwise Reynolds stress.

have been encountered and errors for global statistics (cf. Tab. 2) are within 10 %. The shape of the velocity profile (Fig. 2a) is also well predicted. However, it should be noted that the peak value for the Reynolds stress (Fig. 2b) cannot be predicted due to the large grid spacing above the roughness crests. The intermediate case n2 features two elements from wall to roughness crest and two elements above. The thickness of the first element above the crest is chosen, so that wall normal gridspacing at the crest is 0.5 wall units, similar to the reference case. Therefore, in case n2 the roughness elements are considered well resolved in each direction, while the bulk flow is not. Refinement near the roughness elements clearly leads to an improvement in global statistics and the velocity profile. However, the prediction for the streamwise Reynolds stress is still poor. Case n3 is similar to n2, but features an additional element in the bulk flow. This reduces maximum gridspacing in wall normal direction to 20 wall units, which is still quite large. Surprisingly, the results for velocity and streamwise Reynolds stress lie within roughly one percent of the values of the reference case. This is the same range as the temporal convergence error. Thus, case n3 is selected for further comparison.

In OpenFOAM more control over placement of cells than in Nek5000 is possible and therefore the three meshes o1, o2 and o3 feature more uniform refinement than the Nek5000 cases, i. e. refinement in each direction. As a result, convergence is strictly monotonous with mesh refinement. Case o1 featuring the coarsest grid shows similar error in friction coefficient as Nek5000 but only half the error for Nusselt number. Shapes for velocity and streamwise Reynolds stress are well predicted. In the intermediate case o2 errors in global statistics as well as in the velocity and Reynolds stress profiles are approximately halved. Regarding global statistics and the velocity profile the finest case o3 is in accordance with the reference case to within one percent. The error in the Reynolds stress profile is again reduced by a factor of two compared to o2. Assuming consistent convergence behaviour with further refinement, one additional refinement step would be necessary to reach the streamwise Reynolds stress distribution of the reference case within one percent. However, as will be shown later, case o3 is already quite expensive in terms of CPU hours. The next refinement step would lead to approximately 50 million cells and therefore

to exceptionally high computational costs. Therefore, case o3 is selected for further comparison.

Using SIMSON the deviation of  $c_f$  and Nu from the reference case is always less than 1.4% and no clear trend is observed when using finer meshes. Velocity profiles are predicted within one percent of the reference values when using the two finer meshes s2 and s3. This is not the case for the streamwise Reynolds stresses, where the differences to the reference case are always larger and mesh refinement does not show a clear trend. Due to lower computational costs, s2 is selected for further comparison.

The results using Xcompact3d suggest that the friction coefficient converges with finer mesh size, however it departs further from the reference case, while increasing the mesh resolution. The same can be observed for the streamwise Reynolds stress profile. For the Nusselt number the values converge to the reference case with finer mesh resolutions. The deviation of the velocity profile from the reference case is less than one percent for each of the cases x1 ,x2 and x3. Below the roughness crests ( $z^+ = 18$ ) velocity and streamwise Reynolds stress profiles of the three cases are hardly distinguishable, all showing a similar underprediction of  $\langle \widehat{u'_1 u'_1} \rangle^+$ . Due to the very small differences in the profiles, but somewhat larger differences for the Nusselt numbers, case x3 is chosen for further comparison.

### 3.2 Differences in the resultant statistics

In the present study OpenFOAM needs the most number of gridpoints, likely due to the low numerical order (cf. Tab. 1). Despite being high order codes, SIMSON and Xcompact3d, also need many gridpoints. Unlike OpenFOAM, gridspacing in wall parallel direction is equidistant in these two codes and dictated by the rather small roughness elements. Nek5000 needs the least number of gridpoints because of the high numerical order and the flexibility to place mesh elements at appropriate positions. However, in order to resolve the circular top surface of the roughness elements, a large number of very small elements is introduced. As the polynomial order cannot be changed from element to element and is equal in each direction, this leads to many additional gridpoints. As a result, grid spacing is small which in turn leads to a severe time step restriction. In order to allow larger time steps the *OIFS* scheme (Maday et al. [16]) with second order time discretization is used allowing higher Courant numbers. Nevertheless, the time step is still small compared to the other codes. Therefore, geometric flexibility in Nek5000 is more restricted than in low order codes like OpenFOAM.

