

Technical University of Munich
Department of Mechanical Engineering
Chair of Aerodynamics and Fluid Mechanics

Streaming Modal Decomposition Approaches for Vehicle Aerodynamics

Marco Kiewat



Technische Universität München
Fakultät für Maschinenwesen
Lehrstuhl für Aerodynamik und Strömungsmechanik

Streaming Modal Decomposition Approaches for Vehicle Aerodynamics

Marco Kiewat

Vollständiger Abdruck der von der Fakultät für Maschinenwesen der Technischen Universität München zur Erlangung des akademischen Grades eines Doktor-Ingenieurs (Dr.-Ing.) genehmigten Dissertation.

Vorsitzender: Prof. Dr. Klaus Bengler

Prüfende der Dissertation:

1. Priv.-Doz. Dr.-Ing. habil. Thomas Indinger
2. Prof. Dr. Makoto Tsubokura

Die Dissertation wurde am 09.04.2019 bei der Technischen Universität München eingereicht und durch die Fakultät für Maschinenwesen am 04.09.2019 angenommen.

Preface

I have worked on this dissertation during the past few years of my time as a research associate and Ph.D. student at the Chair of Aerodynamics and Fluid Mechanics of the Technical University of Munich.

First, I want to express my appreciation to my supervisor PD Dr.-Ing. habil. Thomas Indinger, for his continuous encouragement and support throughout the entire project. Under his excellent supervision and insightful guidance, I was able to dive deep into a broad research spectrum and produce meaningful research results. I am especially thankful for multiple opportunities to discuss my research work with renowned scientists from all over the world during the conferences I attended in Japan and the USA.

My gratitude also goes to my valued colleagues from the industry partner Audi AG, Dr. Moni Islam and Dr. Vincent Zander. Thank you for encouraging me to choose modal analysis as the main topic for my Ph.D. and enlightening me with your industry perspective on vehicle aerodynamics engineering.

Special thanks go to my colleagues, Daiki Matsumoto, Lukas Haag, Christopher Collin, Aleksandra Pachalieva and Lu Miao from the research group for vehicle aerodynamics at the Chair of Aerodynamics and Fluid Mechanics of the Technical University of Munich. Your positive energy, team spirit and healthy competitiveness have made long office hours feel a lot shorter and more fun. I enjoyed working with all of you.

I want to thank my parents, Peter and Heidi, for their never-ending motivational aid in everything I do.

Last, I am truly thankful to my wife, Tatyana. She is the most important person in my life, and without her support, this work would not have been possible.

Contents

Preface.....	i
Contents.....	ii
Nomenclature	iv
Dimensionless Numbers.....	iv
Latin Letters.....	iv
Greek Letters, Symbols, Operators	v
Abbreviations	vi
List of Figures	viii
List of Tables	xiv
Abstract	xv
1 Introduction.....	1
1.1 Motivation	1
1.2 Numerical Simulations.....	4
1.2.1 Governing Equations of Navier-Stokes	5
1.2.2 Turbulence Modeling	6
1.2.3 Discretization and Solving	10
1.3 Modal Analysis State of the Art.....	11
1.4 Objective.....	17
1.5 Outline	18

2	Modal Decomposition Methods for Fluid Mechanics	20
2.1	Idea behind Data-Driven Modal Reduction.....	20
2.2	Generic Test Case: Two-Dimensional Cylinder.....	22
2.3	Low-Memory Proper Orthogonal Decomposition	25
2.3.1	Batch-Processed POD	25
2.3.2	Singular Value Truncation.....	27
2.3.3	Incremental Data Processing.....	28
2.3.4	POD Results	31
2.4	Low-Memory Dynamic Mode Decomposition.....	39
2.4.1	Batch-Processed DMD	40
2.4.2	Streaming Total Dynamic Mode Decomposition	41
2.4.3	Frequencies, Amplification and Damping Rates.....	44
2.4.4	Reconstruction and Amplitudes.....	44
2.4.5	Algorithmic Perspective on STDMD	46
2.4.6	Modified Eigenvalue-Weighted Amplitude Scaling.....	51
2.4.7	DMD Results.....	52
2.5	DMD vs. POD vs. FFT	62
3	Application on a Full Vehicle Aerodynamics Simulation.....	65
3.1	DrivAer Reference Configuration	65
3.2	Simulation Setup and Results.....	68
3.3	Guidelines for Modal Analysis for Vehicle Aerodynamics.....	77
3.4	Modal Analysis Results.....	84
3.4.1	Incremental POD Modes.....	84
3.4.2	Streaming Total DMD Modes.....	89
4	Conclusion and Outlook	97
4.1	Conclusion.....	97
4.2	Outlook	98
	Appendix.....	100
A.	Reconstructions of Incremental POD Modes.....	101
B.	Reconstructions of Streaming Total DMD Modes.....	103
	References	109

Nomenclature

Non-bold symbols are scalars, lower case bold symbols are vectors, and upper case bold symbols are matrices.

Dimensionless Numbers

Co	Courant number, Courant-Friedrichs-Lewy condition
Kn	Knudsen number
Ma	Mach number
Re	Reynolds number
St	Strouhal number

Latin Letters

a	Speed of sound
\mathbf{A}	DMD operator
C	Correlation function
C_D	Drag coefficient
C_L	Lift coefficient
C_P	Pressure coefficient
D	Diameter
\mathbf{D}	Rate-of-strain tensor
\mathbf{D}_φ	Diagonal matrix of vector elements from φ
\tilde{d}	DES length scale
d_\perp	Wall-normal distance
e	Energy
E	Turbulence kinetic energy per unit mass
f	Frequency
f_d	RANS model region protection function for DDES
F	Force
H	Frequency response function
\mathbf{I}	Identity matrix
k	Turbulence kinetic energy
K	Eigenvalue weighted mode ordering criterion
K_{stable}	Modified eigenvalue weighted mode ordering criterion
L	Length scale
\mathbf{L}	Velocity gradient tensor
l_0	Initial length scale of turbulence production
m	Number of snapshots
n	Number of observables
n_0	Cell count of a mesh
p	Static pressure or POD expansion criterion as indicated
Q	Q-criterion
\mathbf{Q}	Orthogonal matrix from QR-decomposition

r	Rank
\mathbf{r}	Residual vector
\mathbf{R}	Upper-right triangular matrix from QR-decomposition
$\tilde{\mathbf{S}}$	Turbulence production source term in DES
t	Time
T	Total time span of the investigated data set $T=t_m-t_1$
\mathbf{T}	Vandermonde matrix
C_{DES}	DES constant
$\mathbf{u}, \mathbf{v}, \mathbf{w}$	Snapshot column vectors of the velocity field components
\mathbf{U}	Left-singular orthonormal vector matrix
u_i, U_x, U_y, U_z	Velocity component
u^+	Dimensionless velocity
U_∞	Free stream velocity
u_{\parallel}	Wall parallel velocity component
u_τ	Wall friction velocity
\mathbf{V}	Right-singular vector matrix
w	Weighting function
x, y, z	Spatial coordinates
\mathbf{X}	Snapshot matrix
x_i	Spatial coordinate
\mathbf{x}_j	Snapshot vector of time step j
\mathbf{y}	Eigenvector
\mathbf{Y}	Eigenvector matrix
y^+	Dimensionless wall distance
y_\perp	Wall-normal coordinate
\mathbf{z}	Augmented snapshot vector
\mathbf{Z}	Augmented snapshot matrix

Greek Letters, Symbols, Operators

$ \varphi $	Magnitude of a variable
$\tilde{\Phi}$	Reconstructed or truncated
$\hat{\varphi}$	Non-dimensionalized or scaled variable as indicated
Φ^T	Transpose of a vector or a matrix
Φ^\dagger	Pseudo-inverse of a matrix
$\ \Phi\ _2$	Euclidean / spectral norm of a matrix
$\ \Phi\ _F$	Frobenius norm of a matrix
α	Vector of complex DMD amplitudes
δ_{ij}	Kronecker delta, $\delta_{ij}=0$ if $i \neq j$ and $\delta_{ij}=1$ if $i=j$
ε	Threshold value
Δ	Discrete difference
ρ	Density

μ	Dynamic viscosity
ν	Kinematic viscosity
ν_T	Turbulent (kinematic) viscosity
$\tilde{\nu}_T$	Modified turbulent (kinematic) viscosity
κ	Wave number or von Karman constant as indicated
λ	Complex eigenvalue
Λ	Eigenvalue matrix
$\boldsymbol{\varphi}$	Complex DMD mode vector
η	Kolmogorov length scale
τ_{ij}	Shear-stress component
τ_w	Wall shear stress
ω	Angular frequency
$\boldsymbol{\Omega}$	Rate of rotation tensor
σ	Singular value
$\boldsymbol{\Sigma}$	Singular value matrix

Abbreviations

AAWK	Audi Aeroacoustic Wind Tunnel
CFD	Computational Fluid Dynamics
CFL	Courant Friedrich Lewy criterion
CGSI	Classical Gram-Schmidt Iteration
CPU	Central Processing Unit
DDES	Delayed Detached-Eddy Simulation
DES	Detached-Eddy Simulation
DFT	Discrete Fourier Transform
DMD	Dynamic Mode Decomposition
DNS	Direct Numerical Simulation
EOF	Empirical Orthogonal Function
exp	Experiment
FFT	Fast Fourier Transform
FKFS	Research Institute of Automotive Engineering and Vehicle Engines Stuttgart
FVM	Finite Volume Method
FOM	Full Order Model
HDD	Hard Drive Disk
HPC	High Performance Computing
ID	Index
incPOD	Incremental Proper Orthogonal Decomposition
LBM	Lattice-Boltzmann Method
LES	Large-Eddy Simulation
mav	Moving Average
max	Maximum

NSE	Navier-Stokes Equations
p	Projected
PCA	Principal Component Analysis
PDE	Partial Differential Equation
POD	Proper Orthogonal Decomposition
PSD	Power Spectral Density
QR	QR-decomposition
RAM	Random Access Memory
RANS	Reynolds Averaged Navier Stokes
rel	Relative
ROM	Reduced Order Model
SA	Spalart-Allmaras
SGS	Subgrid-scale
sim	Simulation
SPH	Smoothed-Particle Hydrodynamics
STDMD	Streaming Total Dynamic Mode Decomposition
stdPOD	Standard Proper Orthogonal Decomposition
SVD	Singular Value Decomposition
temp	Temporary
TI	Turbulence Intensity
trunc	Truncated
TUM	Technical University of Munich
URANS	Unsteady Reynolds-Averaged-Navier-Stokes
WTA	Wind Tunnel A of TUM

List of Figures

Figure 1: Energy cascade in turbulent flow, depicting the turbulent kinetic energy E vs. wave number κ	3
Figure 2: Two-dimensional cylinder CFD mesh of hexahedral cells. The outer mesh has a cell size of $\Delta x=D/16$, and the 15 cell layer O-grid around the cylinder expands with a ratio of 1.6 in the radial direction. No free-hanging nodes are present.	23
Figure 3: Snapshot of the velocity field components and the velocity magnitude at t_{1000} . Von Karman vortex shedding is in a purely periodic state.....	23
Figure 4: Temporal evolution of the y-component of the velocity field probed at position $x=2D$ and $y=0$. The ramp-up phase toward reaching the periodic vortex-shedding state takes around half of the total simulation time.....	24
Figure 5: Power spectral density of the velocity field y-component at position $x=2D$ and $y=0$ vs. Strouhal number. The frequency spectrum is computed using the Hanning window function with a minimum of 50% overlap in the period of t_{1000} to t_{2000}	25
Figure 6: Flowchart of the incremental POD algorithm. Matrix indexing is done using MATLAB® notation. Green arrows denote a positive evaluation of an if statement, and red arrows indicate negative.....	30
Figure 7: Performance loss charts for part one (top) and part two (bottom) of the two-dimensional cylinder test case for varying thresholds in the incremental POD algorithm.....	32
Figure 8: ROM execution time for incremental POD analysis. The computation time increases with the rank of the resulting ROM.	33
Figure 9: Incremental POD performance loss behavior for different maximum-rank truncation parameters for part one. Maximum-rank truncation is conducted using singular value truncation after the full update procedure in the algorithm.....	34
Figure 10: Incremental POD execution time for part one, using the maximum-rank singular value truncation after the update procedure. All execution times are in similar orders of magnitude.	35
Figure 11: Singular values of incrementally computed POD decompositions (left) and ROM performance loss curves (right) for part one and part two of the two-dimensional cylinder test case. Mode pairs for the first few shedding modes can be identified in the singular values.	36
Figure 12: Temporal evolution of the velocity y-component probed at position $x=2D$ and $y=0$ in part one of the simulation. Snapshots $t_{100}=10.91$ and $t_{800}=87.28$ of the original data (left) and reconstructed data (right) are marked with a blue line in the upper graph. POD modes one, two and three are used for the reconstruction (mean flow mode and first mode pair).	37

Figure 13: Temporal evolution of the velocity y-component probed at position $x=2D$ and $y=0$ in part two of the simulation. Snapshot $t_{700}=185.5$ of the original data (left) and reconstructed data (right) is marked with a blue line in the upper graph. POD modes one, two and three are used for the reconstruction (mean flow mode and first mode pair).....	38
Figure 14: Mean FFT spectra of all cells of the original versus the reconstructed velocity field of part two of the simulation. Reconstruction is done using incrementally computed POD modes two and three (i.e., the first oscillating mode pair).....	38
Figure 15: STDMD incremental orthonormal basis update. Green arrows denote a positive evaluation of an if statement, and red arrows indicate negative.	48
Figure 16: STDMD mode evaluation. Flow chart following the last time increment $j=m$ from incremental orthonormal basis computation depicted in Figure 15. Green arrows denote a positive evaluation of an if statement, and red arrows indicate negative.	50
Figure 17: Weighting function of the modified eigenvalue-weighted amplitude ordering criterion K_{stable} for a data series with $m=1,000$. The sine has a period of $m-1=999$ and an amplitude of 1.	52
Figure 18: Left (top to bottom): (1) DMD frequency spectrum with three modes. Mean flow mode one at $f_1=0\text{Hz}$ and $-\alpha_1=-20$, mode two at $f_2=0.11\text{Hz}$ and $-\alpha_2=-2.5$, and mode three at $f_3=0.10\text{Hz}$ and $-\alpha_3=-2.1$. Eigenvalues for all three modes are stable (i.e., $-\lambda=1$). Complex conjugate counterparts of the fluctuating modes (negative frequency) are not shown. (2) Reconstruction of the mean flow mode. (3) Reconstruction of the second mode, including its complex conjugate counterpart. (4) Reconstruction of the third mode, including its complex conjugate counterpart. Right (top to bottom): (1) Temporal evolution of the input data. (2) Reconstruction of an ROM with $r_{\text{ROM}}=1$ (i.e., only the mean flow mode one). (3) Reconstruction of an ROM with $r_{\text{ROM}}=2$ (i.e., modes one and two). (4) Reconstruction of an ROM with $r_{\text{ROM}}=3$ (i.e., all modes combined).	55
Figure 19: Performance loss charts for part one (top) and part two (bottom) of the two-dimensional cylinder test case for varying thresholds in the STDMD algorithm. The last point signifies r_{max} (i.e., the rank of orthonormal basis \mathbf{Q}_z).	56
Figure 20: ROM execution time for STDMD analysis. The computation time increases with the rank of the resulting ROM and the increase in the QR-decomposition threshold value ϵ_{QR}	57
Figure 21: Left: STDMD performance loss behavior for different maximum rank truncation parameters in part one of the two-dimensional cylinder test case. Maximum rank truncation is conducted using orthonormal basis compression in the update procedure in the algorithm. Right: STDMD execution time for the same STDMD runs as depicted left.	58

Figure 22: Frequency spectra of the STDMD analysis (left) and ROM performance loss curves (right) for part one and part two of the two-dimensional cylinder test case. Negative frequency counterparts in the frequency spectra are not shown.....	59
Figure 23: Temporal evolution of the velocity y-component probed at position $x=2D$ and $y=0$ in part one of the simulation. Snapshot $t_{100}=10.91$ and $t_{800}=87.28$ of the original data (left) and reconstructed data (right) is marked with a blue line in the upper graph. DMD modes one and two and the complex conjugate mode of mode two are used for the reconstruction (mean flow mode and first oscillating mode including its negative frequency counterpart).....	60
Figure 24: Temporal evolution of the velocity y-component probed at position $x=2D$ and $y=0$ in part two of the simulation. Snapshot $t_{700}=185.5$ of the original data (left) and reconstructed data (right) is marked with a blue line in the upper graph. DMD modes one and two and the complex conjugate mode of mode two are used for the reconstruction (mean flow mode and first oscillating mode including its negative frequency counterpart).	61
Figure 25: Temporal evolution of the velocity y-component probed at position $x=2D$ and $y=0$ in for STDMD execution on a fraction of part two of the simulation. Snapshot $t_5=185.5$ of the original data (top) and reconstructed data (bottom) is marked with a blue line in the left graph. DMD modes one and two and the complex conjugate mode of mode two are used for the reconstruction (mean flow mode and first oscillating mode including its negative frequency counterpart).	62
Figure 26: Dimensions of the DrivAer reference body in 1:2.5 scale. The lower part of the top sting is included for a better approximation of the wind tunnel measurement results.....	67
Figure 27: DrivAer simulation setup including all no-slip wall geometries, street boundaries, and second- and third-finest grid level mesh regions. Dark patches on the bottom (street) are modeled using translating velocity boundary condition walls. Bright gray patches are solid, no-slip, steady walls.	69
Figure 28: Surface-integrated force coefficient convergence over time for simulation and experiment. The graphs show the convergence to the temporal mean ($\Delta C=C-C_{\text{mean}}$) of the respective force coefficient. Dotted horizontal lines represent a deviation from the temporally averaged value by one count....	70
Figure 29: Rolling resistance evaluation of the DrivAer in the wheels-on configuration in the WTA of TUM.	72
Figure 30: PSD force spectra from FFT of the drag force coefficient during a measurement of the wheels-on configuration with and without wind. The rolling road system is running in both cases. Strouhal numbers are computed using a free stream velocity of 42.64m/s for both cases.	73

Figure 31: PSD frequency spectra of the DDES simulation. The integrated drag, front lift, and rear lift force coefficient spectra are generated using FFT with minimum 50% overlap Hanning windowing for the time frame from 2.2s to 3s simulated time. The lift and drag forces do not contain contributions from the wheels' lift forces.....	74
Figure 32: Slices through the temporally averaged mean flow field around the DrivAer body. Averaging period is one second physical time. Top: Slice through $y=0\text{m}$. Middle: slice through $z=0.15\text{m}$. Bottom: Slice through $z=-0.1\text{m}$. Left: Magnitude of the velocity field. Middle: Pressure coefficient. Right: Total pressure coefficient.....	76
Figure 33: Iso-surfaces of the Q-criterion with $Q=2\cdot 10^6$ colored by the instantaneous velocity magnitude depicted on the original CFD mesh.	76
Figure 34: Averaged FFT spectrum of the velocity components of the flow field from an interpolated sampling box in the vicinity of the DrivAer model. FFT is executed using the same parameters as for the integrated force coefficient spectra in Figure 31.	77
Figure 35: coarseMesh interpolation region definition. The data is mapped to this second region during the simulation before it is passed to the online STDMD algorithm. The lower bound of the interpolation mesh is 0.4mm above the rolling road belt patch.	78
Figure 36: Slices through the $z=0\text{m}$ plane of the DrivAer simulation. Top: Hexahedral equidistant interpolation mesh coarseMesh versus CFD mesh. Bottom: Velocity magnitude field. Visualized are the cell center values (i.e., not interpolated). The cell size of the coarseMesh interpolation grid is around $\Delta x=10^{-2}\cdot L_{\text{wheelBase}}$	80
Figure 37: Top: Temporal correlation function for the first 145 consecutive time steps recorded in a subdomain and two interpolated meshes. Bottom: Temporal correlation function for the first 2,500 consecutive time steps recorded on the two interpolated meshes.	81
Figure 38: Frequency response of a moving average temporal filter. Window sizes are in units of the number of simulated time steps.	82
Figure 39: Convergence of the residual norm divided by the number of snapshots used for the DrivAer DDES simulation.	84
Figure 40: Singular values and ROM performance loss convergence over the number of modes used for reconstruction for the DrivAer DDES simulation.	85
Figure 41: Left: Temporal evolution of the leading right-singular values separately for each mode. Right: Respective PSD distribution plotted over the frequency range between 0Hz and 80Hz. PSDs are computed using three Hanning windows with size 512 and averaging the resulting distributions.	86
Figure 42: First oscillating POD mode pair. Top: slice through the $z=-0.1\text{m}$ plane. Bottom: slice through the $y=0$ plane. The reconstructed first time step of a combination of POD mode two and mode three. Left: x-component. Middle: y-component. Right: z-component.....	87

Figure 43: First oscillating POD mode pair. Slice through the $z=-0.1\text{m}$ plane. Seventeen reconstructed time steps are shown from a combination of POD mode two and mode three.....	88
Figure 44: Frequency spectrum of the STDMD analysis of the DrivAer DDES simulation. Modified eigenvalue-weighted amplitude as a function of the Strouhal number. Higher frequencies are not shown.	90
Figure 45: Eigenvalue distribution of the STDMD modes. The angle of the vector in the complex number plane represents the non-dimensionalized frequency. The vector length represents the damping value for each mode. The marker size is linearly proportional to the mode ordering value K_{stable} . Left: full spectrum. Middle and Right: enlarged ranges of positive frequency modes, including the unit circle drawn as a solid line.	91
Figure 46: Reconstruction of a full period of DMD mode number two y -component field. Two probe locations in the wake of the wheel/wheel housing area show the strong phase connection between the left and right front wheels. The reconstructed time steps are located exactly at physical times $2.2s$, $2.2s+1/(6f_2)$, $2.2s+2/(6f_2)$, $2.2s+3/(6f_2)$, $2.2s+4/(6f_2)$ and $2.2s+5/(6f_2)$, which do not coincide with simulated time steps.	92
Figure 47: STDMD frequency spectra results using different memory-saving measures. Maximum rank settings of $r_{\text{max}}=600$, $r_{\text{max}}=400$ and $r_{\text{max}}=200$ are compared to the full rank decomposition (top). STDMD spectra for data interpolated to an even coarser grid (“coarsestMesh”) with a cell size of $\Delta x=2\cdot 10^{-2}\cdot L_{\text{wheelBase}}$ is shown for comparison (bottom).	96
Figure 48: First oscillating POD mode pair. Iso-surfaces of the reconstructed first time step of a combination of POD mode two and mode three. Top: x -component iso-surface $u_x=\pm 1.1\text{m/s}$. Middle: y -component iso-surface $u_y=\pm 1.2\text{m/s}$. Bottom: z -component iso-surface $u_z=\pm 0.8\text{m/s}$	101
Figure 49: Second oscillating POD mode pair. Iso-surfaces of the reconstructed first time step of a combination of POD mode four and mode five. Top: x -component iso-surface $u_x=\pm 1.4\text{m/s}$. Middle: y -component iso-surface $u_y=\pm 1.2\text{m/s}$. Bottom: z -component iso-surface $u_z=\pm 1.2\text{m/s}$	102
Figure 50: Mode two from the STDMD spectrum at $St_2=4.58$ and $f_2=185\text{Hz}$. Iso-surfaces of the reconstructed first time step. Top: x -component iso-surface $u_x=\pm 12.9\text{m/s}$. Middle: y -component iso-surface $u_y=\pm 13.8\text{m/s}$. Bottom: z -component iso-surface $u_z=\pm 12.6\text{m/s}$	103
Figure 51: Mode eight from the STDMD spectrum at $St_8=7.4$ and $f_8=296\text{Hz}$. Iso-surfaces of the reconstructed first time step. Top: x -component iso-surface $u_x=\pm 10.7\text{m/s}$. Middle: y -component iso-surface $u_y=\pm 11.5\text{m/s}$. Bottom: z -component iso-surface $u_z=\pm 11.2\text{m/s}$	104
Figure 52: Mode three from the STDMD spectrum at $St_3=0.72$ and $f_3=29\text{Hz}$. Iso-surfaces of the reconstructed first time step. Top: x -component iso-surface $u_x=\pm 7.4\text{m/s}$. Middle: y -component iso-surface $u_y=\pm 4.8\text{m/s}$. Bottom: z -component iso-surface $u_z=\pm 5.7\text{m/s}$	105

Figure 53: Mode seven from the STDMD spectrum at $St_7=0.99$ and $f_7=40\text{Hz}$. Iso-surfaces of the reconstructed first time step. Left: x-component iso-surface $u_x=\pm 14.0\text{m/s}$. Middle: y-component iso-surface $u_y=\pm 11.0\text{m/s}$. Right: z-component iso-surface $u_z=\pm 10.5\text{m/s}$ 106

Figure 54: Mode 15 from the STDMD spectrum at $St_{15}=2.68$ and $f_{15}=108\text{Hz}$. Iso-surfaces of the reconstructed first time step. Left: x-component iso-surface $u_x=\pm 9.5\text{m/s}$. Middle: y-component iso-surface $u_y=\pm 9.7\text{m/s}$. Right: z-component iso-surface $u_z=\pm 8.8\text{m/s}$ 107

Figure 55: Mode 16 from the STDMD spectrum at $St_{16}=1.18$ and $f_{16}=48\text{Hz}$. Iso-surfaces of the reconstructed first time step. Left: x-component iso-surface $u_x=\pm 4.1\text{m/s}$. Middle: y-component iso-surface $u_y=\pm 3.9\text{m/s}$. Right: z-component iso-surface $u_z=\pm 3.4\text{m/s}$ 108

List of Tables

Table 1: Experimental wind tunnel testing results from the WTA of TUM versus simulation results. Experimental drag values are rolling resistance corrected. Lift forces are without the lift contributions of the wheels but include the drag contributions of the wheels through the momentum balance.	70
Table 2: Unsteady force coefficient standard deviation from experiment and simulation and maximum absolute deviation.	75
Table 3: Execution times and performance loss parameters for computation of incremental POD using different maximum rank settings and interpolation grids.....	89
Table 4: Execution times and performance loss parameters for computation of STDMD using different maximum rank settings and interpolation grids...	95

Abstract

The aerodynamics of passenger cars and light commercial vehicles is vital as government entities mandate increasingly stringent exhaust gas emission standards and manufacturers release more electric cars into the consumer market. Various tools can be employed to study the flow field around the vehicle and understand dominant aerodynamic effects. The instationary nature of flows at high Reynolds numbers complicates the study of flow structures due to the presence of a wide range of temporal and spatial scales. One can gain a better understanding of such complex flow fields by using data-driven modal analysis tools to separate energetic structures from the randomness. This work builds upon modal decomposition methods that are based on Proper Orthogonal Decomposition (POD) and Dynamic Mode Decomposition (DMD). It introduces respective incremental variants for both methods to prove their applicability for analyzing large data sets in an industrial environment with low memory requirements. It also identifies additional adjustments and extensions to the incremental methods and their respective effects on the resulting modal decomposition. In addition, this work publishes two modifications of the incremental DMD variant with a feature to rank modes by a new mode-ordering criterion and a new way of selectively computing modes in a mode-by-mode approach. A simple, two-dimensional cylinder test case is used for studying the methods' central features; namely, the implementation, required computational resources and the properties of the results. Both methods provide the means of reducing the amount of data of a Full Order Model (FOM), an exact representation of the original flow field, to a Reduced Order Model (ROM), which can be seen as an approximation or a compressed version of the FOM. This work reviews methods for validation of the ROM, data sampling strategies and data pre-processing requirements to ensure applicability to different flow configurations. The vehicle aerodynamics application presents results of modal analysis of the unsteady flow field simulation around the DrivAer reference body. The benefit of modal analysis methods is demonstrated by extraction and analysis of the most dominant flow structures. Reconstruction in time allows observers to track back structures to their respective excitation locations with the purpose of gaining additional insight into underlying dynamical flow field processes.

“Data is the new oil. It’s valuable, but if unrefined it cannot really be used. It has to be changed into gas, plastic, chemicals, etc. to create a valuable entity that drives profitable activity; so must data be broken down, analyzed for it to have value.”

Clive Humby in 2006 [1].

1 Introduction

Ground vehicle aerodynamics has been an integral part of vehicle design since the first modern trains and cars were developed in the late 19th century. In early 1920s, automotive engineers discovered that streamlining vehicle shape reduced drag values at high speed [2]. This finding led to farther driving ranges, higher maximum velocities, and decreased fuel consumption. Recently, exhaust gas emission regulations in the European Union and growing customer awareness of the environmental benefits of fuel-efficient cars are pushing manufacturers to further improve vehicle aerodynamics.

Along with improving gas mileage, the essential task of aerodynamics engineering in the automotive world is to enhance aeroacoustics, soiling, and ride comfort. For some applications, a costly wind tunnel test can be effectively replaced with numerical simulations. In the case of unsteady flow field analysis, simulation results can offer much more details than any flow field visualization technique available for wind tunnel experiments, which are limited to either recording data of a small number of probes at the same time or capturing snapshots of a plane in the area of interest. Thanks to the availability and depth of data from simulations, smart data processing tools enable accessibility of flow field features to better evaluate simulation results. The sheer amount of flow field data that can be generated requires efficient methods for analysis. Those data sets are too large and complex to be handled with traditional data processing utilities. Advanced data filtering, reduction, sampling, analysis and visualization are required to gain insight from such large data sets and make relevant information accessible and understandable. The objective of this research work is to stress the importance of new algorithms for data analysis when such challenges are met and to propose a new method of evaluation that builds on the existing methodology for modal decomposition. This work also shows the applicability of this method to typical industrial engineering infrastructures while using limited computational resources. The method is then applied to vehicle aerodynamics simulation to demonstrate its usability and effectiveness.

1.1 Motivation

Computational fluid dynamics simulations provide a useful tool for reducing costs in the aerodynamic development cycle of a vehicle. Along with the increase in computational resources, there is a continuous shift from steady-state simulations employing Reynolds-Averaged Navier-Stokes (RANS) simulations toward scale-resolving simulations with large eddy simulation (LES) characteristics. Unsteady simulations offer more precise approximations of the integrated forces and allow for in-depth analysis of the unstable nature of the flow. Unsteady flow field analysis, however, is not self-explanatory when investigating flows around complex geometries at a large Reynolds number Re . The definition of

dimensionless Re (1) uses the free stream velocity U_∞ , a characteristic length scale L , and the fluid properties density ρ and dynamic viscosity μ as

$$Re = \frac{\rho U_\infty L}{\mu}. \quad (1)$$

The Reynolds number describes the ratio of inertial forces to viscous forces and can be used to predict the behavior of a flow. Small Reynolds number flows are flows where viscous effects with small velocity gradients dominate the flow. Turbulent structures, high-velocity gradients, intensified mixing and momentum transport perpendicular to the flow direction govern large Reynolds number flows. Typical Reynolds numbers for the flow around vehicles at highway cruising speed are in the range of $Re=10^6$, covering a wide range of spatial and temporal scales. Various vortex identification methods are available to visualize those structures. An iso-surfaces of the Q-criterion as defined in equation (101), one of the most popular parameters, is depicted in Figure 33 and allows areas of detachment to be localized. This way, dominant flow structures can easily be visualized for low- Re flows and simple geometries. If geometries are more complex and turbulent production happens at different scales, identifying the relevance of vortex structures from iso-surfaces of the Q-criterion becomes unfeasible. The overlapping nature of flow field fluctuations results in wrinkly Q-criterion iso-surfaces that can no longer be assigned to a specific excitation mechanism. When this occurs, it is not possible to track coherent structures through the flow field; they cannot be separated from low-energy flow field fluctuations that otherwise can be regarded as noise when studying the most dominant processes. Cars are considered to be complex geometries regarding aerodynamics due to the existence of a vast amount of vortex detachment processes, involving turbulence production at a wide range of spatial scales. The turbulence production initial length scale l_0 for various geometries is in the same order as the characteristic length scale L of a geometry. Example geometries are the width of a wheel or the height of the backward-facing step at the interface between the engine hood and the front window. Dominant flow structures in the vehicle wake can be of the size of the vehicle, while the flow structures detaching from the side view mirror support are smaller by several orders of magnitude. Thus, the energy transfer from the mean flow into turbulent kinetic energy (i.e., turbulence production) happens at various scales. Characteristic scales and associated energies in turbulent flows are depicted in Figure 1. See Pope [3] for a detailed reference on the properties of the turbulent energy cascade. After formation of initial structures, vortices break down into smaller structures, and energetic dissipation due to friction starts. On a much smaller length scale, the Kolmogorov length scale η , the structures finally dissipate entirely into heat. The length scales L and η can be several orders of magnitude apart, depending on the Reynolds number, which is also referred to as separation of scales.

$$\frac{L}{\eta} \sim Re^{3/4}. \quad (2)$$

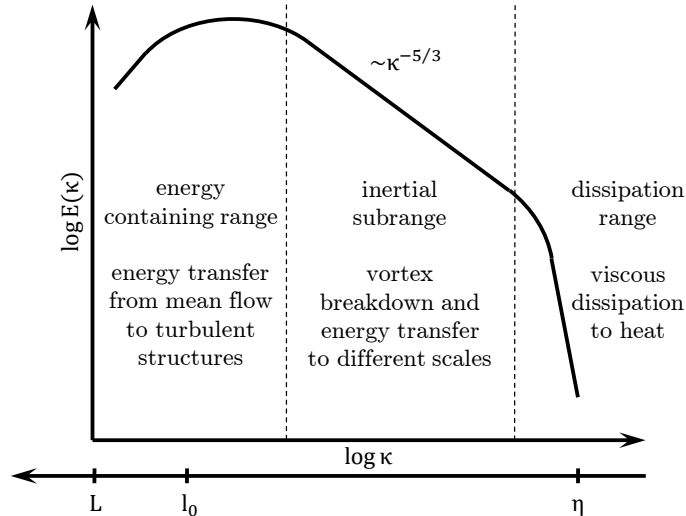


Figure 1: Energy cascade in turbulent flow, depicting the turbulent kinetic energy E vs. wave number κ .

Due to the complex nature of the vortex-shedding mechanisms involved, it is critical to be able to extract the most dominant flow structures to study their respective detachment processes, their convection paths, and their interactions with other coherent structures and vehicle parts. A deeper understanding of those processes will lead to better engineering results. Modal analysis, which is also referred to as modal decomposition or reduced order modeling (ROM), can be employed to obtain this knowledge, as will be shown in this work.

Modal decomposition methods can be applied to vehicle aerodynamics for extraction of coherent flow structures, data compression, and exact frequency filtering or noise reduction. Identifying the globally most dominant modes that induce the most massive oscillations in the flow field is vital for investigating structures. Turbulent structures detaching from the trailing edges of a car are possibly triggered by small-scale disturbances in the flow field far upstream. Changing the behavior of those structures with geometrical modifications or employing flow control devices could lead to decreased overall aerodynamic drag.

If energetic aerodynamic oscillations in the force coefficients are of concern (e.g., oscillations of the rear axle lift coefficient), driving stability and passenger comfort are influenced. A deeper understanding of their generation mechanism is invaluable to modify the aerodynamic processes inducing the force fluctuations of interest. For meaningful insight into the temporal evolution of disturbances, DMD modes can be reconstructed in time.

Aeroacoustic phenomena are rarely well understood and lead to massive damping measures, increasing vehicle weight. If the underlying processes can be simulated, the dominant mechanism for a specific frequency range can be extracted and countermeasures can be designed to prevent those acoustic waves. For example, an aeroacoustic feedback loop at the side mirror that tends to induce a whistling sound. To understand the shedding processes involved, one must extract the most dominant flow structures, but the sheer amount of unsteady data that can be produced during a single simulation is overwhelming.

The generation of a ROM that contains most of the fluctuation energy in a set of modes smaller than those in the original data (i.e., the full-order model (FOM)) can help to reduce the amount of data. A central target is to maintain a majority of unsteady flow field characteristics. Conversion of thousands of time steps into just tens of modes, which can reproduce the time-resolved flow field up to a small error, is possible in some cases.

Complex vortex-vortex interactions lead to the problem that some vehicle add-on parts introduce low drag deficiencies mounted on one car while leading to larger drag increase mounted on another. The side-view mirror installed on vehicles with different rear end geometries is a good example. Interactions between the structures shed by the mirrors and the rear end might lead to unfavorable conditions. The convection paths and temporal evolution could be used to understand these interactions.

Apart from vehicle aerodynamics, there are many application areas of ROM in other fields of fluid mechanics. A good overview of such applications is given in a 2017 review paper by Rowley & Dawson [4], such as noise generating cavity flow and transitional boundary layer flow studies. Other areas of application outside the fluid mechanics community include image processing, data compression, signal analysis, modeling and control of chemical reaction systems, electrical power grid modeling, and financial forecasting for trading strategies.

1.2 Numerical Simulations

Various equations and methods for describing fluid motion in CFD are available. For unsteady simulations in engineering, smoothed-particle hydrodynamics (SPH) approaches, the Boltzmann equation, and the Navier-Stokes equations (NSE) are most commonly used to predict fluid flow behavior.

SPH relies on the representation of fluids by Lagrangian particles in a mesh-free approach. This method is frequently used for multi-phase simulations, especially if free surfaces are involved, which are hard to capture with mesh-dependent methods. The simulation of oil distribution in a gearbox is an excellent example of an engineering application that was predestined for the usage of SPH.

The Boltzmann equation can be solved using a numerical grid, which is referred to as lattice in this context, using the lattice-Boltzmann method (LBM). Originating from a microscopic model called lattice gas automata, which is used for simulating molecular dynamics (i.e., collisions of single gas particles), the LBM was developed to take into account a probability density function of particles and their velocities. This way, the lattice spacing can be much larger than the order of the mean free path length of an individual gas particle. The most crucial benefit of LBM compared to other CFD methods is the local nature of operations to be performed in the collision and streaming step that is executed in each time step. This characteristic make the approach outperform others regarding parallel scaling.

