


CHALMERS



Numerical simulation of flow around a pyramid using Large Eddy Simulation

Master's thesis in Applied Mechanics

ANDERS BENGTSSON

Department of Applied Mechanics
Division of Fluid Mechanics
CHALMERS UNIVERSITY OF TECHNOLOGY
Göteborg, Sweden 2014
Master's thesis 2014:01

MASTER'S THESIS IN APPLIED MECHANICS

Numerical simulation of flow around a pyramid using Large Eddy Simulation

ANDERS BENGTTSSON

Department of Applied Mechanics
Division of Fluid Mechanics
CHALMERS UNIVERSITY OF TECHNOLOGY
Göteborg, Sweden 2014

Numerical simulation of flow around a pyramid using Large Eddy Simulation
ANDERS BENGTSSON

© ANDERS BENGTSSON, 2014

Master's thesis 2014:01
ISSN 1652-8557
Department of Applied Mechanics
Division of Fluid Mechanics
Chalmers University of Technology
SE-412 96 Göteborg
Sweden
Telephone: +46 (0)31-772 1000

Cover:
Isosurface of U velocity, colored with pressure. Data is taken from the 10th POD mode.

Chalmers Reproservice
Göteborg, Sweden 2014

Numerical simulation of flow around a pyramid using Large Eddy Simulation
Master's thesis in Applied Mechanics
ANDERS BENGTSSON
Department of Applied Mechanics
Division of Fluid Mechanics
Chalmers University of Technology

ABSTRACT

In this Master Thesis, flow around a pyramid with 60 apex angle at a Reynolds number of $Re = 27000$ was studied numerically using Large eddy simulation. The pyramid was mounted on a plate, with the base of the pyramid parallel to the plate edges. The pyramid base width was 0.045 m, and the free-stream inlet velocity was $U_\infty = 10$ m/s.

Computations were performed using the software AVL Fire. The Smagorinsky model was used for the sub-grid-scale model. Simulations were compared with experiments using an identical setup which allowed for effective comparisons.

A proper orthogonal decomposition was performed on the computed flow to create a set of modes to describe the flow. The computed modes were correlated to surface pressure at different locations on the pyramid surface in order to find a set of sensor location that can best be used to express the entire flow. The procedure used can then be used to find optimal measurement positions in situations where only limited measurements can be taken, such as experimental setups.

It was concluded that Large Eddy Simulation can be effective in describing the flow around a pyramid, but that this requires a high spatial resolution. It was also shown that proper orthogonal decomposition can be used to identify large scale structures in the flow, as well as finding time-dependent behaviour that can be linked to surface pressures.

Keywords: LES, Pyramid

PREFACE

In this Master thesis LES has been used to perform simulations of turbulent flow around a pyramid. The simulational data has been used to perform a POD analysis of the flow. This data was then used to correlate the flow with surface pressure measurements in order to predict optimal pressure sensor locations. The work was carried out from april 2013 to december 2013 at the Department of Applied Mechanics, Division of Fluid Mechanics at Chalmers University of Technology, Sweden, with Professor Sinisa Krajnovic as supervisor and examiner.

ACKNOWLEDGEMENTS

It should go without saying that this thesis would not be possible if not for the guidance of professor Sinisa Krajnovic, not only during the thesis work but for years preceding. Without him I would know neither how hard or how useful LES-simulations can be.

The experiments by carried out by Zahra Hosseini and Robert Martinuzzi at the University of Calgary were also a must for the thesis. Without those we would have no idea of how our own studies fared.

In addition, Bernd Noack from the University of Potiers deserves thanks for introducing me to POD in a way I could understand.

At the department of applied mechanics at Chalmers, thanks should go to Professor Lars Davidsson and administrator Ulla Lindberg Thieme for making problems with paperwork go away, and to PhD-students Jan Östh and Eystein Helgasson for making problems with computers go away, by science or magic.

Thanks to Dr. Branislav Basara at AVL for providing software licenses and support. Computational resources at SNIC (Swedish National Infrastructure for Computing) at the Center for Scientific Computing at Chalmers (C3SE), and at the national supercomputing centre in Linköping (NSC) are also gratefully acknowledged.

Finally thanks go to the entire student body at Chalmers for making an education incredibly fun and making graduating equally bittersweet.

Göteborg, February 2014
Anders Bengtsson

CONTENTS

Abstract	i
Preface	iii
Acknowledgements	iii
Contents	v
1 Introduction	1
1.1 Limitations	1
2 Theory	1
2.1 Large eddy simulation equations	1
2.2 Numerical calculation	2
2.3 Proper orthogonal decomposition	2
2.4 Correlation analysis	3
2.4.1 Least square regression	4
2.4.2 Monte Carlo optimization	4
3 Method	5
3.1 Numerical simulations	5
3.1.1 Boundary conditions	5
3.1.2 Numerical meshes	6
3.2 Post processing	7
3.3 Sought results	7
4 Results	8
4.1 Velocity cross sections	8
4.2 Flow visualisation	9
4.3 Proper orthogonal decomposition	10
4.4 Monte Carlo Optimization	10
5 Discussion	10
5.1 Validation of the simulations	13
5.2 The flow field	17
5.3 Proper orthogonal decomposition	17
5.4 The optimal sensor positions	18
6 Conclusions	19
6.1 Further studies	20
References	20

1 Introduction

Pyramids are a simple shape present in countless situations on scales ranging from small machine parts to large architectural features. The flow around pyramids have however only been studied in a few experiments and simple simulational setups, and a description of the entire flow field does not exist where good agreement can be found with experiments.

Previous studies have shown some promise, but have failed to properly describe wake flow. As pyramid width decreases with the height, the vortices formed at the leading edge of the pyramid, the effect of the top of the pyramid, and the large vortices formed in the wake of the pyramid all meet behind the pyramid top[1]. This region is therefore needs a high resolution simulation to be able to properly describe the interactions[4].

The aim of this master thesis is to improve the understanding of these flows by producing a large eddy simulation solution which is compared with experiments. To further analyze the flow, a proper orthogonal decomposition of the solution is performed in order to identify prevailing flow patterns. This analysis allows for comparison with surface pressure measurements along the pyramid surface, which may be used for optimization of pressure sensor placement in further experimental studies.

1.1 Limitations

As Large eddy simulation techniques are transient and require a high spatial resolution, the simulations become very large. This puts some practical limitations on the project. This master thesis concerns only a single flow situation with a pyramid with an apex angle of 60 and the base edges parallel to the incoming flow. A Reynolds number of $Re = 27000$ was used. Two simulations were done using different computational grids in order to test for grid independence.

All physical parameters of the simulations were set as to closely mimic experimental studies being carried out concurrently by another team at the University of Calgary in Alberta, Canada. This provides a method to validate the simulations with high certainty.

2 Theory

Viscous flows follow the Navier-Stokes equations, which for incompressible flow of a Newtonian fluid with no external forces read

$$\rho \left(\frac{\partial u_i}{\partial t} + u_j \frac{\partial u_i}{\partial x_j} \right) = - \frac{\partial P}{\partial x_i} + \mu \frac{\partial^2 u_i}{\partial x_j \partial x_j} \quad (2.1)$$

along with a continuity equation.

$$\frac{\partial u_i}{\partial x_i} = 0$$

In general cases these require a numerical solution. The straightforward approach would be to discretize the equation using a finite volume method, and assuring numerical fidelity by using a fine enough discretization to resolve all relevant scales.

