

Comparison of CFD calculations of flows in Hard Disk Drives with Experiments: Part I

Sujit Kirpekar and David B. Bogy

Computer Mechanics Laboratory

Department of Mechanical Engineering

University of California at Berkeley

Berkeley, CA 94720

Telephone: (510)642-4975

Fax: (510)643-9786

`kirpekar@newton.berkeley.edu`

January 23, 2006

Contents

1	Introduction	1
2	Modeling	3
2.1	Experimental Setup	3
2.2	Computational Model Setup	4
2.2.1	Numerical Methods	5
2.2.2	Parametric Grid Generation	6
2.2.3	Approximations for Boundary Conditions	7
2.2.4	Initial Conditions and Statistical Steadiness	8
3	Grid Dependency Studies	11
3.1	Kinetic Energy and Windage	12
3.2	Off-Track and On-Track Drag	15
4	Experimental Validation	17
4.1	Measurements along a line	17
4.2	Measurements on the area	19
4.2.1	Frequency contribution to RMS	20
4.2.2	Length and time scales	21
5	Conclusions and Future Work	23
6	Tables	27
7	Figures	29

List of Figures

1	Top View of Experimental Setup	29
2	Top View of Computational Model	29
3	Wireframe Isometric View of Computational Model	29
4	Typical domain decomposition for parallel computation	29
5	Top view of the computational grid (Grid 1 in Table 1)	30
6	Closeup showing regions 1, 3 and 4 (Grid 1 in Table 1)	30
7	Closeup showing region 3 (Grid 1 in Table 1)	30
8	Global energy-related quantities for Grid 5	31
9	Global energy-related quantities for Grid 5	32
10	Time history of the drag coefficients for Grids 1, 2 and 3.	33
11	Time history of the drag coefficients for Grids 4, 2 and 5.	33
12	Contribution to C_D Off Track RMS from different frequency bands, shown for Grids 1, 2 and 3	34
13	Contribution to C_D Off Track RMS from different frequency bands, shown for Grids 4, 2 and 5	34
14	Contribution to C_D On Track RMS from different frequency bands, shown for Grids 1, 2 and 3	34
15	Contribution to C_D On Track RMS from different frequency bands, shown for Grids 4, 2 and 5	34
16	Comparison of the RMS contributions to C_D Off Track from the finest grid and the extrapolated contributions. Also shown is the GCI_z^{23} across different frequency bands	35

17	Comparison of the RMS contributions to C_D On Track from the finest grid and the extrapolated contributions. Also shown is the GCI_z^{23} across different frequency bands	35
18	Location of measurement line for hot-wire experimental data, from Gross (2003)	36
19	Location of measurement area for hot-wire experimental data, from Gross (2003)	36
20	Mean flow speed along measurement line. LES data is plotted along with error bars representative of GCI	37
21	RMS of flow speed fluctuation along measurement line. LES data is plotted along with error bars representative of GCI	37
22	Turbulence intensity along measurement line. LES data is plotted along with error bars representative of GCI	38
23	Frequency spectrum of velocity fluctuations at different locations, from Grid 5	38
24	Frequency spectrum of velocity fluctuations at different measurement locations, reproduced from Gross (2003)	38
25	Mean flow speed in m/s, plotted for various x-positions	39
26	RMS flow fluctuations in m/s, plotted for various x-positions.	40
27	Mean flow velocity over entire measurement area, from LES	41
28	Mean flow velocity over entire measurement area, from experiments of Gross (2003)	41
29	RMS flow fluctuations over entire measurement area, from LES	41
30	RMS flow fluctuations over entire measurement area, from experiments of Gross (2003)	41
31	0-2 kHz contribution to RMS flow fluctuations, from Gross (2003)	42
32	0-2 kHz contribution to RMS flow fluctuations, from LES	42
33	2-20 kHz contribution to RMS flow fluctuations, from Gross (2003)	42

34	2-20 kHz contribution to RMS flow fluctuations, from LES	42
35	2-6 kHz contribution to RMS flow fluctuations, from Gross (2003)	43
36	2-6 kHz contribution to RMS flow fluctuations, from LES	43
37	6-10 kHz contribution to RMS flow fluctuations, from Gross (2003)	43
38	6-10 kHz contribution to RMS flow fluctuations, from LES	43
39	10-20 kHz contribution to RMS flow fluctuations, from Gross (2003)	44
40	10-20 kHz contribution to RMS flow fluctuations, from LES	44
41	Integral time scale of the flow (in number of disk rotations)	44
42	Integral length scale of the flow (mm)	44
43	Cross term of the time-averaged Reynolds stress tensor	45

List of Tables

1	Grids with variable out-of-plane (z) resolution	27
2	Grids with variable in-plane ($r\phi$) resolution	27
3	Grid Convergence results for global quantities	28

Abstract

In this two part study, comprehensive validation results are presented for the simulation of turbulent flows in model hard disk drives. In this report, which is the first part of the study, we compare simulation results with the hot-wire experimental data from the thesis of Gross (2003). In the second part, contained in a separate report, comparison is drawn against the results of Barbier (2006). Large eddy simulation (LES) results of the airflow between two co-rotating disks with a single e-block arm as an obstruction are reported here. The computational model closely mimics the simulation conditions. In this report, all results are presented in the context of grid convergence and estimates of the associated uncertainty are reported. In the second report, numerical dissipation and iteration errors are reported. We found that grids in the 2-2.5 Million cells range (for a 3 inch drive) are in the asymptotic range. By varying the grid independently in the z - and $r - \phi$ directions, the sensitivity and convergence characteristics of the in-plane and out-of-plane resolutions are reported. It is found that the simulations show monotonic convergence in z and oscillatory convergence in $r - \phi$, with GCI values approximately 20-30 % for most quantities. In validating our LES results with the hot wire experiments, we found good agreement in the mean quantities but larger discrepancies in the RMS quantities. Our findings show that LES results tend to over predict the fluctuations compared to the experiments in almost all frequency bands, and spectra converge to solutions that do not qualitatively agree with the experiments.

1 Introduction

With the rapid proliferation of hard disk drives into non-traditional applications such as music players, cell phones and digital cameras, new demands are placed on the size, speed and reliability of these drives. Correspondingly, there has been a strong demand for higher areal density, faster data transfer rates and better reliability. To achieve higher track densities, the

available area to position the read-write head is reduced, thereby reducing the tolerance for track-misregistration (TMR). On the other hand, higher data transfer rates result in higher speeds of disk rotation, which in turn increase the Reynolds number of the air flow.

Flows in hard disk drives have been investigated by several authors, including the current authors. For a recent summary of both the experimental and the computational progress so far, we refer the reader to Kirpekar and Bogy (2005b). However, in most numerical studies of such complex flows, little attention is paid to reporting the numerical errors and uncertainties of the results. While it has become relatively easy to calculate such flows using commercial CFD software, the accuracy of these results is questionable at best. In the current work, we hope to shed some light on the sources of discrepancies between numerical and experimental results on disk drive flows.

Errors (i.e. the difference between the simulation result and the actual physical value) may be divided into two broad parts: modeling errors and numerical errors. Modeling errors are due to mathematical assumptions of the physical problem itself; e.g. the assumption of incompressibility, the application of simplified boundary conditions, the use of a sub-grid scale turbulence model, the assumption of isothermal flow, etc. Numerical errors are those due to the technique of solving the mathematical problem; e.g. discretization in space and time, grid convergence, artificial dissipation and dispersion, truncation of the iteration process in every time step, computer round off, etc. By providing an estimate for each of these sources of errors, simulation results may be corrected and accurate results may be reported. Estimation of such errors will allow the placement of an accurate “error bar” on any simulation data reported. While it may be next to impossible to account for all of these errors, our work focuses on determining the error introduced by the grid (i.e. the discretization error). In the second part of this work, we also consider the iteration error and the errors introduced by artificial dissipation.

The specific motivations of this work are as follows:

1. To demonstrate typical grid resolutions needed to obtain solutions in the asymptotic range.
2. To quantify the numerical errors and uncertainties of disk drive turbulence simulations which can be computed using current computational resources.
3. To validate our computational results against experimental data sets that investigated realistic drive configurations, and finally,
4. To provide insights into certain physical aspects of the flow that may not be readily understood from experiments

In the first of this two-part paper, we validate our results against the experimental data of Gross (2003), while in the second paper, we discuss the results of Barbier (2006).

2 Modeling

2.1 Experimental Setup

The experimental setup (with which we propose to benchmark our calculations) is described in detail in Gross (2003). For clarity, a schematic of the setup is shown in Figure 1 (reproduced directly from Gross (2003)).

The setup consists of two co-rotating glass disks of 84 mm diameter. The disk spacing is 2.0 mm and the shroud gap is 1 mm wide. A single e-block arm of 1.0 mm thickness was placed between them without the use of any suspension or slider assembly. The e-block arm was not actuated but it could be fixed in three positions to experiment with inner-, middle- and outer-diameter configurations. Additionally, the thickness of the e-block arm was also varied from 1.0 to 1.6 mm. A constant-temperature hot-wire anemometer was used

for velocity measurements. The hot-wire probe was oriented axially at the midplane of the setup, which made it “*most sensitive to the in-plane flow speed component*”.

It is important to note that in the experimental setup, the disks are shrouded for only 250 degrees of their circumferential arc-length. The remaining shroud is cut away to allow for the insertion of the e-block arm and the hot-wire probe. This region is essentially open to the atmosphere and poses some difficulty in computational modeling.

2.2 Computational Model Setup

Our computational model tries to closely follow the experimental setup of Gross (2003). The same geometrical dimensions are used for the disks and the e-block arm. The computational model (without the grid) is shown in Figure 2. An isometric wireframe view is also shown in Figure 3, which shows the smaller out-of-plane (z) dimension as compared to the in-plane ($r - \phi$) dimensions.

The Reynolds number of the flow based on the disk-to-disk spacing is 5,533, while based on the disk outer radius it is 116,197. Traditionally, the former method of reporting the Reynolds number is more prevalent, because it accounts for the axial length scale. In any case, the presence of a blunt body obstruction breaks the azimuthal symmetry and makes the flow turbulent, requiring the use of a turbulence model for simulation.