The NSE describe the conservation of mass and momentum from a macroscopic point of view. CFD simulations using the NSE also employ spatial discretization methods on a numerical grid. Applicability of the NSE is limited

to systems with a sufficiently small Knudsen number $Kn \ll 1$ (i.e. if the physical length scale of interest is much larger than the mean free path of molecules). This condition is met for describing flows around passenger car vehicles so that the continuum hypothesis can be applied and macroscopic mean quantities can be treated as continuous functions.

This work is built on the former NSE-based method using the OpenFOAM® library. OpenFOAM® is a C++-based open source software library for CFD [5]. The following subsections present the NSE, the applied turbulence closure model, the discretization schemes and the solver.

1.2.1 Governing Equations of Navier-Stokes

The mass conservation equation states that the change in density ρ within a fluid element is obtained by the sum of mass flux across its borders through the velocity component u_i . In differential form, it writes

$$\frac{\partial \rho}{\partial t} + \frac{\partial \rho u_i}{\partial x_i} = 0. \quad (3)$$

Passenger car aerodynamics is investigated at free stream flow velocities around $U_\infty = 140 \text{ km/h}$. The dimensionless number to judge if a flow can be considered as incompressible is the Mach number Ma , defined with speed of sound a

$$Ma = \frac{U_\infty}{a}. \quad (4)$$

For $Ma < 0.3$, which is the case for passenger car aerodynamics, compressibility effects play a minor role, and the total derivative of the density becomes zero

$$\frac{d\rho}{dt} = 0. \quad (5)$$

Using the incompressibility simplification, the mass conservation equation simplifies to

$$\frac{\partial u_i}{\partial x_i} = 0. \quad (6)$$

The momentum conservation equations (three partial differential equations) describe the balance of momentum in a fluid volume and the forces acting on it. Momentum can only be added by momentum flux across the borders and stresses acting on the surface. For incompressible flows without any active body forces, the momentum conservation equations read

$$\frac{\partial u_i}{\partial t} + \frac{\partial u_j u_i}{\partial x_j} = -\frac{1}{\rho} \frac{\partial p}{\partial x_i} + \frac{\partial \tau_{ij}}{\partial x_j}, \quad (7)$$

where the shear-stress tensor τ_{ij} (due to molecular viscosity) is

$$\tau_{ij} = \nu \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) = 2\nu D_{ij}, \quad (8)$$

for which the velocity gradient tensor $\mathbf{L} = \partial u_i / \partial x_j$ can be decomposed into a rate-of-rotation tensor $\mathbf{\Omega}$ and the deviatoric rate-of-strain tensor \mathbf{D}

$$\begin{aligned}
\mathbf{L} &= \begin{pmatrix} \frac{\partial u_1}{\partial x_1} & \frac{\partial u_1}{\partial x_2} & \frac{\partial u_1}{\partial x_3} \\ \frac{\partial u_2}{\partial x_1} & \frac{\partial u_2}{\partial x_2} & \frac{\partial u_2}{\partial x_3} \\ \frac{\partial u_3}{\partial x_1} & \frac{\partial u_3}{\partial x_2} & \frac{\partial u_3}{\partial x_3} \end{pmatrix} \\
&= \frac{1}{2}(\mathbf{L} - \mathbf{L}^T) + \frac{1}{2}(\mathbf{L} + \mathbf{L}^T) = \frac{1}{2}\left(\frac{\partial u_i}{\partial x_j} - \frac{\partial u_j}{\partial x_i}\right) + \frac{1}{2}\left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}\right) \\
&= \mathbf{\Omega} + \mathbf{D} = \mathbf{\Omega} + \mathbf{D}.
\end{aligned} \tag{9}$$

Tensor \mathbf{D} also is often referred to as deformation rate tensor or strain rate tensor. It is symmetric and consists of six independent entries. $\mathbf{\Omega}$ is an asymmetric tensor, contains three independent elements and is known as the rotation velocity tensor.

The most straightforward way to simulate vehicle aerodynamics would be to discretize and solve this set of equations directly (i.e., doing direct numerical simulation (DNS)). Due to the separation of scales in high Reynolds number applications as described in section 1.1, it is not possible to discretize the equations properly in space and time due to computational resources. Physical flow fields can only be obtained if all turbulent scales can be adequately resolved in space and time. In this case, the cells would have to be of the order of the Kolmogorov scale, and an exorbitant amount of cells would be required. That is prohibitive regarding memory and would lead to incredibly small simulation time steps to obtain stable and accurate simulation results with a Courant number Co less than or equal to 1. The Courant number, known as the Courant-Friedrichs-Lewy condition, for a discrete cell of side length Δx , local flow velocity U and time step size Δt is defined as

$$Co = \frac{U \cdot \Delta t}{\Delta x}. \tag{10}$$

This law immediately shows the correlation between the spatial resolution and the required time step size to keep $Co \leq 1$. A mesh refinement by a factor of two leads to a reduction of the time step size by a factor of two and an increase in the number of cells by eight, for three dimensions.

1.2.2 Turbulence Modeling

Turbulence modeling approaches are required for vehicle aerodynamics analysis due to the prohibitive computational efforts of DNS for complex 3D geometries and high Reynolds numbers. The amount of turbulence modeling needed is directly linked to the requirements for depth and precision of the results, as well as to available resources (e.g., the maximum time to obtain a result or computational power). For internal flow through pipe systems, it might be enough to compute a steady-state solution to estimate the temporally averaged pressure drop. For vehicle aerodynamics, key results are the temporally averaged integrated drag and lift forces. Even though this information can be obtained

from steady-state simulation approaches, it has been shown that unsteady simulations can improve the results drastically. This is because unsteady flow field effects (e.g., detachment processes) can be captured appropriately and because the shape and pressure drop in large separation regions are approximated better, see Islam et al. [6]. The following section introduces the state-of-the-art simulation approaches from the class of detached-eddy simulation (DES) used in current vehicle aerodynamics engineering that prevent from using DNS while still resolving the large flow structures in time and space. For a more comprehensive overview of DES methods, including applications on different geometries, see the Ph.D. work of Mocket [7].

DES models turbulence with Unsteady Reynolds-Averaged-Navier-Stokes (URANS) in wall proximity and LES in all other regions. The vehicle aerodynamics simulation in this work uses the Spalart-Allmaras (SA) turbulence model for calculating the turbulent viscosity in the RANS region and a Smagorinsky-like turbulence model behavior for the subgrid-scale turbulent viscosity in the LES region. This combination of turbulence models leads to robust complex geometry simulations.

A combination of LES and RANS for such complex geometries at high Reynolds numbers is going to remain state of the art in the foreseeable future, as is explained by Spalart [8]. The complexity of a wall-resolved LES with spatial discretization requirements in all three dimensions leads to the need for RANS models in this region because only the wall-normal spacing is crucial for a sound reproduction of the law of the wall. Additional wall models that can be used in RANS models can then be applied to reduce further the amount of wall-normal refinement to prevent conflict with the $y^+ < 1$ requirement and instead allow for y^+ values of up to 200 (e.g., with the first cell center in the outer log-law region of the boundary layer (y^+ as defined in equation (18))).

Solely relying on RANS, however, is also not suggested, since large separation regions behind bluff bodies cannot be adequately modeled. Also, some engineering tasks rely on knowledge about the unsteady behavior of the flow, so temporally resolved data is needed. For the derivations of the URANS and the LES equations, the Navier-Stokes equations are either Reynolds-Averaged (for URANS) or filtered (for LES). The resulting equations show a structural similarity in that they can both be formulated as seen in equation (11), where the overline stands for the specific operation performed on the variable in either RANS or LES. Both transport equations pose a closure problem of similar nature for which a turbulent shear stress tensor, represented by $\tau_{ij}^{\text{modelled}}$ (also referred to as Reynolds-stresses in context with RANS simulation) needs to be computed:

$$\frac{\partial \bar{u}_i}{\partial t} + \frac{\partial \bar{u}_i \bar{u}_j}{\partial x_j} = -\frac{1}{\rho} \frac{\partial \bar{p}}{\partial x_i} + \frac{\partial \bar{\tau}_{ij}}{\partial x_j} + \frac{\partial \tau_{ij}^{\text{modelled}}}{\partial x_j}. \quad (11)$$

For a detailed derivation of the RANS and LES equations and a possible closure method for each approach separately, see Pope [3]. For both approaches, several closure options exist that are described by turbulence models. A common way of closure is the turbulent-viscosity hypothesis, which claims that the effects

of unresolved turbulent flow can be modeled using a turbulent viscosity ν_T . This increases the local fluid viscosity by a certain amount due to turbulence because turbulence increases momentum exchange in all directions, as shown here:

$$\tau_{ij}^{\text{modelled}} = -\frac{2}{3}k\delta_{ij} + \nu_T \left(\frac{\partial u_j}{\partial x_i} + \frac{\partial u_i}{\partial x_j} \right). \quad (12)$$

For DES based on the SA turbulence model, a single transport equation that must be solved for the evaluation of the modified turbulent viscosity $\tilde{\nu}_T = \nu_T / f_{v1}$ is written:

$$\begin{aligned} \frac{\partial \tilde{\nu}_T}{\partial t} + \frac{\partial u_i \tilde{\nu}_T}{\partial x_i} &= C_{b1} \tilde{S} \tilde{\nu}_T \\ + \frac{1}{C_\sigma} \left[\frac{\partial}{\partial x_i} \left((v + \tilde{\nu}_T) \frac{\partial \tilde{\nu}_T}{\partial x_i} \right) + C_{b2} \left(\frac{\partial \tilde{\nu}_T}{\partial x_i} \right)^2 \right] &- C_{w1} f_w \left(\frac{\tilde{\nu}_T}{\tilde{d}} \right)^2 \end{aligned} \quad (13)$$

with

$$\begin{aligned} f_{v1} &= \frac{\chi^3}{\chi^3 + C_{v1}^3}, \chi = \frac{\tilde{\nu}_T}{\nu}, \tilde{S} = f_{v3} \sqrt{2\boldsymbol{\Omega}_{ij}\boldsymbol{\Omega}_{ij}} + \frac{\tilde{\nu}_T}{\kappa^2 \tilde{d}^2} f_{v2} \\ f_w &= g \left(\frac{1 + C_{w3}^6}{g^6 + C_{w3}^6} \right)^{1/6}, g = r + C_{w2}(r^6 - r), r = \frac{\tilde{\nu}_T}{\tilde{S} \kappa^2 \tilde{d}^2} \\ C_\sigma &= \frac{2}{3}, C_{b1} = 0.1355, C_{b2} = 0.622, C_{v1} = 7.1 \end{aligned} \quad (14)$$

$$C_{v2} = 5, C_{w1} = 3.239, C_{w2} = 0.3, C_{w3} = 2, \kappa = 0.41$$

as defined by Spalart and Allmaras [9]. Furthermore, the implemented model includes corrections from Ashford [10] that protect the source term \tilde{S} from becoming negative, as shown here:

$$f_{v2} = \left(1 + \frac{\chi}{C_{v2}} \right)^{-3}, f_{v3} = \frac{(1 + \chi f_{v1})(1 - f_{v2})}{\chi}. \quad (15)$$

A significant downside of DES is the gray area (i.e., the region between cells modeled by RANS and LES) in which the switching of the modeled turbulent viscosity occurs. The original formulation of DES from Spalart et al. [11] proposes a grid dependent length scale to determine the switch from RANS to LES. That was later found to be problematic because the LES region can be active inside thick boundary layers in which turbulent boundary layer structures cannot yet be adequately modeled with LES due to the mesh requirements in all three dimensions. If LES mode is active inside such a thick boundary layer, the much lower LES-like modeled turbulent viscosity cannot be evened out by resolved Reynolds stresses due to missing resolution. This effect is also referred to as Modeled-Stress Depletion (MSD), which can lead to Grid-Induced Separation (GIS), as described by Menter et al. [12]. To overcome the problem of GIS, Spalart et al. [13] introduced delayed detached-eddy simulation (DDES), which considers the current local flow solution in the RANS-LES switching function. This generic shielding function approach takes into account the local eddy viscosity, the wall distance, and the velocity gradients. The shielding function is

controlled using the turbulent length scale \tilde{d} , which needs to be adjusted to fit the turbulence model region:

$$\begin{aligned} \tilde{d} &= d_{\perp} - f_d \cdot \max(0, d_{\perp} - C_{\text{DES}} \cdot \Delta_{\text{DES}}) \\ f_d &= 1 - \tanh([8r_d]^3), \quad r_d = \frac{v_T + \nu}{\sqrt{\frac{\partial u_i}{\partial x_j} \frac{\partial u_i}{\partial x_j} \kappa^2 d_{\perp}^2}}. \end{aligned} \quad (16)$$

This way, the turbulent length scale returns URANS values (i.e., the wall distance d_{\perp} for computing the SA-based turbulent viscosity) and the LES length scale $C_{\text{DES}} \cdot \Delta_{\text{DES}}$ for LES-like behavior. In this work, $C_{\text{DES}}=0.65$ and $\Delta_{\text{DES}}=\text{cellVolume}^{1/3}$ are chosen. The usage of the cube root volume of the cell as filter size works perfectly fine, even though other investigators recommend different filter definitions (e.g., the maximum cell dimension). The mesh generation in this work is such that prism layers on the surface are extruded and in this area the URANS behavior is active. Outside the prism layers, the mesh is dominated by cubic hexahedral cells, for which the cube root volume and the maximum cell dimension definition are equal. In this study, the LES model region is almost exclusively active in this region of the mesh (i.e., there are mostly cubic cells in areas of $f_d=1$).

Compared to the pressure equation, this transport equation can be solved quite fast and does not need to be coupled with the other equations to be solved for a stable solution. The downside of this model using a single transport equation for both the URANS and the LES regions of the flow field is that there is no clear border between the regions. Only the amount of turbulent viscosity computed indicates whether the model returns RANS values or smaller subgrid-scale values for LES content. The big advantage over zonal methods, in which the boundaries for RANS and LES are defined explicitly, is the missing expense for zone definition and the adaption to possibly temporally changing flow field states. The zone-free model can potentially resolve turbulent structures using LES-like behavior in one time step but automatically switches to RANS if required by a change in local flow properties.

As mentioned above, by employing URANS simulation behavior close to the wall, wall models can be used to approximate the turbulent viscosity in the first cell on the wall surface for $y^+ < 300$. The dimensionless wall distance y^+ can be computed using the wall-normal distance d_{\perp} measured in meters, the wall friction velocity

$$u_{\tau} = \sqrt{\frac{\tau_w}{\rho}} \quad (17)$$

and the kinematic viscosity ν

$$y^+ = \frac{d_{\perp} u_{\tau}}{\nu}, \quad (18)$$

where the wall shear stress is defined as

$$\tau_w = \mu \frac{\partial u}{\partial d_\perp}. \quad (19)$$

That allows evaluation of the magnitude of the wall friction velocity following equation (17) as

$$u_\tau = \sqrt{(v_T + \nu) \left| \frac{\partial u_i}{\partial d_\perp} \right|}. \quad (20)$$

Equation (20) can then be used to compute the turbulent viscosity in the first cell if the usage of a wall function approximates the dimensionless wall distance. Similarly to the dimensionless wall distance, the dimensionless velocity u^+ can be defined with the wall parallel velocity component u_\parallel as

$$u^+ = \frac{u_\parallel}{u_\tau}. \quad (21)$$

In this work, the continuous wall function formulation of Spalding’s publication on “[a] single formula for the law of the wall” [14] is used, returning good results for the entire range of y^+ values. Those results are achieved by employing a smooth blending function between the viscous sublayer and the log-law region, yielding a smooth transition in the buffer layer. This function is crucial for the application of the wall functions to flows around complex geometries with various boundary layer thicknesses and automated mesh generation. It uses local velocity values and the wall distance to approximate y^+ to evaluate the turbulent viscosity in the first cell as

$$y^+ = u^+ + \frac{1}{9.8} \left(e^{\kappa u^+} - 1 - \kappa u^+ - \frac{1}{2} (\kappa u^+)^2 - \frac{1}{6} (\kappa u^+)^3 \right). \quad (22)$$

For more information about the iterative evaluation of equation (22), the reader is referred to the master’s thesis work of Gestrich [15].

1.2.3 Discretization and Solving

Spatial discretization can be obtained by defining a numerical grid of finite elements, finite differences or finite volumes. Each of the discretization methods comes with its pros and cons. The finite volume method (FVM) is the method of choice in this work and is also used in most CFD solver packages. One of the most significant benefits of this approach is the local mass conservation property on each cell (i.e., the sum of all inflows is automatically equal to the sum of all outflows). That leads to excellent global mass conservation. Another advantage over other finite elements is the possibility to build unstructured grids that can be fitted to complex geometries (e.g., with automated meshing routines). The discretization on the FVM mesh is based on the integral form of the underlying system of equations. This representation can be obtained from the differential form using the Gauss theorem or directly from balancing of momentum change and flux across the borders of an infinitesimally small fluid control volume. See Ferziger & Peric [16] for a detailed description of discretization schemes and properties of different representation of the equations. Gradients, appearing in

divergence and Laplacian terms and interpolations to the cell faces (i.e., for integration), are approximated using second-order central differencing. Excluded from this are the advective transport of turbulent viscosity across faces and the non-linear term of the momentum transport equation in cells that have high Courant numbers. Gradients for those two exceptions are treated with linear upwind differencing, which is also second-order accurate but more diffusive and thus introduces more stability.

Before the first iteration of the resulting set of equations can take place, the flow field has to be initialized using an appropriate method. For complex geometry cases, computing the potential flow field solution (i.e., rotation- and viscosity-free flow) and adding a generic boundary layer profile on no-slip wall patches can help to speed up convergence of the instationary simulation drastically and decrease total simulation runtime.

Because the pressure does not appear in the momentum transport equations (11), special care needs to be taken to make sure that the pressure field is set in a way that the velocity field satisfies both the momentum equations and the continuity equation. One way to reach mass conservation is the PISO (pressure-implicit with splitting of operators) algorithm. In a predictor step, the momentum transport equations are solved for the velocity components using pressure values of the previous time step. The first corrector step uses a Poisson equation for the pressure (linking the momentum equations with the mass conservation equation) to make sure the pressure and velocity components satisfy the mass conservation equation. In a second corrector step, the velocity fields are adjusted using the latest pressure field. The corrector steps are repeated until the pressure correction for the time step is negligible. After the pressure and velocity fields are computed, the transport equation for the modified turbulent kinematic viscosity (13) is solved. This solver works well on low Courant numbers because the non-linear coupling between the velocity components is small compared to the pressure-velocity coupling. Time advancing is done using implicit second-order backward differencing, which is also referred to as the three-time level method in Ferziger and Peric [16]. Therefore, the overall resulting numerical order of accuracy for a majority of cells is of second order.

1.3 Modal Analysis State of the Art

Analysis of unsteady flow fields is a broad research area, giving insights into the most complex nature of turbulent flows. Applications in the fluid dynamics community include, among others, optimization of the fluid dynamics properties of a particular product to the calibration of turbulence models. This section presents a short survey of the history of modal analysis methods that build the basis for the methods used in this work. Furthermore, this section synthesizes recent studies on methods for flow field analysis with large data sets and applications on bluff body aerodynamics.

Lumley introduced POD to the fluid mechanics community in 1970 [17]. Sirovich later improved the computation of POD modes by introducing the

“method of snapshots” for a more memory-efficient analysis of well-conditioned data. In this approach, the matrix inner product of the snapshot matrix with its transpose is no longer needed, and left singular vectors can be evaluated incrementally (i.e., mode by mode).

Later, Moore [18] introduced the concept of balanced truncation, which makes it possible to process modes from a control theory point of view. He developed parameters to quantify the controllability and observability of modes to balance them for both parameters before truncation of modes.

Rowley further advanced this concept in 2005 [19] with the introduction of balanced POD, which makes it possible to find a good approximation to Moore’s truncation but is computationally more tractable for large data sets.

A highly important development of POD methods in context with this work is the incremental POD algorithm proposed by Brand in 2002 [20]. He suggested a single-pass snapshot algorithm that can be used to reduce the memory requirements of computing POD modes to a predefined maximum rank limit (e.g., hardware-given). He then used the algorithm for image processing. For the application on simulated flow field data, this advance makes it possible to evaluate POD modes in parallel during simulation without the need to store data, which streamlines analysis by avoiding excessive hard drive disk accessing time.

Oxberry et al. [21] later used a similar definition for computing incremental POD modes and provided useful error analysis for different algorithm control parameters. The same research group also published a parallel incremental POD code library including documentation [22] that was used for their previous work.

One of the downsides of POD is the observation of multiple research groups in the field that the most energetic POD modes might not be the most dynamically relevant ones. Even low-energy POD modes can have a significant impact on the global dynamics of a system, as described by Noack et al. [23] in 2008. Alternative methods that reveal the dynamically relevant flow structures need to be considered, such as Koopman operator analysis.

In 2005, Mezic [24] proposed a method that decomposes the Koopman operator for a dynamical system (e.g., time-resolved data) and results in periodic to almost periodic components (e.g., modes that oscillate at distinct frequencies) and a continuous spectrum.

The first mention of the term DMD was made in 2008 by Schmid & Sesterhenn [25], who were led by the need for data-driven analysis tools that can be applied to experimental particle-image velocimetry (PIV) data sets and to simulated data.

In 2009, Rowley et al. [26] showed that the Arnoldi-type algorithm that was proposed by Schmid & Sesterhenn can be used to produce an approximation of Koopman modes. They further applied the method on a simulated data set from direct numerical simulation (DNS) of a jet in crossflow to study the differences between DMD and POD modes. They showed that DMD can successfully extract globally dominant flow structures at distinct frequencies. Whereas several frequencies pollute POD modes that capture the most energetic structures, DMD

modes are by construction distinct in frequency and can thus isolate and decouple a specific process more effectively for further analysis.

In 2010, Schmid [27] published the Arnoldi-type algorithm he had proposed earlier. The extracted modes were described as a generalization of global stability modes that can be used to capture the underlying physics of a highly dynamical system using fewer degrees of freedom. He explained that the application of his method to linear flow leads to modes equivalent to modes obtained from global stability analysis. In case of non-linear flow, the resulting structures are a linear tangent approximation to the underlying flow. This method is the basis of a Koopman analysis of nonlinear dynamical systems. This work will use the same limitation of data analysis to a subdomain of the flow field as done in Schmid's work to reduce the amount of required memory. This data reduction technique is also shown to be a principal advantage of both POD and DMD in comparison to other global modal analysis techniques, which may require the inclusion of the boundary conditions and the underlying equations that have been used to compute the flow field.

This leads to the conclusion that POD and DMD are perfectly fitted data-driven tools that can be used to generate comparable insights into experimentally and numerically obtained data. The first estimation in Schmid's work pointed out that the requirements for the sampling of higher frequency oscillations are more demanding than in FFT. Schmid suggested a sampling frequency of three times the Nyquist criterion to generate sound DMD results. He also showed applications of the method to the simulated flow over a square cavity, the time-resolved PIV measurement of the flow in the wake of a flexible membrane, and the jet between two cylinders.

Schmid and his coworkers further investigated the application of DMD on the Schlieren image video of a helium jet and on the time-resolved PIV data of the response of an axisymmetric jet to external forcing [28]. They found that the activation of forcing on the jet is well represented in the resulting DMD mode spectrum and the spatial distribution. They concluded that acoustic forcing most actively changes the flow behavior near the nozzle and that DMD might be an exceptionally well-qualified tool to lead to such findings.

In 2012, Schmid et al. [29] applied DMD to three-dimensional tomographic PIV data of a transitional water jet. They suggested the use of transformed eigenvalues for the visualization of stability in the mode spectrum, in contrast to the eigenvalues on the unit circle. They used 200 snapshots in time for their analysis and found that most of the eigenvalues of the DMD operator are close to the unit circle. They forecasted that using more time steps would result in a continuous approaching of eigenvalues toward the unit circle as the modal basis converges toward a linear representation of a saturated nonlinear process. They also constructed spatial snapshots, with time being an independent variable and the axial coordinate becoming the new sampling variable. Spatial DMD modes contain a temporal dependency, and the mode spectrum shows the amplitude versus the streamwise wavenumber instead of a temporal frequency.

Duke et al. [30] conducted an extensive study on the error analysis of the DMD in 2012. Their work is essential for adequately setting the sampling frequency to capture damped modes from a flow field. If highly nonlinear processes are present and highly damped modes are required for generating an ROM using DMD modes, they suggest a minimum of 40 sampling points per one period of an oscillation (i.e., 20 times the Nyquist criterion). This is quite demanding compared to the requirements of FFT analysis. On the lower frequency limit, DMD outperforms FFT by any means. For a linear instability, only a quarter of a full period of oscillations needs to be contained in the data to be able to compute the correct eigenvalues. For non-linear instabilities with growth rates depending on the phase of the waveform, only a single period of oscillation is required.

The need for a physical ordering of modes and the difficulties in obtaining good ROMs from DMD led Jovanovic and coworkers to develop sparsity-promoting algorithms for DMD [31], [32]. Their suggested optimization procedure computes DMD mode amplitudes to obtain the best possible ROM approximation in an energy-norm sense with a minimal amount of modes. The optimization algorithm is computationally inexpensive because it uses the right singular vectors and the singular-value matrix from their underlying POD modes as an input, instead of the full DMD mode matrix plus the snapshot matrix itself. However, it requires the evaluation of the POD.

Tu, in the research group of Prof. Rowley, developed a comprehensive theoretical framework of DMD in his Ph.D. thesis in 2013 [33] and in a journal publication in 2014 [34]. He advanced several extensions of DMD, including strategies for compressed sensing and construction of snapshots that do not necessarily need to be in consecutive order and that can be enhanced to introduce more noise robustness to the analysis.

Hemati from the same research group later encountered problems using DMD related to computational resources and noisy data. The sheer size of the matrices involved leads to considerable memory requirements for the evaluation of DMD modes. He introduced a memory-efficient way to compute DMD modes that makes it possible to evaluate DMD online (i.e., during the simulation) without the need to store simulated data. A single-pass incremental variant of DMD, named streaming DMD (SDMD) [35], uses Gram-Schmidt orthogonalization of snapshots to compute an incrementally updating orthonormal basis. After each update, the method allows for online compression to limit the amount of memory required. Unfortunately, the method does not perform as well as conventional DMD on the same data set; spatial distribution of modes is less smooth, and eigenvalues show higher damping rates.

Hemati then tackled the noise sensitivity of DMD methods by publishing the concept of augmented snapshots in 2017 [36]. The total DMD (TDMD) method makes use of augmented snapshots (i.e., vertically concatenated regular snapshots of preceding time steps) and introduces a preceding singular-value truncation step after POD evaluation of the augmented snapshot matrix. This step is pre-filtering the snapshot matrix for further conventional DMD analysis,

removing noise from the data and leading to increased noise robustness. The concept of augmented snapshots leads to smoother mode distributions, fewer and less-damped modes, and better ROM performance. The method was merged with SDMD to create streaming total DMD (STDMD) in 2016 [37], which includes the memory efficiency from SDMD and the noise robustness from TDMD in one method.

Because the sparsity-promoting approaches are not applicable for the SDMD- and STDMD-based approaches, it is essential to find reliable methods for ordering modes other than the conventional ordering by amplitude, which can be biased by large damping rates. Kou et al. [38] and Alenius [39] each present an approach for considering the temporal evolution of modes during the period sampled. A better mode ordering strategy is obtained through the integration of the amplitude development, including the respective eigenvalue. This method can compete with or even outperform sparsity-promoting DMD in some cases, as shown by Kou et al. This DMD mode post-processing approach will be referred to as eigenvalue-weighted amplitude ordering.

The publications presented up to here form the basis for the methods in this work. Further methodologically relevant developments have been made recently, of which good overviews are given by Rowley & Dawson [4] and Taira et al. [40].

Applications of modal analysis of vehicle aerodynamics are limited, but the method is increasing in popularity because recently published studies show its potential for various purposes. For example, studies on unsteady train aerodynamics were conducted by Muld et al. in a series of publications [41], [42], [43], [44], [45], [46]. They first developed a suitable DES simulation environment, which is validated using a wall-mounted cube case and later a realistic train model. They found that modeling of a tripwire for boundary layer transition at a specific location is not possible using DES. They extracted modes using POD and DMD and found a close resemblance between them. The leading modes are associated with the bending of two counter-rotating vortices in the train wake and vortex shedding. For the visualization of modes, reconstruction in time is found to be most useful. They further investigated differences in flow structures between two different high-speed train geometries and concluded that, despite the large differences in the mean flow field, the unsteady flow is governed by convective vortex shedding. For varying train lengths, it was found that the vortex-shedding frequencies in general decrease as train length increases, which is most likely a result of the increased boundary layer thickness at the separation point. Presumably, a thicker boundary layer increases the characteristic length scale of a geometry. With the assumption of a constant leading Strouhal number, the frequency must decrease.

Frank and coworkers used compressible LES simulations with high-order discontinuous Galerkin spectral element method to study tonal noise generation mechanisms on side-view mirrors [47], [48], [49]. Aeroacoustic flow field effects like these are particularly relevant to understand, predict and prevent, because tonal noise can be perceived as a disturbing whistling sound, decreasing passenger comfort. Frank uses DMD for the extraction of globally dominant modes on the

diffusor side of the mirror and in the area close to the trailing edge at the top of the mirror. Using 200 temporal snapshots, selected unstable modes are visualized; these are the driving modes for an aeroacoustic feedback loop phenomenon that triggers aeroacoustic noise generation due to laminar boundary layer separation. The separating laminar boundary layer introduces a shear layer in which disturbances rapidly grow toward the trailing edge. The structures are then scattered at the trailing edge, reinforcing the dominant frequencies in the shear layer close to the separation point through receptivity. The understanding of the underlying mechanism has led to various possible modifications to prevent side-view mirror whistling.

Recent research in the group of vehicle aerodynamics at the Technical University of Munich (TUM) includes analysis of temporally resolved flow fields from numerical simulation of passenger cars by Peichl, Matsumoto and Kiewat.

Peichl et al. published their first application of conventional DMD analysis on the notchback configuration of the DrivAer geometry [50] in 2014, on which this work builds. They found that the mean flow mode is the most dominant of all modes and that further dominant processes represent the movement of wake structures and vortex shedding from the rear wheels. The application of DMD for the case was limited to the near wake of the vehicle due to memory constraints of the batch-processed conventional DMD algorithm. To obtain reasonable results, they argued that temporal filtering using a Butterworth filter is necessary for preprocessing the data, which was sampled on the CFD mesh. To eliminate chances for phase shift, the temporal filter was applied forward and backward in time. They found that the tremendous amount of memory and hard drive disk space limited the applicability of the method in an industrial engineering environment and that smarter ways of judging a mode's significance are needed.

Matsumoto from the same research group and coworkers were the first to apply an on-the-fly type DMD algorithm to vehicle aerodynamics results [51]. They revealed the difference of the aerodynamic characteristics caused by geometrical modification of the engine bay compartment. As an input for the analysis, surface force vector data was used. Matsumoto and Indinger further developed an alternative of the on-the-fly DMD algorithm by using incremental singular-value decomposition [52]. This approach enables the use of sparsity-promoting DMD after the on-the-fly DMD computation, which results in a better estimation of the DMD amplitudes. They then applied the method to the flow field data of a two-dimensional square cylinder case. The group then applied the newly developed incremental DMD on simulated surface forces of the DrivAer model [53]. This work shows incremental DMD's usefulness and the close resemblance of obtained DMD spectra with FFT spectra. Furthermore they investigated the applicability of incremental POD for compression of very large transient data sets of 901 snapshots from 645 million observables [54]. They found a good reproduction of the original data series by comparing the reconstructed time signal in sampling points of the flow field using only 100 incrementally computed POD modes.

Kiewat et al. concluded from Peichl's work on modal analysis methods that memory-efficient DMD variants are needed to be able to use DMD for engineering purposes. We implemented STDMD to validate its applicability on complex geometries and applied it to the velocity volume vector field of an isolated rotating wheel in comparison to batch-processed DMD [55]. The incremental variant was shown to outperform conventional DMD due to the augmented snapshot strategy. Further memory savings were obtained by spatial cell volume weighted interpolation to a coarser grid. The STDMD method was extended to include eigenvalue-weighted amplitude ordering. The advantages of the new mode-ordering criterion were shown by direct comparison to conventional mode ordering by first-snapshot amplitude, in which the mode spectra and leading mode shapes of two wheel geometries were compared. The first wheel was a generic closed wheel, and the second wheel had the open rim of an Audi Q5 production vehicle, which was simulated using a sliding mesh approach. The leading modes of both geometries were at similar frequencies. The leading modes were classical von Karman vortex shedding, with the wheel width in y-direction being the characteristic length, and horseshoe and shoulder vortex shedding processes. In a second publication [56], the STDMD method was then applied to a full DrivAer reference body vehicle geometry in notchback configuration, including open rims, structured underbody and engine bay flow. As in the first publication, large data sets formed of velocity flow field components were used. Temporal correlation of snapshots was found to be an excellent tool for the determination of the required sampling frequency of an existing data set. The most dominant mode was connected to oscillations observed in the drag forces. Through reconstruction, it was possible to identify the initiation mechanisms of the leading mode. Further modes were attributable to the bluff body vortex shedding due to the height of the vehicle, which can be connected to oscillations in the rear lift force at the same frequency.

1.4 Objective

This dissertation develops a framework of modal decomposition methods for vehicle aerodynamics simulations. The goal is to derive a method that is capable of offering insights into the transient behavior of the flow around the vehicle that is described by large data sets. Central challenges are to overcome the limitations of existing algorithms regarding memory requirements, overall processing time and noise sensitivity. This work investigates incremental algorithms with the possibility for online compression regarding these properties. Furthermore, the usefulness of the resulting decompositions needs to be evaluated in the context of vehicle aerodynamics engineering.

A quantitative parameter has to be chosen that can judge the precision of the ROM, and the strategy presented here is to minimize the amount of data while maintaining most of the energy in the system. For this, mode-ordering criteria and their properties have to be discussed to specify a suitable mode selection parameter for vehicle aerodynamics results.

In addition, the usefulness of the results has to be judged regarding the ability to extract dominant flow structures to gain a deeper understanding of unsteady flow physics. The major differences between modal analysis methods in this respect are the ability to clearly identify a flow field effect by temporal reconstruction offering the ability to track a structure back to its excitation mechanism.

Besides the derivation and the fundamental properties of the suggested algorithms, guidelines for the application on vehicle aerodynamics simulation results have to be defined. The flow field effect under investigation can be analyzed only if the data basis that is used to generate an ROM is conditioned well and if the spatial and temporal resolution of the incoming data is chosen properly.

1.5 Outline

Chapter 1 emphasizes the need for modal decomposition methods for vehicle aerodynamics applications and contains an introduction to the world of unsteady CFD simulations.

In Chapter 2, after the theoretical framework for conducting such simulations, a two-dimensional cylinder simulation is introduced in section 2.2 as a test case for the validation of the methods and to demonstrate their properties. This study aims to show differences between the capabilities and performance of streaming modal analysis tools and related methods. Variants of those basic methods are discussed and compared while always keeping the challenges for applying this method in an industrial context in mind. Sections 2.3 and 2.4 present two classes of modal decomposition algorithms and their properties concerning the goals described in chapter 1.4. Additionally, both algorithms are investigated from a computational cost perspective, which is often the limiting factor when it comes to application to large data sets. After the underlying basic algorithms are presented, already published incremental versions of the algorithms are derived. For the DMD variant chosen in this work, significant changes in the computation of modes are developed and presented here for the first time. These extensions are built on the basis of existing incremental DMD algorithms and lead to further savings in memory and computation time.

Chapter 3 applies the developed analysis tools to vehicle aerodynamics data. Sections 3.1 and 3.2 introduce the DrivAer reference geometry setup and present the numerical flow simulation setup using DDES. Results are discussed and compared to wind tunnel measurements. Reasons are highlighted for the incomparableness of unsteady data from the experiment in the investigated frequency range. For the application of modal analysis on simulations of vehicle aerodynamics, special care has to be given to data preparation, which is a core topic in section 3.3. Spatial and temporal data reduction and filtering for the data that will be used as an input for the mode computation are considered. A strategy is presented for the selection of the sampling period and the sampling frequency. The modal decomposition algorithms are then applied to the aerodynamic simulation of the generic DrivAer reference model in fastback

configuration in section 3.4. Extraction of various flow field phenomena is achieved, their influence on the flow field is investigated, and interactions between different vortex-shedding mechanisms are found to play an influential role for the temporal evolution of the unsteady flow field.

Finally, in Chapter 4, a conclusion for the methods obtained and achieved results is given, and an outlook for possible further research is provided.

2 Modal Decomposition Methods for Fluid Mechanics

Modal decomposition methods provide means of decomposing large sets of data into modes. In fluid mechanics, spatio-temporal data is often represented with spatial modes comprising information about their temporal evolution in additional variables. Dominant large-scale coherent flow structures can be extracted, which allows for an in-depth analysis of flow field phenomena. The goal is to find a suitable representation of the flow field, which can represent the behavior of the flow field well. This also enables condensing all flow field effects to the most dominant phenomena, which can then be analyzed independently or can describe the flow field well while reducing the amount of data to be stored in a ROM.

The decomposition methods discussed in this work are entirely data-driven and do not require any further information about the underlying system being analyzed. Thus, the data analysis does not require any knowledge about the governing equations of the target field of observables.

2.1 Idea behind Data-Driven Modal Reduction

In this work, the data sets are built of snapshots of the velocity field components \mathbf{u} , \mathbf{v} and \mathbf{w} , collected from the flow field with a cell count of n_0 . For a single time step j , a snapshot is constructed as

$$\mathbf{x}_j = \begin{bmatrix} u_{1,j} \\ \vdots \\ u_{n_0,j} \\ v_{1,j} \\ \vdots \\ v_{n_0,j} \\ w_{1,j} \\ \vdots \\ w_{n_0,j} \end{bmatrix}. \quad (23)$$

The number of observables for all investigations in this work is thus $n=3n_0$. To fully capture the dominant physics of a flow field, a reasonable amount of m snapshots is obtained and can then be rearranged into a snapshot matrix $\mathbf{X} \in \mathbb{R}^{[n \times m]}$, shown here:

$$\mathbf{X} = \begin{bmatrix} | & & | \\ \mathbf{x}_1 & \cdots & \mathbf{x}_m \\ | & & | \end{bmatrix}. \quad (24)$$

The amount of samples required in space and time heavily depends on the targeted flow field processes to be investigated or captured in the modal decomposition. Focusing on a very narrow, high-frequency band of the flow field requires a relatively short sampling time if lower frequency oscillations do not modulate the process. Similarly, the spatial sampling can be restricted to a small subdomain in the flow field if a specific part is to be investigated in depth. In

this work, the primary target is to capture the energetically most dominant flow features on a global scale.