In turbulent flows, this quickly becomes unfeasible as the difference between the large and small scales increase exponentially with the Reynolds number. In order to make computations on scales comparable with experiments, some kind of modeling needs to be introduced to reduce the amount of calculations needed.

This study uses Large eddy simulation (LES), which seeks to ease the computational load by explicitly resolving the large scales, while introducing a model only for the small scales. Thus, large scale features should give good agreement with experiments as long as the the dissipation in the small scales is properly modeled[7].

2.1 Large eddy simulation equations

The process of setting up the LES-equations begins by separating the solution field of the sought variables into a filtered field and a small scale fluctuating field.

$$\varphi(\vec{x}, t) = \bar{\varphi}(\vec{x}, t) + \varphi'(\vec{x}, t)$$

Introducing these into the Navier stokes momentum equations (2.1), it can be shown that this reduces to

$$\rho \left(\frac{\partial \bar{u}_i}{\partial t} + \bar{u}_j \frac{\partial \bar{u}_i}{\partial x_j} \right) = -\frac{\partial \bar{P}}{\partial x_i} + \mu \frac{\partial^2 \bar{u}_i}{\partial x_j \partial x_j} + \frac{\partial \tau_{ij}}{\partial x_j}$$

with the residual stress tensor τ_{ij} representing the small scale contribution, which will have to be modelled [2].

Different methods exist for modeling the sub-grid contribution τ_{ij} . This study uses the Smagorinsky model, which assumes that the stress can be modeled as an additional viscosity term, such that

$$\tau_{ij} - \frac{1}{3} \tau_{kk} \delta_{ij} = -2\nu_{sgs} \bar{S}_{ij}$$

with the sub-grid-scale viscosity being further dependent on the grid and large scale rate of strain by

$$\nu_{sgs} = (C_S f \Delta_g)^2 \sqrt{2\bar{S}_{ij}\bar{S}_{ij}}$$

[7] The AVL Fire solver uses the value $C_S = 0.1$, and uses the cubic root of the cell volume for the filter width, $\Delta_g = \sqrt[3]{V}$. As the eddy viscosity should be zero near walls the damping function f is included as

$$f = 1 - \exp \frac{y^+}{25}$$

[2]

2.2 Numerical calculation

The filtered Navier-Stokes equations need to be solved numerically. For computational fluid dynamics software, finite volume methods are currently standard, and is used in this study. The details of the solution procedure goes beyond the scope of this report, however, a few details concerning simulational accuracy needs to be defined to explain choices used when specifying the computational mesh.

As seen in the above section, the equations to be solved are dependent on the cell size for the calculation of sub-grid-scale viscosities. As the sub-grid-scale contributions are modelled, the goal should be to create a numerical grid that resolves as much as possible of the turbulence in the explicitly calculated large scales.

As the residual stress tensor is dependent on the rate of strain tensor, we deduce that fine grid resolution is needed in regions where the rate of strain is expected to be large. These primarily include near wall regions, and in the wake.

A common measure of near wall grid resolution is to express the cell distance in dimensionless wall distance units, with the wall distance weighted by.

$$y^+ = \frac{\sqrt{\tau_w} / \rho}{\nu}$$

where τ_w is the wall shear stress. For LES calculations, proper near wall resolution is widely held to require a dimensionless wall distance below $y^+ < 1$.

To make sure that the calculation converges the Courant-Friedrichs-Levy condition holds that for all computational volumes

$$\text{CFL} = \frac{|\vec{u}| \Delta t}{\Delta x} < 1$$

where Δt is the timestep and Δx the cell side length. This condition is absolute only in fully explicit solvers, and for the software used convergent solution can be still be achieved if this is exceeded only locally [2].

2.3 Proper orthogonal decomposition

The proper orthogonal decomposition (POD) procedure is a method of data analysis which provides a set of orthogonal basis functions that allows reconstruction of the flow. Roughly analogous to performing an eigenvector analysis on a matrix, the procedure expresses the flow field as a combination of a set of coefficients and modes.

$$\vec{\phi}(\vec{x}, t) = \sum_{k=1}^{\infty} a_k(t) \vec{\psi}_k(\vec{x})$$

With $\vec{\phi}$ being any studied flow variable, a_k being the time-dependent modal coefficients, and $\vec{\psi}$ being time-independent modes, we have separated the time and space dependencies from each other. For a case with discrete measurements, such as our simulation, the number of modes becomes finite.

As both $a_k(t)$ and $\vec{v}(x)$ can be seen as vectors, the choice of whether to perform the analysis in the temporal or spatial dimension is dependent on the sought data. For transient flow data, a method known as snapshot POD, calculating $a_k(t)$ first is proper, and used in this study. Data then is captured as a series of snapshots of the studied field providing temporal discretization. Each snapshot is the flow field at a specified time.

Assuming a spatial discretization $\phi(\vec{x}, t) \approx \vec{u}_i(x, t)$ the continuous flow field can be expressed as a vector with elements dependent on the snapshot. With the same reasoning the modes can now be expressed as ordinary vectors \vec{v}_i and the modal coefficients being a discrete vector in time.

$$\vec{u}_i(\vec{x}, t) = \sum_{k=1}^{N_k} a_k(t) \vec{v}_{ik}(x) \quad (2.2)$$

By using the same number of modes N_k as the number of snapshots, the above equation is exact, and numerically the right hand side actually needs more data to describe the same flow, meaning the new form is slightly more inefficient.

The advantage of the POD procedure arises when the modes can be chosen in a way such that some modes contain more information than other. By using only a subset of modes chosen to express as much as possible of the flow a good approximation of the time-dependent structure can be achieved.

To find the modes $a_k(t)$, a correlation matrix is constructed by calculating the inner product of the vectors representing different snapshots t_i and t_j

$$R_{ij} = \frac{1}{N_k} (\vec{u}(\vec{x}, t_i), \vec{u}(\vec{x}, t_j)) \quad (2.3)$$

The choice of inner product space is here taken to be one corresponding to the turbulent kinetic energy of the vector, meaning that the diagonal terms in the correlation matrix is the turbulent kinetic energy for the snapshot.

This also means that element R_{ij} of correlation matrix can be interpreted as the amount of turbulent kinetic energy of time step t_i that can be expressed by time step t_j .

By definition, the correlation matrix becomes symmetrical and positive semi-definite. This means that the eigenvalues and vectors can be computed, and that all eigenvalues are non-negative. The eigenvectors become a time series expressing how the snapshots should be weighted.

With our choice of inner product space this also means that the eigenvalues λ_i correspond to the turbulent kinetic energy resolved by each eigenvector $\vec{\alpha}_i$, and that the first eigenvector is the time series that expresses the most of the turbulent kinetic energy. The second eigenvector is then the time series that describes as much as possible of the turbulent energy that is not expressed by the first eigenvector, and so forth.

From the eigenvectors, the spatial modes can then be computed from equation (2.2), which gives the expression

$$\vec{v}_k(x) = \frac{1}{\sqrt{N_k \lambda_k}} \sum_{m=1}^{N_k} a_k(t_m) \vec{u}(x, t_m)$$

As eigenvectors are orthogonal, these can be sorted in any order. By assigning the highest eigenvalue to index $k = 1$, and so forth, we get the modes sorted in order of the amount of turbulent kinetic energy they describe. We have thus arrived at a set of modes that can be used to describe the flow, as well as a classification method to compare these modes [3]. s

2.4 Correlation analysis

In this master thesis we seek correlations between surface pressure measurements and the computed POD modes. To do this, we may use a similar approach to the calculation of the correlation matrix in (2.3), but using multiple input signals. We now seek to express the POD modal coefficients in terms of a combination of surface pressure measurements.