In respect to global statistics, it is shown above that all codes predict the reference values within roughly one percent if an appropriate mesh is used (Tab. 2). The more interesting differences between the codes become visible by looking at the one dimensional profiles and two-dimensional slices of the flow field. The former are shown in Fig. 3 for different quantities. Velocity and temperature profiles across the whole channel height lie within a band of one percent of the reference case for Nek5000 and OpenFOAM (Fig. 3a,b). The IBM codes SIMSON and Xcompact3d also show excellent agreement above the roughness elements, but closer to the wall larger differences are visible. While Xcompact3d stays within 3 %, SIMSON deviates by more than 5 % from the reference case. The Maximum deviations in peak values for Reynolds stresses (Fig. 3c) are within 4 % for all codes. The agreement for the streamwise component  $\langle \widehat{u'_1 u'_1} \rangle$  is excellent for Nek5000. OpenFOAM and SIMSON slightly overpredict the peak values while Xcompact3d predicts slightly lower values. The spanwise component  $\langle \widehat{u'_2 u'_2} \rangle$  is somewhat underpredicted by all codes with OpenFOAM showing the largest deviation of up to 4 %. Excellent agreement for the wall normal component  $\langle \widehat{u'_3 u'_3} \rangle$  is achieved by the IBM codes and Nek5000 while the values predicted by