The Kolmogorov’s microscale can be estimated from the Reynolds number associated with the largest eddies of the flow:

$$\eta = l \left(\frac{u'l}{\nu} \right)^{-3/4} \quad (1)$$

Here we may estimate the size of the largest eddies (l) to be equal to the disk-to-disk thickness, 2mm. And assuming that the velocity associated with the large eddies (u') is 10% of the maximum linear disk velocity, the Kolmogorov’s scale (η) is approximated to 0.0175

mm. This is a valid *apriori* assumption, based on the experimental data of Gross (2003). Also, η is in good agreement with earlier estimates of Kazemi (2004) and Kirpekar and Bogy (2004).

Given the Reynolds number of the flow, the Kolmogorov dissipation scale and the geometric volume of interest, a true direct numerical simulation would require more than 200 Million cells – which is the reason why most flows in disk drives are addressed using Large Eddy Simulation (LES).

2.2.1 Numerical Methods

Our large eddy simulations are performed using a commercial CFD code, CFD-ACE. The Algebraic Dynamic Model (Germano et al., 1991) was used for the sub-grid scale modeling. Our prior experience with using commercial codes to compute LES solutions (Kirpekar and Bogy, 2005a) indicates that this is the most suitable SGS model.

CFD-ACE uses a box-filter for both the grid filtering and test filtering. The grid filter is taken to be the grid itself, while the width of the test filter is twice that of the grid filter. In addition to spatial averaging (smoothing), the dynamic coefficient is also artificially truncated to prevent instabilities.

CFD-ACE employs a segregated solver using the SIMPLEC technique. The elliptic pressure-coupling equation is solved using a multigrid technique. The maximum number of iterations are limited to 50 per time step. For most time steps, the solver would converge within 40 iterations. The criteria for convergence was that the residuals for each quantity (in the velocity or pressure equations) be less than 10^{-4} .

Central differencing was used in space with a first order implicit method in time. A constant time step of 10^{-5} seconds was used. A second order semi-implicit method (Crank Nicholson) would lead to unstable calculations for all our grids, hence the use of the first order method. Regarding the artificial dissipation associated with the first order time integration,

very little difference was noticed between solutions using a first order and a second order method for the LES of a blunt body flow (see Kirpekar and Bogy, 2005a). This is most likely due to the small time step size.

2.2.2 Parametric Grid Generation

To study the grid dependency of our LES solutions, simulations were conducted for several different grids. To ensure a close geometric relationship between the different grids the mesh generation was parametrized. The grid was completely generated by specifying the number of nodes (and their distribution) along the edges. By changing the number of nodes uniformly (say in geometric progression) very similar (but refined) grids could be generated. The meshing strategy was to completely specify the grid parameters in the plane of the disks and then extrude the entire domain axially.

The in-plane region of the grid was divided into four distinct regions (See Figure 5 and 6):

1. Coarse structured grid: which accounts for a major part of the flow domain and does not contain any obstructions
2. Fine structured grid in the shroud gap: Here the grid is refined to resolve the streamline curvature near the shroud. However, this refinement is only sufficient to resolve the main features of the flow in the shroud, but not the boundary layer adjacent to the shroud
3. Upstream and downstream structured grid refinement: The grid is refined in the region immediately upstream and downstream of the arm. This allows the accurate placement of the first node downstream from the solid wall of the arm. This helps us resolve (or not resolve, depending on the grid) the separated shear layer and the associated small

turbulent structures close to the wall of a blunt body. (See Figure 7 for a close up view of this region)

4. Upstream and downstream unstructured grid relaxation: To interface the fine grid near the arm with the coarser grid in the rest of the domain, an unstructured grid was used. The meshing tool for unstructured grids produces quadrilateral-dominant cells (90% quads, 10% triangular cells) which drastically reduces the number of cells needed compared to a purely triangular mesh (See Figure 6 for close up view of this region)

Our initial attempts at grid dependency studies showed that solutions changed quite differently due to in-plane refinement as compared to out-of-plane refinement. For this reason, the grid was refined independently along the two orthogonal directions and convergence of the solutions is reported accordingly.

2.2.3 Approximations for Boundary Conditions

The boundary conditions for the computational domain are implemented as follows:

1. The top and bottom disks (along with the central hub) are modeled as rigid rotating walls. Effects such as run-out (especially NRRO), clamping distortions and disk vibrations cannot be accounted for in this model.
2. The actuator (with only one e-block arm) is also modeled as a fixed obstruction to the flow, with no-slip boundary conditions. Vibrations of the arm (which are of the order of a few nanometers) are not used as boundary conditions to the flow field.
3. Relaxation zone: The original experiments of Gross (2003) are “open” to the atmosphere in the region downstream of the arm. Similar boundary conditions are applied in our computational domain by radially extending the domain 5 mm from the shroud. Atmospheric pressure boundary conditions allowing the inflow and outflow of air are

then applied to the edge of this extended region. If the computational domain had not been extended, atmospheric pressure boundary conditions would have to be applied at the edge of the rotating disk. This would, however, not be physical, since we do not expect the pressure to be atmospheric immediately close to the edge of the disk. By extending the domain outwards, a “relaxation zone” is created where the pressure in the drive may adjust to the ambient conditions.

4. In the computational domain the shroud gaps are modeled as symmetric boundaries (slip wall boundary conditions). This ensures that in the gap, airflow is permitted only in the plane of the disks, but not perpendicular to them. Since the addition of cells to the top and bottom of the current domain is not computationally feasible, this is a good approximation to the narrow shroud gap. Alternately, a periodic boundary condition may be enforced between the top and bottom shroud gaps, such that the flow leaving the domain at the top reenter the domain at the bottom. This boundary condition, however, led to unphysical travelling waves in the velocity solutions, and hence was not used.

2.2.4 Initial Conditions and Statistical Steadiness

All of our LES calculations are initialized from steady state $k - \epsilon$ solutions to the flow field. CFD-ACE uses the original $k - \epsilon$ implementation of Launder and Sharma (1974) with $C_\mu = 0.09$, $C_{\epsilon_1} = 1.44$, $C_{\epsilon_2} = 1.92$, $\sigma_k = 1.0$ and $\sigma_\epsilon = 1.3$. Given the empirical nature of the ϵ equation, and the use of coefficients based on simple turbulent shear flow, we did not expect the ϵ solution to be accurate. This is manifested in the high residuals for the ϵ variable, which do not reduce even with very large number of iterations (10,000). However, velocity and pressure values at various points in the domain remained constant (within 10% of the mean) after about 250 iterations. Using this as a guideline, each $k - \epsilon$ solution was computed for 2000 iterations and the resulting solution was used as the initial conditions for

the LES calculation.

From prior experience, it is understood that instantaneous solutions of an LES are quite different (qualitatively) from the steady $k - \epsilon$ solutions. On integration in time, the LES solutions change rapidly from the predicted initial conditions and gradually achieve statistical steadiness. However, since the flow is highly turbulent, a local measure of steadiness (e.g. based on the convergence of the mean velocity at one point) is generally inappropriate, and a more global metric needs to be defined. For our simulations, we compute the (filtered) kinetic energy and the Windage ¹ and use an energy balance argument to claim statistical steadiness.

To illustrate this technique, we consider the following definitions and equations. Let U be the three-dimensional velocity vector. The kinetic energy of the flow may be defined as:

$$E(\mathbf{x}, t) = \frac{1}{2} U \cdot U \quad (2)$$

while the filtered kinetic energy can be obtained by filtering the kinetic energy field,

$$\overline{E}(\mathbf{x}, t) = \frac{1}{2} \overline{U \cdot U} = E_f(\mathbf{x}, t) + k_R(\mathbf{x}, t) \quad (3)$$

where the kinetic energy of the filtered velocity field is defined as,

$$E_f = \frac{1}{2} \overline{U} \cdot \overline{U} \quad (4)$$

and the residual kinetic energy is defined as,

$$k_R = \frac{1}{2} \overline{U \cdot U} - \frac{1}{2} \overline{U} \cdot \overline{U} \quad (5)$$

It is easy to derive the conservation equation (see Pope (2003) or Kundu (1990)) for E_f ,

¹Windage is defined as the power supplied by the disk motor to the flow domain through the rotating disks and hub

which is,

$$\frac{\partial E_f}{\partial t} + \bar{U} \cdot \nabla E_f = \frac{\partial}{\partial x_i} \left\{ \bar{U}_j \left(2\nu \bar{S}_{ij} - \tau_{ij}^r - \frac{\bar{p}}{\rho} \delta_{ij} \right) \right\} - 2\nu \bar{S}_{ij} \bar{S}_{ij} + \tau_{ij}^r \bar{S}_{ij} \quad (6)$$

where the filtered rate-of-strain tensor is given by,

$$\bar{S}_{ij} = \frac{1}{2} \left(\frac{\partial \bar{U}_i}{\partial x_j} + \frac{\partial \bar{U}_j}{\partial x_i} \right) \quad (7)$$

and the sub-grid scale (residual) stress tensor τ_{ij}^r is algebraically determined from S_{ij} on applying the grid and test filters (see Germano et al. (1991) and Lilly (1992)).

Let $\mathcal{V}(\mathbf{x})$ be the volume and $\mathcal{A}(\mathbf{x})$ be the surface area of our computational domain. \mathcal{A} may be subdivided into, $\mathcal{A} = \mathcal{A}_w + \mathcal{A}_d + \mathcal{A}_o$, where the subscripts refer to stationary “walls” (both no-slip and symmetry planes), rotating “disks” and flow “outlets”.