2.4.1 Least square regression

Using a set of N_m pressure measurements \vec{p}_j , each a time series with the same temporal resolution as the snapshots, we seek a linear combination of these which best expresses the modal coefficients $a_k(t)$ in a least square sense. We thus seek an approximation

$$a_k(t) \approx \beta_{ij} \vec{p}_j(t)$$

where we may calculate the square residual

$$R^2 = (\beta_{ij} \vec{p}_j(t) - a_k(t)) \cdot (\beta_{ij} \vec{p}_j(t) - a_k(t)) \quad (2.4)$$

which becomes a scalar measure of how well the data match. The problem can then be specified as finding the weights β_{ij} which minimizes R^2 . This then becomes a simple least square minimization problem, as each set of mode coefficients $\vec{\alpha}_i$ can be considered independent, we get a set of matrix equations

$$\beta_{ij} = \left(\frac{1}{N_m} \sum_j \vec{p}_j \vec{p}_j^T \right)^{-1} \left(\frac{1}{N_m} \sum_j \vec{p}_j \vec{\alpha}_i \right)$$

which we may compute numerically with ease [5].

In order to increase the accuracy without having to sample additional data, we expand our set of pressure measurements with a time shifted copy of the pressure series. By shifting the variable in time the current modal value can be matched with a previous surface pressure value. The time shift is chosen as the first positive zero of the autocorrelation function of the modal time series, which gives

$$R(t_j) = \sum_n \alpha_i(t_n) \alpha_i(t_n - t_j) \quad (2.5)$$

This can be deduced to be symmetrical, and the index of the first zero with $t_j > 0$ gives us our time shift.

By sampling the pressure signals slightly earlier than the start of the measurement window used for the calculation of the POD-modes, the time shifted signals can be seen as a second set of pressure measurements, effectively doubling the size of \vec{p}_j without needing more data from the simulation.

2.4.2 Monte Carlo optimization

We now seek the set of pressure measurements p_j which minimizes R^2 in (2.4). While using simulational data, we may easily increase the number of pressure measurements included in the computation. For experimental setups, there are often several restrictions on what measurements can be taken, and a small number of measurements N_{meas} will have to express the flow. A suitable set of measurement positions needs to be found, and we seek to use the large amount of data available in a LES simulation to predict these positions. While the residuals can be easily calculated for a given set of pressure measurements, the problem of finding the set of measurements that minimizes R^2 is significantly harder. As there is no guaranteed connection between the pressure measurements and the modes, the only way to find a solution is to compute the residual R^2 for every unique combination of pressure sensors. This calculation quickly becomes very large, as the number of possible sets N_{set} of N_{meas} chosen measurements from a pool of N_{tot} possible measurements is given by

$$N_{set} = \frac{N_{tot}!}{N_{meas}!(N_{tot} - N_{meas})!}$$

In this study, with a set of 6 measurements taken from a pool of 1521 positions this equals around $N_{set} \approx 1.7 \cdot 10^{17}$ combinations [6]. This can be considered numerically unfeasible, and a method to find a reasonable guess is needed.

We therefore use a random subset of combinations, to find sets of sensors which provide a statistically significant lowering of the residual. While this does not guarantee the optimal solution, it gives a set of data that is relevant enough for further analysis. We then seek to find trends in the sensor placements which yields low residuals, and from these seek to provide a best guess for finding the lowest residual.

The idea is to perform the least square regression procedure outlined in section 2.4.1 on a randomly chosen set of sensors. The calculated residual R^2 is stored along with the chosen sensors, and the process is repeated

a large amount of times. With a large enough amount of calculated residuals, we get an approximation of how the residuals are distributed. By choosing the the sensors that give the lowest residual we get a sensor setting which we cannot tell if it is optimal, but we can with a large enough sample of different sensor location combinations say that the combination is within the top percentile with high certainty. The exact degree of certainty depends on the number of samples used and how these are distributed.

From the sensor sets in the random sample we can also study whether there are any trends in the sensor placement that can be used to produce a good guess for a low residual. The best found set of sensors can then be used as a basis for a local search to find even better matches with the experimental data, but since those measurements cannot be assumed to be randomly distributed, it is important not to use those measurements when comparing for statistically significant residuals. We can however deduce that these sensor positions express the flow modes better than a set of sensors we know in the top percentile.

3 Method

The process of performing the simulation begins with specifying the geometry and the simulational domain. The domain is then discretized to create a computational mesh. The mesh is then used by a software package, that produces a solution to the flow field using specified models and boundary conditions. The solution is then transferred to another software package which computes the proper orthogonal decomposition.

3.1 Numerical simulations

The geometry and simulational domain was specified to be identical to the experimental setup used by the collaborating team. The experiments were carried out in an open test section of a wind tunnel, with the pyramid being affixed to a plate measuring 1100 mm in the flow (x-) direction and 760 mm in the spanwise (y-) direction. For the simulational domain, the plate size was used as the basis for the length. To avoid spending computational resources on areas not affected by the pyramid. The domain extended only 500 mm in the y-direction, and extended in the z- direction to a height of 405 mm.

The pyramid chosen had a base length of 45 mm and a height of 39 mm. This gives a pyramid apex angle of 60. The pyramid was placed with the apex 200 mm from the leading edge of the simulational domain, and centered in the spanwise direction. The pyramid was angled so that all the base sides were parallel to the edges of the base plate. The origin of the simulational domain was then chosen as the point on the plate directly under the apex of the pyramid.

3.1.1 Boundary conditions

The boundary conditions were specified in the AVL Fire software package. For the ground plate and pyramid surface, a no-slip wall condition was used. To recreate the open test section, slip-wall conditions were used on the sides and top of the numerical domain. Due to the small size of the projected pyramid surface are in relation to the simulation domain cross-section (less than 1%), any potential errors introduced by this boundary condition is considered negligible.

For the outlet of the domain, a zero gradient condition was specified, with the outlet assumed to be far enough downstream not to influence the result. For the inlet, a wall-normal velocity of $U_0 = 10$ m/s was specified expect for near the ground plate, where a specified velocity profile was used in order to achieve the same boundary layer thickness observed in the experimental setup. This was applied at below $z = 0.01$ using an expression gained by adjusting experimental data to an exponential expression.

$$U(z) = \begin{cases} 10, & z \geq 0.01 \\ 10 \left(\frac{z}{0.01}\right)^{\frac{2}{9}}, & z < 0.01 \end{cases}$$

The initial conditions used a uniform initialization with a U-velocity of 10 m/s. The pressure was set at 92000 Pa to match the atmospheric pressure used in the experiments.

3.1.2 Numerical meshes

The geometrical domain was specified in the ICEM software which was also used to discretize the domain into the computational mesh. Two meshes were made, with the major difference being the cell sizes in the wake region. To accomplish smooth transitions between cell sizes (a necessity for a properly converging simulation) in the denser mesh, a fundamentally different meshing strategy was needed. The two meshes are therefore considered different enough to show mesh-independence of the solution, given that the results of the simulations are in agreement.