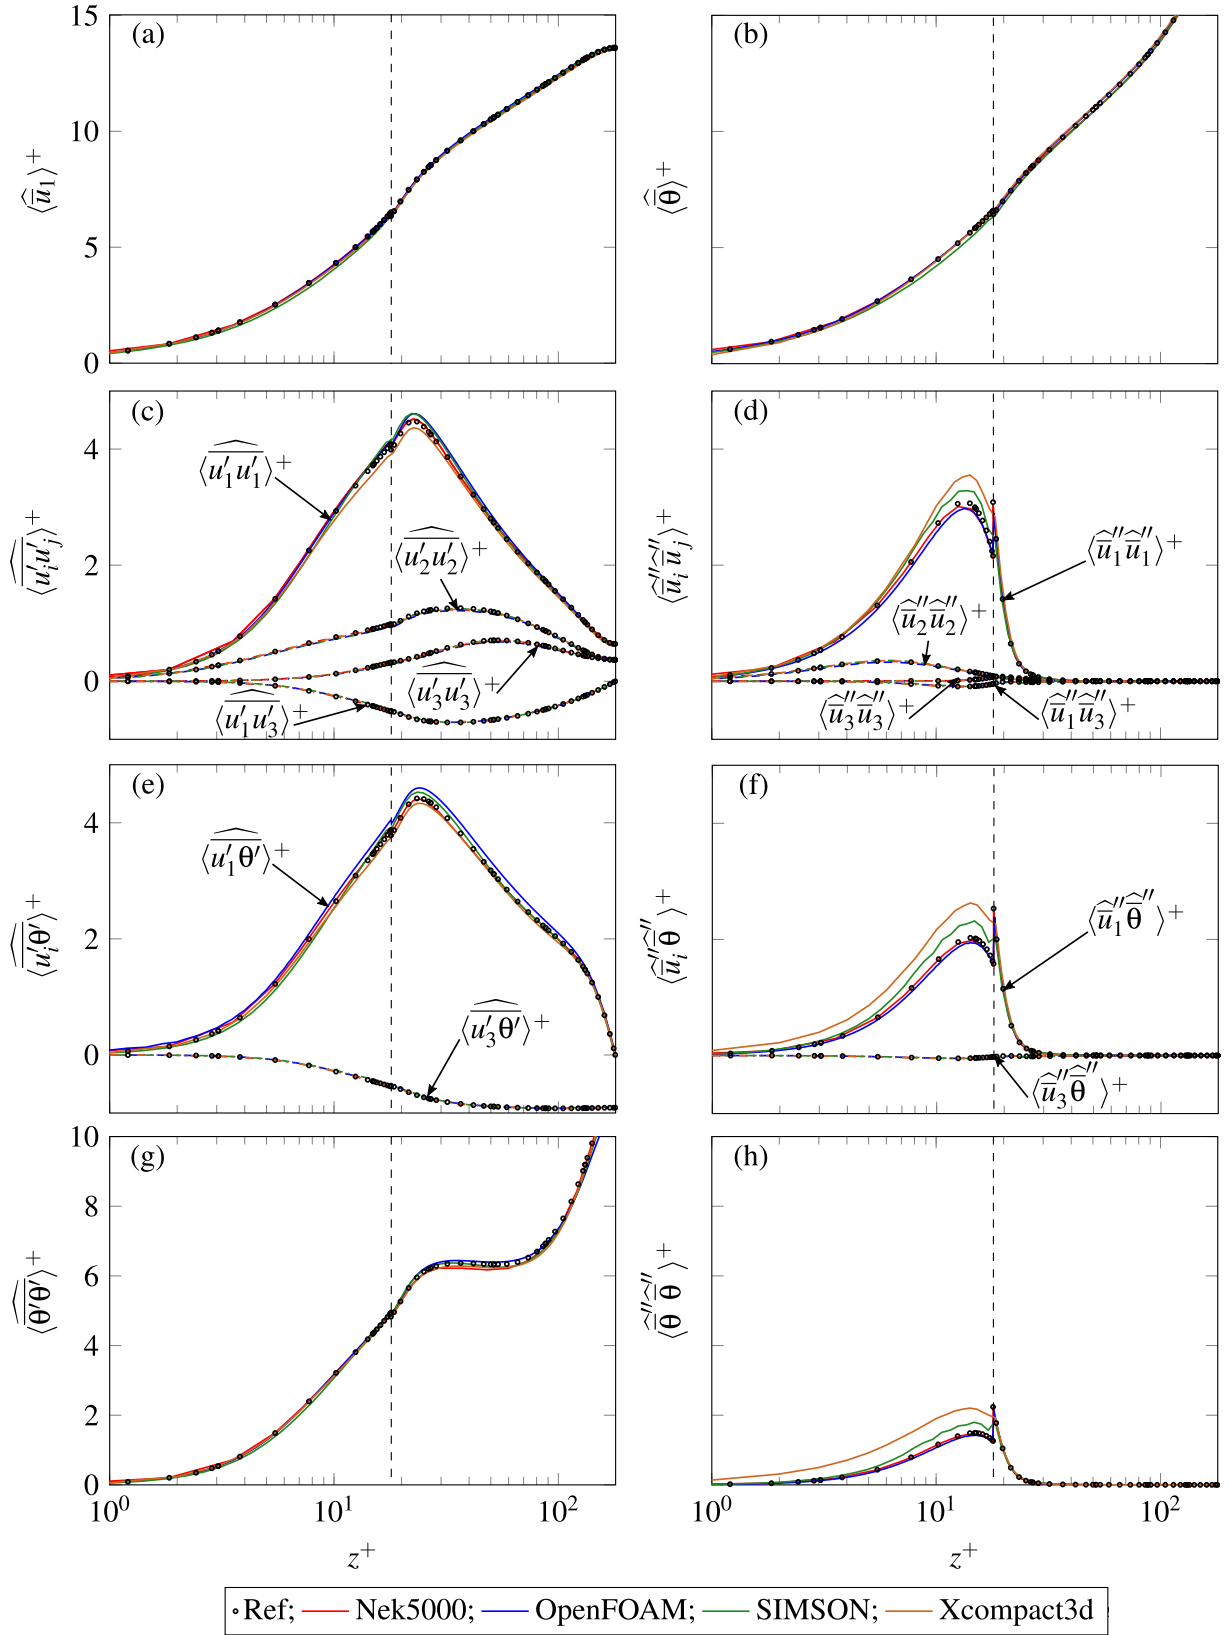


Figure 3: One dimensional profiles in wall normal direction: (a) velocity, (b) temperature, (c) Reynolds stresses, (d) dispersive stresses, (e) turbulent heat fluxes, (f) dispersive heat fluxes, (g) temperature variance, (h) dispersive temperature variance