Integrating Eqn. 6 over \mathcal{V} and converting the divergences into surface integrals over \mathcal{A} , we obtain the following energy balance,

$$\begin{aligned} & \underbrace{\frac{\partial}{\partial t} \int_{\mathcal{V}} E_f d\mathcal{V}}_{\text{Rate of change of KE}} - \underbrace{\int_{\mathcal{A}_d} (2\nu \bar{U}_j \bar{S}_{ij} - \bar{U}_j \tau_{ij}^r) d\mathcal{A}}_{\text{Windage}} = \quad (8) \\ & \underbrace{- \int_{\mathcal{A}_o} E_f \bar{U}_j d\mathcal{A}}_{\text{Flux of KE}} + \underbrace{\int_{\mathcal{A}_o} (2\nu \bar{U}_j \bar{S}_{ij} - \bar{U}_j \tau_{ij}^r) d\mathcal{A}}_{\text{Net work by stresses at outflow}} - \int_{\mathcal{A}_o} \frac{\bar{p}}{\rho} \bar{U}_j d\mathcal{A} \\ & + \underbrace{\int_{\mathcal{V}} (-2\nu \bar{S}_{ij} \bar{S}_{ij} + \tau_{ij}^r \bar{S}_{ij}) d\mathcal{V}}_{\text{Viscous and SGS dissipation}} \end{aligned}$$

In this equation, the *Flux of kinetic energy* is the net kinetic energy produced or destroyed due to the flow of air outside our domain. The *Net work by stresses at outflow* is the work done by the surface forces (arising from the shear stress and SGS stress) on the computational

volume at the boundary.

To achieve a statistical steady state it is important that an energetic balance is achieved, i.e. the energy production and dissipation balance each other, and that the net rate of change of kinetic energy be small. Since the velocity is solenoidal, our calculations are mass conserving, and we do not expect very large contributions to the kinetic energy from the outflow/inflow. A dominant balance is therefore expected between the Windage and the (combined viscous and SGS) dissipation.

Based on the explanation above, we computed the kinetic energy and the Windage of each simulation as the calculation progressed. Statistics of the flow (such as means, r.m.s. and higher moments) are then calculated only after the kinetic energy has “settled down”, i.e. did not change by more than 5% of its mean value. This provided us a systematic method for estimating statistics of the flow based on global quantities rather than on a point by point basis. The initial transients typically lasted for about 2-3 revolutions (1200-1800 time steps) of the disk. Our calculations are continued until 8 revolutions – giving us 6 revolutions (3600 time steps) of useful data.

3 Grid Dependency Studies

Grids in the 0.5 Million cell range, which may be computed on a single desktop machine showed very poor convergence and hence the resolution was increased to approximately 2.5 Million cells. Any more refinement would have been impractical as the LES would require very long computation times. Each of the five simulations reported here was run for 2-3 weeks on a clustered Linux system using 8-32 CPUs to gather data for 8 disk revolutions. (An example of a parallel domain decomposition is shown in Figure 4) While the current authors have simulated cases in excess of 15 Million cells, such computations cannot be practically included in the kind of grid refinement studies presented here.

The various grids used in our work (labeled: Grid 1 to Grid 5) are described in Tables 1 and 2. Grid 1, 2 and 3 denote increasing z - resolution (i.e out-of plane resolution), while Grids 4, 2 and 5 represent increasing $r - \phi$ (in-plane) resolution. Since each grid was generated by completely specifying the grid in one $r - \phi$ plane, and extruding it axially, the in-plane and out-of-plane resolutions could be varied independently. The average resolution of the grid may be computed from the volume or area of the domain and the number of cells. In these tables, representative grid resolutions, h , h_z and $h_{r\phi}$ are determined using the following definitions:

$$h = \left(\frac{\text{Volume}}{N} \right)^{1/3} \quad (9)$$

$$h_z = \frac{\text{Axial dimension}}{N_z} \quad (10)$$

$$h_{r\phi} = \left(\frac{\text{In-plane area}}{N_{r\phi}} \right)^{1/2} \quad (11)$$

In presenting the convergence results for two orthogonal directions, we often notice that the two sets of grids (1-2-3 and 4-2-5) are converging to different results when extrapolated to $h = 0$. Nevertheless, the actual value of the result at $h = 0$ is not of much consequence to us, since it is significantly affected by several factors other than the grid (as discussed briefly earlier). However, extrapolated error and the grid convergence index (GCI) are very useful in quantifying the uncertainty of the results.

3.1 Kinetic Energy and Windage

We start by discussing the convergence of global quantities such as the kinetic energy, windage and drag on the arm. The windage and the drag on the arm are especially important to the disk drive community, because they refer to the power required by the motor

to run the disks and the force on the actuator, respectively. These quantities are referred to as “global”, because they are obtained by integration in space, and the integrand is dependant on the properties of the flow at several locations. Global quantities are expected to show better behavior in convergence than local estimates, since the integration should smooth out local errors and present an average estimate of the rates of convergence.

The evolution of the kinetic energy and windage are shown in Figure 8. In this Figure, results are presented for the finest grid (Grid 5) and error bars are included based on the Grid Convergence Index (GCI) of the mean kinetic energy and windage. In obtaining the GCI, we have followed the guidelines of the ASME Journal of Fluids Engineering, policy statement on the control of numerical accuracy. In this figure, and all subsequent figures, error bars are applied to data from the finest grid itself, instead of the more customary practice of using the extrapolated data. Usually, the extrapolated solutions are close enough to the finest grid calculations to be included in the uncertainty error bars. Nonetheless, if these computations were to serve as a benchmark for future validation efforts, the fine grid data would be more useful than the extrapolated solutions.

In Figure 8 the quantities are non-dimensionalized using the following definitions: Let, $U_o = \Omega r_o$ be the disk edge velocity, where Ω is the rotation speed and r_o is the disk outer radius. Let \mathbb{V} be the volume of the domain and the \mathbb{A}_d be the area of the disks. Then, the non-dimensional kinetic energy and windage may be defined as,

$$k^* = \frac{\frac{1}{2} \int_{\mathcal{V}} \mathbf{u} \cdot \mathbf{u} d\mathcal{V}}{\frac{1}{2} U_o^2 \mathbb{V}} \quad (12)$$

$$W^* = \frac{\int_{\mathcal{A}_d} (2\nu \overline{u_j} \overline{S_{ij}} - \overline{u_j} \tau_{ij}^r) d\mathcal{A}}{\left[\frac{1}{2} U_o^2\right] [U_o] \mathbb{A}_d} \quad (13)$$

In the same Figure 8, the third sub-figure shows the rate of change of k^* . Finally, in the

fourth sub-figure, the difference between the rate of change of kinetic energy and windage (i.e. right-hand-side of Eqn 8) is plotted, which is mainly the combined viscous, SGS and numerical dissipation.

From the figure, we notice that the rate of change of kinetic energy decreases almost to zero after about 2 revolutions of the disk. The kinetic energy *decays* from its steady value, indicating the $k - \epsilon$ solutions tend to overpredict the kinetic energy of the flow. Interestingly, the decay in kinetic energy is very close to an exponential function, and a direct comparison of an exponential curve with the kinetic energy is plotted in Figure 9. The rate of decay was found to have a time constant of 0.737 revolutions, suggesting, that the kinetic energy will achieve 5% of it's mean value in 2.209 revolutions. In reporting the rest of our results, our statistical averaging is started after the kinetic energy is within 5% of its converged mean value. In Figure 8, this is a little after 2 revolutions. At about 3 revolutions, the change in kinetic energy is less than 1% of the mean.

The error bars in the Figure 8 are based on the data from Table 3, by using the higher value of GCI. The table also reports the absolute error in the solutions (e_a) and the error in the extrapolated solution (e_{exp}). The GCI is computed separately in the in-plane and the axial directions and is reported in Table 3. Since we are dealing with global quantities, the GCI calculations are based on the global grid size h , and not on directional resolutions such as h_z and $h_{r\phi}$. The kinetic energy and windage both show monotonic convergence in both the z - and $r - \phi$ directions. In Table 3 the calculated order of convergence ranges from 1.02 to 2.71, which is in good agreement with the formal order of accuracy, 2. This is also an indication that the chosen grids are in the asymptotic range. In general, increasing the resolution causes both the mean kinetic energy and windage to decrease. From this one may infer that increasing the number of cells allows the resolution of smaller flow structures associated with smaller kinetic energies. The energy cascade from the larger to the smaller eddies is thus responsible for lowering the total kinetic energy of the domain. Interestingly, in case of both

the kinetic energy and windage, we observe a higher sensitivity to the z-resolution than the $r - \phi$ resolution. This also leads to the result that $GCI_z^{23} > GCI_{r\phi}^{25}$, which implies a higher uncertainty due to the resolution in the z direction. It is known that the velocity profiles in a disk drive are similar to turbulent Couette flow (interdisk velocity profiles are reported extensively in Kirpekar and Bogoy (2006)). Since the principal mechanism to generate kinetic energy is from the rotating disks, the interdisk resolution plays a vital role in the kinetic energy of the flow. The momentum being “pumped” into the domain is highly dependent on the resolution in the boundary layer. On the other hand, the in-plane resolution (especially in the wake of the arm) determines the rate of loss of kinetic energy to the viscous and SGS sinks. The overall result is that the energetics of the flow domain are more sensitive to the z-resolution than the $r - \phi$ resolution in the range considered.

3.2 Off-Track and On-Track Drag

A similar time history of the coefficient of drag, C_D , on the actuator is plotted in the Figures 10 and 11. The time history is shown for the final 6 revolutions of the computation. The coefficients are further decomposed into Off-Track and On-track directions, where Off-Track is the direction perpendicular to the axis of the e-block arm, and On-track is the direction parallel to the axis of the e-block arm. In computing these coefficients the projected areas of the arm and the disk edge speed U_o are used.

From Figure 10 we observe that C_D Off-track is almost twice as large as C_D On-track, which is due to the orientation of the arm in the rotating flow. Figure 10 also shows that increasing the z-resolution increases the mean Off-Track drag but decreases the mean On-track drag. Interestingly, the RMS values of both the Off-Track and On-track drag reduce. This suggests that under-resolved simulations, which are dominated by the large scale motions, tend to over predict the fluctuations of pressure acting on the arm. Increasing the resolution allows the cascade to (slightly) smaller scales than before, resulting in smaller fluctuations

at the large eddy level.

There is little difference in the convergence results for the z and $r - \phi$ directions for the drag, given in Table 3. Our results indicate that the GCI is high (20-30%) for the mean and RMS values of drag coefficients. The RMS values of the drag coefficients show oscillatory convergence in the $r - \phi$ direction and hence the GCI is not reported.