The coordinate system was set with the origin at the ground plate under the pyramid apex. The x-direction was the streamwise direction, the y-direction was chosen as the spanwise direction, and the z-direction was the height wise direction. To assure that the near wall behavior of the flow was properly resolved, both meshes were designed to produce a dimensionless wall distance of $y^+ < 1$ near all walls (the y^+ designation is slightly confusing as the wall normal direction is more closely aligned with the z-direction). Furthermore the cell sizes in the spanwise and flow wise directions were kept at under $x^+ < 30$ in order to properly resolve the surface. As a consequence of the mesh structure, the cell sizes near the top of the pyramid were even smaller.




Figure 3.1: *x-y plane of the fine mesh at 0.8 times the pyramid height (upper) and x-z plane of the fine mesh through the center of the pyramid (lower)*

The smaller mesh used a total of 8.0 cells, and the finer mesh used 14.1 cells. As the near wall cells are by far the smallest, and were roughly the same in both simulations, the constraints set by the CFL-number were almost identical, the same time step of $\Delta t = 10^{-5}$ s could be used while calculating on both meshes.

Simulations were carried out using computer resources from Chalmers Centre for Computer Science and Engineering (C3SE) and from the National Supercomputer Centre in Linköping. The simulations on the coarse mesh used 64 cores for around 670 hours, and the simulations on the fine mesh needed 128 cores used for around 500 hours.

3.2 Post processing

The post processing consisted first of calculating the POD-modes, then implementing the correlation analysis, and finally to visualize the solution. The POD calculation was done using the xAMC package, which is a set of command line programs designed specifically for POD calculation on flows.

For the correlation analysis, a Matlab code was written to perform the calculation and perform the Monte-Carlo search. The primary focus of this was to get a working code that could natively use the data from xAMC and the simulations. As such the program is likely highly ineffective, and if other cases are to be studied using the same approach, a dedicated application should be written in a compilable language such as Fortran or C.

Plots of data were mainly carried out in Matlab, and the Enight software was used to visualize 2-d planes and 3-d structures, as well as for computation of vortex cores and flow properties.

3.3 Sought results

In order to perform the proper orthogonal decomposition a series of snapshots of the flow need to be established. While it would be possible to save the entire domain, due to the large amount of data generated, 1000 snapshots were taken of a subdomain around the pyramid extending 0.054 m on both sides of the pyramid, extending 0.0675 m in front of the pyramid, and 0.2475 m behind the pyramid apex. The domain extended 0.0675 from the group, as shown in figure 3.2 expressed in terms of the pyramid base width $d = 0.045\text{ m}$.

The snapshots saved the three dimensional velocity field along with the pressure in $120 \times 76 \times 31$ discrete positions in the domain. Denser spacing was used in the center of the domain and the lower part of the domain, as well as some specific planes included to ease comparison with measurements taken in the experimental study, as shown in figure 3.3. Snapshots were saved with a frequency of 1000 Hz , over one simulated second. The first snapshot was saved 0.5 simulated seconds from the beginning of the simulation, when the flow was observed to be independent of the initial conditions.




Figure 3.2: Size of the domain used for the POD snapshots

In addition to the snapshots, mean values of the entire simulation domain taken over the snapshot period was saved.

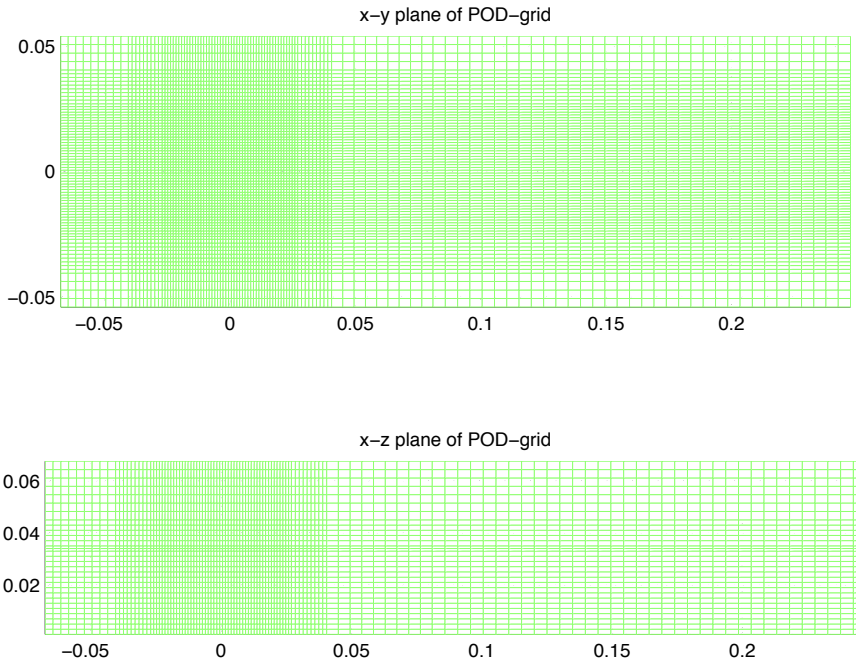


Figure 3.3: *Discretization used for the POD snapshots*

The surface pressure was measured in a set of positions on the pyramid and on the plate behind the pyramid. For the pyramid 1521 positions equally space over the pyramid in a 39x39 square pattern was used, and for the wake nine lines were chosen, with 39 equally spaced sensors. Measurements for these sensors were saved with a frequency of 10000 Hz.

For the Monte-carlo search of optimal sensor positions, a search using 100000 tested combinations were performed using all possible sensor positions on the pyramid. As one purpose of these searches was to suggest proper sensor positions for experiments two other searches were done using subsets of the possible sensor positions. Both of these excluded positions near the base of the pyramid where pressure sensors would be hard to fit, and one of these also excluded positions near the corner edges of the pyramid for the same reasons. These two searches were also done using 100000 tested combinations.

4 Results

As one of the advantages of CFD methods is the availability of large amounts of data, this report can only cover a subset of all data. We first seek to demonstrate that the LES simulations are in agreement with the experiments, for which we compare velocities at specific cross-sections of the domain.

After agreement between the experiments and simulations have been discussed, we present the POD-modes along with the mean flow to present an analysis of the flow features.

Finally we show the results of the correlation analysis, and present the sets of sensor positions found by the optimization algorithm.

4.1 Velocity cross sections

From the POD-snapshot domains we can extract data along any line. We first want to compare the two simulations to see whether the mesh resolution has any effect on the simulation. This is done by comparing the

U-velocities at two x-normal planes at different distances from the pyramid. These are shown in figure 4.1. For the two planes behind the pyramid there are significant disagreements between the two simulations, with the coarser mesh giving higher velocities for cross-section higher up in the wake.

To compare with the experiments, we chose to study mean velocities along two planes at different elevation. Along the planes velocities are plotted at five different distances from the pyramid. These are shown in figures 4.2 and 4.3. The curves from the finer mesh show much better agreement with the experiments, though the smoothness might suggest that a larger sample is needed for the average.