OpenFOAM deviate by up to 4 % from the reference case. Agreement for the shear stress  $\widehat{\langle u_1' u_3' \rangle}$ , which is an important quantity in terms of modelling, is excellent for all codes. Qualitatively, the findings for the turbulent heat fluxes (Fig. 3e) are similar to those found for the Reynolds stresses (Fig. 3c). The same is true for the temperature variance that all codes predicted within 4 % of the reference case, except SIMSON that shows a slightly larger deviation in proximity to the wall (Fig. 3g).

In comparison to the Reynolds stresses the dispersive stresses show considerable larger deviations from the reference case (Fig. 3d). The most obvious differences are clearly visible for the streamwise component. The deviation of the peak value predicted by Nek5000 and OpenFOAM is 2 % and 3 %, respectively. Values determined by SIMSON deviate by up to 7 % and even 16 % for Xcompact3d. A possible reason for the larger deviation of the IBM codes is given at the end of this section, where 2D slices are discussed. The behavior of the four codes concerning the dispersive heat fluxes (Fig. 3f) are much about the same as for the dispersive stresses. However, the deviation in the peak value of the streamwise component for the IBM codes SIMSON and Xcompact3d is now roughly doubled and reaches 15 % and 30 %, respectively. Like the dispersive stresses and dispersive heat fluxes the dispersive temperature variance shows a similar overprediction by the IBM codes, with deviations up to 20 % and 50 % for SIMSON and Xcompact3d respectively (Fig. 3h).

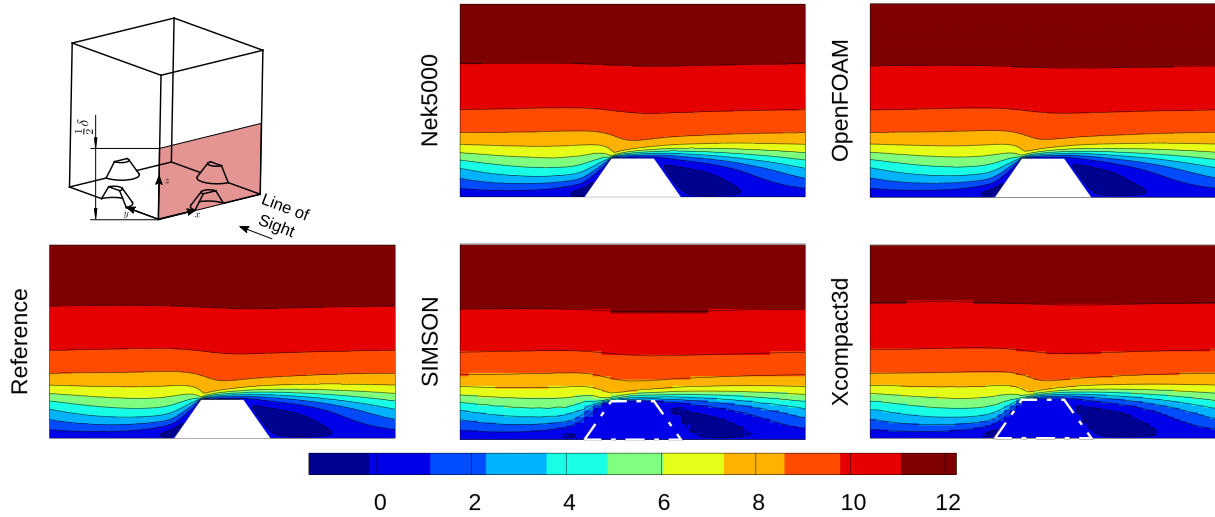


Figure 4: Two dimensional slices of streamwise velocity  $\widehat{u}_1^+$

In order to gain more insight into the phenomena near the roughness elements, two dimensional slices are compared. The excellent agreement of the different codes in predicting the one dimensional velocity profile can be seen in the two dimensional slices as well (Fig. 4). The only noteworthy difference can be seen in the recirculation zone which is slightly smaller in the SIMSON simulation. With respect to the very small differences in the two dimensional slices, it seems inconsistent that profiles of the dispersive stresses for the IBM codes are not in agreement with the reference case. By definition, the dispersive stresses are largest at the fluid solid interface and then rapidly decrease with wall distance. Due to the non sharp distinction between fluid and solid region uncertainties in post processing arise, which might explain the discrepancy.