Two principal goals can be targeted using decomposition algorithms. First, this class of data processing algorithms can be used to gain more insight into the physics of a temporally resolved flow field. Second, the representation of a time series in the modes can be used to generate an ROM. While the original data set, which is also referred to as an FOM, of the flow field using a snapshot matrix with rank r_{FOM} might require massive amounts of data storage, an ROM can potentially maintain the most dominant flow features but use a smaller rank $r_{\text{ROM}} < r_{\text{FOM}}$. From this perspective, modal decomposition methods can thus also be classified as data-compression algorithms. The two resulting approaches that are used in this work are based on previously published methods of POD and DMD variants. Both of them are chosen to guarantee limited memory consumption and realize incremental processing of flow field data during the simulation run. The following sections introduce those methods and suggest modifications of the original formulations that make the methods even more applicable for industrial usage. These modifications are published in this work for the first time.

As the modal decomposition algorithms can be used for data compression and structure extraction purposes, two ways of judging the usefulness of a method are discussed.

First, for data compression, the principal target is to minimize the amount of data while maintaining most of the energy in the system. It is found that comparing the reconstructed flow field of an ROM $\tilde{\mathbf{X}}$, with the original flow field \mathbf{X} , which is used to build the ROM, is an excellent indicator for applicability to large and complex data sets. This way, the performance of the algorithm applied to a specific data set can be boiled down to a single number for one ROM rank. The performance loss indicator is defined as

$$\Pi_{\text{loss}} = 100 \frac{\|\mathbf{X} - \tilde{\mathbf{X}}\|_{\text{F}}}{\|\mathbf{X}\|_{\text{F}}}, \quad (25)$$

where the Frobenius norm of a matrix \mathbf{A} is defined as

$$\|\mathbf{A}\|_{\text{F}} = \sqrt{\text{tr}(\mathbf{A}\mathbf{A}^{\text{H}})}, \quad (26)$$

with \mathbf{A}^{H} being the conjugate transpose of \mathbf{A} . The performance loss indicator and proportional versions have been used in several previous investigations, including Jovanovic et al. [31], Oxberry et al. [21] and Kiewat et al. [55]. Due to the memory restrictions related to modal decomposition algorithms for large data sets, it is worth mentioning that the Frobenius norm is a vector norm and can be evaluated incrementally by computing

$$\|\mathbf{A}\|_{\text{F}} = \sqrt{\sum_{i=1}^n \sum_{j=1}^m |a_{ij}|^2}. \quad (27)$$

This analysis is computationally much more inefficient but has proven to be a valuable workaround for large data sets, for which $\|\mathbf{X} - \tilde{\mathbf{X}}\|_F$ in equation (25) cannot be computed at once.

Second, besides the quantitative performance-loss indicator, the interpretability of a mode or of combined modes is found to be a crucial factor for determining an algorithm's usefulness for simulations. That target is more subjective and can only be decided on a case-by-case basis.

2.2 Generic Test Case: Two-Dimensional Cylinder

The two-dimensional flow around a circular cylinder is used to describe properties of the methods employed here. Their differences are outlined using the obtained results. The flow around circular cylinders is among the most widely studied configurations at a wide range of Reynolds numbers and serves well for pointing out properties of the developed methods in this work.

A $z=0$ slice through the full two-dimensional CFD mesh is shown in Figure 2. The mesh is built of blocks using the OpenFOAM® blockMesh utility. The flow is solved for in x-direction (streamwise) and y-direction (spanwise). Due to the symmetry of the mesh to the $y=0$ plane, an initial disturbance is required to seed a fluctuation in the spanwise direction. If the flow would be initialized using a symmetric flow field, symmetric steady-state flow is observed until numerical errors lead to a small asymmetry and oscillations begin. Such processes can often take longer by several orders of magnitude than the characteristic period of oscillations in the flow. The asymmetry can be forcefully triggered using an initial condition of a uniform internal flow field with the streamwise velocity component of $U_x=U_\infty$ and spanwise velocity component of $U_y=U_\infty$. At the cylinder surface, a no-slip boundary condition is imposed for the velocity while the upper and lower side walls of the domain are of boundary condition-type symmetry (free slip condition). The inlet boundary condition (on the left) is a fixed-value velocity boundary with zero-pressure gradient. A fixed-value pressure boundary condition represents the outlet (on the right). The same pressure value is also used as a uniform initialization for the internal field.




Figure 2: Two-dimensional cylinder CFD mesh of hexahedral cells. The outer mesh has a cell size of $\Delta x=D/16$, and the 15 cell layer O-grid around the cylinder expands with a ratio of 1.6 in the radial direction. No free-hanging nodes are present.

The non-dimensionalized time and velocities for this case are defined as

$$\hat{t} = \frac{U_\infty}{D} t \quad (28)$$

and

$$\hat{u}_k = \frac{u_k}{U_\infty}. \quad (29)$$

A fixed time step size of $\Delta \hat{t}_{\text{CFD}}=0.01s \cdot U_\infty/D$ is used for the simulation. For modal analysis, every hundredth time step is used, $\Delta \hat{t}=1s \cdot U_\infty/D$. Subscript time indexing is from here on reserved for modal analysis time step indices (e.g., \hat{t}_{1000} is snapshot number 1,000, which equals the flow field of the hundred-thousandth CFD time step). Right after the asymmetric initialization at $\hat{t}_0=0$, the flow quickly becomes unstable and reaches a periodic vortex shedding state after a non-dimensionalized time of $\hat{t}_{1000} \approx 110$, for which velocity field plots are shown in Figure 3. That commonly known periodic vortex-shedding phenomenon is also referred to as a Karman vortex street. The process occurs in the range between $50 < \text{Re} < 10^7$.




Figure 3: Snapshot of the velocity field components and the velocity magnitude at \hat{t}_{1000} . Von Karman vortex shedding is in a purely periodic state.

Starting from \hat{t}_{1001} , around 19 full shedding cycles are simulated until \hat{t}_{2000} is reached at the end of the simulation. The evolution of the velocity field at a point in the wake of the cylinder is depicted in Figure 4.




Figure 4: Temporal evolution of the y-component of the velocity field probed at position $x=2D$ and $y=0$. The ramp-up phase toward reaching the periodic vortex-shedding state takes around half of the total simulation time.

A total of 2,000 snapshots are recorded for modal analysis. For a detailed comparison of modal analysis algorithm performance properties, the simulated case is split into two parts. Part one covers the linear growth of the instabilities ranging from \hat{t}_1 to \hat{t}_{1000} , and part two includes pure oscillatory von Karman vortex-shedding phenomena from \hat{t}_{1001} to \hat{t}_{2000} .

A commonly used dimensionless parameter for analysis of periodic behavior is the Strouhal number St , giving the ratio of oscillation effects and mean flow speed, defined as

$$St = \frac{f \cdot D}{U}. \quad (30)$$

Applying FFT on the time series in Figure 4 from \hat{t}_{1000} to \hat{t}_{2000} , the dominating Strouhal number around $St_{\text{FFT}}=0.179$ of the von Karman vortex street can be extracted, as can be seen in Figure 5. For better frequency resolution and more certainty in the extraction of the dominant frequency value, longer simulation runtimes would be necessary, as FFT heavily relies on a sufficiently long time series. It can only be assumed that the exact dominant frequency lies somewhere between the dominant peak in Figure 5 and the adjacent frequency to the left, due to the asymmetry of neighboring values around the maximum. Computing the frequency by directly counting the number of oscillations in Figure 4 also gives a slightly lower value of $St_{\text{simple}}=0.178$. Experimental investigations by Roshko [57] for a long cylinder in a closed test section wind tunnel resulted in a best-fit law for the Reynolds number range $50 < Re < 150$ of

$$St = 0.212 \left(1 - \frac{21.2}{Re} \right), \quad (31)$$

resulting in $St_{\text{exp}}=0.167$. A three-dimensional simulation with a better approximation of real-world boundary conditions would be necessary to match experimental results better. Since the two-dimension cylinder test case is only meant to provide flow field data for the description of modal analysis algorithm behavior and result properties, the uncertain 8% offset in Strouhal number is insignificant for the usefulness of the data.




Figure 5: Power spectral density of the velocity field y-component at position $x=2D$ and $y=0$ vs. Strouhal number. The frequency spectrum is computed using the Hanning window function with a minimum of 50% overlap in the period of t_{1000} to t_{2000} .

2.3 Low-Memory Proper Orthogonal Decomposition

POD is also commonly referred to as Karhunen-Loeve decomposition, Hotelling transform, principal component analysis (PCA), empirical orthogonal function (EOF) analysis or singular value decomposition (SVD). POD offers a way to generate a set of spatially orthogonal real-valued modes, which can be used for entirely data-driven ROMs.

2.3.1 Batch-Processed POD

A snapshot matrix \mathbf{X} is decomposed into a real-valued matrix of scaled modes $\hat{\mathbf{U}}$ that are orthogonal in space and into a real-valued orthonormal matrix \mathbf{V} describing the temporal evolution of the modes, as shown here:

$$\mathbf{X} = \hat{\mathbf{U}}\mathbf{V}^T. \quad (32)$$

That way, \mathbf{X} is separated into variables of spatial ($\hat{\mathbf{U}}$) and temporal (\mathbf{V}) dependency. Furthermore, $\hat{\mathbf{U}}$ is split into an orthonormal mode matrix \mathbf{U} and its weighting factors $\mathbf{\Sigma}$, as shown here:

$$\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T. \quad (33)$$

The matrix \mathbf{U} is column orthonormal for $n>m$, and \mathbf{V} is orthonormal in both dimensions if it is not truncated, such that their inner product is the identity matrix

$$\mathbf{U}^T \mathbf{U} = \mathbf{V}^T \mathbf{V} = \mathbf{V} \mathbf{V}^T = \mathbf{I}. \quad (34)$$

In general, taking the SVD of \mathbf{X} (equation (33)) results in an orthonormal basis of $\mathbf{U} \in \mathbb{R}^{[n \times n]}$, singular values $\boldsymbol{\Sigma} \in \mathbb{R}^{[m \times m]}$ and orthonormal matrix $\mathbf{V} \in \mathbb{R}^{[m \times m]}$. Most real-world fluid problems are typically described with numbers of observables larger than 10^6 and $n \gg m$. Such matrices are also referred to as tall and skinny matrices. In this case, the computation of a full SVD becomes prohibitively expensive even on an up-to-date supercomputer cluster. That is mainly due to the large amount of RAM required and the number of necessary flops, which is of the order $O(n^2 m + m^3)$ (e.g., when using the ‘‘R-SVD’’ procedure described in described by Golub & Loan [58]). That is also the case if only the first r modes are required. As a way out of this quadratic dependence on the number of observables, Sirovich [59] shows the application of the so-called method of snapshots for computing the POD of a system that is represented by discrete time steps. The method of snapshots is a valid approach for most applications in fluid mechanics, where only the most dominant r modes are to be computed. The concept of the method of snapshots is outlined below. The key benefit of the method is that it only requires the computation of the matrix product

$$\mathbf{R} = \mathbf{X}^T \mathbf{X}, \quad (35)$$

with $\mathbf{R} \in \mathbb{R}^{[m \times m]}$. Eigenvalue decomposition of \mathbf{R} gives the approximated first m eigenvalues and eigenvectors of \mathbf{X} ,

$$\mathbf{R} \mathbf{V} = \mathbf{V} \boldsymbol{\Sigma}^2. \quad (36)$$

The spatial distribution of POD modes \mathbf{U} can then be evaluated by rearranging equation (33) to

$$\mathbf{U} = \mathbf{X} \mathbf{V} \boldsymbol{\Sigma}^{-1}. \quad (37)$$

A considerable advantage of this concept is the possibility of evaluating the POD modes incrementally, which reduces the amount of memory required for the evaluation of columns of \mathbf{U} by a factor of almost two, as shown here:

$$\mathbf{u}_j = \sum_{k=1}^m \mathbf{x}_k \mathbf{v}_{k,j} \sigma_j^{-1}. \quad (38)$$

If all j columns of \mathbf{U} are evaluated at the same time using equation (37), the additional memory required increases by the size of the snapshot matrix \mathbf{X} in case of full rank POD. It is worth noting that special care needs to be taken on the condition of the data before applying the method of snapshots. If \mathbf{X} is ill-conditioned (i.e., $\text{cond}_2(\mathbf{X}) > 1$) then \mathbf{R} is even more ill-conditioned:

$$\begin{aligned} \text{cond}_2(\mathbf{R}) &= \text{cond}_2(\mathbf{X}^T \mathbf{X}) = \|\mathbf{X}^T \mathbf{X}\|_2 \|\mathbf{X}^T \mathbf{X}\|_2^{-1} \\ &= (\text{cond}_2(\mathbf{X}))^2. \end{aligned} \quad (39)$$

Several ways of implementing the outlined batch-processed methods exist. These algorithms are implemented in publicly available libraries, and their functionality is widely accepted and well-validated in the fluid mechanics community.

2.3.2 Singular Value Truncation

POD modes are sorted by energy content using their respective singular-value σ_{jj} . When computed, the orthonormal basis can be used to represent multidimensional data in the least-squares-optimal sense. Thus it can be handy for obtaining an ROM by truncating the mode matrix using singular value truncation. By setting a certain threshold for the amount of energy to be conserved in the system, an optimal ROM representation can be obtained. That way, POD can also be categorized as a data-compression routine, resulting in a low-dimensional approximation with rank r_{ROM} of the full state data with r_{FOM} . For other applications, singular-value truncation is also used as a filtering operation for removing low-energy noise and oscillations from data. The FOM with r_{FOM} can be smaller than m in case of linearly dependent snapshots, which can appear for periodic flows. The truncated SVD of \mathbf{X} into $\tilde{\mathbf{U}}=\mathbf{U}_{:,1:r_{\text{ROM}}}$, $\tilde{\mathbf{\Sigma}}=\mathbf{\Sigma}_{1:r_{\text{ROM}},1:r_{\text{ROM}}}$ and $\tilde{\mathbf{V}}=\mathbf{V}_{:,1:r_{\text{ROM}}}$ is depicted in equation (40).

$$\begin{matrix} n \\ \downarrow \\ \mathbf{X} \\ \uparrow \\ m \end{matrix} \approx \begin{matrix} r_{\text{ROM}} \\ \downarrow \\ \tilde{\mathbf{U}} \\ \uparrow \\ m \\ n \\ \downarrow \\ \end{matrix} \cdot \begin{matrix} r_{\text{ROM}} \\ \downarrow \\ \tilde{\mathbf{\Sigma}} \\ \uparrow \\ m \\ r_{\text{ROM}} \\ \downarrow \\ \end{matrix} \cdot \begin{matrix} r_{\text{ROM}} \\ \downarrow \\ \tilde{\mathbf{V}}^T \\ \uparrow \\ m \\ r_{\text{ROM}} \\ \downarrow \\ \end{matrix} \quad (40)$$

Truncation of the modes is based on the descending order of singular values on the diagonal of $\mathbf{\Sigma}$. The least-squares-optimal sense mentioned for equation (33) also holds for any chosen r_{ROM} . That means that any ROM based on SVD represents the original data \mathbf{X} better than any other possible matrix in the Frobenius norm sense (see equations (25) and (26)). So, for any chosen r_{ROM} , it is guaranteed to get an optimal ROM (see also Chatterjee [60]). In many cases of bluff body aerodynamics at Reynolds number of the order of $\text{Re} \approx 10^6$, the first POD mode corresponds to the mean flow mode if the data has not been mean-subtracted before applying the POD algorithm. Each following POD mode contains a part of the remaining oscillations around the mean value in the flow field. The total energy content in a flow field can be calculated by summation of squares of the singular values as

$$e_{\text{total}} = \sum_{j=1}^m \sigma_{jj}^2. \quad (41)$$

Several methods for finding reasonable truncation values for specific applications exist. One approach is to compare a single mode's energy content to the total energy in the system

$$e_{\text{rel},j} = \frac{e_j}{e_{\text{total}}} = \frac{\sigma_{j,j}^2}{\sum_{j=1}^m \sigma_{j,j}^2}. \quad (42)$$

By setting a fixed minimum value for $e_{\text{rel},j}$, modes with a contribution to the total energy of the system of less than $\varepsilon_{\text{POD, trunc}}$ can be truncated from the ROM representation. The modes to be kept must fulfill the criteria

$$e_{\text{rel},j} > \varepsilon_{\text{POD, trunc}}. \quad (43)$$

A good practice value of $\varepsilon_{\text{POD, trunc}}=10^{-5}$ has been proven to work well for high-Re cases for low-energy filtering applications of POD from parameter studies.

Another sensible criterion for singular value truncation, especially for ROM-oriented applications, is the summation of relative energy contributions starting from mode one to mode number r_{ROM} and comparing that sum to the total energy of the model. An ROM preserving a certain percentage of the total energy can then be generated

$$e_{\text{preserved}} \geq \frac{\sum_{j=1}^{r_{\text{ROM}}} \sigma_{j,j}^2}{\sum_{j=1}^m \sigma_{j,j}^2}. \quad (44)$$

Depending on the application, values of $e_{\text{preserved}} \approx 0.99$ are sensible. The choice of $e_{\text{preserved}}$ heavily depends on the energy contained in the mean flow mode. Sometimes, the mean flow mode already makes up over 90% of the total flow field energy. It is found that covering 90% of the fluctuation energy creates reasonable ROMs that contain most of the dominant unsteady flow structures (see also recommendations in Sirovich [59]). Since POD application is usually employed to investigate the temporally resolved flow field, mean-subtracted POD with $e_{\text{preserved}}=0.9$ is a reasonable practice. For visualization purposes, the mean flow can be added to the reconstructed flow field after computing the POD.

2.3.3 Incremental Data Processing

Due to the high complexity of turbulent flow fields around vehicles, an incremental approach that doesn't require the full snapshot matrix in memory is desirable, since memory limitations are otherwise quickly reached. Brand [20] introduces an algorithm to compute an SVD incrementally based on the method of snapshots and applies it to the example of video data processing. Later, Oxberry et al. [21] adopted his algorithm to apply it to CFD data analysis. This section outlines the incremental computation of the POD modes from an algorithmic point of view. A flowchart of the entire algorithm is sketched in Figure 6.

The first time step is normalized and used as the first column of the orthonormal basis \mathbf{U} . The singular value matrix $\mathbf{\Sigma}$ is initialized with the two-norm of the first snapshot and the right singular vector matrix \mathbf{V} is initialized with identity. After initialization of the POD decomposition, each new snapshot

vector is checked for relevance compared to the existing orthonormal basis. By projecting the vector onto the left singular vectors, one can judge the orthogonality of the current snapshot vector with respect to previously processed data. If the resulting orthogonality vector \mathbf{I} shows a small variance (many elements are close to zero), and if this variance is small compared to the variance of the snapshot itself, then the new snapshot vector contains relevant new data that is considered for a full update of the orthonormal basis (large value for p). In this case, the POD basis is expanded. If p is smaller than a user-defined threshold, p is set to zero, and the orthonormal basis is not extended. A matrix similar to an incremental covariance matrix \mathbf{Q} is defined by the existing singular value matrix, which is expanded by one rank with p on the diagonal and \mathbf{I} in the last column above the diagonal. The quadratic matrix \mathbf{Q} is then diagonalized using SVD. In case of a reduced update routine, the resulting \mathbf{U}_{temp} , $\mathbf{\Sigma}_{\text{temp}}$ and \mathbf{V}_{temp} are then sorted by singular value and truncated by one rank in both dimensions. In the case of a full update procedure, the incrementally computed left singular vector matrix is expanded using the new snapshot vector, which is orthonormalized with respect to the existing left singular vectors. The right singular vectors are expanded by a one-rank unity column and row. In any case, the left and right singular vector matrices are then updated by matrix product with \mathbf{U}_{temp} and \mathbf{V}_{temp} respectively.

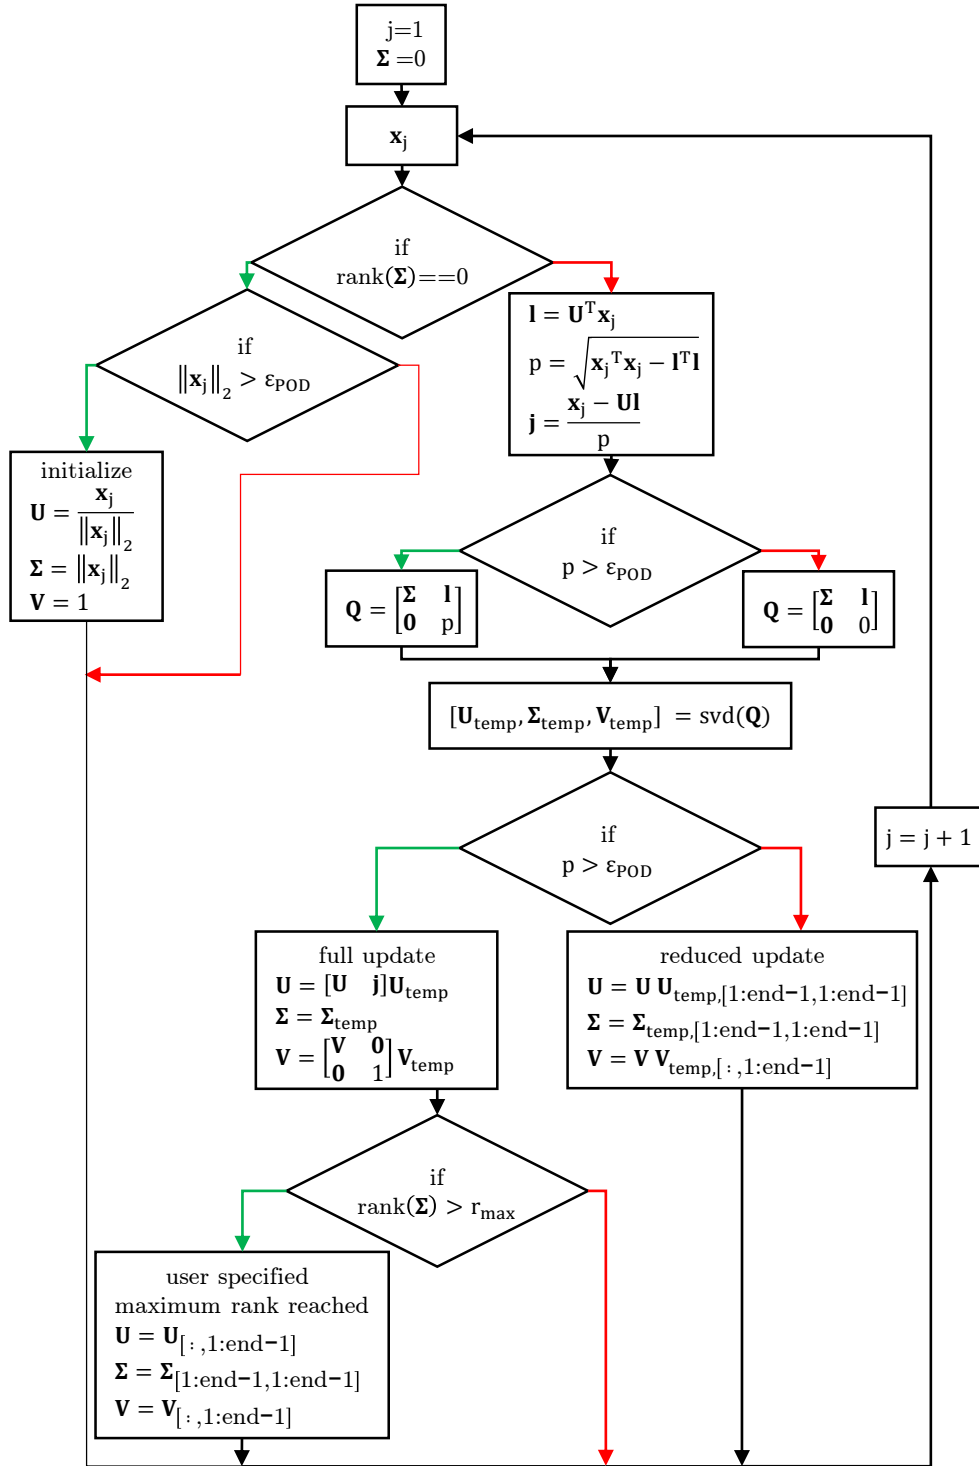


Figure 6: Flowchart of the incremental POD algorithm. Matrix indexing is done using MATLAB® notation. Green arrows denote a positive evaluation of an if statement, and red arrows indicate negative.

2.3.4 POD Results

A detailed comparison of conventional batch-processed POD and incremental POD for the two-dimensional cylinder in the periodic shedding cycle has been conducted by Langer [61] and will, therefore, be omitted from this presentation of incrementally computed POD analysis results of the two-dimensional cylinder case. The focus of the following section is on the algorithm performance concerning memory consumption and execution time.

The real-value matrix multiplication obtains the reconstruction of the snapshot matrix based on an incremental POD ROM following equation (40) with

$$\begin{aligned} & \tilde{\mathbf{X}}(\text{modeIDs})_{\text{observableIDs}, \text{timeIDs}} \\ &= \tilde{\mathbf{U}}_{\text{observableIDs}, \text{modeIDs}} \tilde{\mathbf{\Sigma}}_{\text{modeIDs}, \text{modeIDs}} (\tilde{\mathbf{V}}_{\text{timeIDs}, \text{modeIDs}})^T. \end{aligned} \quad (45)$$

The parameters `modeIDs`, `observableIDs` and `timeIDs` can be lists of indices to select specific modes, observables (i.e., cell indices if only one flow field quantity is used) and time steps for reconstruction. For ROM analysis and performance loss computation, `modeIDs=1:rROM`, `observableIDs=1:n` and `timeIDs=1:m`. For the analysis of a single physical flow phenomenon, modes can be reconstructed in combination with other modes or on their own, which can lead to more insight rather than observation of only the ensemble of all modes in an ROM. Because the modal basis only contains real values, a single mode can only contain the physical flow field phenomena happening in phase. Processes occurring with a phase offset will always require modes to be added to the modal basis. That can lead to disadvantages for the application on very complex flow fields. The reconstructed snapshot matrix generally becomes a better approximation of the original data snapshot matrix by increasing `rROM` if enough care is taken in adjusting the incremental POD computation, as will be highlighted in the discussion of tuning parameters below.

The first algorithm-tuning parameter discussed is a truncation criterion, which affects the expansion of the memory-wise most problematic orthonormal matrix \mathbf{U} . Significant memory and execution time savings can be obtained if expansion can be prevented without decreasing the quality of the resulting ROM. Due to the incremental nature of the algorithm, the expansion of the orthonormal basis to a high rank in the order of $O(n)$ (e.g., by using many snapshots and a very low threshold) can even lead to larger ROM modeling errors than if an increased singular value truncation criterion is used. That is mainly due to the propagation and increase of errors due to machine-precision errors in the involved matrix operations. The decay in ROM performance with a decreasing number of modes used for reconstruction is shown in Figure 7 for several threshold values. For each curve, the last data point marks the rank of the resulting ROM (e.g., `rROM=10` for `threshold=10-1` in the ROM of part two of the time series).




Figure 7: Performance loss charts for part one (top) and part two (bottom) of the two-dimensional cylinder test case for varying thresholds in the incremental POD algorithm.

The threshold criterion limits the number of modes in the orthonormal basis and leads to very compressed representations of the full data, even for lower ranks. An ROM with a rank of only 35 can reproduce the original thousand-time-step data series of part one with only 2.2% error using a threshold of 10^{-1} . Using only 10 modes for the same threshold in part two leads to an error of just 0.9%. For part one, decreasing the threshold to a final value of 10^{-6} also steadily decreases the ROM error. The most accurate approximation of the flow field is reached using an ROM with rank 1,000, leading to $\Pi_{\text{loss}}=3 \cdot 10^{-11}\%$ (not shown in Figure 7 for axis-scaling reasons). The same tendency can also be observed for the ROM analysis of part two down to a threshold of $\epsilon_{\text{POD}}=10^{-5}$. For a threshold of $\epsilon_{\text{POD}}=10^{-6}$, the purely periodic shedding nature of the flow eventually becomes a problem if adding additional time steps beyond \hat{t}_{1100} , when further time steps have to be orthogonalized. This is because the time steps still fulfill the threshold criterion but cannot be orthogonalized well due to the fact that two full shedding periods already have passed. In the process of the complete update routine and orthogonalization, numerical double precision errors add up, and a poorly orthogonalized basis results. An additional orthogonality check after adding a new time step could potentially prevent such errors but would require too much computation time. Keeping that in mind, it is legitimate to require a threshold criterion of $\epsilon_{\text{POD}}>0$ for all applications with the given formulation of the update procedure. A reasonable, well-chosen threshold criterion can then also limit the required memory to compute the ROM, which scales with the rank of the

resulting orthonormal basis. Using a threshold value of $\varepsilon_{\text{POD}}=10^{-4}$ is a good tradeoff between ROM precision and memory usage for both time series, with errors of only $\Pi_{\text{loss}}=2\cdot 10^{-3}\%$ in part one and $\Pi_{\text{loss}}=6\cdot 10^{-4}\%$ in part two. That level of accuracy equates to memory savings with a factor of almost three in part one and of 13 for part two, compared to a threshold value of $\varepsilon_{\text{POD}}=10^{-6}$. Interestingly, all ROM executions lead to similar performance losses for small ranks (i.e., the error curve of ROM with threshold $\varepsilon_{\text{POD}}=10^{-a}$ equals the error curve of ROM with a threshold $\varepsilon_{\text{POD}}=10^{-(a+1)}$ when only using $r_{\text{ROM}}\approx 1/3r_{\text{max}}$). Even the threshold $\varepsilon_{\text{POD}}=10^{-1}$ case with $r_{\text{max}}=10$ can capture the first three modes perfectly fine and can reproduce modes four and five, at least concerning mode distribution. Thus, the computational cost can be reduced drastically by extracting only dominant modes instead of building a precise ROM of the full state data. Beyond the need for significant memory when using current ROM algorithms in engineering applications, the requirement for high computational time can also be a limiting factor. Figure 8 shows the computational time for the evaluation of the ROM for different threshold values. The threshold value of $\varepsilon_{\text{POD}}=10^{-4}$ offers a good compromise between execution time and ROM performance.

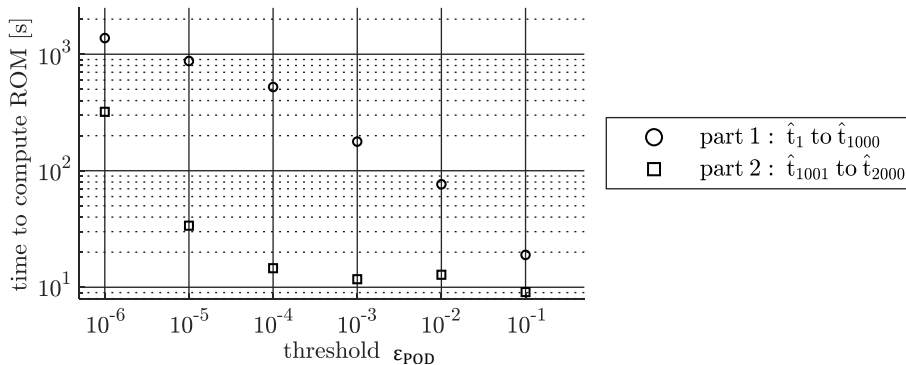


Figure 8: ROM execution time for incremental POD analysis. The computation time increases with the rank of the resulting ROM.

Besides the above thresholding criterion, which can be set independently of the computer system requirements, it is also possible to apply a fixed maximum rank for the creation of an ROM by conventional singular value truncation as in batch-processed POD; however, this occurs at the end of each full update iteration. A fixed maximum rank employed at the beginning of the update routine is not advisable (i.e., where the above threshold criterion is used), because highly relevant snapshot information could be neglected. If the process investigated contains a ramp-up phase or linear growth (e.g., in the two-dimensional cylinder part one time series), snapshots that are added later might be more relevant than snapshots that are already represented in the orthonormal basis. A fixed maximum-rank criterion should thus only be implemented in a singular value truncation after updating the POD basis, even if it increases execution time. That type of criterion can then be tuned to adjust the algorithm

to comply with maximum available memory or to limit the execution time of the algorithm.

While the above tuning parameters must be used universally, the definition of a fixed maximum rank for the creation of an incremental POD basis should be adjusted to account for available computer resources. Due to the large number of observables and snapshots for the application to high-Re production vehicle aerodynamics, incremental POD offers a good workaround when computer system memory limitations are an issue. Due to the problematic reorthonormalization of snapshots in part two after only two shedding cycles (i.e., around 200 snapshots), only results for part one are presented here, as shown in Figure 9. Relatively small deviations are observed for all rank limitations compared to the full incremental POD without the application of any truncation. The performance loss curve of the full incremental POD with a rank of 1,000 is pictured in Figure 7. The first few modes are found well even for minimal maximum-rank settings. A limitation of the POD basis to 10% of the full rank orthonormal basis size with a maximum rank of 100 leads to, at most, a 0.003% performance loss increase compared to an ROM built from the first 100 modes of the full incremental POD. The full incremental POD comes with a number of 1,000 modes in the update routine of the last time step during the incremental POD basis computation and thus also requires 10 times the memory. The performance loss of an ROM with rank 100 for part one is only 0.0175%, computed with a maximum rank setting of 100. Similar ROM performance charts can also be obtained for lower ranks of decompositions of time series one.




Figure 9: Incremental POD performance loss behavior for different maximum-rank truncation parameters for part one. Maximum-rank truncation is conducted using singular value truncation after the full update procedure in the algorithm.

The restriction of the POD basis size does not influence the ROM performance much but does come with a downside. Due to the recurring full update procedure, even for time steps that would be sorted out by a wisely chosen threshold criterion, much computational effort needs to be spent to generate a small-rank POD basis that keeps the memory consumption in a certain limit. A maximum-rank incremental POD with $r_{\text{max}}=100$ takes approximately twice the execution

time as the incremental POD with a threshold criterion of 10^{-2} , leading to an ROM with rank 107 and a high performance loss of $\Pi_{\text{loss}}=0.201\%$. If only a maximum-rank truncation is used, preallocation of memory leads to a further decrease in the execution time for higher ranks. If execution time is not of concern but memory limitations play a role, maximum-rank singular value truncation should always be chosen over a threshold criterion due to better precision in generating ROMs with the same amount of modes. If only the first few leading modes are to be investigated, the threshold criterion offers a good way of drastically decreasing computation time, compared to the maximum-rank criterion. While the resulting ROM can deviate from the FOM representation, leading modes are still in good agreement with the full-size incremental POD modes.




Figure 10: Incremental POD execution time for part one, using the maximum-rank singular value truncation after the update procedure. All execution times are in similar orders of magnitude.

The POD results for a threshold value of 10^{-4} are shown in the following section. The left chart in Figure 11 shows the singular values for both parts of the test case. It quickly becomes evident that many more modes are required to capture the instability growth process in part one and the decrease in singular values is much slower than in part two. That tendency is also reflected in the ROM performance curve in the right chart of Figure 11. Because POD modes are optimal in a least-squares sense, this is not surprising at all. Addition of POD modes to an ROM leads to a better approximation of the FOM.




Figure 11: Singular values of incrementally computed POD decompositions (left) and ROM performance loss curves (right) for part one and part two of the two-dimensional cylinder test case. Mode pairs for the first few shedding modes can be identified in the singular values.

An ROM with $\Pi_{\text{loss}} < 1\%$ can be obtained for both parts of the simulation with a relatively small rank. For part one, $\Pi_{\text{loss}}(r_{\text{ROM}}=18)=0.99\%$, and for part two, $\Pi_{\text{loss}}(r_{\text{ROM}}=7)=0.58\%$. The mean flow performance loss (computed by using the temporally averaged flow field as an approximation) is almost identical to the performance loss of the first POD mode (also referred to as mean flow mode) for both parts. The increase in oscillation energy in the ramp-up phase of part one leads to a slightly bigger contribution of the mean flow to the total energy in the system with $\Pi_{\text{loss,meanFlow}}=16.14\%$ and $\Pi_{\text{loss,meanFlow}}=16.84\%$ in part two. Depending on the exact volume that is analyzed and the amount of energy contained in the fluctuations around a bluff body, the mean flow mode usually contains the most energy of all modes. Confining the investigated fluid volume to the wake or recirculation region right behind a bluff body leads to smaller mean flow mode contributions. The appearance of clear mode pairs in the singular value spectrum is a hint to periodic processes appearing in the flow field. By reconstructing the first mode pair in combination with the mean flow mode, the cyclic shedding behavior of the cylinder can already be modeled well, with $\Pi_{\text{loss}}(r_{\text{ROM}}=3)=7.3\%$ in part one. This can be seen in the time series of a probe and the comparison of the y-components of the original and reconstructed velocity fields in Figure 12. The reconstructed time series in the cylinder wake shows a large relative offset compared to the original data series in the ramp-up phase of part one, but the main flow characteristics are captured well.