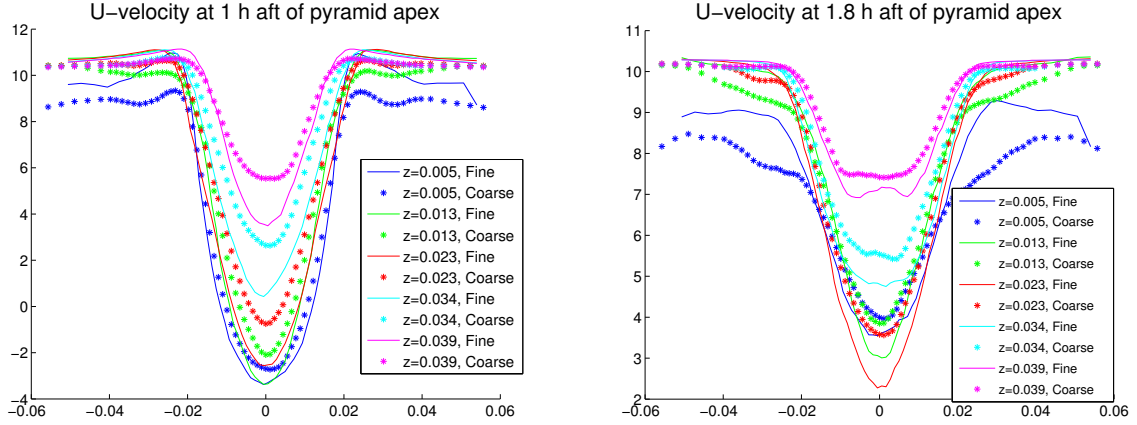


Figure 4.1: U -velocity component in the two simulations, measured at h and $1.8 \cdot h$ behind the pyramid apex.




Figure 4.2: U and V velocity components behind the pyramid at $z = 0.012$ m

Along with these plots, comparisons were done on profiles between these, as well as the W -velocity component and turbulent stresses. For space reasons, these are excluded, as they provide little more detail relevant for the further discussion of the data.

4.2 Flow visualisation

Due to the large amount of available data, any visualization of transient flow needs to include a lot of simplification. As the POD analysis is performed to provide an expression of transient flow expressed in terms of combined time constant modes, visualization of transient flow is not done here.

Instead we provide a set of mean velocity profiles at z -planes spread along the pyramid height. The U and V components are shown in figure 4.4. As these planes are taken from the snapshots, the pyramid appear

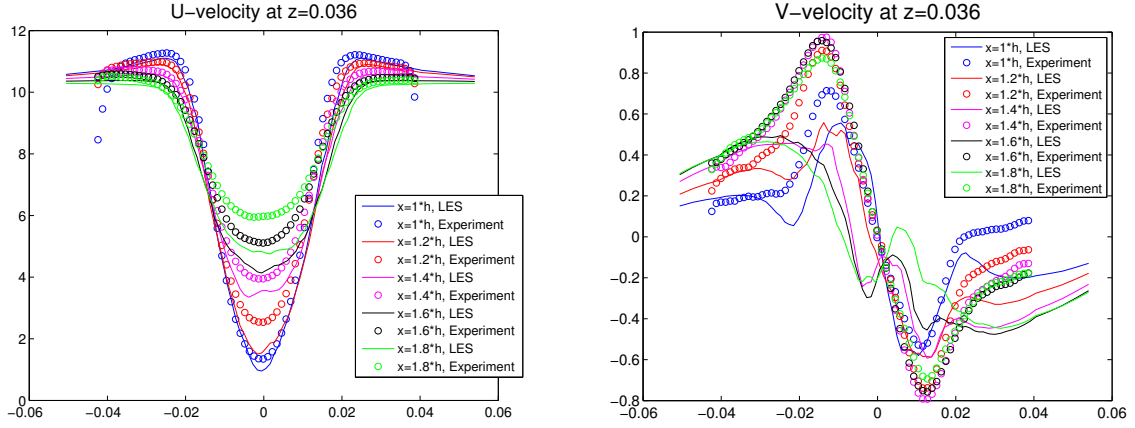


Figure 4.3: U and V velocity components behind the pyramid at $z = 0.036$ m

as an area of zero velocity. The near wall behavior should be taken with a grain of salt, as the interpolation includes these zero velocity measurements. The average surface pressure and mean square pressure fluctuation taken from the surface pressure measurements on the pyramid is shown in figure 4.5. Finally, to illustrate the prevailing vortices, a 3-d picture of identified vortex cores using an eigenvalue scheme on the mean flow is shown. This should be seen primarily as an illustration of the locations of the main structures.

4.3 Proper orthogonal decomposition

The proper orthogonal decomposition performed by the xAMC program yielded at set of 1000 modes. The relative energies of which are shown in figure 4.7. The amplitude is weighted so that the energy of the mean flow corresponds to an energy of 1. As can be seen from the figure, one requires 44 out of 1000 modes to express half of the POD-resolved-energies, and using half of the modes expresses 92 % of the energy.

The POD-modes are shown by tomographies of the pressure of the the mean flow and first five modes at half the pyramid height, and the full pyramid elevation. These are identically scaled to allow for comparison of mode amplitudes. Pressure is measured from -125 Pa to 125 Pa, relative to the atmospheric pressure. These are given in figure 4.8

4.4 Monte Carlo Optimization

Three searches were performed, one including all sensors on the pyramid, one excluding sensors closer to the pyramid base edge than 5 mm in either x - or y - direction, and one additionally excluding sensors near the diagonal edges. The calculated sum of square residuals were distributed according to the histograms in figure 4.9.

The top four sensor position sets were given by these searches as shown in figure 4.10 to 4.12 along with their calculated sum of square residual.

5 Discussion

In this section, as in the preceding ones, we will first discuss whether the simulational data agrees well enough with the experiments to make any worthwhile conclusions. We will then comment on the solution of the flow field as well as the computed POD modes, and finally we will discuss the implications of the correlation analysis with the pressure sensors.




Figure 4.4: Mean U (left) and V (right) velocity fields at planes located (from up to down) at elevation of $z = 0.009\text{ m}$, $z = 0.015\text{ m}$, $z = 0.021\text{ m}$, $z = 0.027\text{ m}$, $z = 0.033\text{ m}$ & $z = 0.039\text{ m}$