The RMS values of the drag coefficient on the arm may be broken down into frequency components using Parseval's theorem. Information regarding the amount of energy associated with different frequency bands is important to the disk drive component designers, who may then design structures with natural frequencies that do not fall in the heavily excited bands.

Figures 12 and 13 show the RMS contribution from different frequency bands to the Off-Track drag coefficient. Similarly, Figures 14 and 15 show the RMS contribution from different frequency bands to the On-Track drag coefficient. Interestingly, some clear trends are demonstrated: By increasing the resolution, the low frequency contribution (0-1 kHz) decreases, while the higher frequency contribution, especially 1-6 kHz, increases. This trend is consistently demonstrated in both the z - and $r - \phi$ directions; However, as seen in Figures 12 and 13, convergence is monotonic in z - but oscillatory in the $r - \phi$ direction. We note that to obtain the resultant RMS due to all frequency bands, algebraic addition is not permitted, but the RMS values should be added geometrically: by summing their squares and taking the square root.

The Figures 12 to 15 display an important trend in the frequency components of the excitation force on the actuator. This data (in the monotonically convergent cases), may be used to obtain the extrapolated solution and the GCI. It is most useful to directly compare the extrapolated values with the values from the finest grid (Grid 3), along with the GCI. This is done in Figure 16 for the Off-Track component and Figure 17 for the On-Track component. Again, the Figures show very interesting results. Firstly, the difference between the extrapolated solution and the solution from the finest grid decreases with increasing

frequency. Generally, there is excess energy in the lower frequencies, but less energies in the higher frequencies. Secondly, the GCI decreases with increasing frequency, indicating that the LES solutions converge much faster in the higher frequency components. In Figure 16, the very high GCI value in the 6-10 kHz range is hard to explain and may be considered spurious. The contributions in the 10-50 kHz range are not analyzed for convergence since the values are very close to each other.

4 Experimental Validation

In this section our numerical results are directly compared with the experimental data of Gross (2003). Two experimental data sets are available: Shown in Figure 18 are measurements along a single line in the wake of the arm, and shown in Figure 19 are measurements in a broader rectangular area, again downstream of the e-block arm. The measurement area in Figure 19 is referenced by an x-y coordinate system

4.1 Measurements along a line

Figure 20 shows the mean velocity along the measurement line, Figure 21 shows the RMS velocity and Figure 22 shows the turbulence intensity (i.e. the ratio of the RMS to the mean velocity). In these figures the distance along the measurement line is non-dimensionalized by the length of the line, so all the plots range from 0 to 1. The 0 end of the plot corresponds to outside the edge of the disk, while the 1 end of the plot corresponds to the inner location (see Figure 18) In all the figures the percentage occurrence of oscillatory convergence is displayed at the top along with the average order of convergence. In plotting the error bars on the figures the GCI was determined using the usual formula, but with the average order p_{avg} , of the method. The error bars were then included in the figures at ten equispaced locations. In all three figures a higher number of points showed oscillatory convergence in

the $r - \phi$ direction than in the z -direction, hence the GCI estimates are from convergence in the z -direction.

In general there is higher agreement in the mean quantities than in the RMS quantities. The spatial variation of the mean velocity along the measurement line is in fairly good agreement with the experimental data. Remarkably, the agreement is very good close to the outer edge of the disk, where we expect the influence of the outflow boundary condition. This indicates that the relaxation region included in our simulations provides a good estimation to the physical outflow boundary. The velocity profiles show a higher local variation compared to the experimental data – which is smoother. The reason for this may be that the experimental data is based on readings taken over several minutes (i.e. several thousand revolutions) while the computational data is averaged for 6 revolutions only. In general, the percentage of points showing oscillatory convergence is higher for the RMS velocity than the mean velocity.

It is the general observation that the LES results tend to over predict the RMS fluctuations of velocity. This is consistent with the drag results outlined previously, and it tends to corroborate the notion that LES simulations on the current grids tend to under resolve the smaller scales of motion, leading to higher fluctuations in the large scales. Figure 22 shows the turbulence intensity along the measurement line, which is the ratio of the RMS to the mean velocities. Again, the turbulence intensity is higher in the simulation compared to the experiment, but the agreement is good close to the outflow boundary condition. In Gross (2003), in addition to the turbulence intensity, the mean and RMS dynamic pressure head is also reported. These quantities can be easily deduced from the mean and RMS velocities, hence we do not report them here.

In Figure 23 the frequency spectrum of the velocity fluctuations is plotted, which may be compared with the experimental results in Figure 24. Several observations can be made with regard to Figure 23. Firstly, the data is more noisy than the experimental results because of

the limited data set available. Secondly, there are no clear peaks corresponding to frequency locking. The orientation of the arm in the rotating flow and the complex geometry of the arm itself, generated a complex wake. The vorticity shed from the arm organizes itself into eddies behind the arm but this phenomenon is not self-selective of any frequency. The fluctuations are contained in the low frequencies (0-6 kHz) and are much smaller at frequencies beyond that. This unsteadiness appears to be mostly random, but the flow structures that are shed are long lived and coherent. These flow structures are carried around by the rotating disk and are dissipated in time.

4.2 Measurements on the area

Figures 25 and 26 show a direct comparison between the LES and experimental data. Again, the LES results are from Grid 3, and the percentage oscillatory convergence and average order of accuracy are included on the top of each sub-figure. The error bars on the LES data are based on the GCI from the average order of accuracy.

In the Figures 25 and 26 the mean and RMS velocities are plotted as functions of the y-coordinate (ranging from -6 to 6). Different x-locations (ranging from 2 to 12) are plotted in different sub-figures. See Figure 19 for the location of the x-y coordinate system.

As shown in these figures, there is better agreement in the mean velocities than in the RMS velocities. The experimental data shows that the mean velocity has a radial gradient and there is a well defined transition from a smaller velocity to a larger velocity when going from $y = -6$ to 6. This is because the flow is blocked immediately downstream of the arm, and is accelerated in the space between the arm and the hub. The LES data also shows a similar trend, but the transition is a little further away from the hub. For $x = 12$ and $x = 10$ the magnitudes are in remarkable agreement.

In terms of RMS, the experiments show a moderate level of fluctuations in the wake, and a slight increase in the fluctuations in the region where the mean velocity transitions, followed

by much smaller fluctuations approaching the hub. The LES results, however, show different qualitative features. They exhibit a higher level of RMS fluctuations and a significantly higher peak in the flow transition region. Additionally, $x = 12$, $x = 10$ and $x = 8$ also show a peak in RMS near the hub too.

Finally, Figures 27 to 30 graphically summarize the results in the rectangular measurement area. While Figures 27 and 28 share the same color scale for the mean velocity, Figures 29 and 30 have different color scales for the RMS as denoted. The figures for mean velocity show the transition of the velocity from the blocked region to the accelerated region. The experimental figure also shows, by a dotted line, the location of the suspension slider assembly if it were to be included in the setup. Figures 29 and 30 show the larger differences between the RMS fluctuations, as discussed earlier.

4.2.1 Frequency contribution to RMS

The frequency contribution to the RMS from different frequency bands is now discussed in detail. Figures 31 and 32 compare the 0-2 kHz frequency contributions to the RMS, while Figures 33 and 34 compare the contributions from 2 to 20 kHz.

In both cases the RMS from the LES is approximately two to three times larger than that predicted from the experiments. While it is unclear what the exact source of discrepancy is, it is well known that simulations tend to over predict some components of the RMS fluctuations. E.g. in a separate study by the authors, the flow across a square cylinder was computed and streamwise Reynolds stresses (i.e. the streamwise velocity RMS fluctuations) were over predicted. This over prediction was due to the nature of the SGS model itself and we may conclude that modeling error contributes significantly to the prediction of the velocity fluctuations.

Both Figures 32 and 34 show the clear stream of shed eddies that contribute to higher fluctuations. The region blocked by the arm has higher fluctuations than the accelerated flow

region. The thesis of Gross (2003) also breaks down the 2-20 kHz contribution to the RMS in to 2-6, 6-10 and 10-20 kHz bands. Figures 35 through 40 provide a direct comparison between the LES and the experiments for these frequency bands. The general trend is that the LES consistently predicts higher fluctuations compared to the experimental data in all frequency bands. With increasing grid resolution, both in the $r - \phi$ and z-directions, the 2-6 and 6-10 kHz contribution to the RMS increases, while the 0-2 kHz contribution decreases. This trend is exactly similar to the trend in the drag coefficients shown in Figures 12-15 and is hence not repeated. This leads to the conclusion that when performing calculations on successively refined grids, LES solutions converge to spectral contents that do not qualitatively agree with the experimental spectra. Thus grid-free LES solutions can never agree perfectly with experiments, which is most likely due to the deficiency in the SGS model, as hinted at earlier. Other factors, such as limited data for averaging LES solutions, influence of boundary conditions and the uncertainty in the hot-wire measurement process may also contribute to the discrepancy between the results.

4.2.2 Length and time scales

The characterization of the turbulent flow is not complete without the specification of a time scale and a length scale. The integral time scale of the flow may be computed using the normalized auto-correlation function,

$$\rho(s) = \frac{\frac{1}{T} \int_0^T u'_\phi(t) u'_\phi(t+s) dt}{\frac{1}{T} \int_0^T u'^2_\phi(t) dt} \quad (14)$$

where $u'_\phi = u_\phi - \overline{u_\phi}$, according to the Reynolds decomposition.

The integral time scale may be then computed as,

$$\tau = \int_0^\infty \rho(s) ds \quad (15)$$

and invoking Taylor’s frozen field hypothesis (see Pope, 2003), the integral length scale may be computed as:

$$\lambda = \tau \overline{u_\phi} \quad (16)$$

The integral time scale and length scale are shown in Figures 41 and 42. The time scale is non-dimensionalized to represent the number of disk rotations. The time scale is the largest in the accelerated part of the area and is relatively small in the region of the wake. This indicates that although the flow is being accelerated in this region, the flow remains largely laminar and fluctuations are well correlated for almost a whole revolution of the disk. In the more turbulent wake, the fluctuations flow remain uncorrelated, and the integral time scales are small. In the laminar region the combined effect of flow acceleration and larger time scales, leads to much larger length scales. The length scales are much smaller in the wake. This indicates that the largest flow structures in the domain are contained in the laminar flow region and the wake is characterized by much smaller flow structures with shorter life spans.