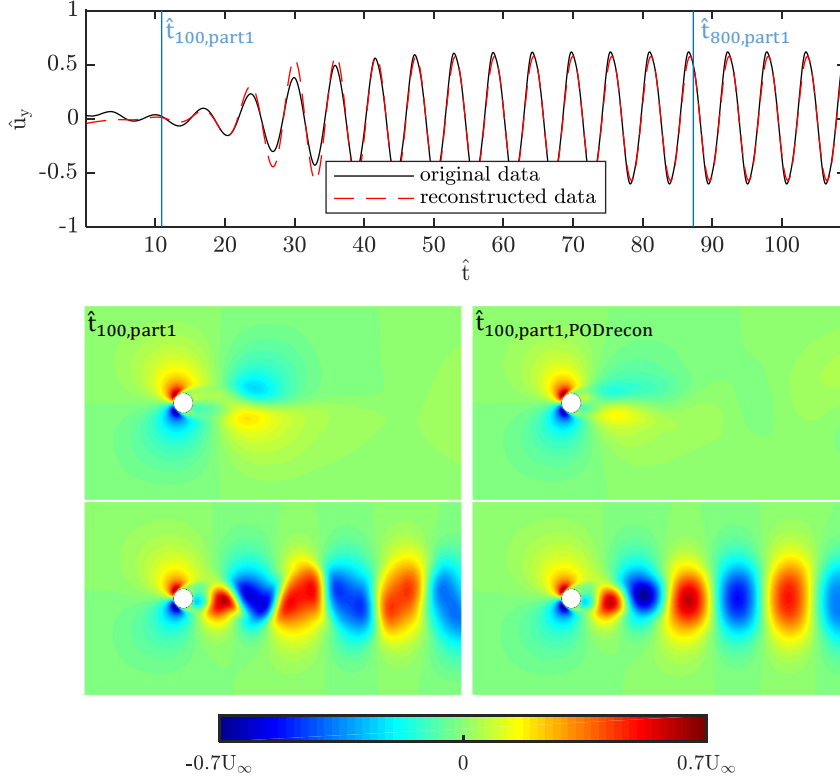


Figure 12: Temporal evolution of the velocity y-component probed at position $x=2D$ and $y=0$ in part one of the simulation. Snapshots $\hat{t}_{100}=10.91$ and $\hat{t}_{800}=87.28$ of the original data (left) and reconstructed data (right) are marked with a blue line in the upper graph. POD modes one, two and three are used for the reconstruction (mean flow mode and first mode pair).

Location and shape of the dominating structures in the cylinder wake are slightly shifted at \hat{t}_{100} . Additional POD modes are necessary to capture better the processes in the ramp-up phase, but the compression performance of the algorithm is still outstanding. The fluctuations contain 16.14% of the systems' kinetic energy, of which 8.84% can be represented using just the first mode pair. The dominating flow structures of 1,000 time steps in the flow field data set can be compressed to three real-valued POD modes that are computed incrementally. After the ramp-up phase, the structures of the first three reconstructed POD modes capture the location, shape and magnitude of the disturbances in the y-direction in very close agreement.

For part two of the simulation, the reconstruction of the first three POD modes even leads to $\Pi_{\text{loss}}(\Gamma_{\text{ROM}}=3)=3.6\%$. A probe sampling of the reconstructed flow, as shown in Figure 13, reveals a very close agreement of the velocity field y-component signal. A slight time delay of the reconstructed signal is observable but does not hinder the extraction and tracking of the dominant coherent structures in the flow field.

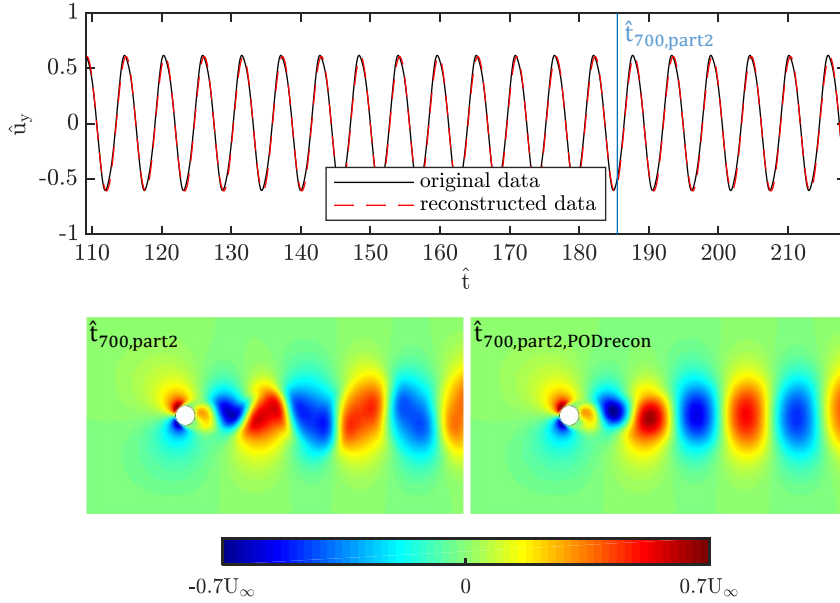


Figure 13: Temporal evolution of the velocity y-component probed at position $x=2D$ and $y=0$ in part two of the simulation. Snapshot $\hat{t}_{700}=185.5$ of the original data (left) and reconstructed data (right) is marked with a blue line in the upper graph. POD modes one, two and three are used for the reconstruction (mean flow mode and first mode pair).

Due to the real-valued nature of POD modes, the entire temporal evolution information of a mode is contained in the related right singular vector. Therefore, application of FFT on the right singular vector is equal to the same analysis applied to all cells of the temporally reconstructed flow field. A comparison between the frequency spectra of the original velocity field and the first POD mode pair for part two of the simulation shows how the first POD mode matches the Strouhal number of the dominant shedding process exactly, as can be seen in Figure 14. The modes contain wideband frequency data.

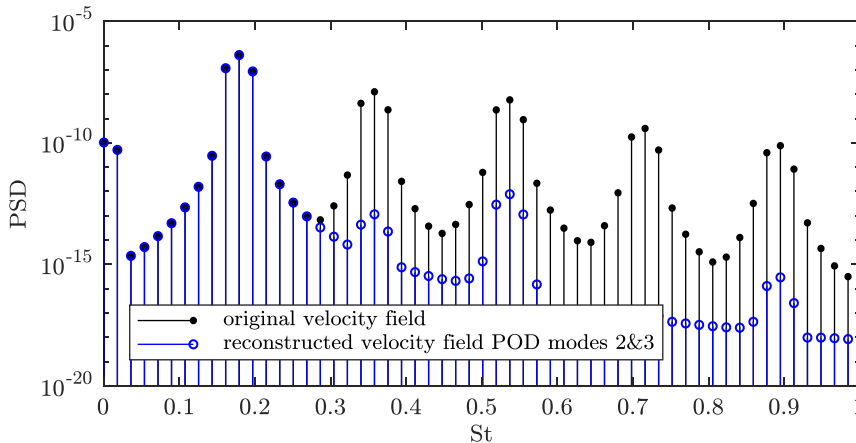


Figure 14: Mean FFT spectra of all cells of the original versus the reconstructed velocity field of part two of the simulation. Reconstruction is done using incrementally computed POD modes two and three (i.e., the first oscillating mode pair).

While POD is very robust and inexpensive for generating ROMs with an optimally small amount of modes, it can also be used to investigate dominant flow features, mode by mode, for simple flow fields. Most of the fluctuation energy is contained in one shedding process, so the first POD modes, which are computed optimally in the least-squares sense, can accurately represent this single flow field process.

2.4 Low-Memory Dynamic Mode Decomposition

Similar to POD, DMD is a data-driven modal analysis method that is found to work exceptionally well for diagnostic purposes in bluff body aerodynamics. DMD modes approximate Koopman modes using a linear model. Upon successful computation of the approximated Koopman modes, one can then analyze the properties of a linear operator instead of relying on the approximated solution of the discretized NSE. In addition to DMD's strengths for periodic or quasi-periodic structure detection and linear growth phenomena, it can represent non-linear flow phenomena using combinations of multiple linear modes in some cases. Correct approximation of modes relies on a smart choice of a basis function of observables. DMD produces useful results in the case of selecting the flow field components as observables, without any higher order combinations. This section presents the underlying fundamental concepts of DMD and shows the evaluation of DMD modes from an algorithmic point of view. That perspective is chosen to demonstrate reasoning behind the specific steps in the algorithm concerning computational efficiency, which is critical to be able to evaluate DMD modes for large data sets in an industrial environment.

A fundamental idea of this modal decomposition approach is that any time step (snapshot) of an instationary flow field can be approximated using a linear combination of previous time steps if a sufficient amount of time steps is used, as shown here:

$$\mathbf{x}_m \approx a_1 \mathbf{x}_1 + a_2 \mathbf{x}_2 + \dots + a_{m-1} \mathbf{x}_{m-1} . \quad (46)$$

A snapshot vector \mathbf{x}_j is formed by the vertical concatenation of the three velocity components in a column vector. Given equation (46), a linear operator \mathbf{A} can be defined that maps many consecutive snapshots to their respective following snapshot:

$$\mathbf{A} \mathbf{X}_1^{m-1} = \mathbf{X}_2^m . \quad (47)$$

The temporal dynamics of the flow field can now be analyzed using eigenvalue analysis of \mathbf{A} . Many different ways of evaluating the properties of \mathbf{A} have been proposed. In the following, a short summary of the relevant algorithms is presented. These algorithms are used to formulate a modified online variant of the existing methods, which allows for an incremental processing of snapshots. This facilitates the introduction of a new way of computing DMD modes incrementally and sorting them by a suitable criterion for bluff body aerodynamics applications.

2.4.1 Batch-Processed DMD

One way of computing \mathbf{A} and its eigenvalues would be to compute the inverse of $\mathbf{A}\mathbf{X}_1^{m-1}=\mathbf{X}_2^m$ and obtain

$$\mathbf{A} = \mathbf{X}_2^m (\mathbf{X}_1^{m-1})^\dagger. \quad (48)$$

Since the computation of this inverse is not feasible in the case of high-dimensional snapshots, the matrix is decomposed using SVD instead:

$$\mathbf{X}_1^{m-1} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T. \quad (49)$$

The SVD returns POD modes, as in equation (33), which represent coherent structures that are orthogonal in space. As for the process of computing POD modes, using the method of snapshots leads to considerably lower computational cost and memory needs. POD modes are columns of \mathbf{U} sorted by their respective singular value in diagonal matrix $\mathbf{\Sigma}$ accordingly (the right singular vectors \mathbf{V} express the temporal evolution of the modes). For the application to very high dimensional data with low-order discretization errors and local Courant numbers larger than one, numerical noise is expected and disturbs the simulated physical flow field signal. Because it was found that this erroneous noise occurs in the high-frequency range and contains low energy fluctuations, singular value truncation of the SVD result shown in equation (49) and as described in equation (40) can and should be applied. This step is also required if physical frequencies that cannot be captured by DMD are contained in the data. Such situations occur if only a selection of time steps are sampled (e.g., every hundredth time step) when concentrating on low-frequency oscillations. As Duke et al. mention in their work, it is difficult to distinguish between numerical noise and high frequency flow field oscillations:

$$\mathbf{X}_1^{m-1} \approx \tilde{\mathbf{U}}\tilde{\mathbf{\Sigma}}\tilde{\mathbf{V}}^T. \quad (50)$$

This step is indispensable for successful evaluation of the DMD modes, which was also found by Peichl et al. [50]. They describe singular value truncation as an essential step in their algorithm as a low-pass filter for removing high-frequency oscillations, which cannot be captured by DMD. Further description of a suitable time step sampling procedure and data pre-processing conditions in case of online modal analysis for vehicle aerodynamics will follow in section 3.3. Inserting into equation (48) and rearranging, the projected DMD operator \mathbf{A}_p can be computed as the original DMD operator \mathbf{A} mapped onto the POD basis of \mathbf{X}_1^{m-1} as

$$\mathbf{A}_p = \mathbf{U}^T \mathbf{A} \mathbf{U} = \mathbf{U}^T \mathbf{X}_2^m \mathbf{V} \mathbf{\Sigma}^{-1}, \quad (51)$$

or

$$\mathbf{A}_p \approx \tilde{\mathbf{U}}^T \mathbf{X}_2^m \tilde{\mathbf{V}} \tilde{\mathbf{\Sigma}}^{-1} \quad (52)$$

respectively for the singular value truncated SVD. The eigenvalues of \mathbf{A}_p approximate the eigenvalues of \mathbf{A} , so the eigenvectors \mathbf{y}_k are evaluated by solving the eigenvalue problem

$$\mathbf{A}_p \mathbf{y}_k = \lambda_i \mathbf{y}_k. \quad (53)$$

Finally, the DMD modes are evaluated as

$$\boldsymbol{\varphi}_k = \mathbf{U}\mathbf{y}_k. \quad (54)$$

That is the most straightforward way of computing DMD modes, which has been applied successfully for a range of applications and is commonly referred to as conventional DMD or Schmid DMD. The computationally most expensive step in this method is the evaluation of the truncated SVD in equation (50). The memory requirement of conventional DMD is in the order of two times the snapshot matrix size.

A modification of this approach is presented in the formulation of a new method by Hemati et al. [36]. TDMD makes use of augmented snapshots that are composed of two consecutive snapshots:

$$\mathbf{z}_j = \begin{bmatrix} \mathbf{x}_j \\ \mathbf{x}_{j+1} \end{bmatrix}. \quad (55)$$

The method takes advantage of a noise-removing pre-processing step via an SVD of the augmented snapshot matrix

$$\mathbf{Z} = \begin{bmatrix} \mathbf{X}_1^{m-1} \\ \mathbf{X}_2^m \end{bmatrix} = \mathbf{U}_Z \boldsymbol{\Sigma}_Z \mathbf{V}_Z^T. \quad (56)$$

As is the case with conventional DMD, truncating this decomposition by removing small energy modes filters noise from the augmented snapshot matrix:

$$\mathbf{Z} \approx \tilde{\mathbf{Z}} = \tilde{\mathbf{U}}_Z \tilde{\boldsymbol{\Sigma}}_Z \tilde{\mathbf{V}}_Z^T. \quad (57)$$

The SVD in equation (56) can be evaluated using the method of snapshots as described in section 2.4.1 for considerable computation time and memory savings. Using the augmented snapshot matrix in this pre-processing step and applying singular value truncation, a de-biased snapshot matrix can be computed as

$$\tilde{\mathbf{X}}_1^{m-1} = \mathbf{X}_1^{m-1} \tilde{\mathbf{V}}_Z \tilde{\mathbf{V}}_Z^T \quad (58)$$

and

$$\tilde{\mathbf{X}}_2^m = \mathbf{X}_2^m \tilde{\mathbf{V}}_Z \tilde{\mathbf{V}}_Z^T \quad (59)$$

accordingly. The explicit evaluation of $\tilde{\mathbf{U}}_Z$ is not required in this approach. Only $\boldsymbol{\Sigma}_Z$ is needed for singular value truncation and truncated $\tilde{\mathbf{V}}_Z$ in equation (58), which means a massive saving in memory. After the de-biased snapshot matrix is evaluated, the conventional DMD process is executed with $\tilde{\mathbf{X}}_1^{m-1}$ and $\tilde{\mathbf{X}}_2^m$ instead of \mathbf{X}_1^{m-1} and \mathbf{X}_2^m following equations (50) to (54). The de-biasing effect of the augmented snapshots procedure and its resulting benefits is demonstrated for noisy experimental PIV data in the original paper by Hemati et al. [36]. A direct comparison to conventional DMD analysis shows much better structure identification and better approximation of eigenvalues, compared to conventional DMD in their work.

2.4.2 Streaming Total Dynamic Mode Decomposition

As is the case for POD, the application of DMD to high-dimensional data sets is limited due to memory constraints. In case of vehicle aerodynamics, hard drive disk writing times and storage costs also are of concern for the application. A modal analysis tool that can be applied online during the simulation (i.e., making use of the data in memory) with minimal sensitivity to numerical noise is desired.

The following description of an online variant of DMD that incorporates the ideas from TDMD is outlined and extended. This work suggests a new way of evaluating the DMD modes in a more time- and memory-efficient way than previously published methods. The properties of those extensions compared to the existing and already published algorithm are highlighted.

The STDMD algorithm is based on the methods shown in Hemati et al.'s work on TDMD [36], SDMD [35], and STDMD [37]. The core of the algorithm consists of the incremental computation of an orthonormal basis, similar to the process in computing incremental POD modes but using the classical Gram-Schmidt iteration (CGSI) process for the generation of an orthonormal basis from augmented snapshots \mathbf{z}_j . Instead of an SVD, STDMD employs a QR-decomposition of augmented snapshot matrix \mathbf{Z} as

$$\mathbf{Z} = \mathbf{Q}_Z \mathbf{R}_Z. \quad (60)$$

In this section, matrix \mathbf{X}_1^{m-1} is used as the first part of the snapshot matrix containing the first $m-1$ snapshots, and \mathbf{X}_2^m contains the last $m-1$ snapshots. The upper part of \mathbf{Z} can also be represented as

$$\mathbf{X}_1^{m-1} = [\mathbf{I} \quad \mathbf{0}] \mathbf{Q}_Z \mathbf{R}_Z \quad (61)$$

$$(\mathbf{X}_1^{m-1})^T = \mathbf{R}_Z^T \mathbf{Q}_Z^T \begin{bmatrix} \mathbf{I} \\ \mathbf{0} \end{bmatrix} \quad (62)$$

and the lower part as

$$\mathbf{X}_2^m = [\mathbf{0} \quad \mathbf{I}] \mathbf{Q}_Z \mathbf{R}_Z, \quad (63)$$

as given in equation (56) above. Because \mathbf{R}_Z is not explicitly required for the evaluation of DMD modes, only

$$\mathbf{G}_Z = \mathbf{R}_Z \mathbf{R}_Z^T \quad (64)$$

is evaluated during the incremental updates. Similar to the derivation of batch-processed DMD in 2.4.1, the expression for the evaluation of the DMD operator from equation (48) is the starting point for the derivation

$$\mathbf{A} = \mathbf{X}_2^m (\mathbf{X}_1^{m-1})^\dagger. \quad (65)$$

Using the identity

$$\mathbf{X}^\dagger = \mathbf{X}^T (\mathbf{X}_1^{m-1} (\mathbf{X}_1^{m-1})^T)^{-1}, \quad (66)$$

equation

$$\mathbf{A} = \mathbf{X}_2^m (\mathbf{X}_1^{m-1})^T (\mathbf{X}_1^{m-1} (\mathbf{X}_1^{m-1})^T)^{-1} \quad (67)$$

is obtained. To arrive at a valid definition of an STDMD operator, the definition of another QR-decomposition

$$\mathbf{Q}_X \mathbf{R}_X = [\mathbf{I} \quad \mathbf{0}] \mathbf{Q}_Z, \quad (68)$$

is needed. \mathbf{Q}_X is a more noise-insensitive representation of the image of the snapshot matrix. It is worth noting that the above QR-decomposition should not be mistaken for the QR-decomposition of snapshot matrix \mathbf{X}_1^{m-1} . Rearranging equation (61) for \mathbf{Q}_Z and inserting into equation (68) leads to

$$\mathbf{Q}_X \mathbf{R}_X = \mathbf{X}_1^{m-1} \mathbf{R}_Z^{-1}. \quad (69)$$

After simplification and rearrangement for the snapshot matrix \mathbf{X}_1^{m-1} ,

$$\mathbf{X}_1^{m-1} = \mathbf{Q}_X \mathbf{R}_X \mathbf{R}_Z \quad (70)$$

$$(\mathbf{X}_1^{m-1})^T = \mathbf{R}_Z^T \mathbf{R}_X^T \mathbf{Q}_X^T \quad (71)$$

is obtained.

Finally, the STDMD operator is found by substituting all matrices on the right-hand side of equation (67) with expressions (63), (62), (70) and (71) as

$$\mathbf{A} = [\mathbf{0} \quad \mathbf{I}] \mathbf{Q}_Z \mathbf{R}_Z \mathbf{R}_Z^T \mathbf{Q}_Z^T \begin{bmatrix} \mathbf{I} \\ \mathbf{0} \end{bmatrix} (\mathbf{Q}_X \mathbf{R}_X \mathbf{R}_Z \mathbf{R}_Z^T \mathbf{R}_X^T \mathbf{Q}_X^T)^{-1}. \quad (72)$$

Another definition is the combination of terms on the right-hand side using definition (64)

$$\mathbf{R}_X \mathbf{R}_Z \mathbf{R}_Z^T \mathbf{R}_X^T = \mathbf{R}_X \mathbf{G}_Z \mathbf{R}_X^T = \mathbf{G}_X. \quad (73)$$

After simplification, the projected STDMD operator is then given as its projection onto the orthonormal basis \mathbf{Q}_X with

$$\mathbf{A}_p = \mathbf{Q}_X^T \mathbf{A} \mathbf{Q}_X = \mathbf{Q}_X^T [\mathbf{0} \quad \mathbf{I}] \mathbf{Q}_Z \mathbf{G}_Z \mathbf{Q}_Z^T \begin{bmatrix} \mathbf{I} \\ \mathbf{0} \end{bmatrix} \mathbf{Q}_X \mathbf{G}_X^{-1}. \quad (74)$$

As is the case with other variants of DMD, projecting the DMD operator onto the orthonormal basis representation of the snapshot matrix turns out to be computationally more manageable due to the lower dimensionality of the orthonormal basis in case of singular value truncation or online orthonormal basis compression respectively. The STDMD complex number eigenvectors and eigenvalues are again computed using the eigendecomposition of \mathbf{A}_p , as in equation (53) for batch-processed DMD variants

$$\begin{aligned} \mathbf{A}_p \mathbf{y}_k &= \lambda_k \mathbf{y}_k \\ \mathbf{A}_p \mathbf{Y} &= \Lambda \mathbf{Y}. \end{aligned} \quad (75)$$

For the evaluation of the complex-valued STDMD modes, the eigenvectors of \mathbf{A}_p are then applied onto the orthonormal basis \mathbf{Q}_X

$$\boldsymbol{\varphi}_k = \mathbf{Q}_X \mathbf{y}_k. \quad (76)$$

Equation (76) is known from the Hemati et al. work on STDMD and produces excellent results, comparable to batch-processed TDMD.

To further reduce the amount of memory required for the evaluation of DMD modes, this work suggests a new formulation of the projected DMD operator. By replacing \mathbf{Q}_X with rearranged equation (68) with

$$\mathbf{Q}_X = [\mathbf{I} \quad \mathbf{0}] \mathbf{Q}_Z \mathbf{R}_X^{-1}, \quad (77)$$

the following equation for the projected STDMD operator can be obtained

$$\mathbf{A}_p = ([\mathbf{I} \quad \mathbf{0}] \mathbf{Q}_Z \mathbf{R}_X^{-1})^T [\mathbf{0} \quad \mathbf{I}] \mathbf{Q}_Z \mathbf{G}_Z \mathbf{Q}_Z^T \begin{bmatrix} \mathbf{I} \\ \mathbf{0} \end{bmatrix} ([\mathbf{I} \quad \mathbf{0}] \mathbf{Q}_Z \mathbf{R}_X^{-1}) \mathbf{G}_X^{-1}. \quad (78)$$

Using the same replacement for equation (76), results in

$$\boldsymbol{\varphi}_k = [\mathbf{I} \quad \mathbf{0}] \mathbf{Q}_Z \mathbf{R}_X^{-1} \mathbf{y}_k. \quad (79)$$

$$\boldsymbol{\Phi} = [\mathbf{I} \quad \mathbf{0}] \mathbf{Q}_Z \mathbf{R}_X^{-1} \mathbf{Y}. \quad (80)$$

This new formulation reduces the total amount of memory required because \mathbf{Q}_X does not need to be computed at all in the QR-decomposition of the upper part of \mathbf{Q}_Z (equation (68)) for the computation of \mathbf{A}_p and $\boldsymbol{\varphi}_k$. Because \mathbf{Q}_X is not explicitly required, a QR-decomposition implementation that does not evaluate and keep all vectors of \mathbf{Q}_X in memory should be used. As will be shown in section 2.4.4, the new formulations in equation (78) and equation (79) also have substantial advantages in the evaluation of amplitudes.

2.4.3 Frequencies, Amplification and Damping Rates

Because the eigenvalues and eigenvectors appear in pairs of negative and positive frequency, the negative frequency counterpart for each positive frequency mode can be obtained by taking the complex conjugate of the positive frequency mode. The frequency of each DMD mode is computed by

$$f_k = \frac{\text{Im}[\log \lambda_k]}{2\pi \Delta t}. \quad (81)$$

The sign of the respective frequency is the same as the sign of the imaginary part of the eigenvalue. Δt represents the fixed time step size between snapshot \mathbf{x}_j and \mathbf{x}_{j+1} . It is found that all runs of STDMD produce identical mode distributions for all frequency pairs down to numerical precision. Therefore, computation of all $m-1$ STDMD modes is not required, since almost half of the modes appear redundant, depending on how many zero Hz mean flow modes are found, even for full rank DMD models. The amount of storage space required to save the full order representation using a modal basis of DMD modes is thus similar to the space needed for a POD. This is because while POD modes are real numbers only, DMD modes are complex numbers (i.e., one DMD mode takes as much storage as two POD modes). In case of a single mean flow mode at zero Hz, the full order STDMD model is represented by $(m-1)/2+1$ complex number modes, if no orthonormal basis compression is applied during STDMD execution. Concerning memory consumption, this is in the same order as the original flow field data with m real number snapshots. It is assumed here that each, the real and imaginary part of the STDMD modes, is stored in the same data type as the real number flow field snapshots. From equation (79) it becomes evident that the memory required for this step can be limited to the size of $[\mathbf{I} \ \mathbf{0}] \mathbf{Q}_z$ (i.e., the upper half of \mathbf{Q}_z) if the STDMD modes are evaluated incrementally, because usually $n \gg m$, so the complex number vectors \mathbf{y}_k are of negligible memory concern.

The magnitude of the complex eigenvalue represents damping and amplification rates of DMD modes. Highly damped modes have eigenvalue magnitudes smaller than one, and amplified modes have eigenvalues greater than one. Stable modes with continuous oscillation behavior come with eigenvalue magnitudes of one, as shown in equations (83). The possibility of having modes that are amplified or damped gives DMD the possibility to represent unsteady flow field effects, like linear growth, exactly with ROMs.

$$\begin{aligned} |\lambda_k| < 1 &\rightarrow \text{damped mode} \\ |\lambda_k| = 1 &\rightarrow \text{stable mode} \\ |\lambda_k| > 1 &\rightarrow \text{amplified mode} \end{aligned} \quad (82)$$

2.4.4 Reconstruction and Amplitudes

As in POD, the DMD modes can be used to generate ROMs. Modes can be reconstructed using the Vandermonde matrix of complex eigenvalues, describing the temporal evolution of modes

$$\mathbf{T} = \begin{bmatrix} \lambda_1^0 & \lambda_1^1 & \dots & \lambda_1^{m-1} \\ \lambda_2^0 & \lambda_2^1 & \dots & \lambda_2^{m-1} \\ \vdots & \vdots & \ddots & \vdots \\ \lambda_r^0 & \lambda_r^1 & \dots & \lambda_r^{m-1} \end{bmatrix}. \quad (83)$$

For approaches using augmented snapshots, the value of m needs to be adjusted to the number of augmented snapshots, which is one snapshot less than conventional DMD analysis on the same set of snapshots. Multiplication of the Vandermonde matrix with the DMD modes and the diagonal matrix of amplitudes \mathbf{D}_α leads to the approximated reconstruction of the snapshot matrix

$$\tilde{\mathbf{X}} = \Phi \mathbf{D}_\alpha \mathbf{T}. \quad (84)$$

Evaluation of this expression for the first time step, using only the first column of \mathbf{T} , leads to $\lambda_k^0=1$ and

$$\mathbf{x}_1 = \Phi \boldsymbol{\alpha}. \quad (85)$$

As for conventional DMD, the STDMD modes' amplitudes can then be evaluated using the method of first snapshot as

$$\boldsymbol{\alpha} = \Phi \setminus \mathbf{x}_1. \quad (86)$$

The complex number amplitudes can be used to scale the unit length STDMD modes accordingly. Other advanced amplitude computation methods exist for the application of results from batch-processed DMD approaches or for wherever the full singular values and right-singular POD vectors of equation (33) are available. In case the mode distribution of a specific mode is to be visualized, its respective negative frequency counterpart is always to be reconstructed with it to obtain the right flow field magnitudes. Due to the redundancy of information, only the number of positive frequency modes is accounted for in this work's discussion of mode selection strategies and the creation of an ROM.

The real part of the reconstructed flow field is expressed in the same units as the observables from the snapshot matrix. That means a reconstructed snapshot from a dominating DMD mode is likely to be visible in the original flow field snapshot if no overlapping modes are present in that specific time step that extinguish the disturbance from the observed mode. In case a specific DMD mode is to be analyzed individually, a reconstruction of an entire oscillation period of the mode of interest can be extremely helpful. For reconstructed flow field visualization the Vandermonde matrix does not necessarily require integer number exponents for the evolution of the eigenvalues. Integer number exponents are crucial only for the comparison to the original snapshot matrix. For visualization on the other side, floating point numbers can be used as exponents. That technique returns much smoother reconstructions, especially in case of modes with frequencies close to the maximum resolvable frequency. Because the entire DMD mode matrix Φ inverse needs to be computed for the evaluation of the mode amplitudes, this step can be quite computationally intensive. This is because the inversion of such a large complex number matrix consumes considerable CPU time and because the entire DMD mode matrix needs to be evaluated and be present in memory at once, with a RAM contribution in the

same order as the orthonormal basis \mathbf{Q}_z . That is undesirable because not all modes are required at once for applications like dominant flow structure detection, frequency-band analysis or generation of an ROM. Knowing the mode amplitudes and frequencies before computing the modes themselves could avoid those problems and make selective and incremental STDMD mode output possible. Mode output could occur by writing to HDD or by passing the modes to another in-memory post-processing tool for visualization of the modal analysis results without the need for expensive and time-consuming HDD writing tasks. For this matter, equation (86) can be rewritten as

$$\boldsymbol{\alpha} = (\mathbf{Q}_x \mathbf{Y})^{-1} \mathbf{x}_1 \quad (87)$$

$$\boldsymbol{\alpha} = (\mathbf{Y}^{-1} \mathbf{Q}_x^{-1}) \mathbf{x}_1. \quad (88)$$

Therefore, using the inversion identity for orthonormal matrices

$$\mathbf{Q}_x^{-1} = \mathbf{Q}_x^T = ([\mathbf{I} \quad \mathbf{0}] \mathbf{Q}_z \mathbf{R}_x^{-1})^T = (\mathbf{R}_x^{-1})^T \mathbf{Q}_z^T \begin{bmatrix} \mathbf{I} \\ \mathbf{0} \end{bmatrix}, \quad (89)$$

the mode amplitudes can be computed as

$$\boldsymbol{\alpha} = \mathbf{Y}^{-1} (\mathbf{R}_x^{-1})^T \mathbf{Q}_z^T \begin{bmatrix} \mathbf{I} \\ \mathbf{0} \end{bmatrix} \mathbf{x}_1. \quad (90)$$

Through this formulation for the mode amplitudes, it is possible to evaluate the amplitudes before computing any of the STDMD modes. Selective mode computation based on amplitude-related selection criteria is now possible. Ordering modes by their physical dominance is critical to be able to extract dominant structures. Considerations about sensible sorting strategies are presented in section 2.4.6.

2.4.5 Algorithmic Perspective on STDMD

With equation (74), it is possible to evaluate the properties of the projected DMD operator based on an incrementally computed orthonormal basis of \mathbf{Q}_z and \mathbf{G}_z . In this work, the algorithm for updating the matrices is the CGSI process [62]. It proves to construct orthogonal vectors after few iterations using matrix multiplications, which can be parallelized. Figure 15 depicts the incremental update procedure in a flow chart including the initialization routine and compression of the orthonormal basis in case a specific maximum rank is reached. As new snapshots become available, augmented snapshots \mathbf{z}_j are formed. The first column of \mathbf{Q}_z is initialized with the normalized first augmented snapshot vector. The first entry of \mathbf{G}_z is formed by the square of the Euclidean norm of the first snapshot, representing the total kinetic energy contained in case velocity components are used as observables. After initialization, each new augmented snapshot vector is orthonormalized with respect to the existing orthonormal basis using CGSI. Upon reaching a user-defined number of CGSI iterations, the algorithm checks for the significance of the orthogonalized vector by computing its Euclidean norm and comparing it to a predefined truncation limit, ε_{QR} . In case this test is passed, \mathbf{Q}_z is expanded by the normalized and orthogonalized augmented snapshot vector, and \mathbf{G}_z is zero-padded in both dimensions. Before compression of the orthonormal basis, \mathbf{G}_z is updated using the orthogonality of

the current augmented snapshot with respect to the orthonormal basis, which is also referred to as non-orthogonality residual in the first iteration of the CGSI process. If the user-specified maximum rank limit r_{\max} is reached, the STDMD algorithm allows for an online compression of the orthonormal basis. That step is crucial for allowing the application of the method to large data sets because it avoids the requirement to store snapshots in memory (i.e., a single-pass operation). As the leading driver for computational CPU time cost is the CGSI process, orthonormal compression steps can also be used to limit the cost of further reorthonormalizations of augmented snapshots with respect to \mathbf{Q}_Z . The STDMD algorithm variant proposed here differs from the SDMD algorithm outline published in [35] because the compression stage is following the update of \mathbf{G}_Z rather than being applied right after the CGSI process (steps three and four in Hemati’s original SDMD publication [35]). Switching these two steps leads to much more consistent results concerning the reconstruction of the original snapshot matrix from DMD modes; therefore, ROMs reconstructed from the computed modes can represent the flow field more precisely (reduced performance loss Π_{loss}). Compression is obtained by eigenvalue decomposition of \mathbf{G}_Z , sorting the eigenvectors in its eigenvector matrix \mathbf{Y}_{G_Z} and eigenvalues $\boldsymbol{\lambda}_{G_Z}$ in descending order, and projecting the oversized $r_{\max}+1$ orthonormal columns of \mathbf{Q}_Z onto the leading r_{\max} eigenvectors. \mathbf{G}_Z is set equal to the diagonal matrix of the leading r_{\max} eigenvalues. At this point, all matrices and vectors involved in the STDMD procedure are real-valued numbers. After the potential compression procedure is done, the time step index j is incremented, and the next available augmented snapshot is constructed for the next orthonormalization. While STDMD and conventional DMD can be applied to any snapshot matrix constellation of \mathbf{X}_1^{m-1} and time incremented snapshots \mathbf{X}_2^m , snapshots two to $m-1$ are contained in both \mathbf{X}_1^{m-1} and \mathbf{X}_2^m , which drastically reduces the amount of memory required for batch-processed DMD variants. For the application of STDMD, reusing \mathbf{x}_{j+1} for the upper part of the augmented snapshot of the next iteration reduces the complexity of the algorithm.

Due to the need for interpolation of CFD-simulated data to a coarser grid before STDMD processing, as will be shown for the application on a vehicle aerodynamics case in section 3.3, the reuse of snapshots from a previous time step also halves CPU time required for spatial cell volume weighted interpolation tasks.

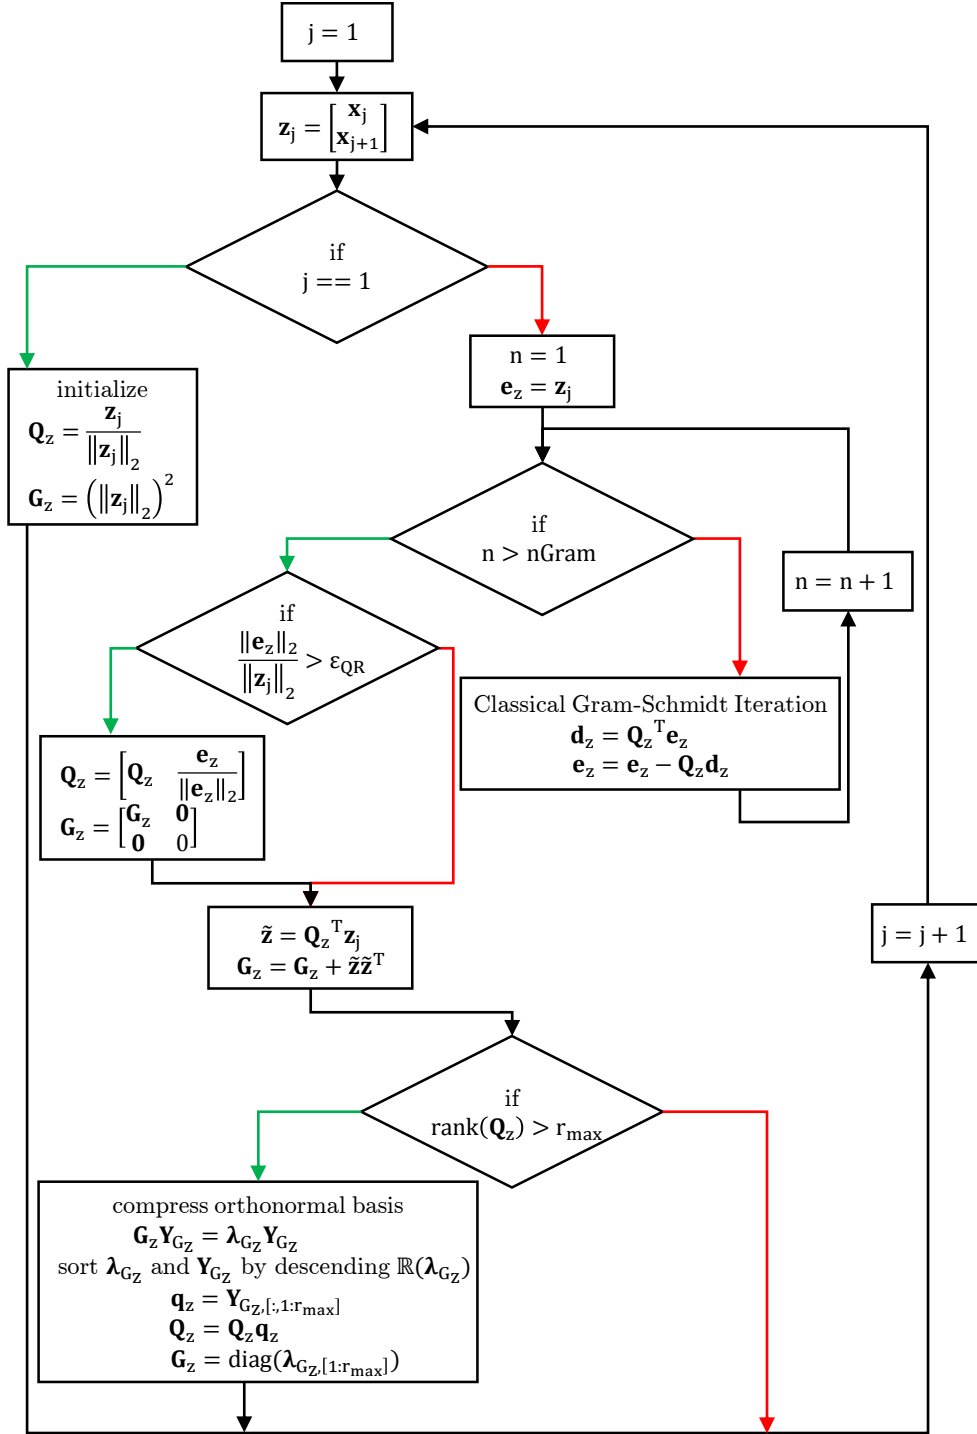


Figure 15: STDMD incremental orthonormal basis update. Green arrows denote a positive evaluation of an if statement, and red arrows indicate negative.