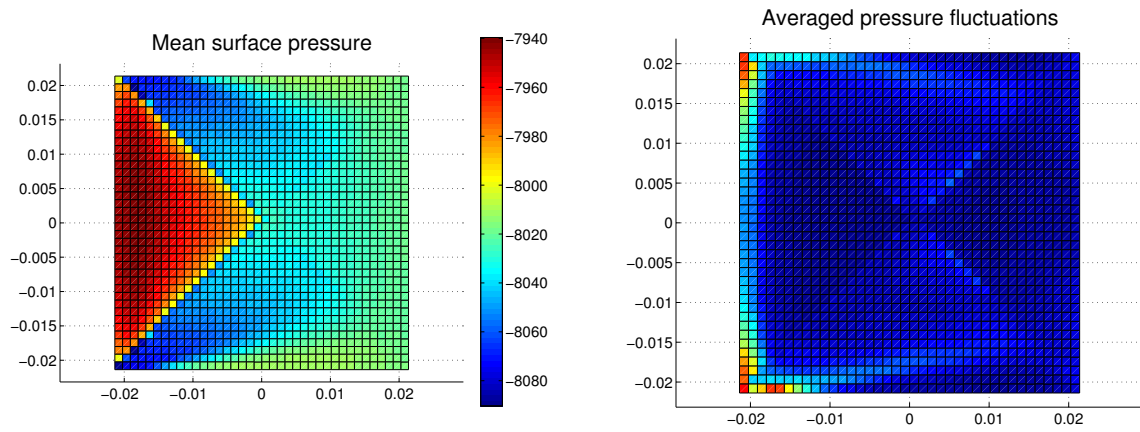


Figure 4.5: *Mean surface pressure and mean squared pressure fluctuation.*




Figure 4.6: *Vortex core visualization*

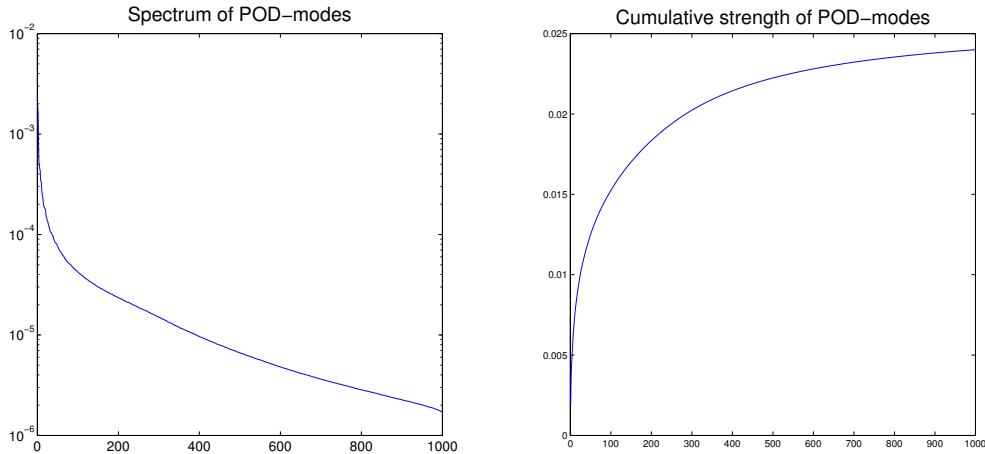


Figure 4.7: *Spectrum (left) of POD-mode intensity, and cumulative intensity of POD-modes (right)*

5.1 Validation of the simulations

The difference between the simulations show that the solution made with the coarser mesh is not resolved enough to be considered independent of the mesh. While the general picture is consistent, the larger velocities in the wake point to insufficient resolution of the wake. This is consistent with the refined mesh giving much better agreement with the experimental setup.

The agreement between the two simulations is large enough to disregard the data from the coarse mesh in the further discussion of the results. The coarse mesh used is simply not dense enough for LES to be used.

Studying the comparison of the fine grid simulation and the experiment give two important clues to the usability of the simulation. When studying the U-profile, we see that apart from the points closest to the pyramid, the experiments give higher velocities than the corresponding point in the simulation. We also see that the velocities also increase with the distance from the pyramid. Therefore, if experiments at one point are compared with a point slightly further away in the simulation, almost perfect agreement is found. This would indicate that the main flow picture in the simulation is sound, but that something shifts the information in the simulation further down in the wake.

Comparing the V-velocity profile, we also see that in the simulation, positions closest to the pyramid show alternating directions of the V-velocity near the center, typically indicative of the presence of a recirculation vortex. Such a zone is known to exist in both experiments and simulations, and is seen in previous studies, but stops being relevant further down the wake. For this comparison, this effect has completely ceased when in the experiments, but is still present in the experiments. Once again, if experiments at one point are compared with a point slightly further away in the simulation, much better agreement is found.

Some of the discrepancies between the simulation and the experiments is also likely to be caused by the measurement technique used in the experiments. As the measurement window in the experimental data begins roughly $0.9 \cdot h$ aft of the pyramid apex, the cross section studied near the pyramid may not contain all information about the wake. The lack of the recirculation zone in the experiment may only point to a small difference in the shape of those vortices, but as their center is located outside the measurement window, the shape can not be studied further.

There is therefore a possible scenario where the only main difference between the simulations and experiments is the extent of the two wake vortices in the stream wise direction. A difference here then influences the entire wake, with a larger vortex in the simulations shifting the velocity profiles slightly rearwards to give the simulations generally larger deviations from bulk flow, as seen in the comparisons. Near the tip where velocity tapers off the most, this difference would be the largest. Any discrepancies between simulations and experiments are then related to a lack of resolution upstream of the place where the differences are seen.

As the shift needed to create a match in the data is not constant, the shift is highly likely caused by a lack of resolution in the wake, which makes the difference between experiments and simulations too large to neglect. While the following analysis can be used to describe the simulated solution, any conclusions drawn can not be




Figure 4.8: Comparison of mean flow (top) and POD-modes 1-5 (from up to down), at elevation of $z = 0.5h$ (left) and $z = h$ (right)

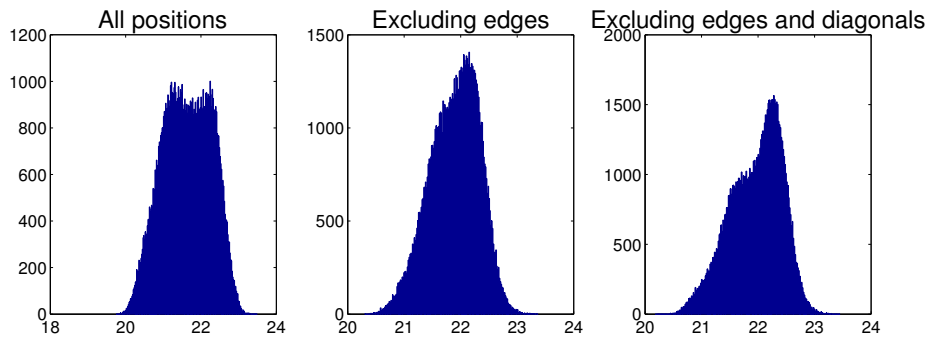


Figure 4.9: *Distribution of the residuals from the three Monte Carlo Optimizations*

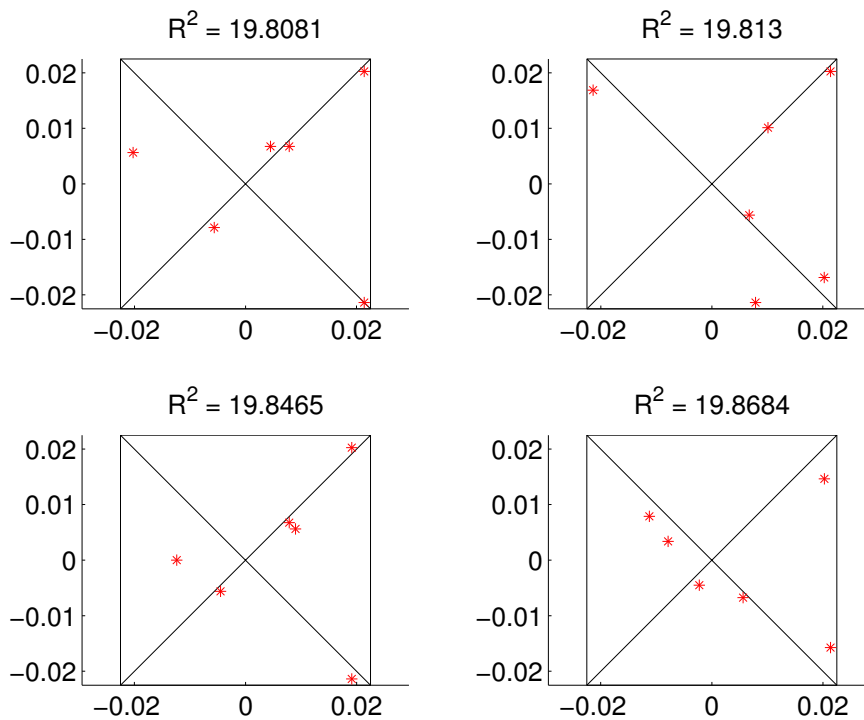


Figure 4.10: *The four sets of sensors with the lowest residuals found while searching all possible combinations*