We also note that most of the eddies in the turbulent wake have a length scale of about 2 mm or less, which is also the disk-to-disk spacing in the model. Hence, estimations of the Kolmogorov’s microscale based on this estimate is valid, as done previously.

Finally, the cross term of the Reynolds stress tensor representing $\overline{u_r u_\phi}$ is plotted in Figure 43. It clearly shows a distinct ridge in the stress component in the region where the flow transitions from the accelerated region to the wake region. This indicates the region of strong production of turbulence and the region where the turbulent field is anisotropic. The negative sign of the stress component is typical of a turbulent shear flow and indicates the production of turbulence from the interaction between the fluctuating field and the mean field.

5 Conclusions and Future Work

To summarize:

1. For the first time, comprehensive grid convergence results have been presented for flows in hard disk drives. We found that grids in the 2-2.5 Million cells range (for a 3 inch drive) are in the asymptotic range. While it is customary to vary the grid uniformly in all three dimensions and report the convergence, such an effort would have missed the independent sensitivity and convergence characteristics of the in-plane and out-of-plane resolutions.
2. In the face of limited computational resources and very long simulation time, we have also outlined a rigorous and novel technique to define the (statistical) steadiness of the flow. This is based on monitoring the kinetic energy and windage of the flow. We found that simulations initiated from steady $k - \epsilon$ solutions decay exponentially to their steady values, which is helpful in deciding an averaging interval for reporting the statistics of the flow.
3. Our grid convergence results mainly show monotonic convergence in z and oscillatory convergence in $r - \phi$, with GCI values approximately 20-30 % for most quantities. More importantly, we noticed higher sensitivity of the quantities to the z -resolution, which indicates the importance of resolving the axial dimension adequately for accurate simulation. The simulation results also show that increasing the grid resolution changes the spectral content of the drag on the arm. Increasing grid resolution decreases the 0-1 kHz content while increasing the higher 2-6 kHz spectral content. Finally, the results of this paper can assist disk drive CFD practitioners to estimate the grid based uncertainty of their simulations and compensate (correct) their results based on the data presented here.

4. In validating our LES results with the hot wire experiments, we found good agreement in the mean quantities but larger discrepancies in the RMS quantities. Generally, statistical quantities reported in an LES do not account for the direct influence of the unresolved scales and hence such comparisons should be made with caution. Our findings show that LES results tend to over predict the fluctuations in almost all frequency bands, and the spectra converge to solutions that do not qualitatively agree with the experiments. While the contribution from the highest frequencies is very small, (e.g. the contribution of 10-20 kHz range to the RMS is only 0.6%) LES results still overpredict the amount of fluctuations arising from this frequency band. It is postulated that the modeling assumptions, mainly the SGS model, are responsible for the lack of agreement.

Acknowledgment

This study was supported by the Computer Mechanics Laboratory (CML) at the University of California, Berkeley.

The authors wish to acknowledge the contribution from Intel Corporation, Hewlett-Packard Corporation, IBM Corporation, and the National Science Foundation grant EIA-0303575 in making hardware and software available for the CITRIS Cluster which was used in producing these research results.

References

- C. Barbier. *Experimental and numerical study of the flow in a simulated hard disk drive*. PhD thesis, University of Virginia, Charlottesville, 2006.
- M. Germano, U. Piomelli, P. Moin, and W. H. Cabot. A dynamic sub-grid scale eddy viscosity model. *Physics of Fluids*, A(3):1760–1765, 1991.
- H. Gross. *Off-Track Vibrations of the Read-Write Heads in Hard Disk Drives*. PhD thesis, University of California, Berkeley, 2003.
- H. Kazemi. *Mathematical Modeling of Flow-Induced Vibrations of Suspension-Head Units in Hard Disk Drives*. PhD thesis, University of Virginia, Charlottesville, 2004.
- S. Kirpekar and David B. Bogy. A comparison of large eddy simulation models for numerical simulation of airflow in hard disk drives. *CML Report*, 04-013, 2004.
- S. Kirpekar and David B. Bogy. Computing bluff body flows using commercial cfd software. *CML Report*, 05-008, 2005a.
- S. Kirpekar and David B. Bogy. A study on the efficacy of flow mitigation devices in hard disk drives. *CML Report*, 05-003, 2005b.
- S. Kirpekar and David B. Bogy. Computing the aeroelastic disk vibrations in a realistic hard disk drive. *CML Report*, 06-001, 2006.
- P. J. Kundu. *Fluid Mechanics*. Academic Press, 1990.
- B. E. Launder and B. I. Sharma. Application of the energy-dissipation model of turbulence to the calculation of flow near a spinning disc. *Letters Heat Mass Transfer*, 1:131–137, December 1974.

D. K. Lilly. A proposed modification of the germano subgrid scale closure method. *Physics of Fluids*, A(4):633–635, 1992.

S. B. Pope. *Turbulent Flows*. Cambridge University Press, 2003.

6 Tables

Table 1: Grids with variable out-of-plane (z) resolution

Grid Name	Volume mm^3	Number of cells, N	Number of out-of-plane cells, N_z	Number of in-plane cells, $N_{r\phi}$	h mm	h_z mm	$h_{r\phi}$ mm
Grid 1	1.043×10^4	1,101,264	16	68,829	0.2116	0.1250	0.2735
Grid 2	1.043×10^4	1,651,896	24	68,829	0.1848	0.0833	0.2735
Grid 3	1.043×10^4	2,202,528	32	68,829	0.1679	0.0625	0.2735

Table 2: Grids with variable in-plane ($r\phi$) resolution

Grid Name	Volume mm^3	Number of cells, N	Number of out-of-plane cells, N_z	Number of in-plane cells, $N_{r\phi}$	h mm	h_z mm	$h_{r\phi}$ mm
Grid 4	1.043×10^4	1,171,632	24	48,818	0.2072	0.0833	0.3248
Grid 2	1.043×10^4	1,651,896	24	68,829	0.1848	0.0833	0.2735
Grid 5	1.043×10^4	2,361,324	24	98,389	0.1640	0.0833	0.2288

Table 3: Grid Convergence results for global quantities

Global Quantity		p_z	$e_{a,z}^{23}$ (%)	$e_{ext,z}^{23}$ (%)	GCI_z^{23} (%)	$p_{r\phi}$	$e_{a,r\phi}^{25}$ (%)	$e_{ext,r\phi}^{25}$ (%)	$GCI_{r\phi}^{25}$ (%)
Kinetic Energy	Mean	1.85	2.66	15.87	17.12	2.71	2.19	6.09	7.17
Windage	Mean	1.45	1.19	8.71	10.01	1.19	1.19	8.54	9.83
Off-track drag	Mean	1.56	2.48	13.28	19.15	1.02	2.39	15.63	23.16
	RMS	1.09	1.26	12.85	14.23	—			
On-track drag	Mean	1.74	1.90	11.69	13.09	1.48	5.02	34.93	32.36
	RMS	1.74	2.38	15.05	16.35	—			

7 Figures




Figure 1: Experimental Setup (from Gross, 2003) (diagram is not to scale)

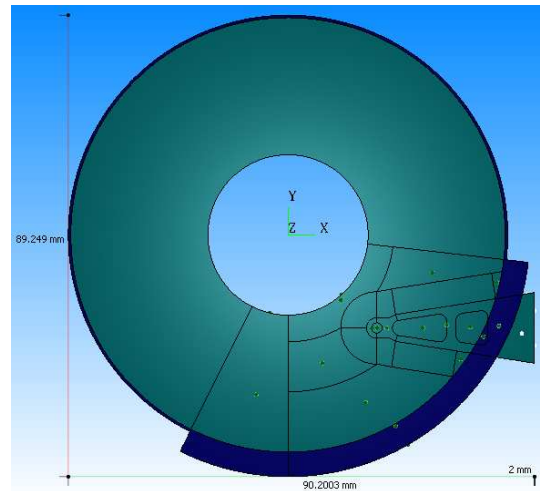


Figure 2: Top View of Computational Model (diagram is to scale)




Figure 3: Wireframe Isometric View of Computational Model

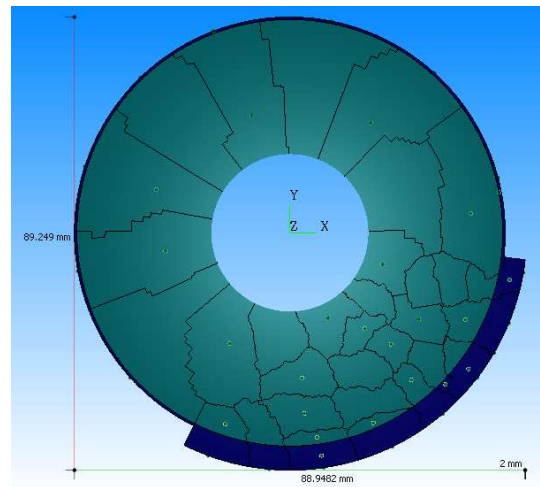


Figure 4: Typical domain decomposition for parallel computation

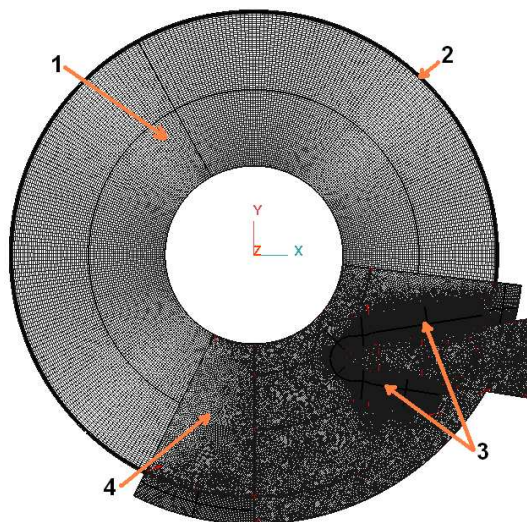


Figure 5: Top view of the computational grid (Grid 1 in Table 1)