Upon reaching the last time step m , the very last augmented snapshot \mathbf{z}_{m-1} is constructed from \mathbf{x}_{m-1} and \mathbf{x}_m , and the orthonormal basis is updated one last time. The description of the orthonormal basis as a POD basis is omitted here because

\mathbf{R}_Z is not computed explicitly and POD was introduced in this work as another means of computing the SVD of a matrix. It is possible to reformulate the algorithm for the evaluation of the regular $\mathbf{Q}_Z\mathbf{R}_Z$ representation, but this does not allow for POD compression as it is presented here, which is crucial for the applicability of the method. Based on \mathbf{Q}_Z and \mathbf{G}_Z , evaluation of equation (68) can return the remaining part of the required matrices for the computation of the projected STDMD operator in equation (74). The most considerable memory-consuming step is found at this stage with a maximum RAM requirement in the order of approximately

$$\text{RAM}=\text{size}(\mathbf{Q}_Z)=(2\cdot n\cdot m)\cdot 8\text{Bytes}, \quad (91)$$

if no online compression of \mathbf{Q}_Z is employed and double precision values are used. For example, a commonly used CFD mesh of 100 million cells with three velocity field components (i.e., 300 million observables n and 1,000 snapshots m for an adequate representation of the flow field) would require 4.8 terabytes of memory. Additional matrices that are being held in memory only require memory of the order of m and m^2 , which is negligible compared to the size of \mathbf{Q}_Z .

After the update procedure is finished and \mathbf{Q}_Z and \mathbf{G}_Z are finalized, the algorithm follows the set of equations from section 2.4.2, which are visualized in the flow chart in Figure 16. The QR-decomposition at the beginning can be a memory-consuming equation if \mathbf{Q}_X is evaluated explicitly. A QR-decomposition variant that does not compute \mathbf{Q}_X should be employed. A popular example of such a method is the highly parallelizable algorithm of Benson et al. [63], which avoids communication and is numerically stable and efficient. Sayadi et al. [64] further employ this algorithm as the basis for conventional DMD evaluation in an application on transition flow with half a billion observables and 101 snapshots. They show linear scaling for up to 1,024 cores using this algorithm.

Next, the projected DMD operator is computed, and its complex eigenvalues λ_k and eigenvectors \mathbf{y}_k are obtained by eigendecomposition. With λ_k and \mathbf{y}_k , the DMD mode amplitudes and frequencies can be computed. As mentioned before, negative frequency modes can be omitted in general as a first mode selection step without losing ROM accuracy, due to redundancy.

Afterward, the frequency band for which the modes are being computed can be narrowed with a user-input frequency range. That option for narrowing modes can be specifically helpful in case of acoustic flow field effect analysis (i.e., if only a specific frequency is targeted in the analysis) or for exact band-pass filtering applications. As is shown in section 1.3, resolving unstable modes can require high-frequency sampling intervals (i.e., 20 to 40 times higher than the frequency to be resolved), according to Duke et al. [30]. Therefore, a wide range of low-amplitude, high-frequency modes is likely to appear in the mode spectrum, which can be truncated before further mode selection steps. Following the frequency selection, the modes (respectively α_k , λ_k and \mathbf{y}_k) are sorted by the modified eigenvalue weighted amplitude ordering method, introduced in section 2.4.6. The modes can be sorted by their physical relevance using this parameter.

Finally, a selected number of modes is computed from previously chosen mode selection steps using equation (79) (i.e., frequency range selection by f_{\min}

and f_{\max} and relevance ordering by K_{stable}). Upon reaching this point, the r_{ROM} modes can be visualized by mode distribution, reconstruction in time or saved to hard drive as required for further modal post-processing steps.

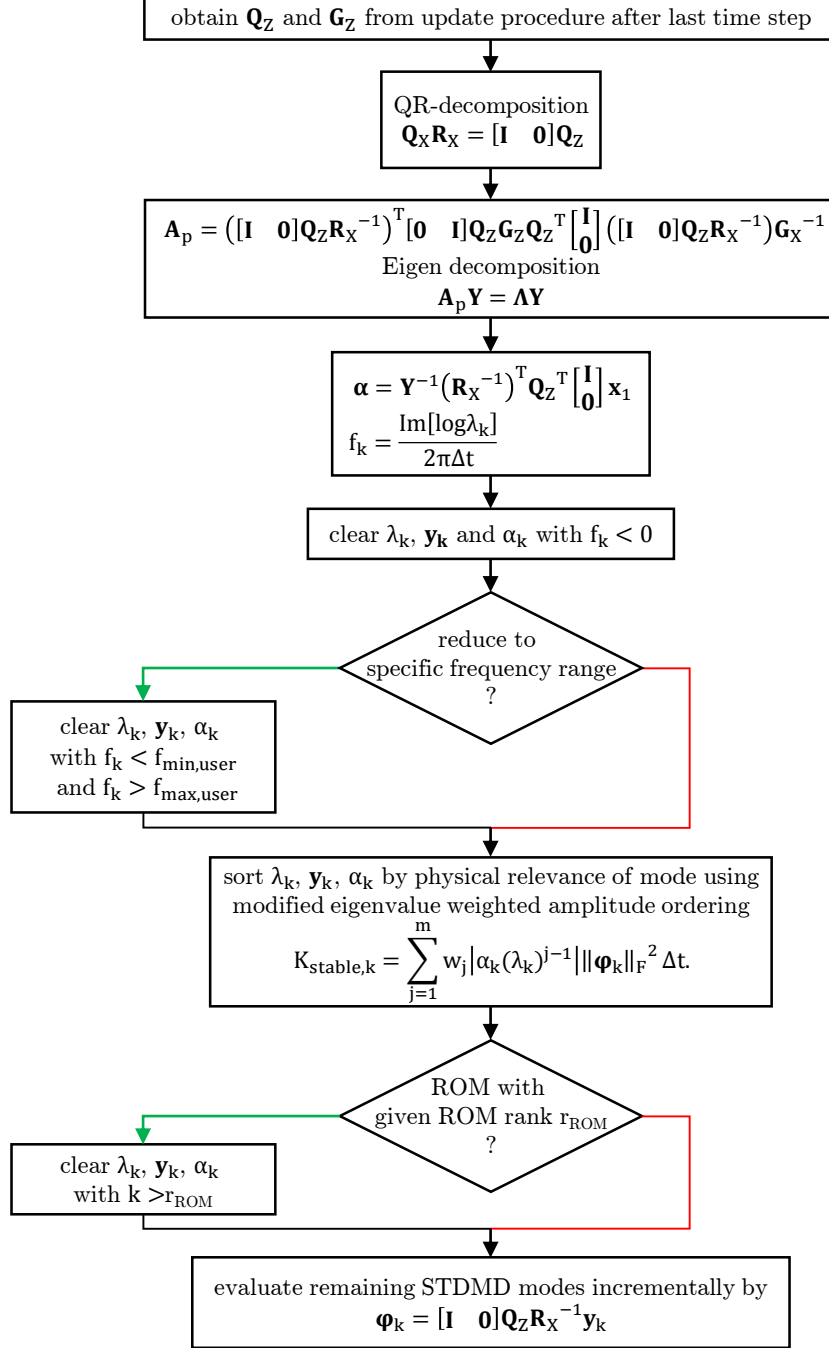


Figure 16: STDMD mode evaluation. Flow chart following the last time increment $j=m$ from incremental orthonormal basis computation depicted in Figure 15. Green arrows denote a positive evaluation of an if statement, and red arrows indicate negative.

2.4.6 Modified Eigenvalue-Weighted Amplitude Scaling

For the extraction of dominant flow field processes, strategies are required for ordering the modes by dynamic importance. While selection of the most dominant POD modes is straightforward, as shown above, ordering DMD modes requires application-specific sorting mechanisms. While there is a significant number of methods for choosing important modes for conventional DMD methods that involve the computation of the POD modes, there is only a small number of choices for incrementally computed DMD modes.

The most conventional mode ordering criterion is ordering by the magnitude of the complex mode amplitude $|\alpha_k|$. Amplitude ordering is used by most researchers in the field of applied modal analysis and can lead to helpful insights. For vehicle aerodynamics investigations, Frank [49] and Peichl et al. [50] have employed this method.

The amplitude ordering approach works well for applications with mostly stable modes, but it does not include the influence of the eigenvalue of the respective modes. For cases with growing or decaying modes on the other side, the temporal evolution of a single mode can change a mode's dynamic importance significantly. While a mode with large amplitude and large damping rate would be considered dominant in case of mode ordering by amplitude, its quick decay makes its influence vanish in the majority of time steps due to the multiplication with higher powers of the eigenvalue. That deficiency is also found by Kou et al. [38], who propose the integration of the temporal evolution of each mode amplitude over time. The suggested mode ordering criterion K_k for each mode k is defined as

$$K_k = \sum_{j=1}^m |\alpha_k(\lambda_k)^{j-1}| \|\boldsymbol{\phi}_k\|_F^2 \Delta t. \quad (92)$$

The definition in equation (92) includes the length of the mode vectors $\boldsymbol{\phi}_k$, which is incorporated here for the sake of completeness but is usually not required, because most DMD algorithms return unit length mode vectors by default. The definition does not include the temporal evolution of all individual modes because the temporal evolution of each cell in a specific mode will always even out due to the oscillation around zero, except possibly for very low-frequency modes.

The appearance of highly damped modes in noisy and complex simulation data can lead to very large amplitudes for those modes, making them appear physically relevant even if the eigenvalue-weighted amplitude ordering is applied. Flow fields that are not well represented by the computed orthonormal basis (e.g., by high compression or truncation criteria used for evaluation of DMD modes) have also shown modes with large damping values. The bright side is that unphysical, high amplitudes appear in combination with large damping ratios. Previous investigations have revealed damped modes with amplitudes in the order of the mean flow mode, for instance. To overcome this problem, the eigenvalue-weighted amplitude ordering is modified by a weighting function w_j , which is a function of the time step as

$$K_{\text{stable},k} = \sum_{j=1}^m w_j |\alpha_k(\lambda_k)^{j-1}| \|\boldsymbol{\varphi}_k\|_F^2 \Delta t. \quad (93)$$

The weighting function chosen in this work is a sine function with the period of the investigated time frame indices $m-1$ and function roots $j=1$ and $j=m$, as shown here:

$$w_j = \sin\left(\left(\frac{2\pi}{m-1}\right)\left(j-1 - \frac{m-1}{4}\right)\right) + 1 \quad (94)$$

For a data series with $m=1,000$, the weighting function is depicted in Figure 17.




Figure 17: Weighting function of the modified eigenvalue-weighted amplitude ordering criterion K_{stable} for a data series with $m=1,000$. The sine has a period of $m-1=999$ and an amplitude of 1.

The weighting function can be adjusted on a case-by-case basis. The function chosen here favors stable modes over damped or amplified modes. The newly introduced method outperforms other DMD mode-ordering strategies if the DMD spectrum contains unphysical, highly damped modes. That can happen for noisy data sets or for DMD executions with either high orthonormal basis compression ratios or small rank truncation limits. Sorting the modes by this criterion in descending order allows for mode selection by physical relevance. Due to the existence of degenerate eigenvalue pairs with modes at similar frequencies, it cannot be guaranteed to return the best possible ROM. However, it returns better results on average than previously used conventional mode selection criteria. Better mode-ordering results for STDMD can only be obtained by prohibitively expensive optimization algorithms, which would also require the original flow field data to be present after the simulation is done. Such an optimization problem can be defined using the definition of the performance loss criterion numerator as

$$\begin{aligned} \|\mathbf{X} - \tilde{\mathbf{X}}\|_F &= 0 \\ \|\mathbf{X} - \boldsymbol{\Phi} \mathbf{D}_\alpha \mathbf{T}\|_F &= 0, \end{aligned} \quad (95)$$

with optimization variable vectors $\boldsymbol{\alpha}$ and possibly $\boldsymbol{\lambda}$, which are both complex value vectors with r_{ROM} elements. Such optimization is rather expensive and should be omitted or simplified. Because the full POD is not available in this method, smarter and faster algorithms like sparsity-promoting DMD, as introduced by Jovanovic et al. [31], [32], are also not applicable in this context.

2.4.7 DMD Results

As for the incrementally computed POD results, STDMD is applied to the two-dimensional cylinder case from section 2.2 for discussing the properties and

influence of different algorithm parameters on the results. As opposed to r_{ROM} model rank in POD, the limiter for restricting maximum memory usage is given by the maximum allowable size of the orthonormal basis \mathbf{Q}_Z with r_{max} . The actual ROM rank of DMD modes, on the other hand, is not explicitly restricting the memory consumption in the regular workflow because the DMD modes can be evaluated incrementally, as discussed in section 2.4.5. Due to the complex number nature of DMD modes, the r_{max} columns in the orthonormal basis reduce to a rank of $r_{\text{ROM}} \leq 1/2 \cdot (r_{\text{max}} - 1)$ DMD modes in case of a single mean flow mode, while the required storage of a single DMD mode is twice the storage of a POD mode. The STDMD executions in the two-dimensional cylinder results employ five Gram-Schmidt iterations for orthonormalization of augmented snapshots, which is found to be sufficient for other flow field decompositions as well.

The first STDMD tuning parameter evaluated here is the incremental QR-decomposition update threshold ε_{QR} for the expansion of the orthonormal basis \mathbf{Q}_Z . The following graphs in Figure 19 depict the performance losses over ROM model ranks r_{ROM} for different given thresholds ε_{QR} . For a direct comparison to POD regarding storage requirement, the complex number nature of DMD modes needs to be considered. All settings for the threshold value lead to a decrease in performance loss, dropping from a mean flow mode performance loss of 16.28% to a value between 10.23% and 12.17% with the addition of the first DMD mode. The first DMD mode can, therefore, model around a third of the total oscillation energy of the system. While the first mode is usually the only mean flow mode in a DMD ROM, the linear growth leads to the appearance of several DMD modes at zero Hz. The actual performance loss computed using the temporal mean of the snapshot matrix is 16.14%, as observed in section 2.3.4, but the first DMD mode at zero Hz (which is also the mode with the overall largest amplitude) only mimics a part of the physically present mean flow. Part one of the data series includes the first snapshot with minimal flow field disturbances and makes the method of first snapshot that is used for STDMD in this work especially unsuitable since the mode amplitudes are computed using a reconstruction of the first snapshot. If oscillation energies are not present or very small, as it is in this case, the exact reconstruction of the first snapshot can be achieved without assigning amplitudes to the oscillating modes. Equation (86) has several solutions close to the optimal solution with oscillating mode amplitudes that are not optimal for the entire data series. Therefore, it can be advised to apply DMD in combination with the method of first snapshot for developed flows or flows that contain a substantial amount of characteristic oscillations in the first snapshot. Amplitude optimization algorithms like the method applied in sparsity-promoting DMD, on the other hand, can overcome this deficiency but are either not applicable to this variant of STDMD or too expensive to use. If the POD basis is not evaluated explicitly, an optimization as executed in sparsity-promoting DMD requires the full snapshot matrix for a very memory- and CPU time-intensive optimization procedure.

The most precise ROM with full mode evaluation can be achieved using a threshold value of 10^{-7} or lower for both parts of the simulation. Using all modes

for reconstruction of part one leads to a performance loss of 0.50% and of $7.70 \cdot 10^{-4}\%$ in part two of the simulation. The worst threshold of 10^{-3} produces unusable ROM representations of the flow field for part one with a performance loss of 10.32% for 84 computed modes and a reasonably good ROM with a performance loss of 0.05% for part two with only 13 modes. Another notable property of DMD ROM results is visible in the upper graph of Figure 19. While the ROM performance loss curves monotonically decrease as more modes are added for part two of the simulation, the addition of modes to the ROM does not necessarily decrease the modeling error in part one. Reconstructing an ROM with the first four modes picked by the modified eigenvalue-weighted amplitude mode ordering method leads to a performance loss of 67.69%. An ROM of DMD does not only fail to produce a best-possible representation of the original data with any given rank but can also produce deteriorated reconstructed flow fields that are far from the original flow field, both in amplitude and phase in any time step. That does not mean that dominant flow field effects cannot be extracted, but special care needs to be taken when analyzing DMD results in this respect using the current performance loss definition. Due to the squaring of differences in the Frobenius norm, the addition of a DMD mode can worsen the performance loss indicator significantly. An illustration of this effect is given in Figure 18 using a generic example of a DMD decomposition of a data set with one observable $n=1$, 15 time steps of $m=15$ and with an ROM rank of $r_{\text{ROM}}=3$. By adding modes ordered by their amplitudes, the best possible ROM approximation with a given rank is obtained in this generic example. The ROM with $r_{\text{ROM}}=1$ results in a performance loss of 7.35%. Adding mode two to the ROM increases the performance loss significantly to 14.5%. The addition of the third mode yields precisely the original data set without error (i.e., the resulting ROM can fully capture the temporal evolution of the data series). Other possible mode orders lead to worse deterioration of the performance loss indicator in this example. Instead of adding modes in the order one, two, three, one could add modes as one, three, two. That would lead to a performance loss model error of 17% with $r_{\text{ROM}}=2$. Any possible combination of two modes out of the three leads to an increase in Π_{loss} , compared to mode addition by amplitude, which is used in the right column of Figure 18.

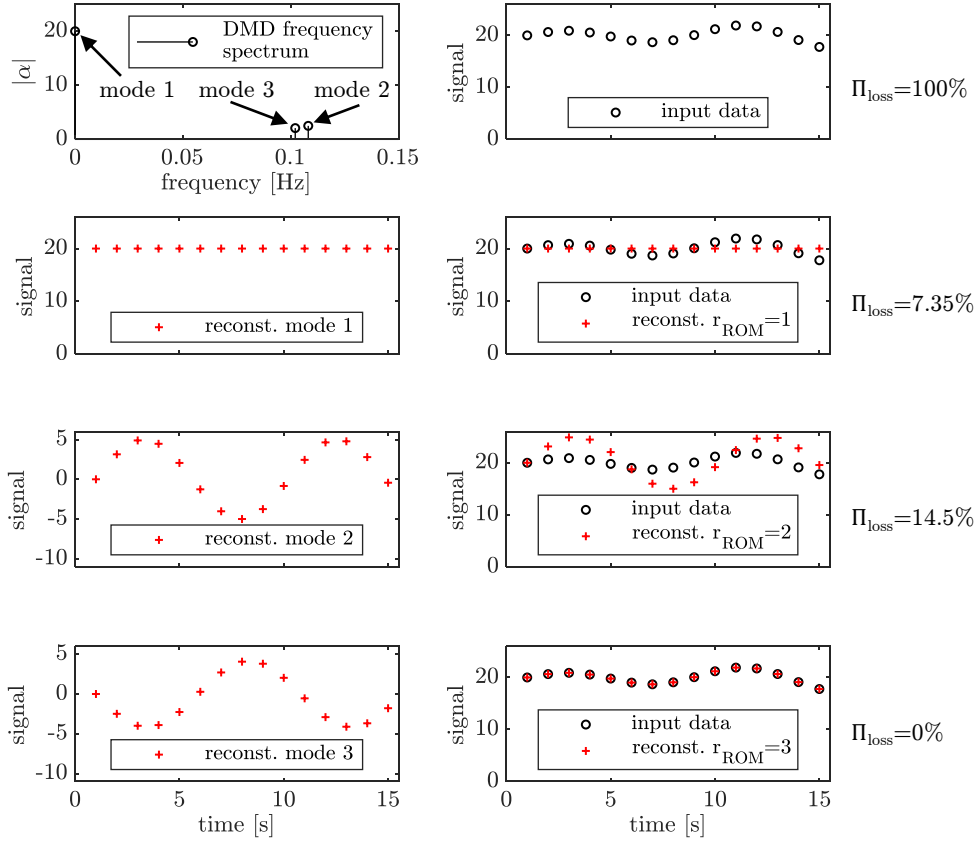


Figure 18: Left (top to bottom): (1) DMD frequency spectrum with three modes. Mean flow mode one at $f_1=0\text{Hz}$ and $|\alpha_1|=20$, mode two at $f_2=0.11\text{Hz}$ and $|\alpha_2|=2.5$, and mode three at $f_3=0.10\text{Hz}$ and $|\alpha_3|=2.1$. Eigenvalues for all three modes are stable (i.e., $|\lambda|=1$). Complex conjugate counterparts of the fluctuating modes (negative frequency) are not shown. (2) Reconstruction of the mean flow mode. (3) Reconstruction of the second mode, including its complex conjugate counterpart. (4) Reconstruction of the third mode, including its complex conjugate counterpart. Right (top to bottom): (1) Temporal evolution of the input data. (2) Reconstruction of an ROM with $r_{\text{ROM}}=1$ (i.e., only the mean flow mode one). (3) Reconstruction of an ROM with $r_{\text{ROM}}=2$ (i.e., modes one and two). (4) Reconstruction of an ROM with $r_{\text{ROM}}=3$ (i.e., all modes combined).

The observations above can be generalized. For a good representation of a complex flow field, DMD ROM might need all of the modes combined, because each mode only consists of a single frequency oscillation governed by a specific amplification or decay rate. Complex flow field effects (e.g., linear growth followed by decay of the same oscillation frequency) require several overlapping DMD modes to be modeled. While this property can be seen as a flaw, it also integrates a strength. The simplistic appearance of DMD modes can still be used for analyzing dominant flow field effects, and due to their physical simplicity, assigning a DMD mode to a specific vortex-shedding phenomenon is much easier than for other related modal reduction techniques. Due to the increases in the performance loss indicator for small numbers of modes, it is highly advisable to

always check the performance loss of the full rank reconstruction of a DMD ROM.

In part two of the two-dimensional cylinder case, the performance loss curves fall monotonically, and the ROM error curves are mostly lower for smaller QR-decomposition thresholds. It is notable that the very first oscillating DMD mode can reduce the performance loss to 3.58%, so a single DMD mode can represent more than three-fourths of the oscillation energy in the system (the mean flow performance loss is 16.14%). While for complex systems, it seems unavoidable to construct a sizeable orthonormal basis in any case for generating a good ROM, a threshold of 10^{-3} for limiting the size of the orthonormal basis is sufficient to get a 0.05% performance loss resulting from the reconstruction of all 13 computed DMD modes.

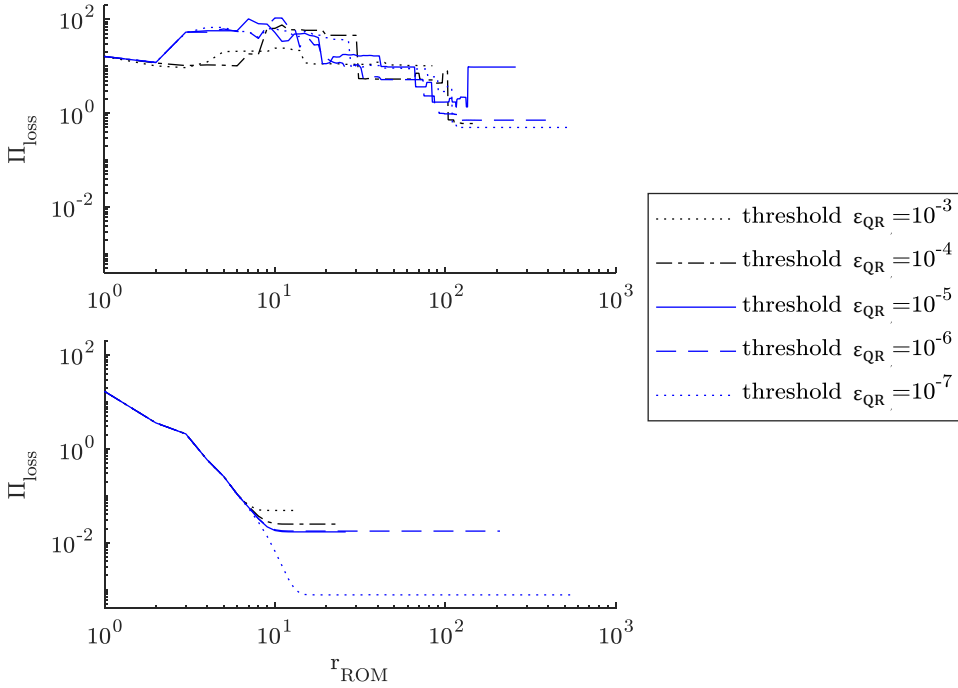


Figure 19: Performance loss charts for part one (top) and part two (bottom) of the two-dimensional cylinder test case for varying thresholds in the STDMD algorithm. The last point signifies r_{\max} (i.e., the rank of orthonormal basis \mathbf{Q}_Z).

Figure 20 gives insight into the computation time required for the execution of the STDMD algorithm using the same QR-decomposition thresholds as discussed above, for which the ROM representation results are shown in Figure 19. While the full evaluation of STDMD without truncation, which can be achieved using a threshold value of 10^{-7} , uses computation times in the order of 10^3 seconds on a 20-core machine, the higher threshold settings using 10^{-3} take one order of magnitude less time for part one and take two orders of magnitude less time for part two. Due to the large ROM error using a threshold value of 10^{-3} for part one, it is recommended to use low threshold values, especially for very complex flow field phenomena, to allow for a robust expansion of the

orthonormal basis. In general, the better the orthonormal basis, the better the evaluated DMD modes. The two-dimensional cylinder flow in this work is suited well for a threshold value of 10^{-4} , for which a reasonable tradeoff can be achieved among ROM accuracy, HDD storage requirements for the modes and STDMD execution time. For simple flow fields with many time steps to be analyzed, threshold values larger than zero should always be used to avoid adding snapshots to the orthonormal basis that are already adequately represented by it.

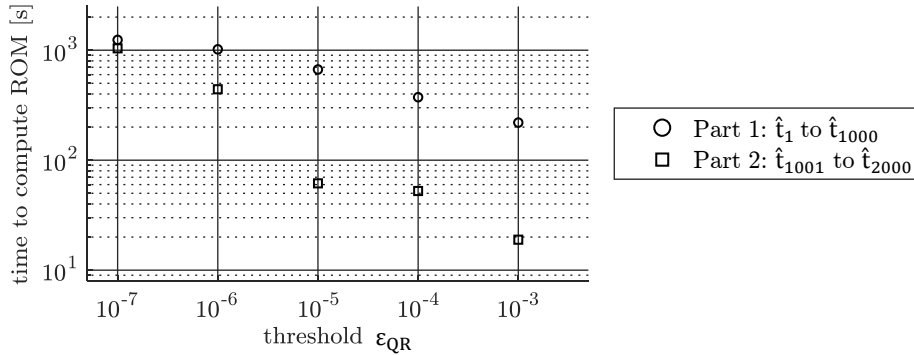


Figure 20: ROM execution time for STDMD analysis. The computation time increases with the rank of the resulting ROM and the increase in the QR-decomposition threshold value ϵ_{QR} .

There is no exact guideline for setting a threshold value universally, but there is always a maximum allowable memory consumption on any computer system. The compression step in the STDMD algorithm from Figure 15 is capable of reducing the size of the orthonormal basis after a predefined rank is exceeded. This way, the total memory consumption of the algorithm can be restricted to fit the available computer system resources. Part two of the simulation is not sensible to be evaluated using a fixed maximum rank without thresholding, but part one is. Figure 21 shows the performance loss behavior of different maximum rank compression criteria. No apparent tendency can be observed for increasing maximum ranks, but the overall trend leans to better ROMs with more massive maximum ranks allowed. All STDMD runs with $r_{\max} \geq 600$ can produce ROMs with performance losses smaller than 1%. The computation time required for STDMD is governed by the number of iterations in the CGSI process, the rank of the orthonormal basis and the number of times the orthonormal basis compression is executed. The orthonormalization process is less time-consuming for smaller orthonormal bases, but the compression procedure needs to be executed more often. For this case, the computational benefit of having a small orthonormal basis rank evens out with the additional compression executions at a maximum rank of 600. By choosing a smaller number of Gram-Schmidt iterations for each update execution, this rank shifts to larger values.




Figure 21: Left: STDMD performance loss behavior for different maximum rank truncation parameters in part one of the two-dimensional cylinder test case. Maximum rank truncation is conducted using orthonormal basis compression in the update procedure in the algorithm. Right: STDMD execution time for the same STDMD runs as depicted left.

As mentioned above, running STDMD with a threshold value of $\varepsilon_{\text{QR}}=10^{-4}$ gives a good tradeoff between the ROM rank, ROM performance and STDMD execution time for both parts of the two-dimensional cylinder simulation test case. The following section presents specific results for the execution of STDMD on using this threshold criterion without a maximum rank criterion for both parts of the simulation. The left graph in Figure 22 shows the frequency spectra of both parts of the simulation. While in the second part of the simulation, only a few modes need to be used to approximate the flow field correctly, the first part requires a wide band of frequencies to capture the increasing instabilities of the flow. For both parts, the most dominant oscillatory modes are at the same frequencies, and the relevant flow field effects can be extracted. The right graph in Figure 22 shows the performance loss curves. As presented in Figure 18, the performance loss increases when adding more modes in case of more complex flow field phenomena (i.e., part one in contrast to the ROM convergence of the second part). Also, the final ROM model error for the first part of the simulation is more significant than for part two by one order of magnitude. While the largest ROM model error appears for the reconstruction of only the mean flow mode in part one, the performance loss in the second part of the simulation peaks at STDMD mode number 11's addition to the ROM, with a value of $\Pi_{\text{loss}}=76.38\%$. The final ROMs have a rank of 144 for the first part and 24 for the second part. The K_{stable} value of the mean flow modes and the most dominant shedding modes are in the same range with $K_{\text{stable,part1,mode1}}=2858.2$ and $K_{\text{stable,part2,mode1}}=2857.6$ for the mean flow, as well as $K_{\text{stable,part1,mode2}}=337.4$ and $K_{\text{stable,part2,mode2}}=313.3$. Steady changes in the mean flow of part one are captured by additional mean flow modes and low-frequency modes with sufficient damping rates, which are not present in part two. All other oscillations of part two are ramped up similarly, using multiple modes at similar frequencies with different damping rates, resulting in mode

addition phenomena as presented in Figure 18. In the case of part one, further modes added to the ROM are in antiphase to the dominant oscillating mode (i.e., the dominant shedding mode at $St=0.2$). Those modes surrounding the $St=0.2$ mode are slightly damped so that their subtraction effect decays toward zero in the reconstruction of later time steps of part one. The data compression performance of DMD results deteriorates as soon as more complex flow field phenomena are involved, as will be shown later for the vehicle aerodynamics application.

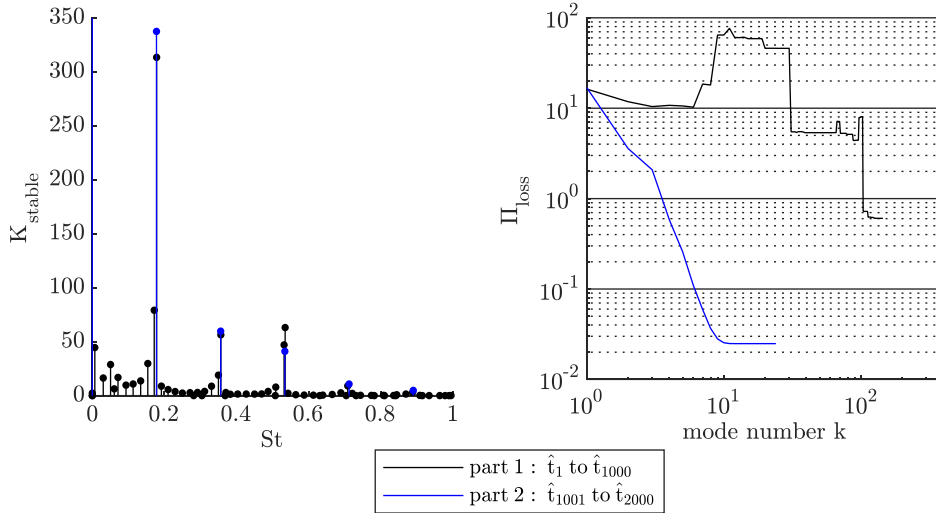


Figure 22: Frequency spectra of the STDMD analysis (left) and ROM performance loss curves (right) for part one and part two of the two-dimensional cylinder test case. Negative frequency counterparts in the frequency spectra are not shown.

The reconstruction of DMD ROMs gives insight into the unsteady flow field effects. In combination with the mean flow mode, an approximation of the temporally resolved flow field is returned. Due to the fixed frequency and fixed eigenvalue of single DMD modes, the reconstruction of the y -direction component of modes one and two leads to the visualization shown in the right columns of Figure 23. The dominant vortex shedding mode at $St=0.2$ is extracted well. There is a larger offset in the first half of part one in comparison to the original data than in the POD results compared to the same data and using the same amount of information (i.e., a single oscillating DMD mode requires the same amount of storage as a POD mode pair, which is used for reconstruction in Figure 12). The simplicity of DMD modes comes in handy when analyzing many time steps, since a single period of reconstruction of a mode's fluctuation is enough to understand its behavior. Information about the mode's actual appearance in the investigated time frame is not retained in a single mode reconstruction and is only available when reconstructing using multiple oscillating modes together.

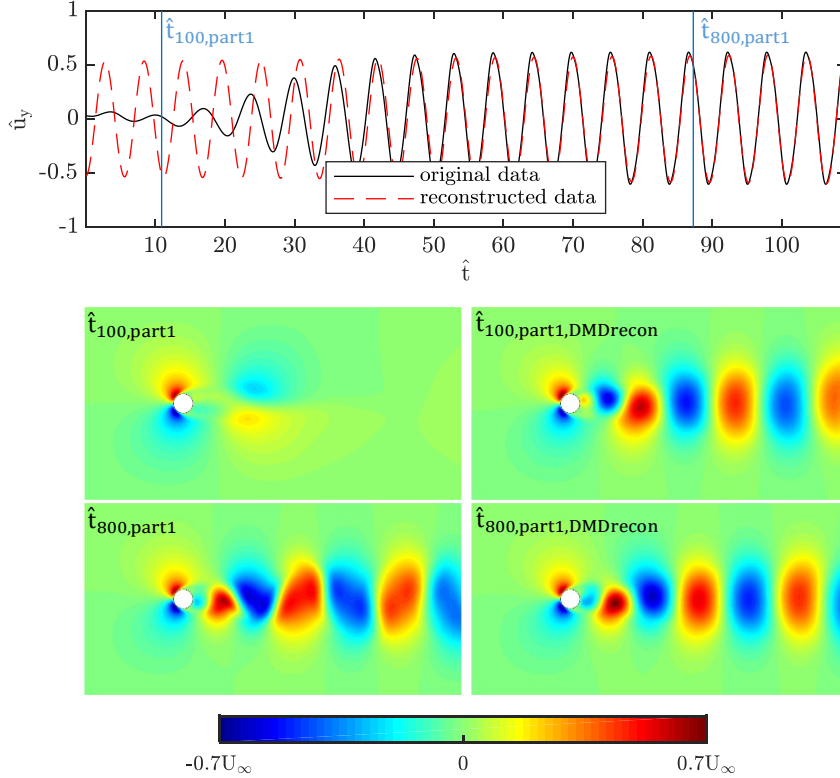


Figure 23: Temporal evolution of the velocity y-component probed at position $x=2D$ and $y=0$ in part one of the simulation. Snapshot $\hat{t}_{100}=10.91$ and $\hat{t}_{800}=87.28$ of the original data (left) and reconstructed data (right) is marked with a blue line in the upper graph. DMD modes one and two and the complex conjugate mode of mode two are used for the reconstruction (mean flow mode and first oscillating mode including its negative frequency counterpart).

The first oscillating mode in part one comes with a slight amplification rate (i.e., $|\lambda_2|=1.00012$), which improves the performance loss in comparison to the addition of an entirely stable mode, as it appears in part two, because the amplified mode can mimic the increasing instability better. Location, shape and energy of the y-component structures are well approximated in the second half of part one.

For part two of the simulation, the reconstruction of the first two DMD modes leads to much better results. The ROM with rank two can approximate the flow field down to a precision of $\Pi_{\text{loss}}(\text{rROM}=2)=3.6\%$, the same as for the rank three model of POD modes. Due to the complex-value nature of DMD modes, the required data storage is in the same range as a POD ROM. Similar to the visualization for part one, Figure 24 depicts the temporal evolution of the y-component in a probe point and shows the comparison between the original flow field and the reconstructed flow field at time step index 700 of part two of the two-dimensional cylinder simulation.