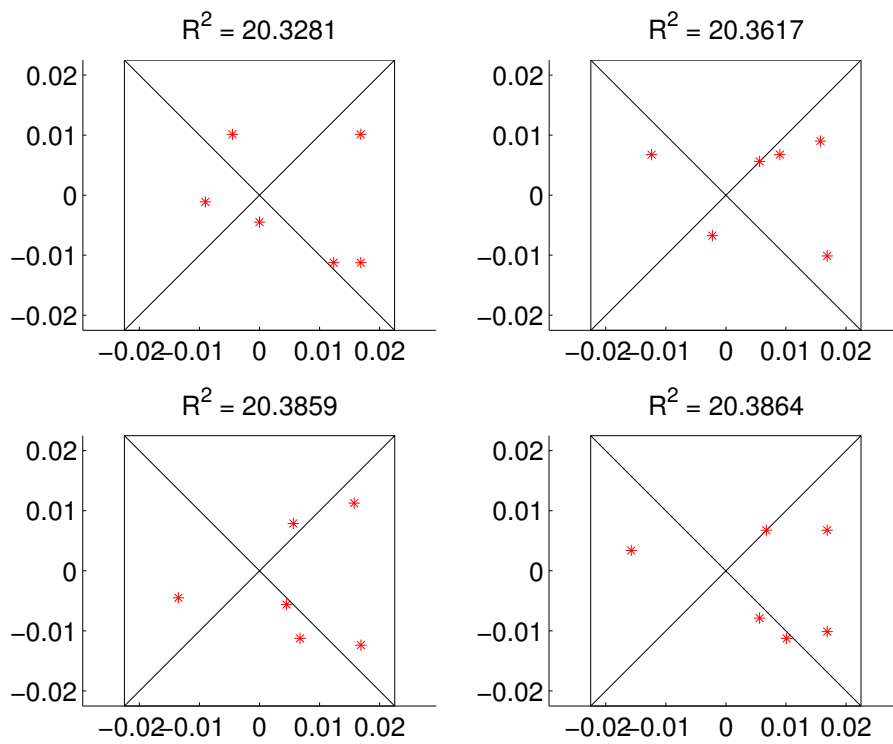


Figure 4.11: *The four sets of sensors with the lowest residuals found while searching in a area excluding sensors near the ground plate*

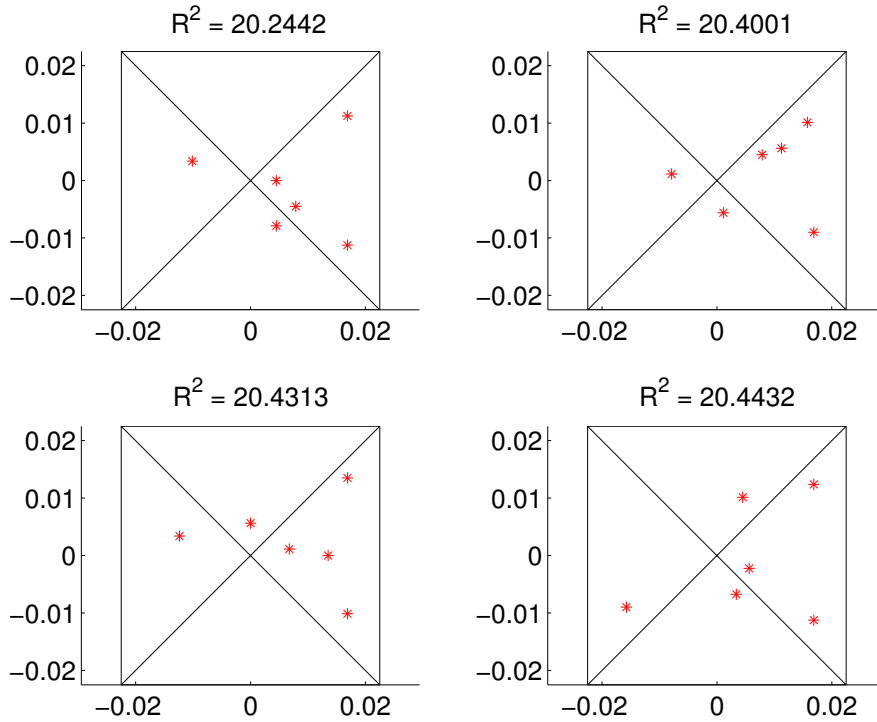


Figure 4.12: *The four sets of sensors with the lowest residuals found while searching only sensors feasible for experiments*

applied to the experiments and real world flows.

5.2 The flow field

The features expected from previous studies of the flow around a pyramid can be seen in the simulations. The formation of a large horseshoe vortex near the base along with two vortices emanating from the front base corners of the pyramid are all expected and seen in the mean flow. Together with two counterrotating vortices meeting behind the tip of the pyramid, these make up the most prevalent turbulent features [1].

More detailed analysis shows that the large horseshoe vortex is accompanied by a smaller one forming near the base of the pyramid rotating in the opposite direction. The vortices emanating near the base of the pyramid are in fact trend differently and seem to alternately form a low vortex near the ground plate and a weak one trending along the leading diagonal edge of the pyramid. This alternates with a stronger vortex forming along the diagonal of the pyramid side, which split up close to the trailing side of the pyramid and either continues to join the rear vortices (as seen in figure 4.6), or trails of upwards to join the wake further up. The choice of directions is highly likely related to the timing of the main wake vortices.

Overall the flow field agrees well in all main features with previous studies suggesting that the flow is well described in behavior, though specific points may still differ.

5.3 Proper orthogonal decomposition

The decomposition of the flow shows a ordinary spectrum of modal strength. One might look at the total mode intensity, totaling 0.023 compared to the mean flow, set at 1, to deduce that the turbulence is very low intensive, but as flow far from the pyramid and ahead of the pyramid is very well described by the mean flow, this is not an indication of turbulent intensity, but rather indicates that the same results could be achieved with a smaller POD domain.

The spektrum show that the modal strength decreases very fast, as the first few modes expresses the large scales, with modes expressing successively smaller scales containing less and less energy. This is further confirmed by looking at the shown pressure modes in figure 4.8.

The first three modes in the $z = 0.5h$ plane resolve the large scale near pyramid wake. The first mode give the two vortices directly behind the pyramid. That the two pressure peaks has alternating signs indicate that the vortices are shed alternately. The second mode basically shifts the first one further down the flow to allow resolution of the shedding of the vortices, this mode also includes the formation of a new vortex, which is in phase with the shedding. The third mode balances the two first modes, effectively resolving effects of any deviation from a constant frequency shedding.

Worth noting is that the first three modes has very little impact on the pressure in the plane located at $z = h$. Instead the fourth and fifth mode give high amplitudes at this plane, higher than the lower plane. This indicates that the lower flow fluctuates at different frequencies than the higher flow. The prescence of multiple superposed flow patterns would also explain why the mode amplitudes has an obvious oscillation, but show no clear frequencies. These frequency spectra were calculated, but as the results gave no obvious conclusion, they were not included for discussion in this report. Doing a POD analysis on a subset of the data, i.e. doing separate calculations on different planes, could strengthen or disprove this hypothesis.

5.4 The optimal sensor positions

Studying the histograms from the three searches show that all of these plots show a bimodal behavior, which is most pronounced in the case with all allowed sensors, where the distance between the two peaks is significantly larger than in the other two cases. The two searches excluding sensors give roughly the same two peaks, which are closer together than for the first search. The bimodal behavior here is less pronounced, as the two peaks roughly merge in to one, but the model of two superimposed normal distributions does still explain the shape of the distribution.