Figure 6: Closeup showing regions 1, 3 and 4 (Grid 1 in Table 1)




Figure 7: Closeup showing region 3 (Grid 1 in Table 1)

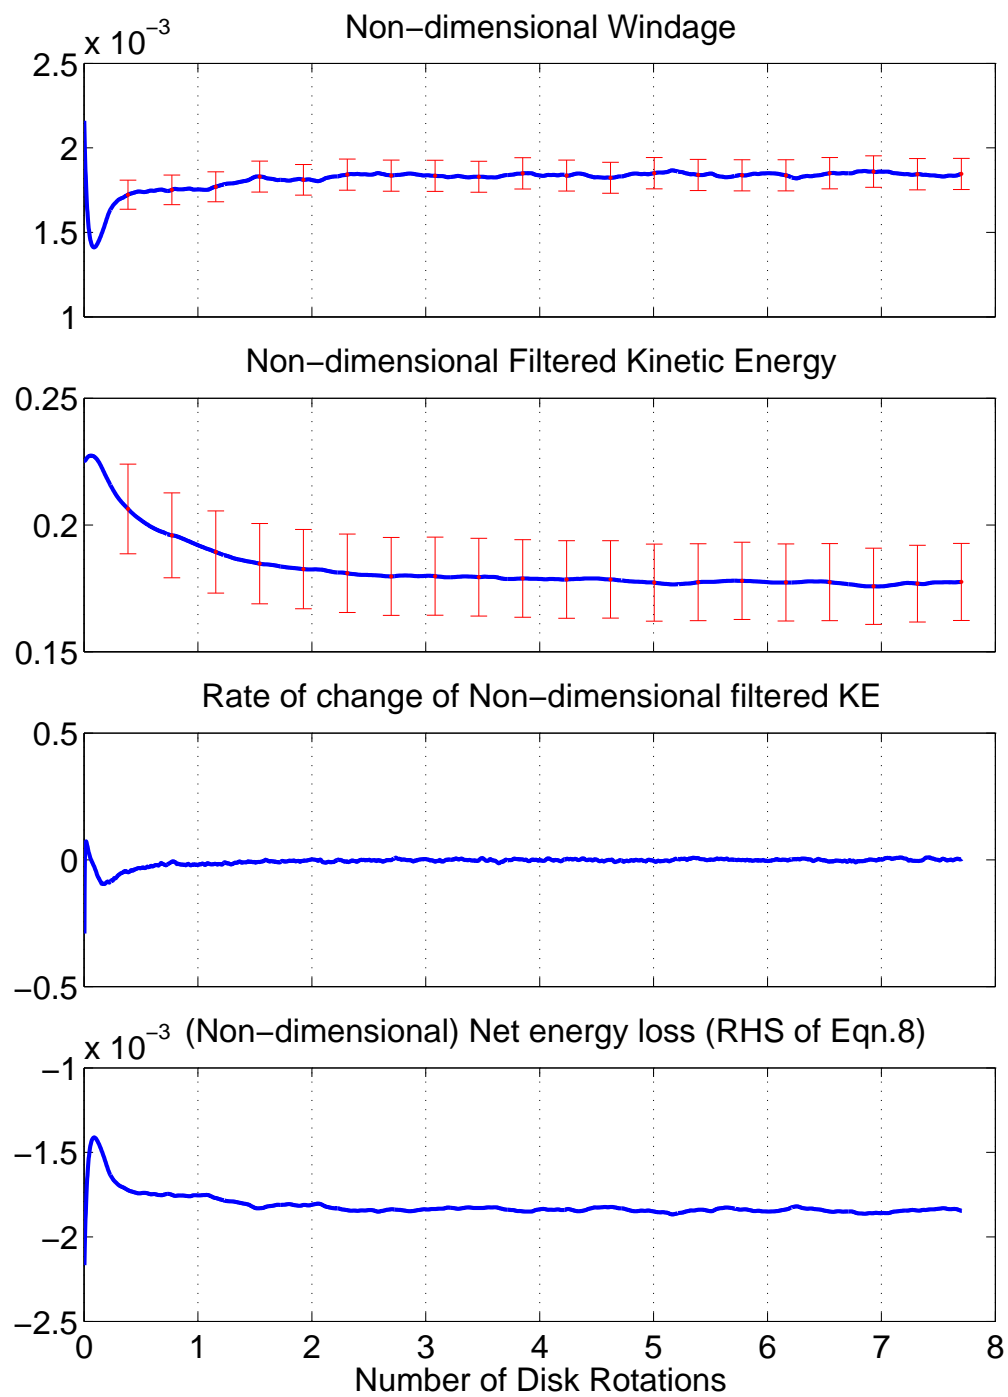


Figure 8: Global energy-related quantities for Grid 5

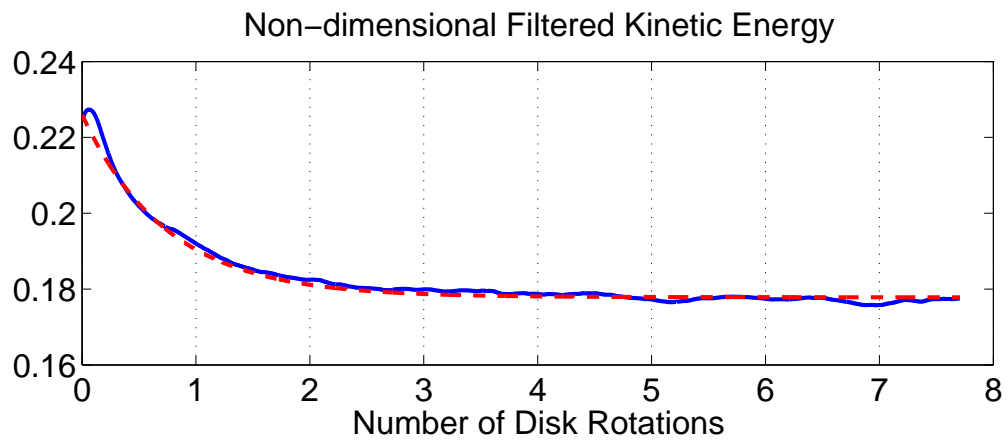


Figure 9: Global energy-related quantities for Grid 5




Figure 10: Time history of the drag coefficients for Grids 1, 2 and 3.

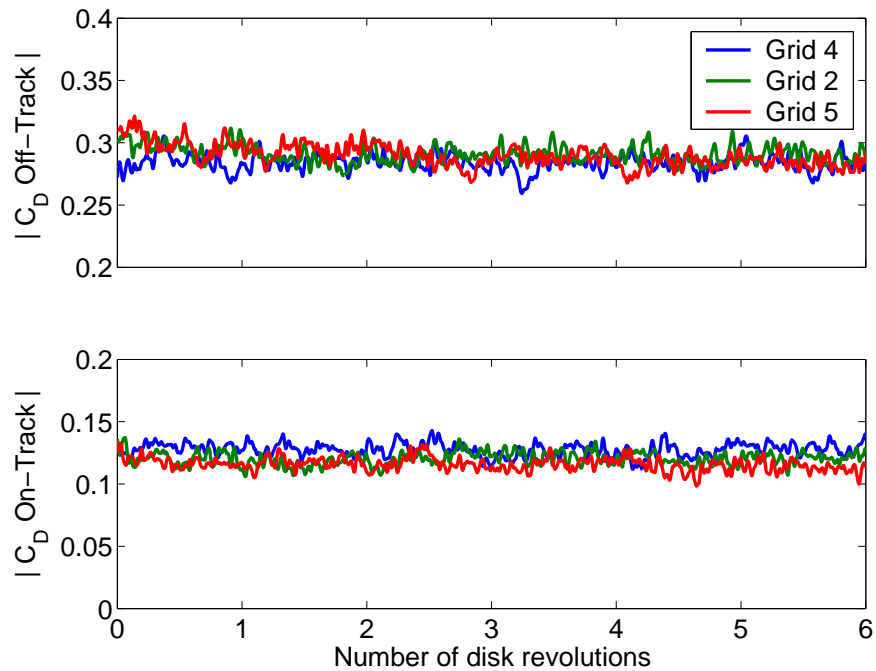


Figure 11: Time history of the drag coefficients for Grids 4, 2 and 5.




Figure 12: Contribution to C_D Off Track RMS from different frequency bands, shown for Grids 1, 2 and 3




Figure 13: Contribution to C_D Off Track RMS from different frequency bands, shown for Grids 4, 2 and 5




Figure 14: Contribution to C_D On Track RMS from different frequency bands, shown for Grids 1, 2 and 3




Figure 15: Contribution to C_D On Track RMS from different frequency bands, shown for Grids 4, 2 and 5

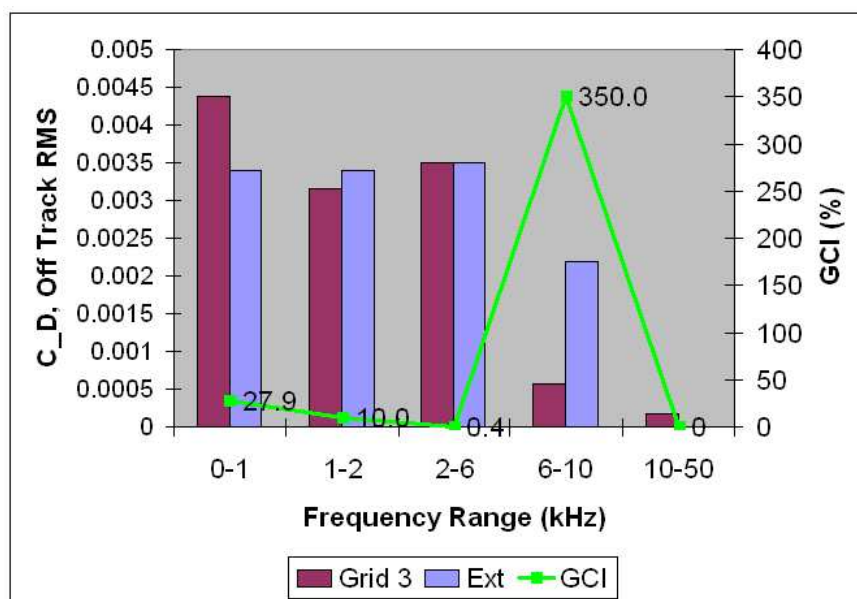


Figure 16: Comparison of the RMS contributions to C_D Off Track from the finest grid and the extrapolated contributions. Also shown is the GCI_z^{23} across different frequency bands