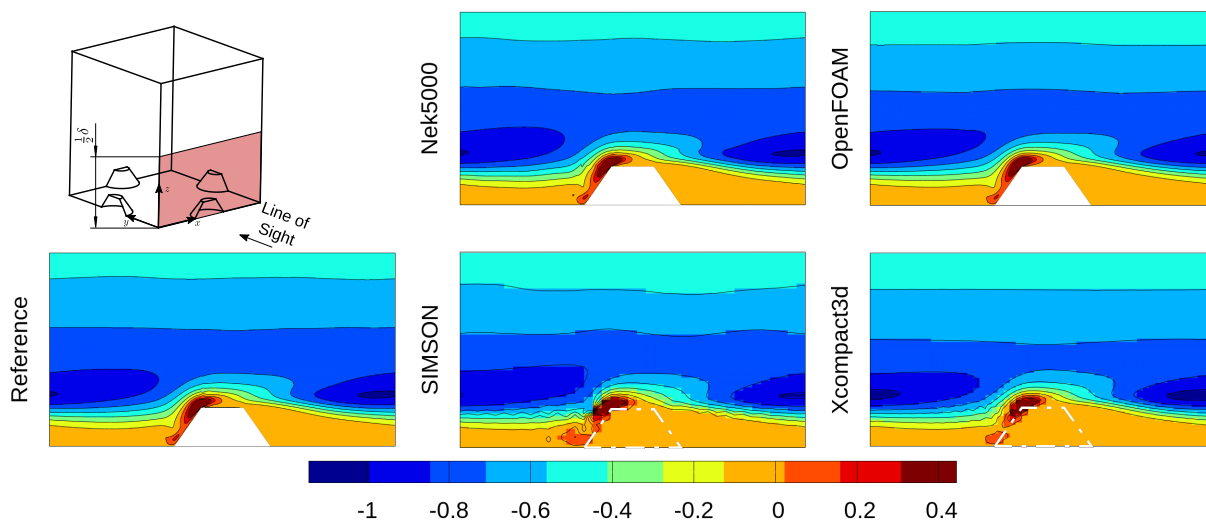


Figure 5: Two dimensional slices of Reynolds shear stress  $\widehat{u'_1 u'_3}$

Unphysical oscillations can be observed in several quantities like the shear stress  $\widehat{u'_1 u'_3}$  in the Nek5000, SIMSON and Xcompact3d simulations (Fig. 5). While those oscillations can be seen in the Nek5000 simulations with coarse grids as well, they disappear with mesh refinement. This is not the case with the IBM codes SIMSON and Xcompact3d where the oscillations persist when grid resolution is increased. These oscillations remain in the vicinity of the cone for Xcompact3d, while for SIMSON they extend further upstream. It should be noted that these oscillations have no negative influence on the 1D profiles which are very well predicted by all codes.

### 3.3 Scalability

Comparison of physical results and the refinement study gave a coarse overview of the necessary problem size in terms of gridpoints and therefore memory requirements. Now, in this last section focus is on how fast results are obtained in terms of CPU hours. For this purpose, a scalability study is conducted on BwUniCluster (2.0) [17] using CPU nodes based on the Intel Broadwell architecture. We consider a strong scaling study, i.e. the total problem size remains fixed. In order to reach maximum performance, hardware optimization flags were used. In addition, for the codes Nek5000 and Xcompact3d proprietary Intel HPC software is compared to the free GNU Compiler and OpenMPI.