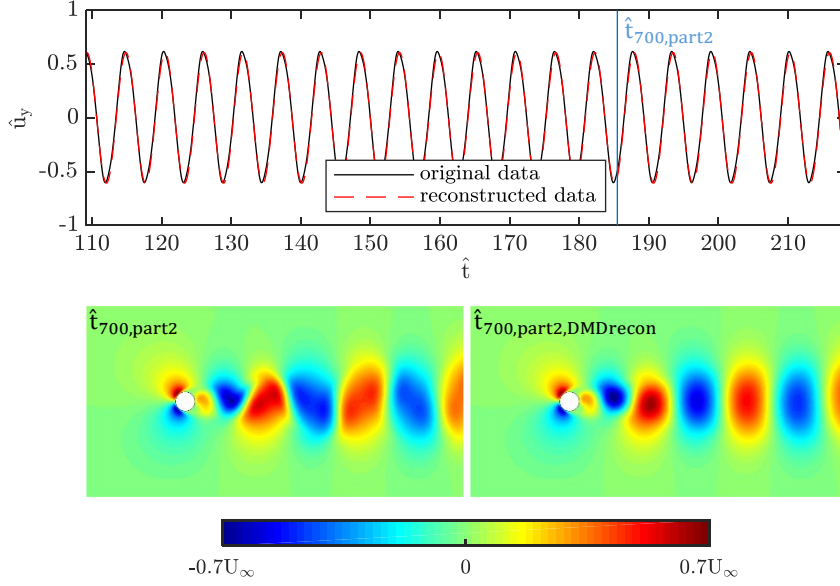


Figure 24: Temporal evolution of the velocity y-component probed at position $x=2D$ and $y=0$ in part two of the simulation. Snapshot $\hat{t}_{700}=185.5$ of the original data (left) and reconstructed data (right) is marked with a blue line in the upper graph. DMD modes one and two and the complex conjugate mode of mode two are used for the reconstruction (mean flow mode and first oscillating mode including its negative frequency counterpart).

Another exciting feature of DMD is the ability to extract eigenvalues of oscillation processes that are only contained in the investigated time frame with less than an entire oscillation. While FFT, with its fixed frequency spacing, fails to do so, DMD can be used to extract comparably low-frequency structures and approximate eigenvalues. For this matter, 10 time steps from \hat{t}_{1001} to \hat{t}_{1010} are used as a separate data set and analyzed using STDMD (i.e., the first 10 time steps of part two). The investigated period contains only a fraction of an entire vortex shedding period. A single shedding cycle takes around 52 time steps, so the data series contains only around a fifth of a total shedding period. The data series, the respective reconstruction using DMD modes one and two in a probe point, and the spatial distribution for a single snapshot and reconstructed snapshot are shown in Figure 25. The ROM of two DMD modes can follow the given data series well, and the STDMD result also gives an excellent approximation for the system's dominant eigenvalue. The extracted shedding frequency for this case of $St_{\text{periodFraction}}=0.22$ is surprisingly close to the effective frequency of the system at $St=0.178$, given the tiny sampling period. For increasingly smaller fractions of an entire period of oscillation, DMD modes tend to form an additional mean flow mode, as the change is observed as a steady increase in the mean flow instead of a flow field fluctuation. By increasing the number of investigated time steps, the extracted eigenvalue converges to the actual system eigenvalue.

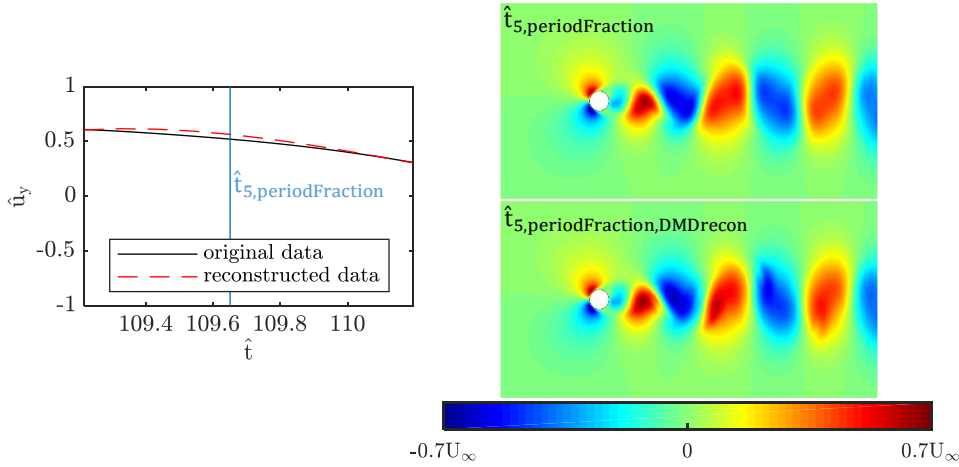


Figure 25: Temporal evolution of the velocity y-component probed at position $x=2D$ and $y=0$ in for STDMD execution on a fraction of part two of the simulation. Snapshot $\hat{t}_5=185.5$ of the original data (top) and reconstructed data (bottom) is marked with a blue line in the left graph. DMD modes one and two and the complex conjugate mode of mode two are used for the reconstruction (mean flow mode and first oscillating mode including its negative frequency counterpart).

2.5 DMD vs. POD vs. FFT

One of the significant disadvantages of applying FFT for modal analysis of flow fields can be observed for the one-dimensional time signal evaluated in the PSD spectrum in Figure 5. While it is possible to extract a dominant frequency, FFT offers poor frequency resolution for short sampling periods. If a time signal is covering only one or a minimal amount of oscillations, it is impossible to extract the exact frequency of oscillation applying FFT. Due to the fixed frequency spacing, the true eigenvalue of the oscillation might be right in the middle of two neighboring peaks, and only longer data sampling times, leading to more appropriate frequency resolution, can improve the result.

Additionally, the necessity of applying a windowing function (i.e., a Hanning type window) leads to an underweighting of data contained at the start and end of a time series. For low-frequency analysis, either very long simulation runs are needed to gain further low-frequency range certainty with FFT or only a single window can be analyzed, which in turn leads to the underweighting problem. For DMD, only small portions of a full oscillation period are necessary to capture the frequency precisely. There is no strict lower frequency limit for DMD, but it is advised to analyze at least a quarter of a period of a signal for stable modes to get a reasonable approximation of eigenvalues.

Another DMD benefit is that time series length is not constrained to sample numbers of an even power of two. Therefore, the sampling period and frequency can be chosen based on physical properties of interest. Additionally, weighting functions are not required in general, so all data in the sampling period is equally represented in the spectral representation. In FFT, the necessity to consider the number of samples and the application of window functions usually leads to

multiple evaluations of FFT on a single data set using partially overlapping windows and then averaging the resulting spectra. The averaging of FFT spectra leads to the loss of phase information, which is needed for possible reconstruction of modes in time.

Even though previous investigations use mean-subtracted data sets, DMD does not require such manipulation, which makes it possible to apply the algorithm during runtime of the simulation before computing an approximation of the temporally averaged flow field. If DMD is applied on mean-subtracted data sets without the presence of amplified or decaying flow field processes, the result equals discrete Fourier transform (DFT) results. If this data manipulation is omitted, DMD can represent the original data sets better than a DFT because there is no fixed frequency spacing constraint and a smaller amount of modes is required to represent the same amount of energy from the original data set.

A prominent benefit of FFT over DMD is the easy parallelization on high performance computing (HPC) systems with many processors since FFT can be applied for each cell individually and no communication between processors is necessary. This computational advantage is diminished though by the necessity to store data in memory or on a hard drive, which can drastically increase simulation time and can become prohibitively costly for large data sets. STDMD, on the other hand, can process snapshots on the fly and can compress the orthogonal basis during the simulation time to keep memory requirements in a certain limit.

While FFT returns a frequency spectrum that represents the average activity of specific frequency components in the flow field, DMD results also contain information about the temporal evolution of oscillation amplitudes. DMD captures linear growth well, for instance, while FFT can only return averaged amplitudes over the whole sampling window. This feature of DMD allows for the extraction of modes that are highly energetic, despite that energy being shown only in a short window during the sampled period, which is not possible using FFT due to the temporal averaging effect. The presence of damped modes that are physically not relevant from noisy data sets requires a modified mode selection criteria like the modified eigenvalue-weighted amplitude scaling suggested here.

POD modes are orthogonal in space, which is an essential property concerning the creation of an ROM. When constructing an ROM with a small number of modes, each added mode decreases the approximation error of the ROM in comparison to the FOM. The mode selection strategy for POD-based ROM is quite simple. Their respective singular value selects modes in decreasing order. In DMD, however, some dominant oscillations might only be represented well by a combination of two modes at a similar or even identical frequency that might oscillate in antiphase. If an ROM is generated using only one of the two modes, the resulting reconstructed flow field can be further away from the FOM than if that single mode was not considered (i.e., increasing performance loss while adding more modes to the ROM). Whereas an explicit, energy-related ordering and ROM building is not strictly converging with DMD, its strength

lies in the creation of modes at a distinct frequency. That offers the possibility to extract and analyze comparably simple but still energetic flow field structures or to investigate flow field phenomena in a selected frequency band of interest. If one was interested in the excitation of a particular acoustic frequency, for instance, one could focus on modes around that specific frequency to find possible sources of the noise in the flow field by tracking back the spatial structures to their excitation mechanism.

Another benefit of having orthogonality in time is that oscillations of the integrated force on the vehicle surface can be investigated using a DMD on the force vectors on the vehicle surface. Surface areas of high DMD mode amplitude at the investigated frequency can be found and tracked back to vortex generation mechanisms in the flow field. That offers a way to identify potential improvements to optimize driving stability and comfort.

In part two of the two-dimensional cylinder simulation, POD modes two and three together represent as much oscillation energy as DMD mode two. Due to the complex number nature of DMD modes, this results in the same ROM performance if the singular values and the right singular vectors are considered negligible in size. That is only the case for this specific flow field because only a single frequency is dominant and active in this especially suitable case. In general, POD always outperforms DMD for ROM generation.

An essential drawback of all data-driven methods is the strong dependence on the choice of observables. In the context of DMD, the choice of observables is particularly crucial in case of strong nonlinear flow field effects that need to be captured. In this work, the full state representation of a flow field is considered to be the velocity components. To find the dominant eigenvalues of a nonlinear system, higher order combinations of the basis function (e.g., $\mathbf{u}\mathbf{w}^2$) could lead to a much better approximation of physics (i.e., computing dominant eigenvalues) than is contained in the velocity components. This inherent property of DMD as a data-driven method was demonstrated on a two-dimensional map system in a talk by Rowley [65]. However, including random higher order combinations of observables would increase the degree of complexity drastically, so no such advanced choices of observables are considered in this work.

A disadvantage of DMD is the missing correlation between a mode's amplitude and its actual energetic relevance in the flow field. Because the structures are not necessarily orthogonal in space by construction, DMD modes can overlap entirely and cancel out each other by being in opposite phase. The appearance of degenerate eigenvalues is not uncommon [55], and possible remedies need to be developed. POD modes, as well as DMD modes, should be thought of as oscillations or coherent structures. If the orthonormal basis for computing DMD modes is highly compressed, DMD modes can also become orthogonal in space but maintain their distinct frequency content.

3 Application on a Full Vehicle Aerodynamics Simulation

The following chapter shows the application of the POD and STDMD methods to a large data set from a vehicle aerodynamics simulation. The simulated vehicle geometry setup is introduced, and a reference is made to wind tunnel measurement results. Central difficulties in the comparability of wind tunnel measurement results to CFD simulations are outlined. Suitable spatial and temporal data coarsening and filtering strategies are then presented in order to succeed in evaluating modal analysis results. Finally, modal analysis results are compared for the two methods before concluding remarks are given.

3.1 DrivAer Reference Configuration

The vehicle geometry used in this work is the DrivAer reference body. The mockup configuration was developed by the vehicle aerodynamics group of the Chair of Fluid Mechanics and Aerodynamics of the Technical University of Munich. The geometry with its three different rear end shapes was first published by Heft et al. in 2011 [66]. Further details on the reference body and the wind tunnel setup can be found in other papers published by group member Mack in 2012 (Mack et al. [67]) and with more details about the wheels-on configuration by Miao et al. in 2015 [68]. The former reference also shows the design and functionality of the suspension system and implications on the measured forces. The reference body comprises the geometrical properties of the Audi A4 and the BMW 3 series production cars. The merger of those car shapes to a publicly available geometry that can be used by anyone makes it possible to use a realistic geometry in order to develop new methods that are relevant for production car engineering. While previously the most often used reference body was the SAE body [69], that option comprises only the most basic vehicle shape features. In contrast, the DrivAer resembles the complexity of an actual production vehicle (i.e., including realistic A-pillar shape and wheelhouses). Additionally, the reference body comes with a variety of optional parts to be added or changed:

- smooth underbody or structured underbody, including gearbox tunnel and asymmetrical exhaust pipe system
- mockup configuration or engine bay flow with heat exchanger setup
- with or without side view mirrors
- realistic or generic closed wheels
- fastback, station wagon, and sedan rear-end shapes

The engine bay flow extension was developed in cooperation with the Research Institute of Automotive Engineering and Vehicle Engines Stuttgart (FKFS), and first results were presented in 2015 [70]. For the purposes in this work, the baseline setup is the fastback mockup configuration with closed wheels, without side view mirrors and with a smooth underbody. This vehicle configuration is among the most straightforward possible configurations that can

be used. The simplicity of the geometry reduces the number of error sources for a direct comparison of wind tunnel measurement data to simulation data. It also eases the assignment of flow structures to specific parts or length scales of the geometry. The addition of a structured underbody, for instance, increases the number of cells needed for discretization and gives space for numerical errors that are transported toward the rear and could alter flow phenomena in the vehicle wake. The model support system from the wind tunnel setup is included in the simulation setup to allow for better comparability to wind tunnel measurements. While most previous publications of the DrivAer reference body from TUM feature a setup with wheels decoupled from the vehicle body (i.e., the wheels-off configuration), this work employs the wheels-on configuration developed by Miao et al. [68]. Usage of the wheels-on setup is intended to minimize oscillations in the measured signals originating from the eigenfrequencies of the model support system. The wheels being in contact with the ground significantly reduces aerodynamically induced movement of the body and the wheels. This behavior is desired to reduce separation processes induced by geometry oscillation.

There are two significant differences here compared to the geometrical setup of Miao et al. [68]. First is the geometry of the wheels, which are generic open wheels in the work of Miao et al. and generic closed wheels in this work. Second, the side-view mirrors are not mounted in the configuration of the current work. Both measures are taken for geometry simplification purposes and to reduce the number of error sources for the simulation setup (e.g., interpolation errors for rotating wheels with sliding mesh method are omitted because closed wheel rotation can be modeled by rotating wall velocity boundary condition). That allows for better reproducibility and decreases total simulation cost. The side-view mirrors of the DrivAer with a drag increase of 16 counts in the fastback configuration [71] hint at a laminar detachment process. Possible laminar detachment from the side-view mirrors cannot be captured properly by the chosen turbulence model type.

For the wheels-on setup, a vehicle-mounted suspension system is added to the geometry. The suspension system does not add to the frontal area of the model but introduces a significant blockage to the wheelhouses. The fastback rear-end shape is chosen to minimize the required simulation period for a good representation of the flow field. While the sedan rear-end shape would have sufficed, the station wagon rear-end shape shows bi-stable wake behavior, for which long simulation runtimes are needed to capture the very low-frequency wake instabilities. Such low-frequency oscillations are also referred to as bi-stable states or asymmetric wake states because they do not appear to be connected to any length scale of the bluff body. Previous research by Grandemange et al. on the Ahmed body ([72], [73]) and Pavia et al. on the Windsor body ([74], [75]) investigates such oscillations in the rear wake. They find a connection between globally dominant vortex shedding modes and this bi-stable behavior. They also find similar vortex structures in the wake between the two reference geometries at Strouhal numbers based on the height of the rear-end H of $St_H=0.07$, $St_H=0.13$

and $St_H=0.19$. A direct comparison to their results is difficult due to the large geometrical differences (e.g. missing wheels and fundamentally different rear-end shapes).

The workflow presented in this work aims at making modal analysis tools available for everyday production vehicle aerodynamics engineering, for which typical simulation runtimes are too short for resolving those low frequencies. Figure 26 shows the DrivAer reference body in the wheels-on configuration with the top sting of the model support system of the wind tunnel A (WTA) of TUM. The top sting is included in this configuration due to the impact on the force coefficients, as previously investigated by Heft [76]. She found the top sting to increase the drag by five counts for a fastback configuration of the DrivAer.




Figure 26: Dimensions of the DrivAer reference body in 1:2.5 scale. The lower part of the top sting is included for a better approximation of the wind tunnel measurement results.

The simulations and measurement are conducted at a Reynolds number of $Re=3.2 \cdot 10^6$ with the wheel base as the characteristic length scale. The Reynolds number is chosen to allow comparability of the results to previously published results from TUM. In the experimental setup, the Reynolds number of $Re=3.2 \cdot 10^6$ is reached at a free flow velocity of $U_\infty=42.6\text{m/s}$ and an air temperature of $T=8.5^\circ\text{C}$ in the wind tunnel measurement during the winter. The maximum Reynolds number that can be achieved in the WTA is limited by the maximum speed of the moving belt in combination with higher air temperatures and thus lower air densities during the summer. The maximum rolling-road system belt speed is reached at 50m/s. The wind tunnel is not equipped with a heat exchanger, so the temperature of the oncoming flow cannot be adjusted. Due to the steel tube wall of the wind tunnel, outside temperatures are quickly reached in the test section of the wind tunnel. For long measurements or multiple measurements in a series, temperature increases of up to 15°C compared to

outside temperatures are to be expected due to turbulent dissipation and fan engine heat losses. For a single experiment lasting 50 seconds, temperature variations can be neglected and are averaged to calculate a mean Reynolds number.

3.2 Simulation Setup and Results

The DrivAer simulation is executed using the OpenFOAM® software package. Discretization schemes and SA-DDES turbulence modeling are implemented as described in sections 1.2.2 and 1.2.3. The 1:2.5 scaled DrivAer model geometry as used in the wind tunnel test is placed in the domain with around 40 vehicle lengths of size in the x-direction, 32 in the y-direction and 20 in the z-direction. These distances are necessary to avoid interferences with the boundaries of the domain due to blockage effects. The mesh is constructed using an unstructured but hex-dominant grid. To capture geometrical details on the surface of the vehicle, the volume mesh near the surface is successively refined to a hexahedron edge length of approximately $\Delta x_{\min} = L_{\text{wheelBase}}/1,000$.

Furthermore, prism layers are extruded on the surface for the adequate boundary layer representation using the nutUSpalding wall function as described in section 1.2.2. The next-coarsest refinement region encloses the entire vehicle geometry in a surface-fitted region; see Figure 27 for the opaque, box-shaped geometry closest to the vehicle. It captures the small-scale fluctuations close to the surface and resolves steep flow field gradients (e.g., the static pressure increase close to the stagnation point in the front). A one-level-coarser mesh region encloses the vehicle and large parts of the most dominant structures of the vehicle wake, which is also depicted as an opaque, box-shaped geometry in Figure 27. The same refinement level is used on all no-slip floor patches and surroundings. Further, incrementally coarser grid refinement levels toward the outer base mesh size are designed as scaled versions of the previous refinement levels.

In addition to the vehicle and the model support system, the solid walls surrounding the moving belt system are modeled as steady, solid, no-slip walls. The rest of the bottom boundary conditions are implemented as a translating wall velocity boundary condition. This way, the removal of the boundary layer with the wind tunnel’s passive boundary layer scoop in the experimental setup is approximated. The oncoming velocity profile in block shape does not produce a boundary layer on the bottom boundary up to the solid wall upstream of the moving belt, where a boundary layer is generated. Further downstream, the boundary layer is again accelerated by the moving belt and washed out, and parts of the momentum deficit are recovered. The velocity increase due to displacement thickness becomes smaller before the flow reaches the model. The exact modeling of the flow field effects in this region is difficult since the entire wind tunnel collector geometry is ignored.

The explicit modeling of the wind tunnel nozzle and collector geometry is omitted here for several reasons. First, the goal of this work is to extract flow

field effects and to gain knowledge about dominant instationary detachment processes that can be transferred to other vehicles. Second, the correct approximation of flow structures of the full wind tunnel nozzle, test section and collector geometry would require an incredibly large number of additional cells, which would make this investigation prohibitively expensive for the use of the SA-DDES turbulence model employed here. Resolving the classic Seiferth wings at the nozzle, the formed free jet in the entire test section and the reattachment of the free jet in the collector would each result in similar grid cell requirements for the vehicle geometry. Last, the developing recirculation regions in the test section driven by the free jet take a very long physical time to develop fully. Modeling this process in a simulation to the full extension would lead to overwhelming computation times.




Figure 27: DrivAer simulation setup including all no-slip wall geometries, street boundaries, and second- and third-finest grid level mesh regions. Dark patches on the bottom (street) are modeled using translating velocity boundary condition walls. Bright gray patches are solid, no-slip, steady walls.

The simulation time step size is chosen as $\Delta t_{\text{CFD}} = 0.0016 \cdot L_{\text{wheelBase}} / U_{\infty}$, which leads to a good compromise between Courant number limitation and simulation execution times. The CFL=1 criterion is exceeded in highly refined cells close to the surface where the turbulence model operates in RANS mode for the majority of cells.

The initial flow field is computed using a potential flow solution (i.e., inviscid flow without friction and irrotational velocity field). Additionally, uniform-thickness boundary layer profiles are added to all solid walls, which further reduces the time for convergence toward a physical flow field in the first instationary time steps. After an initiation of two seconds physical flow field time, the turbulent flow field in the wake is properly and fully developed. After two seconds, the flow field is then averaged in time, and integrated forces are recorded for the creation of averaged force coefficients. The drag force coefficient acting in the x-direction C_D is defined through the surface-integrated drag force F_D and the projected surface area of the vehicle in an x-plane A_x as

$$C_D = \frac{2F_D}{\rho A_x U_{\infty}^2}. \quad (96)$$

Similarly, the front and rear lift coefficients are defined as

$$C_{L,\text{front}} = \frac{2F_{L,\text{front}}}{\rho A_x U_\infty^2} \quad (97)$$

and

$$C_{L,\text{rear}} = \frac{2F_{L,\text{rear}}}{\rho A_x U_\infty^2}. \quad (98)$$

The lift forces F_L on the front and the rear are evaluated through momentum balancing of all surface-integrated forces around a line parallel to the y-axis through the mid-point between the wheel-ground contact patches. Figure 28 shows the convergence of the force coefficients toward the averaged force coefficient. For both the simulation and the experiment, the rear lift quantities show the most significant fluctuations. The converging average of the rear lift coefficient from the simulation takes 0.5 seconds before it stays within one drag count bound from the mean value. In the experiment, it takes almost 19 seconds to reach this state. The actively changing coefficients in the experiment are also shown later in the frequency spectrum of Figure 31.




Figure 28: Surface-integrated force coefficient convergence over time for simulation and experiment. The graphs show the convergence to the temporal mean ($\Delta C = C - C_{\text{mean}}$) of the respective force coefficient. Dotted horizontal lines represent a deviation from the temporally averaged value by one count.

A comparison of the temporally averaged surface-integrated forces from simulation and experiment is shown in Table 1. The force coefficients from the simulation contain the surface-integrated forces of all cells. The experimental coefficients are collected from the internal balance inside the DrivAer model.

	C_D	$C_{L,\text{front}}$	$C_{L,\text{rear}}$
Simulation DDES	0.245	-0.085	0.047
Experiment WTA	0.231	-0.037	0.050

Table 1: Experimental wind tunnel testing results from the WTA of TUM versus simulation results. Experimental drag values are rolling resistance corrected. Lift forces are without the lift contributions of the wheels but include the drag contributions of the wheels through the momentum balance.

There are multiple uncertainties on the experimental side that need to be addressed to judge the comparability between temporally averaged integrated forces from simulation and experiment. The topics can be separated into drag- and lift force-component related issues and uncertainties related to boundary conditions of the experiment.

The missing collector in the simulated geometry is the most significant contributor to the offset of the drag value. It has been shown by Collin et al. [77] that a steep pressure gradient upstream of the collector of the WTA increases the base pressure behind the DrivAer body significantly when compared to measurements of the same 1:2.5 scale model, measured in the full-scale Audi Aeroacoustic Wind Tunnel (AAWK), which has a much longer test section. This effect leads to an under-prediction of drag values in WTA measurements. The investigations by Collin et al. show a difference in drag values of seven drag counts between the wind tunnels and an almost exact match between CFD simulations and AAWK measurements. For this study, a DrivAer fastback configuration almost identical to the present setup is used. The evaluation of the velocity is based on the plenum method. Compared to the nozzle method, which is biased by nozzle blockage effects, the plenum method returns force values that are approximately 2% lower. It is worth noting that massive volumetric extension of the wheel suspension acts like a volumetric blocking of the wheelhouses, modifying the flow in the wheelhouse from what is seen in the wheels-off configuration.

The evaluation of rolling resistance in the wind tunnel measurement leads to additional uncertainty, which can be quantified and ruled out. The rolling road resistance is measured by doing a force measurement with different belt speed without any wind. The belt drags air with it and thus induces aerodynamic forces at higher speeds with a quadratic increase of the measured force with the belt velocity. The ventilation movement is another contributor to the quadratic part in the measured forces in the moving belt without wind setting. As shown in Figure 29, by least-squares fitting a second order polynomial curve onto the rolling road resistance measurement points for different velocities, the quadratic term (i.e., aerodynamic forces) can be removed, and the rolling road resistance remains. Once the wind tunnel is turned on, the linear and constant part of the rolling road resistance measurement can be subtracted from the measured drag force, with the assumption that the rolling road resistance is still the same. That is not the case, since the wheels themselves observe a lift force, changing the loading on the contact patch between wheel and belt. This lift force is unfortunately not measurable in the current WTA setup. In the simulation, however, the front wheels observe nine lift counts upwards, leading to a vertical force of 1.8N. With a pre-loading of the contact patch around 20N (including the wheels weight force), this decreases the rolling resistance by 9% because the normal force on the contact patch is in linear relation to the resulting rolling resistance. The measured drag of 84N is increased by a negligible 0.2N. This deficiency in the wheels-on setup is also present in the wheels-off configuration, but even more pronounced, because the symmetric wheel arm airfoil add a vertical force on the wheels that are also not measured. Changes in the rolling resistance and aerodynamic drag on the wheels result in changed lift forces through momentum balancing around the midpoint of the wheel base on the ground. Horizontal aerodynamic forces acting on the wheels are accounted for in the lift balance of the vehicle, but lift forces on the wheels are not. Due to the

small relative error introduced by these effects, this uncertainty can also be neglected.




Figure 29: Rolling resistance evaluation of the DrivAer in the wheels-on configuration in the WTA of TUM.

Comparison of unsteady fluctuations to the present wind tunnel geometry setup is difficult to obtain because the top sting mounting system behaves like a free-swinging pendulum, and solid body eigenfrequencies overlay the aerodynamically induced forces. The mounting system through the top sting is not perfectly stiff. While the boundaries of the vehicle in the simulation stay fixed, external forcing leads to an excitation of the entire experimental setup, which makes the DrivAer surface move. Without exact values measured, the maximum amplitude of displacement in the y-direction of the DrivAer surface close to the front end and rear end due to torsion around the z-axis is in the range of 4mm for the wheels-off setup and around 2mm for the wheels-on setup.

Also, the speed control of the rolling road system is not able to hold the belt speed constant at all times, especially for higher velocities. Therefore, the belt speed fluctuates with up to ± 0.1 m/s (read from the control display). The ever-changing belt speed leads to irreproducible time-dependent boundary conditions that affect both, unsteady lift and drag values. The belt speed control and nonuniformities in the belt (e.g., at the location of the belt joint) introduce fluctuations in the x-direction.

The imbalanced wheels introduce fluctuations in the x-direction as well through the suspension system. The frequency of wheels rotation is at 52.56 Hz and $St=1.37$ respectively and is observable as the most significant peak in the frequency spectrum of the experimentally measured force coefficient.

Very low-frequency fluctuations in the y-direction occur due to the belt track control system, and these forces are also transferred to the internal balance from all four wheels. Belt control-induced fluctuations of the RRS are directly transmitted through the wheels to the model balance in phase in the wheels-on configuration because all four wheels are connected to the body in y-direction through the suspension system.

All previous solid body fluctuations induce oscillations, preferably at the vehicle mounting system eigenfrequencies, which are larger than aerodynamic forces by several orders of magnitude. While those strong, solid body forces can be measured in the internal balance due to the spatial displacement, flow

structures might not be able to overcome the inertia of the vehicle to induce measureable oscillations at higher frequencies.

The wheels-on setup is not as prone to spatial displacement because it allows for less movement of the DrivAer body relative to the ground in the y-direction. The downside of the wheels-on setup is that forces acting on the wheels are not recorded separately but are contained in the internal balance of the DrivAer body. Overall, the experimental wheels-off configuration produces similar magnitudes of oscillation as the wheels-on configuration due to eigenfrequencies of the measurement system. The standard deviations of the drag values for the wheels-on and wheels-off configurations are both in the same order with a value of around 0.07. The drag force coefficient oscillations of two wind tunnel measurements are evaluated using FFT in Figure 30.




Figure 30: PSD force spectra from FFT of the drag force coefficient during a measurement of the wheels-on configuration with and without wind. The rolling road system is running in both cases. Strouhal numbers are computed using a free stream velocity of 42.64m/s for both cases.

Because the current work is not explicitly targeted at the validation of unsteady simulation data, a match between experimental and numerical results is not aimed for here. Comparing the two spectra, it is apparent right away that the experimentally obtained forces are dominated by solid-body vibrations, since the peaks are at the same frequencies for the experiment with and without wind. The oscillations up to $St_{L,wheelBase}=0.2$ seem to be supported by aerodynamic forces, leading to an increase in the solid-body eigenfrequency oscillations in this area. Even though an assignment of aerodynamic forces to peaks is not visible in this comparison, Reiß [78] was able to show a shift in a narrow peak that can be assigned to vortices from the wheels when thickening the wheels in the y-direction.

Figure 31 shows the frequency spectra of integrated forces from the simulation results. The most considerable oscillations in the PSD spectrum of experimental data is larger than in the simulation data by three orders of magnitude. The most dominant oscillations occur in the range up to $St=4$ for

the simulated data. The most dominant frequency in the simulation for the drag and the front lift occurs around $St=0.2$. Comparing to the volume FFT from Figure 34, it is noticeable that there is no steep decrease in the frequency spectrum in the integrated forces after $St=1.5$. Some aerodynamic processes that are relevant for the temporal evolution of the surface-integrated forces are not revealed explicitly as peaks in the averaged frequency spectrum of the velocity field.




Figure 31: PSD frequency spectra of the DDES simulation. The integrated drag, front lift, and rear lift force coefficient spectra are generated using FFT with minimum 50% overlap Hanning windowing for the time frame from 2.2s to 3s simulated time. The lift and drag forces do not contain contributions from the wheels' lift forces.

Even though there is no literature reference with the same geometrical configuration and the same Reynolds number, Table 2 shows a comparison of the statistical values of the integrated force fluctuations with data from other publications. Strangfeld et al. use a 25% scale DrivAer model in the wind tunnel of Technical University of Berlin for their time-resolved force and pressure measurements [79]. Their fastback model is in mockup configuration including DrivAer rims, side-view mirrors, steady wheels and fixed ground plate. They do not find any dominant frequencies in their force and pressure measurement data, which is recorded at a sufficient rate of 2500Hz. They also argue that the inertia of the vehicle might be too large to detect flow frequencies. Their measured maximum deviation from the mean drag value is below 5%, which is connected to the much lower Reynolds number of $Re_{L,wheelBase}=1.7 \cdot 10^6$ compared to the

$Re_{L,wheelBase}=3.2 \cdot 10^6$ in the present study. Simmonds et al. [80] obtain slightly larger standard deviations in their simulations with an also somewhat more complex DrivAer geometry than in the current work (i.e., with DrivAer rims, structured underbody and side-view mirrors). Additionally, their cases are run at a larger Reynolds number of approximately $Re_{L,wheelBase}=4.9 \cdot 10^6$. These complexities in the model and increased Re are expected to add oscillations in the force coefficients, so an increased standard deviation is not surprising. The maximum difference from the mean drag value is in close agreement to the ones obtained from Simmonds et al..

	DDES	Strangfeld et al. [79] experiment	Simmonds et al. [80] simulation	Ruettgers et al. [81] LES simulation
Standard deviation $\sqrt{\frac{1}{N-1} \sum_{i=1}^N (C_{D,i} - \overline{C_D})^2}$	0.0058	-	IDDES: 0.0064 LBM: 0.0071	0.005
Maximum absolute deviation of C_D $\max\left(\frac{ C_{D,i} - \overline{C_D} }{\overline{C_D}}\right)$	7.37%	<5%	IDDES: 5.97% LBM: 7.92%	-

Table 2: Unsteady force coefficient standard deviation from experiment and simulation and maximum absolute deviation.

Figure 32 shows several mean flow fields of the DDES simulation in representative slices. Even though the top sting is a symmetric airfoil profile, the momentum loss induced is visible in the $y=0m$ plane and in the shape of the vehicle wake in the $z=0.15m$ plane. The dimensionless static pressure coefficient for the visualization is evaluated using the pressure at the domain outlet as a reference with

$$C_p = \frac{p - p_{outlet}}{\frac{1}{2} \rho U_{inlet}^2}. \quad (99)$$

The total pressure coefficient is defined as the non-dimensionalized sum of the same static pressure difference and the dynamic pressure,

$$C_{p,total} = \frac{(p - p_{outlet}) + \frac{1}{2} \rho |\mathbf{u}|^2}{\frac{1}{2} \rho U_{inlet}^2}. \quad (100)$$

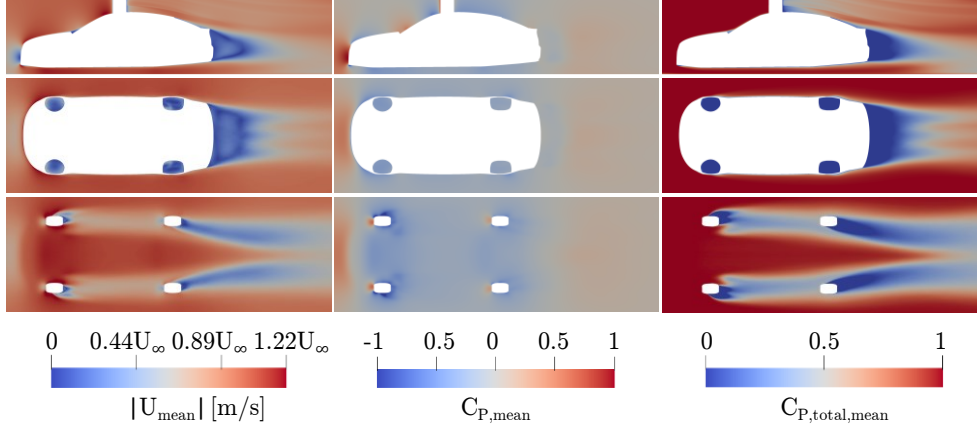


Figure 32: Slices through the temporally averaged mean flow field around the DrivAer body. Averaging period is one second physical time. Top: Slice through $y=0\text{m}$. Middle: slice through $z=0.15\text{m}$. Bottom: Slice through $z=-0.1\text{m}$. Left: Magnitude of the velocity field. Middle: Pressure coefficient. Right: Total pressure coefficient.

The complexity of the temporally resolved flow field of the last simulated time step at $t=3\text{s}$ is visualized using the Q -criterion in Figure 33. The Q -criterion is an often-used criterion for the visualization of turbulent structures. The criterion is defined as the second invariant of the velocity gradient tensor by Hunt et al. [82],

$$Q = \frac{1}{2}(\mathbf{\Omega}\mathbf{\Omega} - \mathbf{D}\mathbf{D}). \quad (101)$$

With this definition, a turbulent structure can be visualized as an iso-surface of positive Q (i.e., the cells in which the strain rate is larger than the rotation rate).

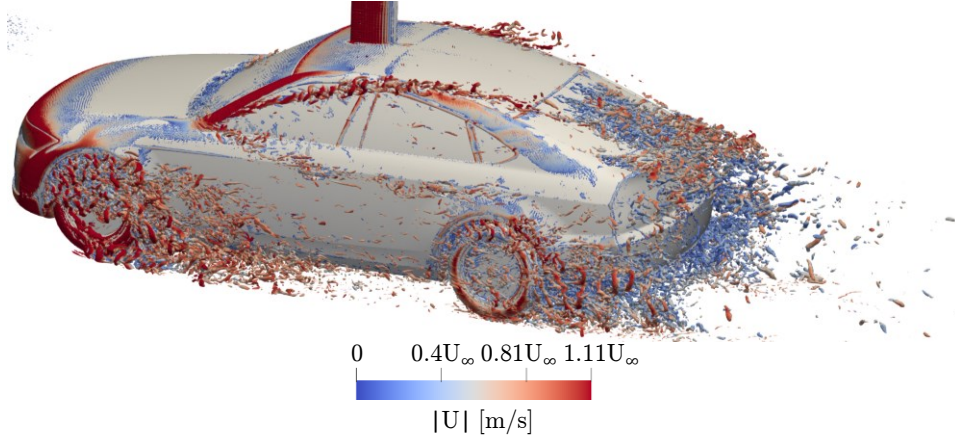


Figure 33: Iso-surfaces of the Q -criterion with $Q=2 \cdot 10^6$ colored by the instantaneous velocity magnitude depicted on the original CFD mesh.

For a first glance at the unsteady properties of the flow field, FFT can be applied to get an overview of the frequency range. For the volumetric flow field data, application of FFT is computationally cheap due to the local type of operations performed (i.e., temporal FFT can be executed on each cell independently and thus scales well in parallelization). To make results

comparable, the x-axis of the frequency spectra is in non-dimensional units of the Strouhal number St . Figure 34 shows the connection between the Strouhal number, based on the wheel base as the characteristic length scale, and the frequencies in Hz obtained from the simulation setup with free stream velocity of 45m/s. The experiment, conducted at the same Reynolds number, is done using a wind speed of 42.64m/s. The FFT results are obtained from an average of spectra by applying three windows of size 512 to each line/observable in the snapshot matrix consisting of 801 columns/snapshots. Each column includes one window at the start, one with its center at time step number 401 and one at the end of the time series. The Hanning windowing function is used in this work. No amplitude or energy correction is applied for FFT results. First, the spectra from three windows of one observable are averaged, and then the spectra of all observables are averaged. A total of 1.6 million cells and three velocity field components from an interpolation grid as defined in section 3.3 are used to form the snapshot matrix. The PSD spectrum shows a decay in energy from low frequencies toward higher frequencies. Oscillations below 150Hz govern the average cell in this subdomain. The fact that there are no apparent peaks visible in the area of 100 Hz or $St=2.5$ is somewhat surprising since bluff body vortex shedding processes from the front wheels are to be expected in this frequency range (i.e., the typical Strouhal number of 0.2 based on the wheel width using the free stream velocity leads to a frequency in that range). For the rear wheels, which observe a much smaller average velocity due to the wake from the front wheels, smaller frequencies in the range of $St=1.5$ are expected.