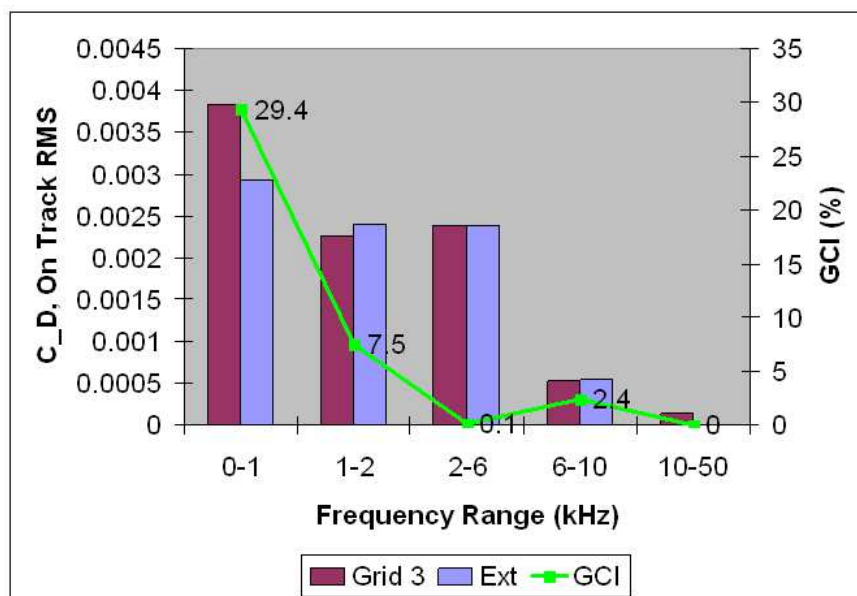


Figure 17: Comparison of the RMS contributions to C_D On Track from the finest grid and the extrapolated contributions. Also shown is the GCI_z^{23} across different frequency bands




Figure 18: Location of measurement **line** for hot-wire experimental data, from Gross (2003)




Figure 19: Location of measurement **area** for hot-wire experimental data, from Gross (2003)

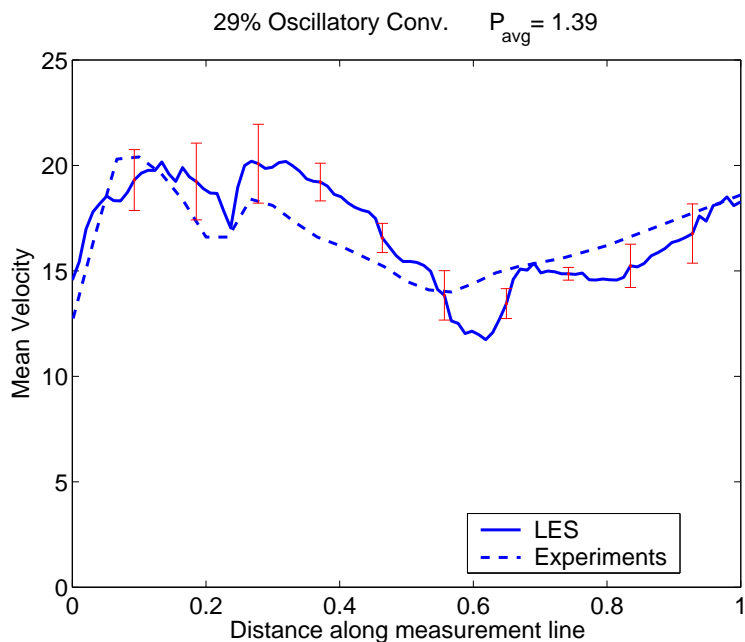


Figure 20: Mean flow speed along measurement line. LES data is plotted along with error bars representative of GCI

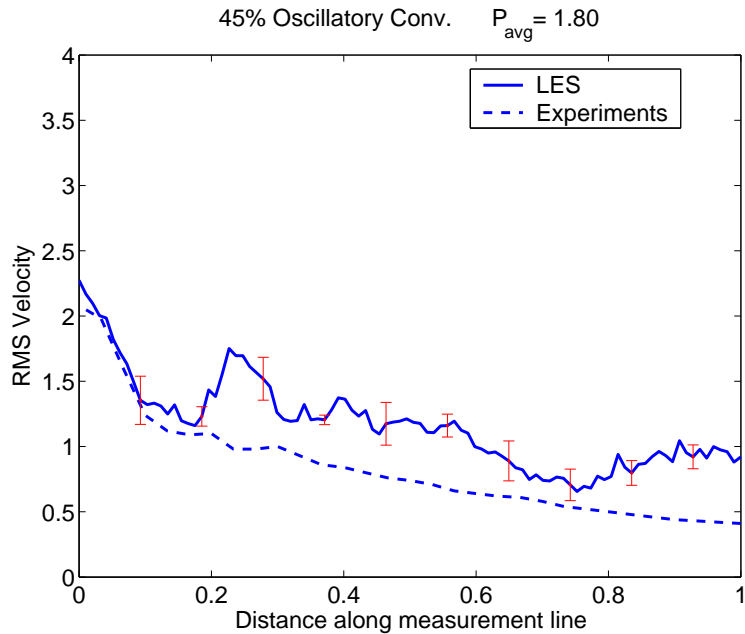


Figure 21: RMS of flow speed fluctuation along measurement line. LES data is plotted along with error bars representative of GCI

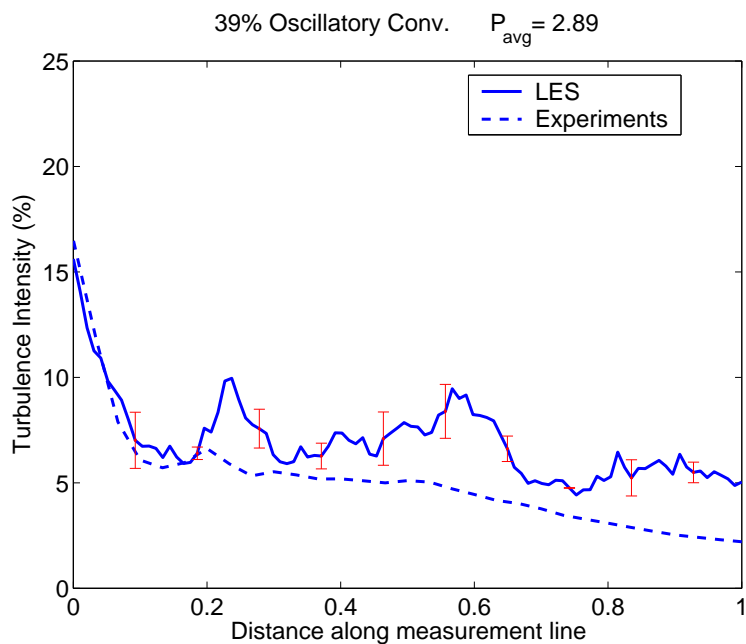


Figure 22: Turbulence intensity along measurement line. LES data is plotted along with error bars representative of GCI




Figure 23: Frequency spectrum of velocity fluctuations at different locations, from Grid 5




Figure 24: Frequency spectrum of velocity fluctuations at different measurement locations, reproduced from Gross (2003)