The bimodal behavior implies that there are at least two statistically significant ways that the set of sensors chosen affects the residual. Had the distribution been normal, there would be no way to prove that the residuals gained were not due to random variations. With the histograms given in figure 4.9 there is a significant effect of the sensor placement on the residual, and we may proceed with our analysis.

The figures 4.10 to 4.12 show that there is a large variation of sensors giving similar residuals, showing that no specific measurement position gives significantly better results. There are however some discernible patterns, which becomes even more apparent when studying more sets of positions giving residuals on the same level.

All sets of sensors giving favorable residuals use two sensors located low on the the leeward side of the pyramid. Most favorable residuals are found when using sensors on the very corners of the pyramid. As the residuals when allowing sensors very close to the ground are discernably lower than the cases excluding such positions, we see that experimental measurements giving good agreement would be hard to perform, as these positions make physical sensor placement hard.

Apart from the pair of sensors, the optimizations give preference to locations near the diagonal edges of the pyramid, which also make experimental sensor placement hard. One concludes that placing sensors near where vortices form and detach from the pyramid give best agreement with the POD modes. Comparing this with the pressure fluctuations in figure 4.5 one sees that these sensors are located in a region of high pressure fluctuations.

Two additional considerations when studying the different allowed sensor positions is that due to the three searches using the same amount of random sets, but having differing amounts of possible combinations, the results should statistically be less converged for the larger amount of allowed sensors. However, for this sample size, this effect is negligible compared to the advantage of having access to more relevant sensors.

Also, one might note that a lot of sensor positions are located close to each other, which may suggest that all six sensor positions might not be needed to give a good estimation.

When comparing the two searches where only a subset of sensor positions are allowed, little difference is found, and in fact the best set of sensors found for the third search has a lower residual than the ones found in the second search. As the third search uses a subset of the sensors included in the second search, that set could have been found in both searches. It is seen that even when excluding sensors on the diagonals, low residuals are gained when using positions near the diagonals.

The overlap also serves to show that the Monte-carlo method gives no absolute results. The difference between the cases is mostly seen in the histograms, which shows that it is easier to find a half-good set of sensor positions when diagonal placement is allowed, but that the results for the very best sets are equal in both cases.

When positions near the ground plate are excluded there are still two sensors located low on the leeward side, separated and roughly symmetrical. That this pair is found in every good combination of sensors is likely related to the two large wake vortices, which are resolved by these sensors. Even when placement on the diagonal is allowed, the optimum position of these is found with the sensors closer to the middle. This, together with the wider spacing when low sensors are allowed, indicates that the best position is near the root of the wake vortices, and that higher up, it is better to place the center closer to the core of the vortex.

On the leeward side, a third sensor is always present which is placed higher up on the pyramid. When diagonal placement is allowed, this is placed on one or both leeward diagonal edges. These are found in proximity to where the vortices growing along the side of the pyramids detach from the surface. When diagonal placement is not allowed, there is instead a single central sensor located high on pyramid, making a triangle of relevant sensors on the leeward side. While an additional position is also found somewhere in the formed triangle, this is most often close to another sensor, and is therefore assumed to be less necessary.

Another sensor is usually placed on the side of the pyramid, but in most cases only a sensor on one of the sides is found. The behavior on the sides of the pyramid is assumed to be symmetric, and therefore information on one side and a time shift can describe flow on both sides. When diagonal placement is possible, diagonals can also be deduced to provide information about both sides.

Finally a position somewhere on the windward side of the pyramid is often included, and found close to the center and slightly below the sensors on the side near the top. This gives information on the large scale vortex shedding, and when sensors close to the ground plate are allowed, also about the near pyramid horseshoe vortex.

Using these conclusions a set of sensors is constructed and tested for a low residual, and produces a lower residual than all others found when sensors near the ground is allowed. This is shown in figure 5.1 and show all patterns explained above. The procedure is repeated allowing all sensor positions and also gives a significantly lower residual than the randomized cases. That these optimizations can be done serves to point out that while




Figure 5.1: *Optimized sensor positions. On the left using full limitations, and on the right allowing all sensor positions.*

the Monte-carlo approach is usable for optimization, but that local searches around the found optimum can increase the quality of the results.

6 Conclusions

The results arrived at in this project give further insight into flow around pyramids. While the results primarily enforces the previously known results about the flow around the pyramid, additional information was given that helps explain why both this and previous studies have had problems resolving the wake with LES.

The main cause of the discrepancies between the simulations and the experiments were the location and shape of the two main rear vortices, which then shifts the entire wake rearwards to cause comparison errors for every compared location. As this study and previous studies has sought to remove the discrepancies by increasing mesh resolution in the far wake, more effort should be directed at the region close to the pyramid, and at the area where the main vortices form. As this area includes interactions between both these large vortices, the smaller vortices along the side, and the effects near the tip, insufficient resolution in this area is a real possibility even in a mesh that is well resolved further along the wake.

The POD analysis also show that this is a powerful technique to gain insight into complex flows such as flow around a pyramid. The correlation with the surface pressure does show that there is a link between the pyramid surface pressure and the wake flow, but this also comes to the somewhat obvious conclusion that in order to describe the flow, it is best to place sensors close to known flow features.

6.1 Further studies

As the two performed LES-simulations are not in agreement, it would be beneficial to show mesh independence of the solution by further refining the mesh. If the near pyramid wake area is refined as discussed above, and a simulation still gives results similar to the present LES, it would cast doubt on these conclusions, and also cast doubt on the fidelity of the experiments. Using an approach using more large scale modelling could give interesting results if agreeing results can be found without needing a refined mesh, and should be done as well.

The important result of the wake having different frequencies at different elevations should be further studied by doing a more detailed analysis of the data. The interaction should be tested to see if a direct relation between the local pyramid width can be established. This does not require additional simulations, but rather doing a set of simpler POD-analyses for each plane, as well as studying the raw pressure measurements for frequency maximas. However, a significant amount of time will have to be spent on analyzing the data.

While the correlation and sensor placement optimization could be further studied, it is not clear if any worthwhile conclusions can be achieved by more data. Varying the amount of sensors to place, as well as studying the effects of changing the time shift on the distribution of the residuals will give a better understanding of how good the conclusions drawn are, but no immediate goals can be set to directly draw additional conclusions. The method does however show some promise and deserves further study when applied to a simulation better matching experimental flows.

References

- [1] M. M. AbuOmar and R. J. Martinuzzi. Vortical structures around a surface-mounted pyramid in a thin boundary layer. *Journal of Wind Engineering and Industrial Aerodynamics* 96 (2008), 769–778.
- [2] *AVL-Fire CFD Solver Users Guide*. AVL. 2013.
- [3] L. Cordier and M. Bergmann. “Proper Orthogonal Decomposition: an overview”. Lecture series 2003-04 on post-processing of experimental and numerical data, von Karman Institute for Fluid Dynamics, 2003.
- [4] S. Krajnovic and R. Lárusson. “Large eddy simulations of the flow around pyramids”. *Proceedings CFD society of canada conference, May 2012, Canmore, Canada*. 2012.
- [5] D. C. Montgomery. *Design and Analysis of Experiments*. 7 th. John Wiley & Sons, 2009.
- [6] L. Råde and B. Westergren. *Mathematics Handbook for Science and Engineering*. 5th. Studentlitteratur, 2004.
- [7] H. Versteeg and W. Malalasekera. *An introduction to computational fluid dynamics*. 2nd. Pearson Ed., 2007.