The simulation costs in terms of CPU hours per bulk time unit  $t_b$  are compared in Fig. 6a.  $\Delta t_{wc}^*$  is the wall clock time per  $t_b$ . It is clear that for the current problem OpenFOAM is by at least one order of magnitude slower than the other codes. Due to the low numerical order OpenFOAM is highly memory bound whereas the higher order counterparts are more efficient in this case. The performance of Nek5000, SIMSON and Xcompact3d depend on the number of processes. While SIMSON and Xcompact3d perform very well on low core numbers, Nek5000 catches up at higher number of processes.

Speedup and efficiency of each code is given in Fig. 6c and Fig. 6d respectively. Speedup  $s$  and efficiency

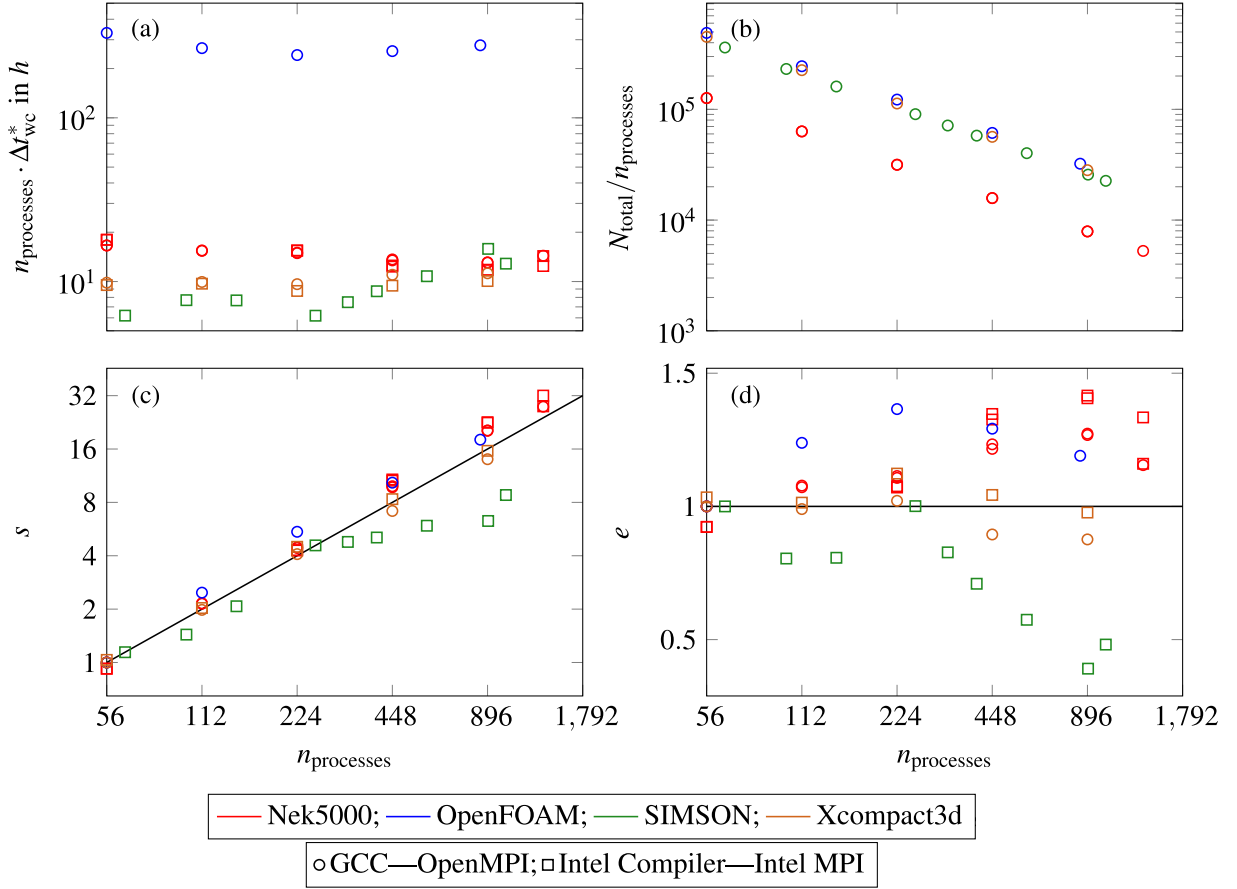


Figure 6: Scalability study: (a) Cost per  $t_b$ , (b) Number of gridpoints per MPI task, (c) Speedup, (d) Efficiency