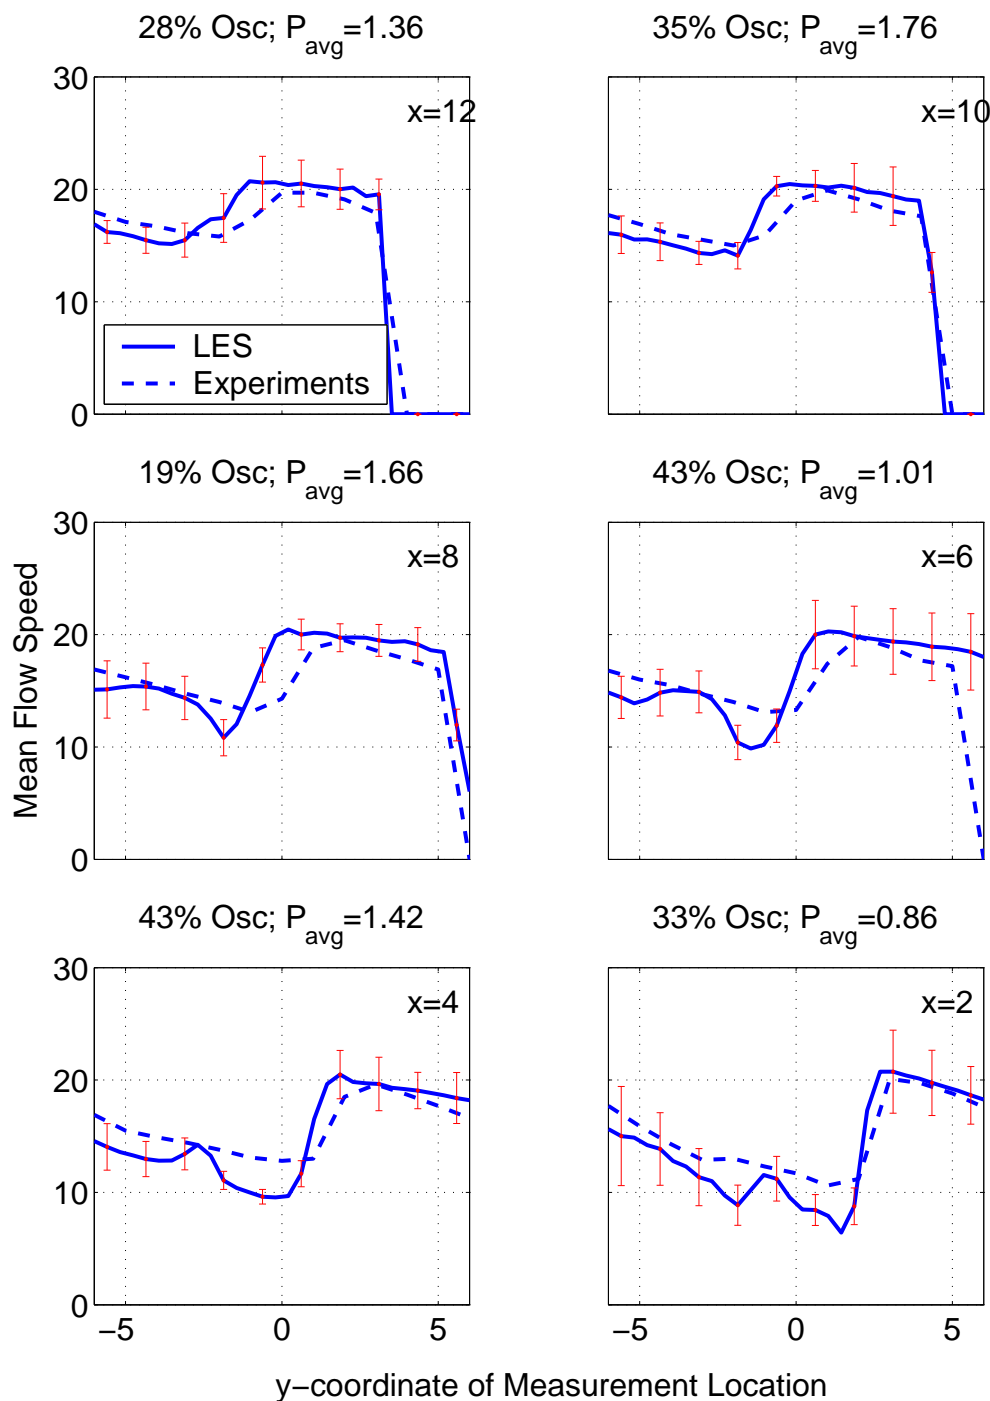


Figure 25: Mean flow speed in m/s, plotted for various x-positions. Both Experimental and LES data is shown, the latter with error-bars based on the GCI. At the top of each sub-figure, the percentage occurrence of oscillatory convergence and the average order of convergence is printed.

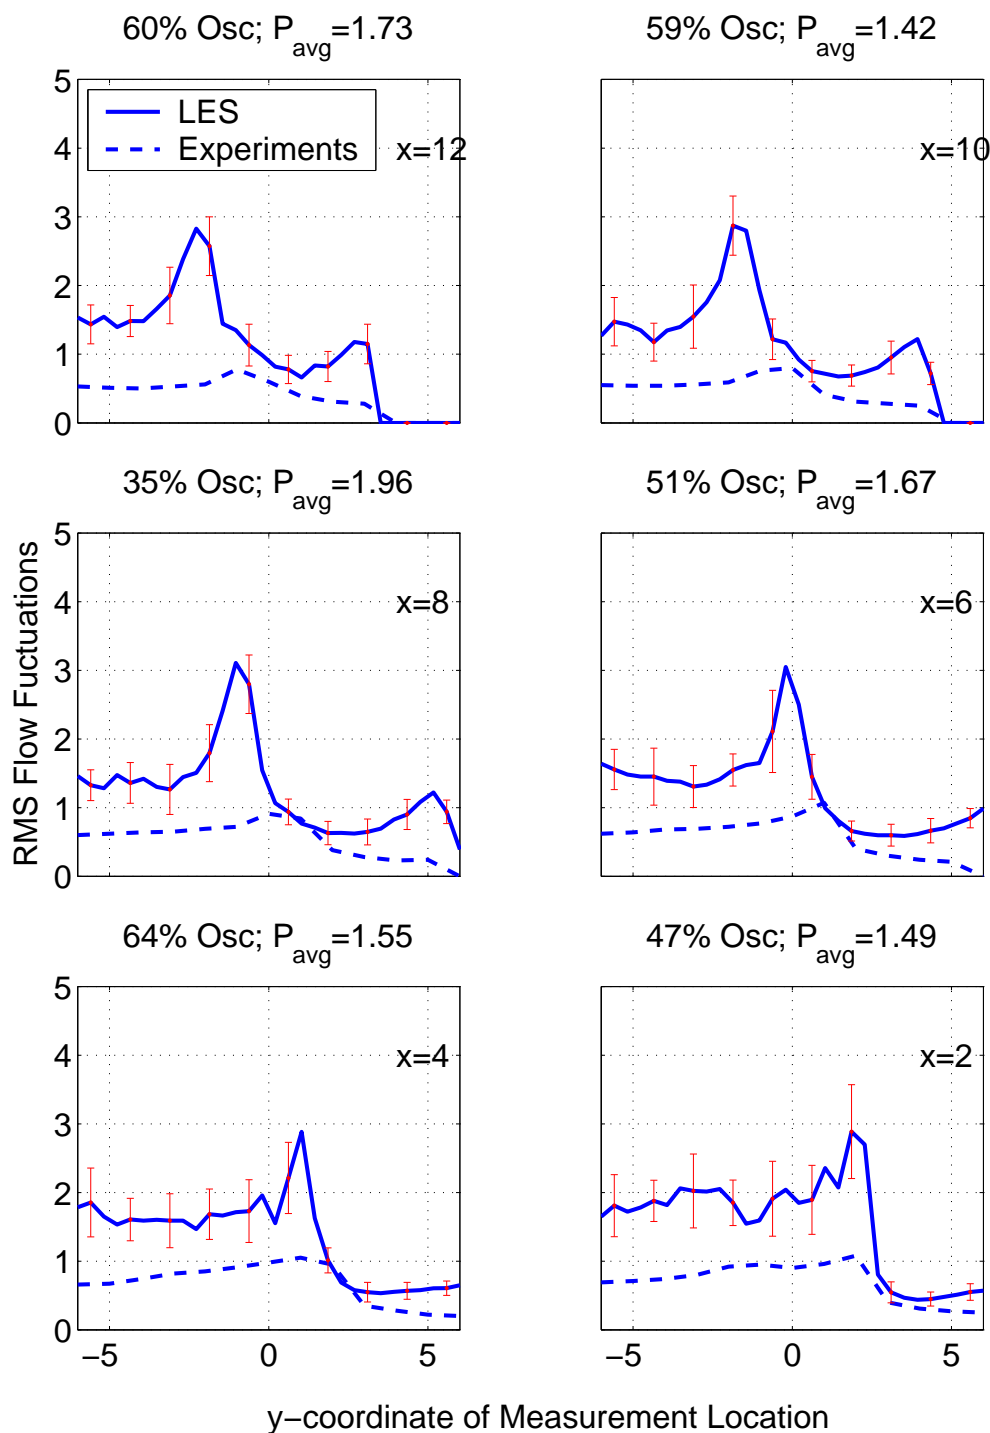


Figure 26: RMS flow fluctuations in m/s, plotted for various x-positions. Both Experimental and LES data is shown, the latter with error-bars based on the GCI. At the top of each sub-figure, the percentage occurrence of oscillatory convergence and the average order of convergence is printed.




Figure 27: Mean flow velocity over entire measurement area, from LES




Figure 28: Mean flow velocity over entire measurement area, from experiments of Gross (2003)




Figure 29: RMS flow fluctuations over entire measurement area, from LES




Figure 30: RMS flow fluctuations over entire measurement area, from experiments of Gross (2003)

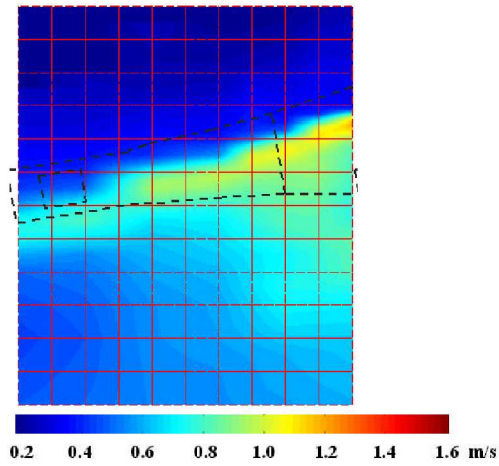


Figure 31: 0-2 kHz contribution to RMS flow fluctuations, from Gross (2003)




Figure 32: 0-2 kHz contribution to RMS flow fluctuations, from LES




Figure 33: 2-20 kHz contribution to RMS flow fluctuations, from Gross (2003)




Figure 34: 2-20 kHz contribution to RMS flow fluctuations, from LES




Figure 35: 2-6 kHz contribution to RMS flow fluctuations, from Gross (2003)




Figure 36: 2-6 kHz contribution to RMS flow fluctuations, from LES




Figure 37: 6-10 kHz contribution to RMS flow fluctuations, from Gross (2003)




Figure 38: 6-10 kHz contribution to RMS flow fluctuations, from LES




Figure 39: 10-20 kHz contribution to RMS flow fluctuations, from Gross (2003)




Figure 40: 10-20 kHz contribution to RMS flow fluctuations, from LES




Figure 41: Integral time scale of the flow (in number of disk rotations)

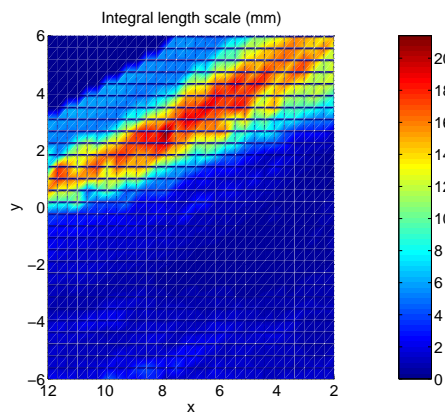


Figure 42: Integral length scale of the flow (mm)




Figure 43: Cross term of the time-averaged Reynolds stress tensor