Figure 34: Averaged FFT spectrum of the velocity components of the flow field from an interpolated sampling box in the vicinity of the DrivAer model. FFT is executed using the same parameters as for the integrated force coefficient spectra in Figure 31.

3.3 Guidelines for Modal Analysis for Vehicle Aerodynamics

This chapter outlines guidelines for the application of modal analysis methods on data sets generated by simulation setups as described in section 3.2. Following these instructions, modal analysis can be employed sensibly as an add-on tool in the aerodynamics engineering process with CFD. The data to be investigated is

limited to areas where deviations from the mean flow field occur. For this matter, only a subdomain around the vehicle surface, including parts of the vehicle wake, is analyzed. The dimensions of the analysis subdomain are defined to contain aerodynamically relevant and energetic structures. The subdomain is therefore not extended all the way into the far downstream wake regions, where the CFD mesh size is increasing, but is instead limited to contain the third-finest grid cells and finer. The entire subdomain is fully submerged inside the larger refinement box depicted in Figure 27. The interpolation mesh on the subdomain for data sampling is shown in Figure 35.




Figure 35: coarseMesh interpolation region definition. The data is mapped to this second region during the simulation before it is passed to the online STDMD algorithm. The lower bound of the interpolation mesh is 0.4mm above the rolling road belt patch.

Spatial interpolation is conducted in parallel during the simulation. After the flow field is solved for on the CFD grid and all equations are converged, the velocity components are mapped to the equidistant grid using cell-volume-weighted interpolation. By heavily parallelizing the interpolation and generating an interpolation map for each processor with fixed cell-volume weighting factors, the interpolation procedure does not add much of an overhead to the actual CFD solution time. This simulation requires approximately 20,000 CPU hours for the investigated time frame and another 82 CPU hours for the execution of the STDMD algorithm. The CPU hours for the simulation include the evaluation of multiple function objects. The majority of CPU hours for all function objects is used for the interpolation to the equidistant grid and writing to HDD with 2,100 CPU hours, of which most time is spent waiting for the HDD writing operations to be finished. The raw data is written to the hard drive to allow evaluation of the performance loss parameter and to visualize the transient flow field data. For parallelization, the interpolation grid is decomposed in the same manner as the CFD grid so that the interpolation tasks are equally distributed among all processors. OpenFOAM® provides the basis for computing these operations efficiently in memory with the built-in multi-region environment. The interpolation grid is also referred to as region coarseMesh.

Spatial interpolation is extremely useful for preconditioning data for conducting modal analysis for vehicle aerodynamics for two critical reasons. First, DMD is very demanding for capturing high-frequency content, especially in the case of the appearance of highly damped or amplified modes. To be on the safe

side, the sampling rate should be 40 times the frequency of interest if extremely unstable processes are involved. Spatial interpolation to an interpolation grid coarser than the CFD grid acts as a spatial filtering operation and introduces smoothing in the temporal domain by invoking cell center information from multiple cells via the cell-volume-weighting method. A simple probing of local flow field values using a single cell center information without any spatial interpolation (e.g., probing by taking the cell center values of the closest cell center) does not suffice for this purpose. To achieve the temporal smoothing effect of the spatial interpolation filter, an interpolation grid coarser than the CFD grid should be used, since each interpolation cell will surely contain information of multiple cells from the CFD grid. The larger the cells are in the interpolation grid, the more significant the impact will be on the temporal information of the data. If only the most energetic, lower frequency, aerodynamically dominant modes are to be investigated and used for ROMs, resolving frequencies that are orders of magnitude higher than the modes of interest would be prohibitively expensive because every single simulated time step would have to be analyzed. This would not be productive. If higher frequency oscillations are to be neglected for generating an ROM, the resulting model surely cannot represent the original data to the full extent, and validation of the ROM is not possible due to the strict sampling conditions. Being able to validate an ROM using a single parameter Π_{loss} is desired, so conditioning the data for the frequency range that can be captured by an ROM is advisable. As the modal analysis only processes data and does not involve the solution of the flow field itself, it should be seen as a way of visualizing instationary data. As long as the most dominant flow field processes that are to be extracted from the simulated data are visible in the interpolated flow field, the respective modes will be present in the ROM that is computed.

For vehicle aerodynamics data, excitation mechanisms of detaching flow structures are often of interest. Because those structures are typically excited from vehicle parts or areas of the flow field that are at least adjacent to the vehicle surface (i.e., the near wake), it is still important to choose a fine enough interpolation grid to be able to observe where structures initially form and detach. For this reason, the cells that are cut by the surface of the vehicle (i.e., wall boundary condition patches) are kept in the interpolation grid, while all cells inside the vehicle geometry are removed. The coarseMesh interpolation grid and a single interpolated time step for a $z=0$ slice are shown in Figure 36.




Figure 36: Slices through the $z=0\text{m}$ plane of the DrivAer simulation. Top: Hexahedral equidistant interpolation mesh `coarseMesh` versus CFD mesh. Bottom: Velocity magnitude field. Visualized are the cell center values (i.e., not interpolated). The cell size of the `coarseMesh` interpolation grid is around $\Delta x=10^{-2}L_{\text{wheelBase}}$.

The direct comparison shows that spatial filtering reduces the number of small-scale fluctuations while maintaining more massive energy-containing structures. Spatial interpolation to a coarser grid results in a low-pass temporal frequency filter (i.e., more significant energy-containing structures are maintained in the data) because in turbulence, small-scale fluctuations are related to low-energy, high-frequency content. Previous investigations show that conventional DMD can only be applied to pre-processed data from DES simulation if sufficient low-pass temporal frequency filtering is used, removing numerical noise and high-frequency content. Because the augmented snapshot-based approach considered in this work is much less prone to errors from noisy data, there is no need for a low-pass temporal cut-off filter as is the case for conventional DMD. Furthermore, the mode-ordering technique proposed here replaces the traditional mode ordering (i.e., ordering by the amplitude computed by the method of the first snapshot), reducing the chance for mistakenly judging highly damped modes for being dominant.

In addition to the spatial and temporal low-pass filtering effect, the interpolation provides a useful tool for reducing the amount of data to be analyzed and thus overall processing time. For the interpolation grid `coarseMesh` used here, the reduction narrows from 106 million cells to 1.6 million. This data reduction allows for fast processing of many time steps and reduces the required storage for the modes that are computed. Finally, the visualization of modal analysis results can be done much faster and with smoother iso-surfaces. This

way, coherent structures are easier to find and follow through reconstruction of modes than if original CFD grid flow field data is used.

If only the most dominant flow structures are to be analyzed using modal analysis, only a selection of simulated time steps need to be processed. The temporal correlation between time steps can be used to make sure that two following time steps are still related to one another and a full state representation of the flow field is possible. For finding the required spacing between time steps to be used in the modal analysis, the flow field inside the interpolation mesh coarseMesh and a set of cells of the CFD mesh that is located inside the outer boundaries of coarseMesh, named subDomainMesh, are written to the hard drive and investigated by computing the correlation function between snapshots. In addition to coarseMesh, a second interpolation grid coarsestMesh is investigated with twice the cell size of coarseMesh. All data sets are recorded using the same time frame of simulation data. Figure 37 shows the normalized correlation function C_i , defined as

$$C_i = \frac{\mathbf{x}_1 \cdot \mathbf{x}_i}{\mathbf{x}_1 \cdot \mathbf{x}_1}.$$

All data is sampled from the DrivAer simulation starting from $t=2.00004s$ to $t=2.1s$ for the CFD mesh data recorded on subDomainMesh (partially) and interpolated data on coarseMesh and coarsestMesh.




Figure 37: Top: Temporal correlation function for the first 145 consecutive time steps recorded in a subdomain and two interpolated meshes. Bottom: Temporal correlation function for the first 2,500 consecutive time steps recorded on the two interpolated meshes.

Online temporal filtering is another option for preparing data for modal decomposition. For conventional DMD, low-pass filtering is an excellent tool to

take out oscillations at frequencies that cannot be adequately resolved by DMD. The Butterworth filter, for instance, shows good cut-off frequency properties. Being able to apply this type of temporal filter relies on the possibility to apply the filter forward and backward in time to eliminate phase shifting. In the incremental variants, only a single time step (or possibly a small range of time steps) is available in memory, and backward filtering with a low-pass filter is not possible. Applying an online Butterworth filter operation in only one direction leads to phase shifting effects and makes an analysis of structure interaction difficult.

For online modal analysis algorithms, in addition to the spatial interpolation, a narrow temporal box filter can be added for preconditioning the data. Temporal filtering has low computational costs due to the local nature of the processor operations involved, so the temporal filter can be applied using time steps of the CFD grid before the temporally filtered field is mapped to the interpolation grid after every Δt_{DMD} . Narrow window temporal box filtering, which is also referred to as moving average filtering, only has a minor effect on the resolvable frequency spectrum of interest if the window size is small enough. The frequency response function H of a moving average filter with window size T_{mav} can be computed as a function of the angular frequency ω , as shown here:

$$H = \frac{1}{T_{\text{mav}}} \frac{1 - e^{-i\omega T_{\text{mav}}}}{1 - e^{-i\omega}}, \text{ with } \omega = [0.. \pi]. \quad (102)$$




Figure 38: Frequency response of a moving average temporal filter. Window sizes are in units of the number of simulated time steps.

For a window size of 25 CFD time steps, the DMD resolvable frequencies don't change much. In this case, the temporal filter only acts as a smoothing function on the flow field data, flattening wrinkles in the resulting mode plot iso-surfaces. Smoother mode distributions can help to identify vortex-shedding mechanisms and simplify the analysis of interaction processes. In some cases, it also is found that applying a moving average filter with window size 25 can reduce the performance loss parameter for generating an ROM with a limited number of modes.

However, if strict low-pass filtering is needed (i.e., because of energetic oscillations in the higher frequency range that cannot be captured by modeling the data basis in a DMD), moving average window filtering is insufficient because of the non-zero frequency response for higher frequencies. Gauss and Butterworth filters are better approaches for this purpose. Low-pass Butterworth filters offer

very steep frequency response curves close to a defined cut-off frequency, rejecting higher frequency content while maintaining the amplitude of the oscillations in the frequency range to be analyzed almost uniformly. The downside of the application of such operations in one direction in time is that phase shifting occurs at different levels for different frequencies. For the analysis of a wide frequency range, the resulting ensemble of structures becomes distorted (i.e., the temporally resolved flow field with vortex-shedding mechanisms at a wide range of temporal and spatial scales changes). Afterward, tasks like data compression through ROM, structure-structure interaction or even causality analysis could become problematic.

For the higher frequencies, the previous considerations present useful tools for fulfilling the requirements of DMD for capturing the full physics of all dominant flow field processes adequately. The lower frequency limit representable by DMD modes on the other side is set by the total period simulated. As demonstrated for the two-dimensional cylinder test case, DMD is out-performing conventional tools like FFT for resolving low-frequency oscillations. As described in Schmid [27], a representative data basis for the complete description of a physical system and sensible evaluation of DMD modes is required. Schmid argues that if the number of snapshots is large enough, the last time step \mathbf{x}_m of a snapshot matrix can be approximated using a linear combination of previous time steps \mathbf{x}_1 to \mathbf{x}_{m-1} as

$$\mathbf{x}_m = a_1 \mathbf{x}_1 + a_2 \mathbf{x}_2 + \dots + a_{m-1} \mathbf{x}_{m-1} + \mathbf{r} . \quad (103)$$

As m increases, the 2-norm of the residual vector \mathbf{r} decreases (i.e., the snapshot \mathbf{x}_m can be approximated better). For very long periods, the columns in the snapshot matrix eventually become linearly dependent, and the residual converges. Following the Ph.D. thesis of Alenius [39], the convergence of the DMD residual can be visualized as displayed in Figure 39. For a more evident convergence behavior, the last time step is held constant for all evaluations of the residual. A total period of 0.8 seconds is sampled for modal analysis in this work. A particle in the outer free stream of the flow at 45m/s traverses the length of the model scale vehicle around 20 times during the sampled period. The lowest dominant frequency to be expected from a simple bluff body aerodynamics view with $St=0.2$ for the vehicle's total length L_{length} is at 4.9Hz, resulting in a recurrence of this possible low-frequency shedding mode of roughly four times. Because no bi-stable wake phenomena are expected from this rear end shape, lower frequencies are not likely. For station wagon rear end shapes on the other side, very low-frequency oscillations with strong amplitudes are possible and should be captured, or at least be kept in mind, when analyzing simulation results.




Figure 39: Convergence of the residual norm divided by the number of snapshots used for the DrivAer DDES simulation.

If the lowest frequencies that are contained in the signal cannot be resolved appropriately, DMD results can generally still produce acceptable ROM results by introducing multiple zero-frequency modes with different amplification or decay rates. Keeping all discussed temporal resolution and period requirements in mind, it is then also possible to use DMD as an exact band-pass frequency filtering tool.

3.4 Modal Analysis Results

3.4.1 Incremental POD Modes

This section presents the results from incrementally computed POD analysis using the algorithm introduced in section 2.3. Figure 40 shows the singular value results from incrementally computed POD modes. The POD computation is carried out without any online compression of the orthogonal basis. Opposite to the two-dimensional cylinder test case, the singular values of all oscillating modes are not several orders of magnitude apart. For complex flow field phenomena, almost all modes are required for a good approximation of the original flow field data. Using all modes, the ROM shows a minimal modeling error of only $3 \cdot 10^{-12}\%$, which is comparable to the results from the two-dimensional cylinder case, but a representation of the flow field with a performance loss of $\Pi_{\text{loss}} < 1\%$ is achieved using 775 out of 801 modes. The proximity of singular values leads to a slow decrease in performance loss as modes are added for reconstruction. The enormous advantage of POD over DMD for generating an ROM with few modes for complex data seems to diminish if a large portion of fluctuation energy is to be maintained; this is because of the almost linear decrease of performance loss over a wide range of modes. If a tool for data compression is wanted, POD is not exceptionally useful (i.e., the compression ratio in comparison to merely storing the instationary flow field data itself is not considerable). In addition to the mode distributions (i.e., the left singular orthonormal vectors), which are the same size as the original flow field, the reconstruction of POD modes requires the singular values and the right-singular vector matrix, describing the temporal evolution of

modes, to be stored. For the most energetic fluctuating POD modes, the singular values reveal the appearance of mode pairs, which have singular values in similar ranges.




Figure 40: Singular values and ROM performance loss convergence over the number of modes used for reconstruction for the DrivAer DDES simulation.

The first POD mode is a non-oscillating mode that is almost identical to the mean flow field. That is the reason why the first POD mode is commonly referred to as the mean flow mode. In this ROM, the mean flow mode includes a slightly changing right singular vector, which makes the reconstruction of the mode a better approximation of the full flow field compared to the arithmetically computed mean flow field. The obvious strategy for the analysis of POD modes is to start with the leading oscillating modes sorted by amplitude. The temporal evolution of modes can be studied using the right singular vectors. Figure 41 shows the temporal evolution of the first nine leading POD modes. Mode pairs are plotted in the same window to show their spectral similarities.

Further fluctuating mode pairs (i.e., modes 10 and 11) show increasing differences in their spectral content. While modes two and three show almost identical frequency spectra, modes eight and nine differ significantly. The differences in frequency spectra are partly due to the application of window function in FFT, but the diverging behavior can also be seen in the absolute values of the right singular vectors over time. The mean flow mode shows negligible changes in time that are smaller than the first genuinely oscillating modes appearing in pairs by six orders of magnitude. All nine modes shown here lie well within the expected active frequency range seen in the FFT analysis in Figure 34. Higher-frequency content is not contained in any of the modes and low-frequency content close to $St_{\text{wheelBase}}=0.2$ is not dominantly visible.

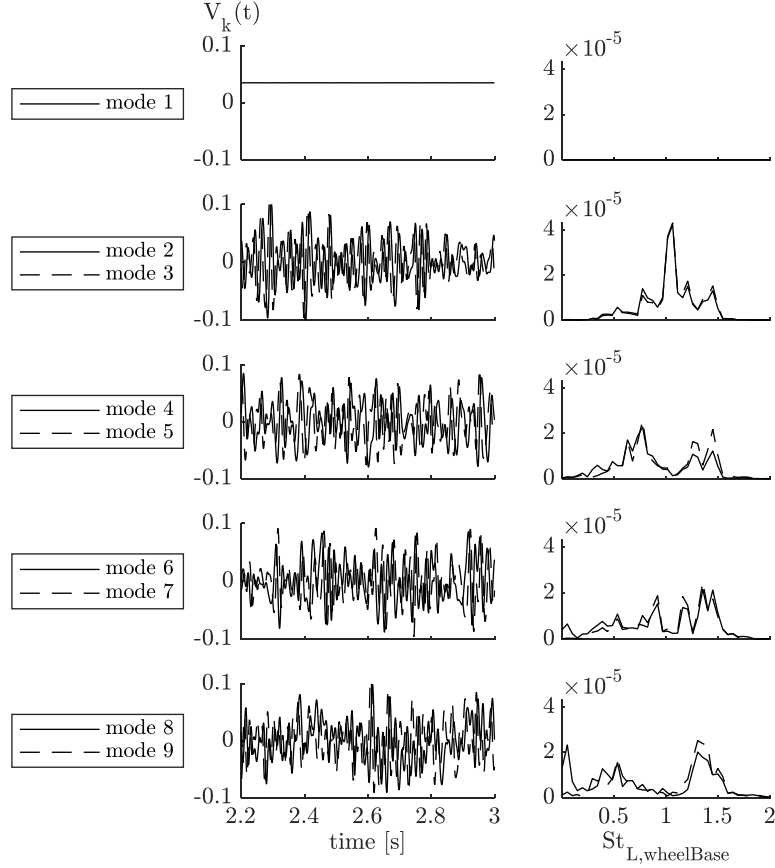


Figure 41: Left: Temporal evolution of the leading right-singular values separately for each mode. Right: Respective PSD distribution plotted over the frequency range between 0Hz and 80Hz. PSDs are computed using three Hanning windows with size 512 and averaging the resulting distributions.

Due to the spatial orthogonality and the real value nature of a single POD mode, reconstruction of a single POD mode does usually not lead to knowledge about the origin of excitation mechanisms. Single POD mode reconstruction results in local oscillations. Convective transport of structures cannot be observed. A combination of at least two POD modes is necessary for meaningful investigation of detachment processes and extraction of structures that are convected downstream. For the analysis of the extracted structures, temporal reconstruction can be used to create iso-surfaces or slices of mode component flow fields (i.e., reconstructed velocity field components u_i , v_i and w_i for time steps $i=[1..m]$). Figure 42 and Figure 48 show such reconstructions for the first time step. The iso-surfaces in this work are chosen to be equal to plus and minus two and a half times the standard deviation of the respective field component in the first time step.

The y-component is the most active for the mode pair. Large coherent structures in the lower part of the wake fluctuate in y-direction similarly to the structures seen in the two-dimensional cylinder case. The x-direction is slightly

less active, and the z -direction is much less active, at least in the first time step. Even though the y -component suggests a correlation to the vehicle width of the excited structure, the excitation mechanism seems to be approximately at the z -location of the axle. Further inspection of the x -component in the bottom view of the vehicle reveals two bluff body-like structures in the distant wake behind the wheels. Even though the structures do not seem to be spatially connected to the rear wheels, the structures that are formed look much like the two-dimensional cylinder structures for each rear wheel.

The most striking finding is that the wheel wake structures seem to be shed precisely in phase, which is the very reason why the y -component structures can connect nicely and move large air packages in the y -direction. If those structures were to shed in anti-phase, the resulting y -component structure would be much smaller in size and possibly also in energy. A similar effect is investigated in a paper about active flow control on a simplified bluff body by Pastoor et al. [83].




Figure 42: First oscillating POD mode pair. Top: slice through the $z = -0.1\text{m}$ plane. Bottom: slice through the $y = 0$ plane. The reconstructed first time step of a combination of POD mode two and mode three. Left: x -component. Middle: y -component. Right: z -component.

The critical issue with the reconstruction of POD modes is the ever-changing temporal behavior of individual modes and mode pairs (i.e., the contribution of modes two and three to the flow field decreases over time, as can be seen in Figure 41). If the interest of investigation lies within the evaluation of temporal evolution of dominant structures, POD modes do not necessarily decrease complexity to a minimum, since their behavior changes over time. The reconstruction of mode pairs can result in spatially connected structures that are convected downstream, so the identification of source mechanisms for dominant shedding processes is possible in some cases. Some mode pair reconstructions, though, show structures that cannot be traced through the flow field, and they appear disconnected from the vehicle surface in the wake. The disconnection from the surface makes it difficult to assign a mode to a specific shedding mechanism, as a production type car has many geometrical details that can cause separation. Because the reconstruction of mode pairs does not lead to ever-recurring flow structures, all snapshots from the original data sets (i.e., 800 snapshots, in this case) need to be reconstructed and analyzed. For the present data set, reconstruction to a video with five frames per second is useful for the first four mode pairs and yields videos 160 seconds long filled with more or less complex reconstructions.

The complexity of the reconstructions depends on the frequency content and temporal evolution of the right singular vectors. If both modes of the mode pair oscillate with a single frequency, or at least with a narrow frequency band across the entire time frame (e.g., as is in the $Re=100$ two-dimensional cylinder case part two), then the reconstructed flow field can be interpreted quickly, as the processes are ever recurring and only one period of oscillation needs to be reconstructed. In the DrivAer setup, though, this is not the case. The temporal evolution with different time steps for a slice through the $z=-0.1m$ plane is shown in Figure 43. The decrease in oscillation amplitudes leads to almost zero disturbance of the flow field for this mode in reconstructed snapshot 600. That leads to a complete disappearance of iso-surfaces as well.

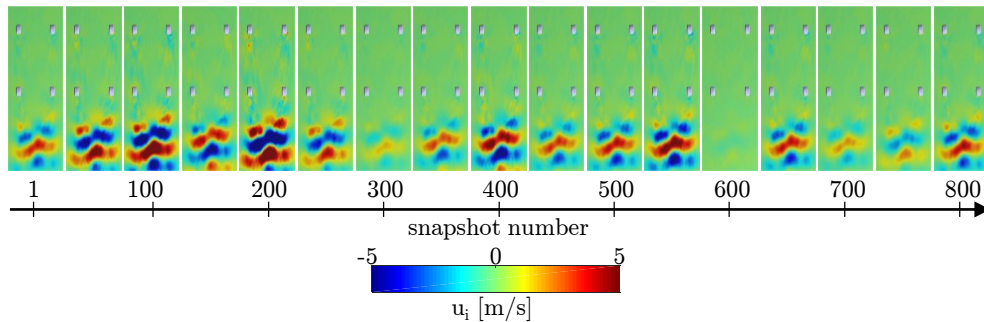


Figure 43: First oscillating POD mode pair. Slice through the $z=-0.1m$ plane. Seventeen reconstructed time steps are shown from a combination of POD mode two and mode three.

The second POD mode pair, which consists of oscillating mode three and mode four, is visualized in Figure 49. The mode pair shows large structures oscillating in the z -direction in the center of the vehicle wake. The structures seem to feed on the underbody airflow. An upward detachment lowers the pressure in the wheel wake, and the air is transported from the wheel wake toward the $y=0$ plane. On downward movement of the air package, the base pressure in the rear wheel wake is increased, and the air is pushed outside. That leads to detachments from the rear wheels that are precisely in anti-phase (i.e., both wheels shed structures toward the $y=0$ plane at the same time).

The third and fourth POD mode pairs show similar behavior to the second mode pair. That is not surprising, since their frequency spectra show dominant peaks at similar frequencies around $St=0.5$ and $St=1.5$. An active region behind the front wheel indicates a connection of the front wheel wake and detachments from the rear wheels, but a clear identification of the shedding process triggering the oscillations in the wake is not possible due to the spatial distance to the geometry surface.

When computational resources do not allow execution of the full rank incremental POD, maximum rank settings should be considered. Beyond the HDD writing and reading time that can be saved by using online algorithms, the maximum rank truncation ability is one of the central benefits of this online-capable algorithm. Table 3 shows the required execution times and resulting performance losses for different memory-saving incremental POD settings.

r_{\max}	incPOD mesh	Execution time [CPU hours]	Π_{loss} [%]
∞	coarseMesh	52	$2.9 \cdot 10^{-12}$
600	coarseMesh	60	3.0
400	coarseMesh	51	4.7
200	coarseMesh	29	6.9
∞	coarsestMesh	7	$3.8 \cdot 10^{-12}$

Table 3: Execution times and performance loss parameters for computation of incremental POD using different maximum rank settings and interpolation grids.

Due to the large computational effort in the compression step, lowering the maximum rank increases the required execution time down to a maximum rank of $r_{\max}=400$. At this point, the benefit of having a smaller orthogonal basis for adding new snapshots evens out the additional effort for the online compression. Furthermore, coarsening the interpolation grid to an even sparser one named `coarsestMesh`, with a cell size twice that of the `coarseMesh` data set, is an excellent method of reducing memory requirements if the flow field processes to be extracted are still visible in the data set. In this case, reducing the number of observables by a factor of eight leads to a decrease in computational time by a factor of eight. These numbers are only exemplary and heavily depend on the extent of parallelization of the algorithm.

3.4.2 Streaming Total DMD Modes

This section presents the results from incrementally computed DMD using the algorithm introduced in section 2.4. Figure 44 shows the frequency spectrum from incrementally computed DMD modes. The STDMD computation is carried out without any online compression of the orthogonal basis. Usage of all 800 DMD complex modes leads to a successful ROM with a total performance loss of $10^{-8}\%$, so the resulting ROM can be used for exact frequency filtering applications. As can be seen in Figure 44, the frequency distribution from DMD is fundamentally different from the FFT spectrum in Figure 34. While FFT offers fixed frequency spacing and a frequency resolution that depends on the window size, DMD can make full use of the entire data set, leading to higher frequency resolution where needed and resolving lower frequency content than FFT. The lack of windowing function usage for DMD gives a much clearer reflection of the actual system behavior because all time steps are considered with equal weight for the decomposition. This way, a single modal DMD basis can represent the entire time window. The resulting frequency spectrum contains two modes with zero imaginary parts (i.e., two mean flow modes). Considering that all oscillating modes are contained in the modal representation redundantly with an adjacent complex conjugate counterpart, the actual modal description boils down to 401

modes (i.e., 399 non-zero frequency modes). One of the mean flow modes has $K_{\text{stable},1}=4.2\cdot 10^7$ and $\lambda_1=1$ and another mode has $K_{\text{stable},401}=2.9\cdot 10^2$ and $\lambda_{401}=0.79$. The appearance of multiple mean flow modes is typical in cases where the data length is not enough for capturing the lowest frequencies. This is usual when the total number of time steps is odd, since the augmented snapshot approach leads to an even amount of orthonormal vectors if online compression is not employed. That leads to the recommendation either to always use an even amount of snapshots for augmented snapshot-based DMD algorithms or to compress the orthonormal basis to an odd amount of orthonormal modes to prevent the appearance of a second zero Hz mode. Even though DMD gives better performance than FFT for capturing very low frequencies (e.g., if only a quarter of a total period of oscillation is contained in the signal), there are cases when low-frequency oscillations are assigned zero Hz frequency with damping or amplification. This way, slowly changing mean flow fields can still be represented correctly by an ROM from DMD modes.

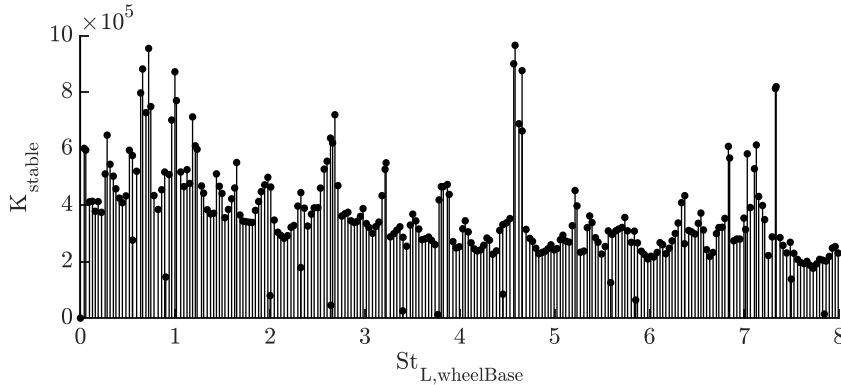


Figure 44: Frequency spectrum of the STDMD analysis of the DrivAer DDES simulation. Modified eigenvalue-weighted amplitude as a function of the Strouhal number. Higher frequencies are not shown.

Another useful visualization of DMD results is the real and complex part of the eigenvalues. The eigenvalue distribution is shown in Figure 45. Most modes are either stable or slightly damped (i.e., with $|\lambda|<1$), but there is a tendency to slightly higher damping values in the higher-frequency range.

Due to the complex-number nature of DMD modes, reconstruction of a single mode in time can yield spatially coherent structures that can be interpreted as clusters of air moving at the same frequency. Often, these clusters are convected downstream with the mean flow and allow the interpretation of a specific DMD mode to be related to a vortex shedding process of a specific part, because well-known von Karman-like vortex streets or shear layer instabilities are formed. Vehicle parts that are elongated in the z -direction, for instance, will usually induce a vortex street, which can most easily be spotted using the y -component of a mode. The same holds for vehicle parts that have a significant extension in the y -direction, which introduces dominant large-scale structures in the z -component of a mode. Figure 50 shows the reconstructed first snapshot of STDMD mode two (i.e., the first oscillating mode picked by the modified

eigenvalue-weighted amplitude ordering method). A von Karman vortex shedding process dominates the mode in the wake of the front wheels. The iso-surface is again selected as two and a half times the standard deviation of the mode distribution in the first time step for each flow field component separately. Taking a look at the resulting iso-values, it quickly becomes evident that the y-component is the most dominant, while the other two components appear in similar magnitude. Even though the rotating wheel in ground contact located in the complex wheel housing geometry is a much more complex bluff body than the two-dimensional cylinder, the dominant detachment process in this reduced order flow model looks relatively similar.




Figure 45: Eigenvalue distribution of the STDMD modes. The angle of the vector in the complex number plane represents the non-dimensionalized frequency. The vector length represents the damping value for each mode. The marker size is linearly proportional to the mode ordering value K_{stable} . Left: full spectrum. Middle and Right: enlarged ranges of positive frequency modes, including the unit circle drawn as a solid line.

In addition to vortex shedding of this Strouhal number dominating the flow around the front wheels, it also affects the global field around the full vehicle. The vortices shed from the front wheels slowly decay toward the rear wheels but partially recover by triggering oscillations at the rear wheel with the same frequency. This behavior is also underlined by a reconstruction of multiple time steps, which shows the front wheel structures passing by the rear wheels, intensifying in energy and then reducing in convection speed (i.e., the structures appear to be more compact and compressed in their x-direction dimension) as the rear wheel wake in the diffuser area slows down the mean flow field. New turbulent energy is generated by the transfer of energy from the mean flow field at the rear wheels at the same time scales as on the front wheels even though the observed mean flow speed around the rear wheels is much lower than in the front. The structures then reach the trailing edge of the diffuser, inducing additional separations at this frequency. Figure 46 shows six reconstructed time steps of the y-component. The possibility to choose non-integer time steps for reconstruction enables reconstruction of a smooth period of oscillation because

the exponents in the Vandermonde matrix can be freely chosen for visualization purposes.

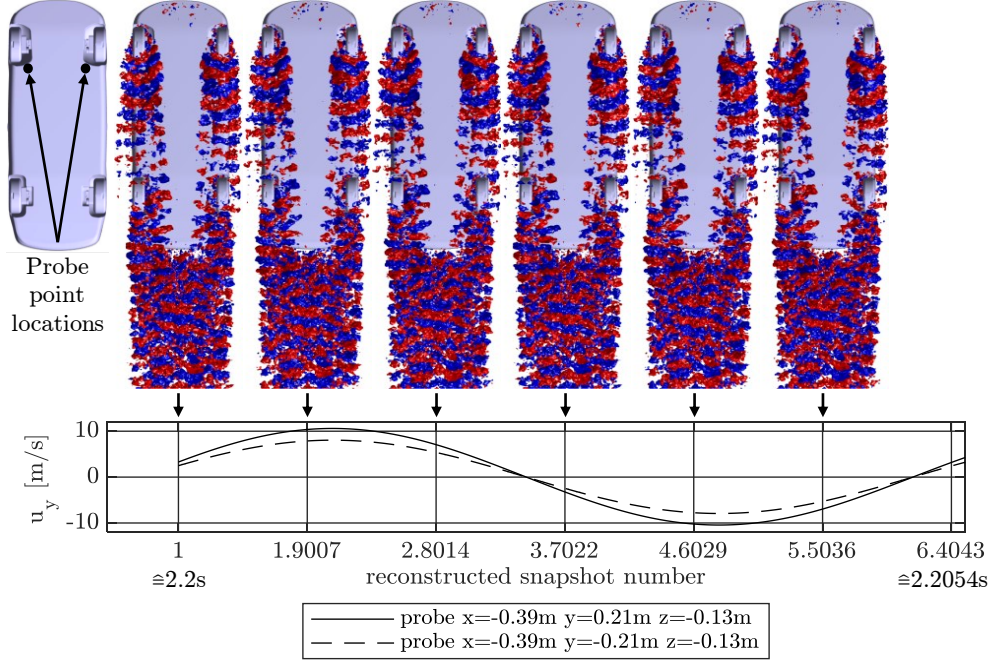


Figure 46: Reconstruction of a full period of DMD mode number two y -component field. Two probe locations in the wake of the wheel/wheel housing area show the strong phase connection between the left and right front wheels. The reconstructed time steps are located exactly at physical times $2.2s$, $2.2s+1/(6f_2)$, $2.2s+2/(6f_2)$, $2.2s+3/(6f_2)$, $2.2s+4/(6f_2)$ and $2.2s+5/(6f_2)$, which do not coincide with simulated time steps.

While the z -component is naturally not well represented in the underbody, except in the wheel housing, the component shows large oscillation emerging from the upper part of the front wheel on the outer side. As is the case for the y -component, the z -component influences the shedding behavior of the rear wheels. Passing by the rear wheels, the structures increase in energy before filling the entire vehicle wake region with oscillations at this frequency. A minor part of the oscillations of the A-pillar vortex also seems to be amplified by this process (i.e., the instabilities that form in the shear layers around the vortex core might be preferentially forming at this specific frequency), which contradicts the widespread opinion that oscillation frequencies of the A-pillar vortices are somewhat erratic.

The first 20 modes contain additional dominant STDMD modes with both a strong influence of von Karman-like y -component vortex structures below the underbody in the wake of the front wheels and a strong z -direction vortex street downstream of the front wheel housings. STDMD modes four with $St_4=5.57$ and $f_4=184\text{Hz}$, mode six with $St_6=4.66$ and $f_6=188\text{Hz}$, and mode 19 with $St_{19}=4.65$ and $f_{19}=187\text{Hz}$ all complement mode two with a slightly offset frequency, showing very similar shedding processes but at lower amplitudes than found in mode two.

Modes with much higher frequency content also appear with similar but smaller structures. These are mode eight with $St_8=7.4$ and $f_8=296\text{Hz}$ and mode

nine with $St_9=7.3$ and $f_9=295\text{Hz}$. The impact on the instabilities in the A-pillar vortex increases at higher frequencies. Interestingly, for the very high-frequency modes, disturbances do not seem to be convected downstream. Instead, they become either almost steady in space or spread upstream. An explanation for this lies in the formation of instabilities that are in a frequency range that cannot be resolved adequately by DMD modes using the current temporal sampling resolution. The structures seem to move upstream, but their convective movement upstream is an artifact of the chosen sampling interval. Similar effects can occur in videos that are recorded with frame rates lower than the Nyquist frequency of the recorded object's movements (e.g., the passing frequency of spokes on a spinning wheel or the flapping frequency of a bird).

The reconstruction of the first time step of a mode with similar shedding characteristics as mode two, but at higher frequencies, is shown in Figure 51 for the reconstruction of mode eight. As the frequency of front wheel driven vortex shedding modes increases, the structures grow smaller (i.e., the size of the y-component bands in the bottom view). Additionally, it can be observed that there is a tendency to earlier (i.e., further upstream) detachment on the front wheel as the frequency increases.

Reconstruction of the STDMD mode three in iso-surfaces is shown in Figure 52. The proximity to the Strouhal numbers contained in POD modes four and five leads to similar, spatially coherent structures, as described for the second POD mode pair above. Large air clusters in the size of the vehicle tail height $L_{\text{tailHeight}}$ move in the vehicle wake. Some of the initial disturbances seem to be sourcing from the front wheel housing area and the upper part of the rear end, but the majority of upstream fluctuation energy comes from the rear wheel housing areas. The frequency of oscillation of $f_2=29\text{Hz}$ can be converted to a different definition of the Strouhal number based on $L_{\text{tailHeight}}$ as the characteristic length scale as $St_{L_{\text{tailHeight}}}=f_2 \cdot L_{\text{tailHeight}}/U_\infty=0.2$. A Strouhal number of 0.2 in bluff body aerodynamics is usually a good indicator that this specific geometry dimension is causing a flow field effect at the respective frequency. The large amplitude in the FFT spectrum of the integrated rear lift forces in the same frequency range indicates that this mode is also to be considered in the context of driving stability. Even if rear lift forces are negative, large scale structures might lead to lift forces that temporally exceed predefined and required bounds for good traction at high velocities. Being able to visualize such structures can help to find possible geometrical modifications for those structures to decrease their impact on the integrated force oscillations.