$e$  is given as:

$$s(\text{code}, n_{\text{processes}}) = \frac{\Delta t_{\text{wc}}^*(\text{code}, n_{\text{processes,ref}})}{\Delta t_{\text{wc}}^*(\text{code}, n_{\text{processes}})}, \quad e(\text{code}, n_{\text{processes}}) = \frac{s(\text{code}, n_{\text{processes}})}{n_{\text{processes}}/n_{\text{processes,ref}}} \quad (5)$$

with  $n_{\text{processes,ref}} = 56$  for Nek5000, OpenFOAM and Xcompact3d and  $n_{\text{processes,ref}} = 64$  for SIMSON. Nek5000, OpenFOAM and Xcompact3d show favourable scalability while SIMSON shows an alternating scalability curve. In SIMSON only 2D decomposition in the two periodic directions is possible and therefore limited. Additionally, the possible decompositions in our case do not fit well to the number of cores in each compute node. Nek5000 is highly efficient down to approximately  $8 \cdot 10^3$  gridpoints per process, which translates to approximately 15 spectral elements with polynomial order of seven. Efficiency of OpenFOAM peaks at approximately  $10^5$  cells per process. Omitting unsuitable decompositions, SIMSON has ideal efficiency at  $4 \cdot 10^5$  and  $10^5$  gridpoints per process. Xcompact3d shows maximum efficiency at approximately  $10^5$  gridpoints per process.

For Nek5000 and Xcompact3d different compilers and MPI implementations are compared. While Xcompact3d performs better with the Intel HPC software throughout the whole study, no clear statement

can be given for Nek5000. At low core counts Intel is a bit slower than the GNU compiler-OpenMPI combination while clearly taking lead at core counts higher than 448.

#### 4 CONCLUSIONS

According to the refinement study Nek5000 needs the least number of gridpoints for the quantities which were compared. OpenFOAM, SIMSON and Xcompact3d need at least three times more gridpoints due to different reasons. Despite their different individual geometry representation strategies all codes show good agreement for integral quantities and one dimensional profiles. The three dimensional fields show also good agreement, but the higher order codes, especially those using the IBM show some local unphysical behaviour in the vicinity of the introduced geometry. While for Nek5000 this unphysical behaviour seems to vanish with mesh refinement, it can not be entirely eliminated for the IBM codes by further increasing mesh resolution.

Good scalability can be achieved with all codes. However, for SIMSON the decomposition must be carefully adjusted to the hardware layout of the cluster. In this specific application, OpenFOAM was significantly slower than the higher order counterparts.

For the problem investigated herein, results achieved with the codes using the IBM were very close to the ones with body conforming approaches at a similar computational cost. However, for the IBM codes larger errors are observed for the dispersive quantities, whose largest magnitudes are found directly at the fluid-solid interface. Therefore, if quantities of interest feature peak values close to the fluid-solid interface, e.g. the dissipation rate of turbulent kinetic energy, one should avoid the usage of the considered IBM approaches. If local three dimensional fields are to be evaluated, the unphysical oscillations generated by the IBM codes, even with very fine grids, should be taken into account.

#### 5 ACKNOWLEDGMENTS

The authors acknowledge support by the state of Baden-Württemberg through bwHPC. The lead author is indebted to the German Federal Ministry of Education and Research for their financial support (grant no. 13FH617IX6).

#### References

- [1] C. S. Peskin, “Flow patterns around heart valves: A numerical method,” *Journal of Computational Physics*, vol. 10, no. 2, pp. 252–271, 1972.
- [2] NEK5000 Version 17.0, 17 Dec 2017, Argonne National Laboratory, Illinois  
Available: <https://nek5000.mcs.anl.gov>.
- [3] A. T. Patera, “A spectral element method for fluid dynamics: Laminar flow in a channel expansion,” *Journal of Computational Physics*, vol. 54, no. 3, pp. 468 – 488, 1984.
- [4] H. G. Weller, G. Tabor, H. Jasak, and C. Fureby, “A tensorial approach to computational continuum mechanics using object-oriented techniques,” *Computers in Physics*, vol. 12, no. 6, pp. 620–631, 1998.
- [5] M. Chevalier, P. Schlatter, A. Lundbladh, and D. S. Henningson, “Simson: A pseudo-spectral solver for incompressible boundary layer flows.” Tech. Rep. TRITA-MEK 2007-07. KTH Stockholm, Stockholm, Sweden.

- [6] P. Forooghi, A. Stroh, P. Schlatter, and B. Frohnäpfel, “Direct numerical simulation of flow over dissimilar, randomly distributed roughness elements: A systematic study on the effect of surface morphology on turbulence,” *Phys. Rev. Fluids*, vol. 3, p. 044605, 4 2018.
- [7] D. Goldstein, R. Handler, and L. Sirovich, “Modeling a no-slip flow boundary with an external force field,” *Journal of Computational Physics*, vol. 105, no. 2, pp. 354 – 366, 1993.
- [8] P. Bartholomew, G. Deskos, R. A. Frantz, F. N. Schuch, E. Lamballais, and S. Laizet, “Xcompact3d: An open-source framework for solving turbulence problems on a cartesian mesh,” *SoftwareX*, vol. 12, p. 100550, 2020.
- [9] R. Gautier, S. Laizet, and E. Lamballais, “A dns study of jet control with microjets using an immersed boundary method,” *International Journal of Computational Fluid Dynamics*, vol. 28, no. 6-10, pp. 393–410, 2014.
- [10] J. Ohlsson, P. Schlatter, C. Mavriplis, and D. S. Henningson, “The spectral-element and pseudo-spectral methods: A comparative study,” in *Spectral and High Order Methods for Partial Differential Equations* (J. S. Hesthaven and E. M. Rønquist, eds.), (Berlin, Heidelberg), pp. 459–467, Springer Berlin Heidelberg, 2011.
- [11] G. L. Kooij, M. A. Botchev, E. M. Frederix, B. J. Geurts, S. Horn, D. Lohse, E. P. van der Poel, O. Shishkina, R. J. Stevens, and R. Verzicco, “Comparison of computational codes for direct numerical simulations of turbulent Rayleigh-Bénard convection,” *Computers & Fluids*, vol. 166, pp. 1 – 8, 2018.
- [12] K. Schäfer, P. Forooghi, S. Straub, B. Frohnäpfel, and A. Stroh, “Direct numerical simulations of a turbulent flow over wall-mounted obstacles—a comparison of different numerical approaches,” in *Direct and Large Eddy Simulation XII* (M. García-Villalba, H. Kuerten, and M. V. Salvetti, eds.), (Cham), pp. 91–96, Springer International Publishing, 2020.
- [13] M. Stripf, A. Schulz, and S. Wittig, “Surface Roughness Effects on External Heat Transfer of a HP Turbine Vane,” *Journal of Turbomachinery*, vol. 127, no. 1, pp. 200–208, 2005.
- [14] W. Reynolds and A. Hussain, “The mechanics of an organized wave in turbulent shear flow. Part 3. Theoretical models and comparisons with experiments,” *J. Fluid Mech.*, vol. 54, no. 2, pp. 263–288, 1972.
- [15] Y. Nagano, H. Hattori, and T. Houra, “DNS of velocity and thermal fields in turbulent channel flow with transverse-rib roughness,” *Int. J. Heat Fluid Fl.*, vol. 25, no. 3, pp. 393–403, 2004.
- [16] Y. Maday, A. T. Patera, and E. M. Rønquist, “An operator-integration-factor splitting method for time-dependent problems: Application to incompressible fluid flow,” *Journal of Scientific Computing*, vol. 5, pp. 263–292, Dec 1990.
- [17] “BwUniCluster 2.0.” [https://wiki.bwhpc.de/e/Category:BwUniCluster\\_2.0](https://wiki.bwhpc.de/e/Category:BwUniCluster_2.0). Accessed: 2021-02-22.