In comparison to the POD modes presented earlier, DMD modes show better connections to small-scale perturbations of the flow upstream and make it possible to track them as they are convected through the flow field. The substantial contribution from the rear wheels hints that synchronized shedding from the wheels again plays a vital role in supporting the creation of such large-scale structures. The y-direction components are patchy, but their overall size is of the order of the vehicle width L_{width} . As is the case for mode two, multiple modes appear as dominant modes at a similar frequency as mode three. Those

are mode five with $St_5=0.64$ and $f_5=26\text{Hz}$, mode 10 with $St_{10}=0.64$ and $f_{10}=26\text{Hz}$, mode 12 with $St_{12}=0.74$ and $f_{12}=30\text{Hz}$, and mode 14 with $St_{14}=0.69$ and $f_{14}=28\text{Hz}$. Such clusters of modes are also found by Schmid for an application of DMD on a jet between two cylinders [27] and Schmid et al. for the application on a helium jet flow [28]. Reconstructed together the modes form coherent packets. Schmid and Schmid et al. find that such mode clusters form arc-like structures in plots of the eigenvalues and consist of fluid elements that show similar support but have slightly offset wavenumbers.

An additional cluster of energetic STDMD modes is in the vicinity of mode seven with $St_7=0.99$ and $f_7=40\text{Hz}$, mode 11 with $St_{11}=1.02$ and $f_{11}=41\text{Hz}$, and mode 18 with $St_{18}=0.97$ and $f_{18}=39\text{Hz}$. Figure 53 shows a reconstruction of the first snapshot of STDMD mode seven in iso-surfaces of the velocity components. The modes in this frequency range have some contribution of the x-component fluctuations from the front wheels in ground proximity. The rear-wheel shedding is not in phase for the y-components, as opposed to the vortex shedding from the front wheels contained in mode two. The phase setting and the lower free stream velocity seen by the rear wheels is likely to be the driving mechanism for the lower amplitude for the dominant modes from the rear wheels. Surely, modes at lower frequencies are more likely to be sourced by the rear wheels due to the lower surrounding velocity, as can be grasped from Figure 32.

Another interesting dominant mode, for which an in-phase vortex-shedding mechanism from the front wheels plays a vital role, is contained in a cluster around mode 15 with $St_{15}=2.68$ and $f_{15}=108\text{Hz}$. Reconstruction of the first time step in iso-surfaces is shown in Figure 54. The dominance of the y-component structures in the wheel wakes is striking. Furthermore, similar to STDMD mode two, the y-component structures are synchronized and shed in phase, resulting in large-scale oscillations in the y-direction in the vehicle wake that stay connected and almost stable throughout the sampling box. There is a connection to oscillations in the drag and front lift force coefficient.

STDMD mode 16 is visualized in Figure 55. There is a coincidence of peaks with the frequency spectra of the surface-integrated force coefficients for drag and rear lift. A strong connection to detachments to the upper upstream part of the front and the rear wheel housing is evident for the x-component. Further downstream, parts of the fluctuation energy from the wheel housings are transferred to the z-component and y-component. The structures are of the same size as the vehicle tail height, similar to modes three and five, but they are smaller in the z-component, which is again connected to the slightly higher frequency of oscillation.

Considering further memory savings in STDMD, three different maximum rank settings and one even coarser sampling grid are evaluated. Upon reducing the maximum rank, the resulting ROM can only partially describe the original data set and the performance loss increases, as seen in Table 4. For the coarsestMesh STDMD computation, the number of cells is reduced by a factor of eight, leading to a drastic decay in execution time. The coarseMesh and coarsestMesh STDMD ROM have similar performance losses. The highest CPU

time usage is for the $r_{\max}=600$ case, due to the expensive compression stages with 600 columns of the orthogonal matrix. Somewhere in between, the excessive cost for compressing large orthogonal matrices evens out with the decrease in computational cost for the CGSI process, where only r_{\max} columns need to be considered for orthonormalization of new incoming augmented snapshots. The $r_{\max}=200$ case, which is computationally only half as expensive as the full rank run, cannot extract the same flow field structures from the full rank case, and the performance loss deteriorates to almost the level of the mean flow performance loss.

r_{\max}	DMD mesh	Execution time [CPU hours]	Π_{loss} [%]
∞	coarseMesh	82	$1.1 \cdot 10^{-8}$
600	coarseMesh	87	5.31
400	coarseMesh	71	7.2
200	coarseMesh	40	10.6
∞	coarsestMesh	9	$1.0 \cdot 10^{-8}$

Table 4: Execution times and performance loss parameters for computation of STDMD using different maximum rank settings and interpolation grids.

Figure 47 shows the STDMD frequency spectra of the executed STDMD runs. Due to the eigenvalue decomposition-based online compression for the maximum rank settings, high frequency content is reduced. The maximum rank cases with $r_{\max}=600$ and $r_{\max}=400$ can still extract the most dominant structures in the lower-frequency range. The spatial distribution of the resulting modes is also well captured, comparable to the full rank decomposition. The higher-frequency content is reduced in amplitude, and peaks disappear as the maximum rank is reduced. The coarsestMesh case allows for very fast computation of STDMD modes and only slightly reduces high-frequency content. Overall, the magnitude of amplitudes is reduced, which is expected from an interpolation to a grid that is twice as coarse, but the locations are maintained for all the most dominant peaks discussed for the coarseMesh data set. If the flow field structures to be analyzed and their detachment positions on the surface are sufficiently visible, it is advisable to choose an interpolation grid as coarse as possible. This leads to lower STDMD execution times, lower mode storage cost and faster visualizations with smoother iso-surfaces.




Figure 47: STDMD frequency spectra results using different memory-saving measures. Maximum rank settings of $r_{\max}=600$, $r_{\max}=400$ and $r_{\max}=200$ are compared to the full rank decomposition (top). STDMD spectra for data interpolated to an even coarser grid (“coarsestMesh”) with a cell size of $\Delta x=2 \cdot 10^{-2} \cdot L_{\text{wheelBase}}$ is shown for comparison (bottom).

The DMD modes presented here show drastically less complex flow phenomena, which can be tracked through the flow field, and allow respective source mechanisms to be identified and distinguished. Different modes show changing convection paths and interactions in the wake. There is a clear tendency toward larger structures for lower frequencies. Synchronization of shedding mechanisms from the wheels and the wheel house regions appears to affect multiple modes. Additional modes can be assigned to dominant vortex-shedding phenomena related to the tail height of the vehicle, but none of the first 20 dominant modes is related to the vehicle length. The possibility to reconstruct periods of oscillations instead of the entire sampled period adds to the usefulness and seamless comprehensibility of DMD modes. Additionally, DMD mode reconstruction can be obtained by reconstruction of interpolated time steps using non-decimal powers of the complex eigenvalues so that smooth reconstructions can be obtained even for high-frequency modes.

4 Conclusion and Outlook

4.1 Conclusion

ROM strategies for vehicle aerodynamics simulations are developed and implemented in a workflow applicable to large CFD meshes with many time steps, which often equates to significant memory and computational resource needs. Analyzing properties of the various modeling methods, however, leads to the conclusion that incremental algorithms with low memory requirements, such as the new STDMD variant described in this work, can produce meaningful results. Additionally, the possibility to execute the algorithm during the simulation saves HDD writing and reading times because flow field data does not need to be stored. The investigated measures for pre-processing the data and minimum sampling requirements ensure applicability in an industrial environment, in which the cost for such an advanced flow field analysis needs to remain lower than for the flow field solver itself.

While FFT can be used to compute the frequency spectrum in each cell for bluff body vortex shedding, the resulting spectrum does not yield results comparable to POD or DMD. In contrast to those methods, windowed FFT cannot extract spatial structures, retain phase information or reconstruct single modes in a direct or simple way. POD, however, extracts energetically dominant mode pairs and can be applied to simulated results from vehicle aerodynamics. The convergence properties of the incremental variant of POD regarding an energy-based norm are proven the most reliable. Also, the data for POD processing does not need to be filtered in any way for this method to generate an ROM and, therefore, yields robust performance in the presence of noise. For the complex flow around a vehicle, creation of a reasonable ROM requires a large number of modes, so the data compression factors are not striking compared to directly storing the raw data. The broad frequency range content of leading POD modes also produces significantly more complex reconstructed flow fields from mode pairs that are not as well understood as the modes computed by DMD. Reconstruction of a single POD does not lead to any additional insight into its temporal evolution since it represents stationary oscillations (i.e., no convective transport of structures is observed, and an excitation mechanism source can hardly be identified). Due to the single frequency and complex number nature of DMD modes, their single mode reconstruction can usually be interpreted as a vortex-shedding mechanism. This mechanism can be either assigned to a specific part of the vehicle or used to visualize dynamic instabilities in the wake of a vehicle that are resulting from shear layer instabilities or other phenomena in the recirculation region of the bluff body wake. A highly complex turbulent flow field with separation at a wide range of length and time scales can be narrowed down to globally dominant DMD modes, which as a result can be analyzed independently. Due to its periodic nature, the reconstruction of a mode for one period of oscillation is sufficient.

The variant of STDMD introduced here does not require inversion of the full DMD mode matrix for amplitude evaluation and, as a result, requires less computational resources and allows for a priori mode selection by relevance and frequency before the actual modes are evaluated. This leads to drastic memory savings. Even though the augmented snapshot approach increases the robustness of the method in the presence of noise, the new STDMD method requires more thoughts on data pre-processing, sampling interval and sampling frequency in comparison to POD and FFT.

The new method uncovers insights into the dominant flow features of vehicle aerodynamics and facilitates a better understanding of simulated flow field processes. The absence of overlapping structures of different length and time scales leads to significant improvement in visibility of the dominant structures. Bluff body vortex-shedding processes are identified and can be assigned to characteristic length scales of the vehicle geometry.

The research on globally dominant vortex-shedding processes in the flow field around the DrivAer body stimulates the development of geometrical modifications that influence structures in specific modes. Those modifications can reduce drag and possibly result in less excitation of acoustic phenomena. Following these findings, three modifications validated in the TUM WKA and CFD simulations were submitted to the German Patent and Trade Mark Office (Deutsches Patent- und Markenamt). By close resemblance of the DrivAer model with actual production cars, the invented parts can be implemented in real-world applications, leading to innovative advances in vehicle aerodynamics.

4.2 Outlook

Clear identification of the excitation mechanism is often impossible in very low-frequency structures. In such cases, one could apply causality analysis methods to investigate which modes are inducing other modes (i.e., redistributing energy between scales). That shows connection between modes from different frequency ranges and helps explain complex structure-structure interactions. Granger causality analysis [84] could be employed for such investigations as proposed by Tissot et al. [85].

The size of structures in the wake of vehicles is heavily dependent on the phase of modes that are shed. Using the phase and frequency information from DMD modes, one can identify locations and frequencies at which active flow control is beneficial for drag reduction. Multiple successful studies on such approaches have been conducted by Wieser et al. [86] on the DrivAer body using fluidic oscillators. As an extension to active flow control at fixed frequencies, closed-loop feedback flow control is a promising approach for reducing drag with minimal excitation energies. Nisugi et al. [87] have tested such an approach on another generic vehicle geometry observing drastic drag reductions. Pastoor et al. [83] further investigated the underlying mechanisms of drag reduction from feedback shear layer control on a simplified bluff body. These findings should be further researched and investigated with specific modal analysis results to

develop suitable active flow control drag reduction tools for realistic vehicle geometry.

Most of the existing research focuses on flow field analysis based on temporal snapshots of the flow field. One alternative way of defining snapshots is the accumulation of temporal information of single x-plane slices through the flow field in the vehicle wake into one snapshot, allowing construction of a full snapshot matrix from consecutive slices located downstream. The received modes return information about the spatial evolution of the bluff body wake in terms of spatial wavelengths instead of temporal DMD frequencies. Due to a strong connection between spatial and temporal frequencies in turbulent flow, the resulting modes could complement the temporal snapshot matrix-based results.

Another critical application of modal analysis for vehicle aerodynamics engineering is the evaluation of DMD modes for aeroacoustics. Often, little-understood aeroacoustics phenomena, which emerge from the underbody or the side window of a vehicle, are eradicated by passive acoustic damping measures to ensure acoustic comfort. Once the main drivers for those oscillations are well understood, and geometrical modifications are identified to suppress the underlying excitation mechanism, aeroacoustics engineering will shift away from acoustic damping concepts and toward acoustic source management. These studies could benefit from the exact band-pass frequency filtering property of DMD, which permits analysis of structures and their excitation mechanisms for specific frequency ranges.

On the method development side of STDMD-based modal analysis, future work could benefit from the treatment of degenerate modes. Performance loss convergence could be improved by taking into account the orthogonality of a mode with respect to other modes and degradation of degenerate or near to degenerate modes in the mode ordering method. Phase information of all highly active cells in the respective modes could be taken into account to find problematic modes that are in anti-phase, canceling out each other. For treating problematic modes, a strategy to reorthogonalize degenerate modes before amplitude computation could be implemented.

For a better ROM approximation of the flow field, DMD ROMs could be drastically improved for low-rank ROM using optimization techniques known from sparsity-promoting algorithms. That could be achieved with acceptable computational effort using sparse in time and space data sampling, which requires storing a sparse subset of the snapshot matrix in memory. From this sparse data set, one could construct a sparse version of the optimization problem (95) and solve for real and imaginary parts of the amplitudes and possibly eigenvalues to obtain better ROM approximation with a given ROM rank. As starting criteria for the optimization solver, conventionally computed amplitudes and eigenvalues could be used.

Appendix

A. Reconstructions of Incremental POD Modes




Figure 48: First oscillating POD mode pair. Iso-surfaces of the reconstructed first time step of a combination of POD mode two and mode three. Top: x-component iso-surface $u_x = \pm 1.1 \text{ m/s}$. Middle: y-component iso-surface $u_y = \pm 1.2 \text{ m/s}$. Bottom: z-component iso-surface $u_z = \pm 0.8 \text{ m/s}$.

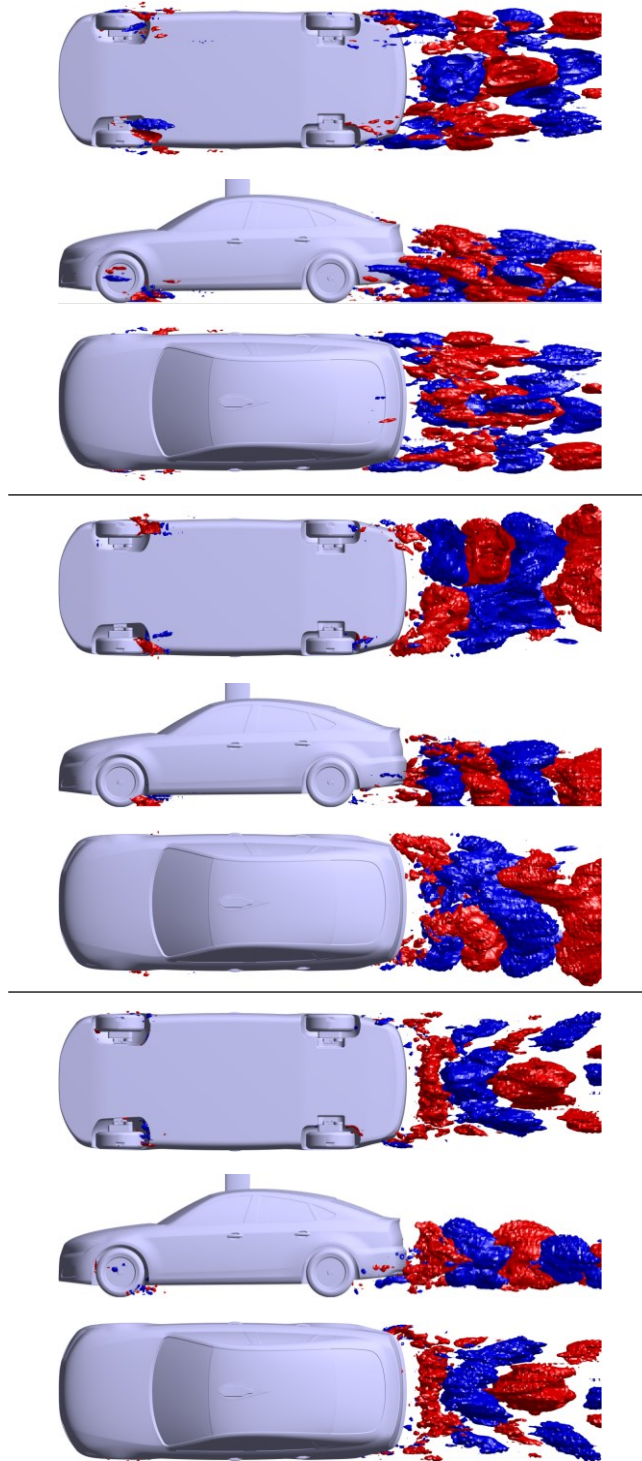


Figure 49: Second oscillating POD mode pair. Iso-surfaces of the reconstructed first time step of a combination of POD mode four and mode five. Top: x-component iso-surface $u_x = \pm 1.4 \text{ m/s}$. Middle: y-component iso-surface $u_y = \pm 1.2 \text{ m/s}$. Bottom: z-component iso-surface $u_z = \pm 1.2 \text{ m/s}$.

B. Reconstructions of Streaming Total DMD Modes




Figure 50: Mode two from the STDMD spectrum at $St_2=4.58$ and $f_2=185\text{Hz}$. Iso-surfaces of the reconstructed first time step. Top: x-component iso-surface $u_x=\pm 12.9\text{m/s}$. Middle: y-component iso-surface $u_y=\pm 13.8\text{m/s}$. Bottom: z-component iso-surface $u_z=\pm 12.6\text{m/s}$.




Figure 51: Mode eight from the STDMD spectrum at $St_s=7.4$ and $f_s=296\text{Hz}$. Iso-surfaces of the reconstructed first time step. Top: x-component iso-surface $u_x=\pm 10.7\text{m/s}$. Middle: y-component iso-surface $u_y=\pm 11.5\text{m/s}$. Bottom: z-component iso-surface $u_z=\pm 11.2\text{m/s}$.

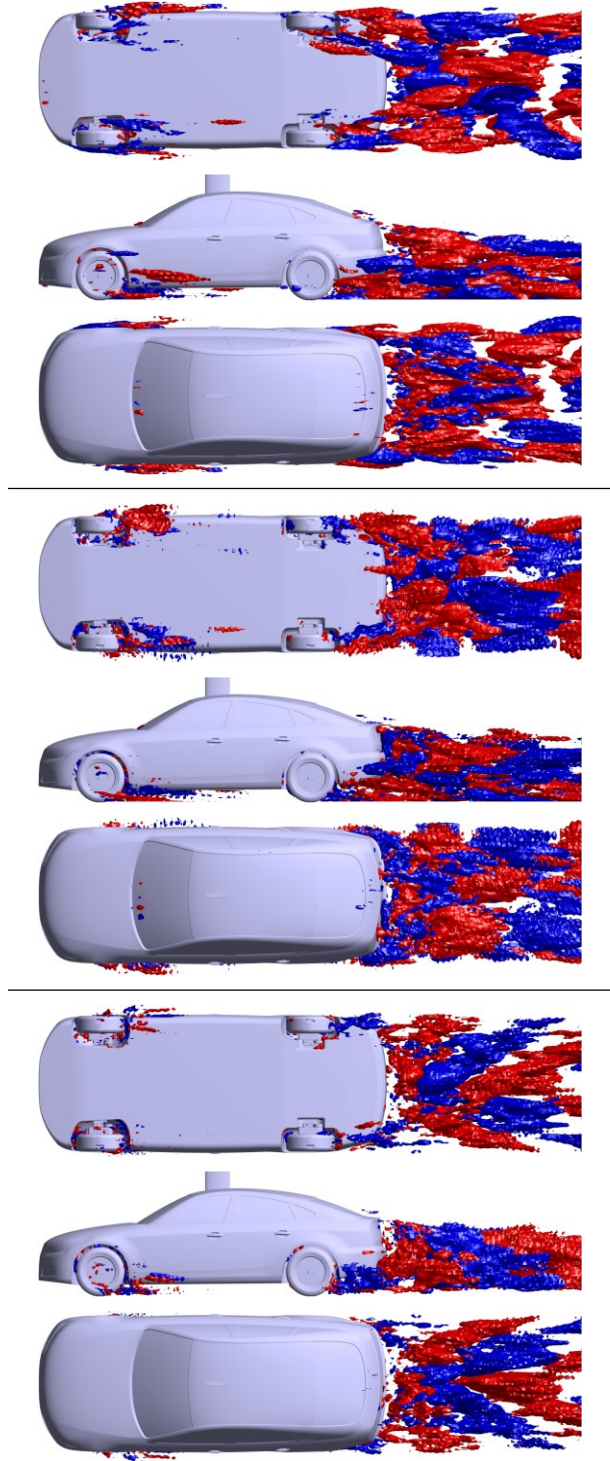


Figure 52: Mode three from the STDMD spectrum at $St_3=0.72$ and $f_3=29\text{Hz}$. Iso-surfaces of the reconstructed first time step. Top: x-component iso-surface $u_x=\pm 7.4\text{m/s}$. Middle: y-component iso-surface $u_y=\pm 4.8\text{m/s}$. Bottom: z-component iso-surface $u_z=\pm 5.7\text{m/s}$.

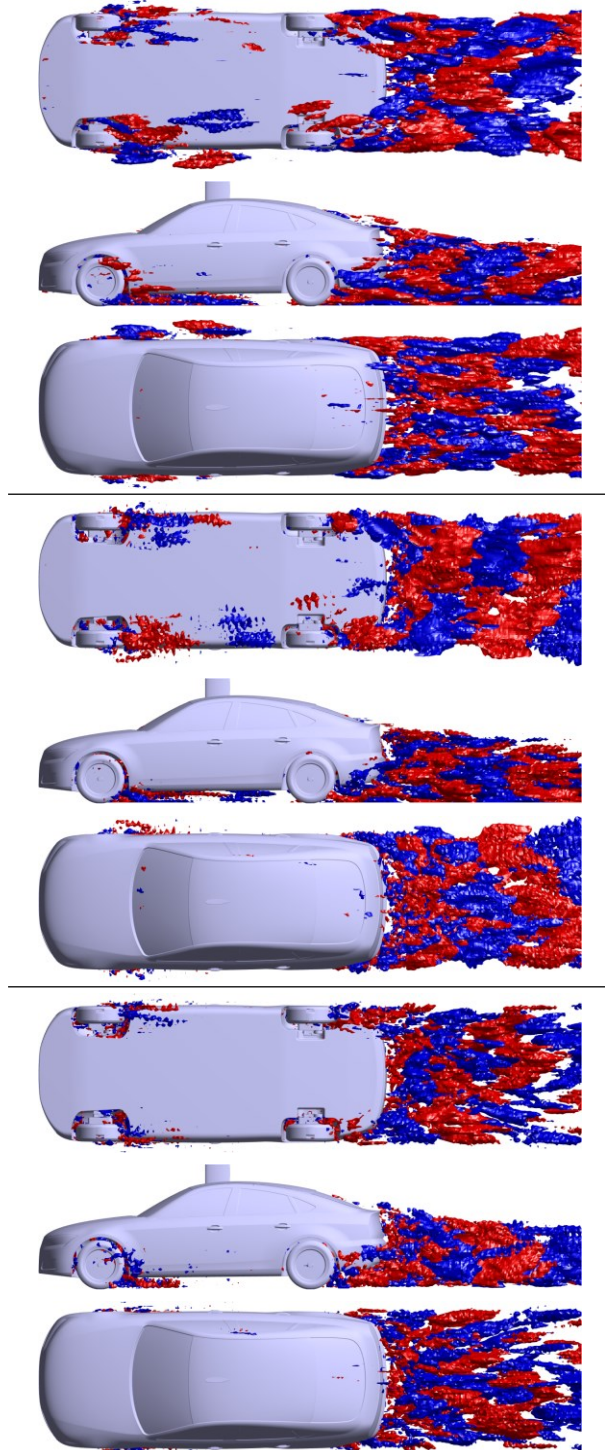


Figure 53: Mode seven from the STDMD spectrum at $St_7=0.99$ and $f_7=40\text{Hz}$. Iso-surfaces of the reconstructed first time step. Left: x-component iso-surface $u_x=\pm 14.0\text{m/s}$. Middle: y-component iso-surface $u_y=\pm 11.0\text{m/s}$. Right: z-component iso-surface $u_z=\pm 10.5\text{m/s}$.




Figure 54: Mode 15 from the STDMD spectrum at $St_{15}=2.68$ and $f_{15}=108\text{Hz}$. Iso-surfaces of the reconstructed first time step. Left: x-component iso-surface $u_x=\pm 9.5\text{m/s}$. Middle: y-component iso-surface $u_y=\pm 9.7\text{m/s}$. Right: z-component iso-surface $u_z=\pm 8.8\text{m/s}$.




Figure 55: Mode 16 from the STDMD spectrum at $St_{16}=1.18$ and $f_{15}=48\text{Hz}$. Iso-surfaces of the reconstructed first time step. Left: x-component iso-surface $u_x=\pm 4.1\text{m/s}$. Middle: y-component iso-surface $u_y=\pm 3.9\text{m/s}$. Right: z-component iso-surface $u_z=\pm 3.4\text{m/s}$.

References

- [1] C. Humby, "ANA Marketing Maestros," 3 11 2006. [Online]. Available: <https://ana.blogs.com/maestros/2006/11/data'is'the'new.html>. [Accessed 8 1 2019].
- [2] T. Schütz, "Hucho - Aerodynamik des Automobils," Springer, 2012.
- [3] S. B. Pope, "Turbulent Flows," Cambridge University Pr., 2000.
- [4] C. W. Rowley and S. T. M. Dawson, "Model Reduction for Flow Analysis and Control," *Annual Review of Fluid Mechanics*, pp. 387-417, 2017.
- [5] H. G. Weller, G. Tabor, H. Jasak and C. Fureby, "A tensorial approach to computational continuum mechanics using object-oriented techniques," American Institute of Physics, 1998.
- [6] M. Islam, F. Decker, E. de Villiers, A. Jackson, J. Gines, T. Grahs, A. Gitt-Gehrke and J. Comas i Font, "Application of Detached-Eddy Simulation for Automotive Aerodynamics Development," in *SAE International*, 2009.
- [7] C. Mocket, "A comprehensive study of detached-eddy simulation," Ph.D. thesis, Berlin, 2009.
- [8] P. R. Spalart, "Detached-Eddy Simulation," *Annual Review of Fluid Mechanics*, pp. 181-202, 2009.
- [9] P. R. Spalart and S. R. Allmaras, "A one-equation turbulence model for aerodynamic flows," *La Recherche Aeronautique*, pp. 5-21, 1994.
- [10] G. A. Ashford, "An Unstructured Grid Generation and Adaptive Solution Technique for High-Reynolds-Number Compressible Flows," Ph.D. thesis, University of Michigan, USA, 1996.
- [11] P. R. Spalart, W.-H. Jou, M. Strelets and S. R. Allmaras, "Comments on the Feasibility of LES for Wings and on a Hybrid RANS/LES Approach," *Advances in DNS/LES*, pp. 137-147, 1997.
- [12] F. Menter and M. Kuntz, "Adaptation of Eddy-Viscosity Turbulence Models to Unsteady Separated Flow Behind Vehicles," Springer, 2002.
- [13] P. Spalart, S. Deck, M. Shur, K. Squires, M. Strelets and A. Travin, "A new version of detached-eddy simulation, resistant to ambiguous grid densities," *Theoretical and Computational Fluid Dynamics*, pp. 181-195, 2006.
- [14] D. B. Spalding, "A single formula for the "law of the wall"," *Journal of Applied Mechanics*, pp. 455-458, 1961.

- [15] M. Gestrich, "Hybride RANS/LES - Methoden in der Fahrzeugaerodynamik," Master's thesis, Technical University of Munich, Germany, 2014.
- [16] J. H. Ferziger and M. Peric, "Computational Methods for Fluid Dynamics," Springer, 2002.
- [17] J. L. Lumley, "Toward a turbulent constitutive relation," *Journal of Fluid Mechanics*, pp. 413-434, 1970.
- [18] B. C. Moore, "Principal Component Analysis in Linear Systems: Controllability, Observability, and Model Reduction.," *IEEE Transactions on Automatic Control*, pp. 17-32, 1981.
- [19] C. W. Rowley, "Model Reduction for Fluids, Using Balanced Proper Orthogonal Decomposition," *International Journal of Bifurcation and Chaos*, 2005.
- [20] M. Brand, "Incremental Singular Value Decomposition of Uncertain Data with Missing Values," in *European Conference on Computer Vision*, 2002.
- [21] G. M. Oxberry, T. Kostova-Vassilevska, B. Arrighi and K. Chand, "Limited-memory adaptive snapshot selection for proper orthogonal decomposition," Lawrence Livermore National Laboratory, 2015.
- [22] W. Arrighi, G. Oxberry, T. Vassilevska and K. Chand, "libROM User Guide and Design," Lawrence Livermore National Laboratory, 2015. [Online]. Available: <https://computation.llnl.gov/projects/librom-pod-based-reduced-order-modeling>. [Accessed 21 12 2018].
- [23] B. R. Noack, M. Schlegel, B. Ahlborn, G. Mutschke, M. Morzynski, P. Comte and G. Tadmor, "A Finite-Time Thermodynamics of Unsteady Fluid Flows," *Journal of Non-Equilibrium Thermodynamics*, pp. 103-148, 2008.
- [24] I. Mezic, "Spectral Properties of Dynamical Systems, Model Reduction and Decompositions," *Nonlinear Dynamics*, pp. 309-325, 2005.
- [25] P. J. Schmid and J. Sesterhenn, "Abstract: MR.00007 : Dynamic Mode Decomposition of numerical and experimental data," in *61st Annual Meeting of the APS Division of Fluid Dynamics*, San Antonio, Texas, USA, 2008.
- [26] C. W. Rowley, I. Mezic, S. Bagheri, P. Schlatter and D. S. Henningson, "Spectral analysis of nonlinear flows," *Journal of Fluid Mechanics*, vol. 641, 2009.
- [27] P. J. Schmid, "Dynamic mode decomposition of numerical and experimental data," *Journal of Fluid Mechanics*, 2010.
- [28] P. J. Schmid, Juniper M. P. and O. Pust, "Applications of the dynamic mode decomposition," *Theoretical and Computational Fluid Dynamics*, pp. 249-259, 2011.

- [29] P. J. Schmid, D. Violato and F. Scarano, "Decomposition of time-resolved tomographic PIV," *Experiments in Fluids*, pp. 1567-1579, 2012.
- [30] D. Duke, J. Soria and D. Honnery, "An error analysis of the dynamic mode decomposition," *Experiments in Fluids*, vol. 52, p. 529-542, 2011.
- [31] M. R. Jovanovic, P. J. Schmid and J. W. Nichols, "Low-rank and sparse dynamic mode decomposition," *Center for Turbulence Research Annual Research Briefs*, vol. 26, pp. 139-152, 2012.
- [32] M. R. Jovanovic, P. J. Schmid and J. W. Nichols, "Sparsity-promoting dynamic mode decomposition," *Physics of Fluids*, 2014.
- [33] J. H. Tu, "Dynamic Mode Decomposition: Theory and Applications," Ph.D. thesis, Princeton University, USA, 2013.
- [34] J. H. Tu, C. W. Rowley, D. M. Luchtenburg, S. L. Brunton and N. J. Kutz, "On Dynamic Mode Decomposition: Theory and Applications," *Journal of Computational Dynamics*, pp. 391-421, 2014.
- [35] M. S. Hemati, M. O. Williams and C. W. Rowley, "Dynamic mode decomposition for large and streaming datasets," *Physics of Fluids*, vol. 26, 2014.
- [36] M. S. Hemati, C. W. Rowley, E. A. Deem and L. N. Cattafesta, "De-biasing the dynamic mode decomposition for applied Koopman spectral analysis of noisy datasets," in *Theoretical and Computational Fluid Dynamics*, Berlin Heidelberg, Springer-Verlag, 2017, pp. 349-368.
- [37] M. S. Hemati, E. A. Deem, M. O. Williams, C. W. Rowley and L. N. Cattafesta, "Improving Separation Control with Noise-Robust Variants of Dynamic Mode Decomposition," in *54th AIAA Aerospace Sciences Meeting*, San Diego, CA, USA, 2016.
- [38] J. Kou and W. Zhang, "An improved criterion to select dominant modes from dynamic mode decomposition," *European Journal of Mechanics B/Fluids*, 2016.
- [39] E. Alenius, "Flow Duct Acoustics - An LES Approach," Ph.D. thesis, KTH Royal Institute of Technology, Stockholm, Sweden, 2012.
- [40] K. Taira, S. L. Brunton, S. T. M. Dawson, C. W. Rowley, T. Colonius, B. J. McKeon, O. T. Schmidt, S. Gordeyev, V. Theofilis and L. S. Ukeiley, "Modal Analysis of Fluid Flows: An Overview," *AIAA Journal*, pp. 4013-4041, 2017.
- [41] T. W. Muld, G. Efraimsson and D. S. Henningson, "Mode Decomposition on Surface-Mounted Cube," *Flow, Turbulence and Combustion*, pp. 279-310, 2012.

- [42] T. W. Muld, G. Efraimsson, D. S. Henningson, A. H. Herbst and A. Orellano, "Analysis of Flow Structures in the Wake of a High-Speed Train," in *Proceedings of Aerodynamics of Heavy Vehicles III, Truck, Buses and Trains*, 2010.
- [43] T. W. Muld, G. Efraimsson and D. S. Henningson, "Flow structures around a high-speed train extracted using Proper Orthogonal Decomposition and Dynamic Mode Decomposition," *Computers & Fluids*, 2012.
- [44] T. W. Muld, G. Efraimsson and D. S. Henningson, "Mode Decomposition and Slipstream Velocities in the Wake of Two High-Speed Trains," *The International Journal of Railway Technology*, 2012.
- [45] T. W. Muld, G. Efraimsson and D. S. Henningson, "Wake Characteristics of High-Speed Trains with Different Lengths," *Journal of Rail and Rapid Transit*, pp. 333-342, 2014.
- [46] T. W. Muld, "Slipstream and Flow Structures in the Near Wake of High-Speed Trains," Ph.D. thesis, , KTH Royal Institute of Technology, Stockholm, Sweden, 2012.
- [47] H. M. Frank and C. D. Munz, "Direct aeroacoustic simulation of acoustic feedback phenomena on a side-view mirror," *Journal of Sound and Vibration*, pp. 132-149, 2016.
- [48] H. M. Frank and C. D. Munz, "Large eddy simulation of tonal noise at a simplified side-view mirror using a high order discontinuous Galerkin method," *AIAA Journal*, 2016.
- [49] H. M. Frank, "High Order Large Eddy Simulation for the Analysis of Tonal Noise Generation via Aeroacoustic Feedback Effects at a Side Mirror," Ph.D. thesis, University of Stuttgart, Germany, Verlag Dr. Hut, 2017.
- [50] M. Peichl, S. Mack, T. Indinger and F. Decker, "Numerical Investigation of the Flow Around a Generic Car using Dynamic Mode Decomposition," in *ASME Fluids Engineering Division Summer Meeting*, Chicago, USA, 2014.
- [51] D. Matsumoto, L. Haag and T. Indinger, "Investigation of the Unsteady External and Underhood Airflow of the DrivAer model by Dynamic Mode Decomposition Methods," *International Journal of Automotive Engineering*, pp. 55-62, 2017.
- [52] D. Matsumoto and T. Indinger, "On-the-fly algorithm for Dynamic Mode Decomposition using Incremental Singular Value Decomposition and Total Least Squares," arXiv:1703.11004 [physics.flu-dyn], 2017.

- [53] D. Matsumoto, M. Kiewat, L. Haag and T. Indinger, "Online Dynamic Mode Decomposition Methods for the Investigation of Unsteady Aerodynamics of the DrivAer Model (First Report) - Validation of DMD Methods for Surface Forces," *International Journal of Automotive Engineering*, pp. 64-71, 2018.
- [54] D. Matsumoto, M. Kiewat, C. A. Niedermeier and T. Indinger, "Application of Incremental Proper Orthogonal Decomposition for the Reduction of Very Large Transient Flow Field Data," *International Journal of Automotive Engineering*, pp. 117-124, 2019.
- [55] M. Kiewat, L. Haag, T. Indinger and V. Zander, "Low Memory Reduced Order Modelling with Dynamic Mode Decomposition Applied on Unsteady Wheel Aerodynamics," in *ASME Fluids Engineering Division Summer Meeting*, Waikoloa, HI, USA, 2017.
- [56] M. Kiewat, D. Matsumoto, L. Haag, V. Zander and T. Indinger, "Online Dynamic Mode Decomposition Methods for the Investigation of Unsteady Aerodynamics of the DrivAer Model (Second Report) - Application on Velocity Fields," *International Journal of Automotive Engineering*, pp. 72-78, 2018.
- [57] A. Roshko, "On the Development of Turbulent Wakes From Vortex Streets," National Advisory Committee for Aeronautics, 1954.
- [58] G. H. Golub and C. F. van Loan, "Matrix Computations," The Johns Hopkins University Press, London, 1996.
- [59] L. Sirovich, "Turbulence and the Dynamics of Coherent Structures Part 1: Coherent Structures," *Quarterly of Applied Mathematics*, pp. 561-571, 1987.
- [60] A. Chatterjee, "An introduction to the proper orthogonal decomposition," *Current Science*, pp. 808-817, 2000.
- [61] S. C. Langer, "Incremental Proper Orthogonal Decomposition For Aerodynamics," Bachelor's thesis, Munich, Germany, 2017.
- [62] A. Björk, "Solving Linear Least Squares Problems by Gram-Schmidt Orthogonalization," *BIT Numerical Mathematics*, vol. 7, no. 1, pp. 1-21, 1967.
- [63] A. R. Benson, D. F. Gleich and J. Demmel, "Direct QR factorizations for tall-and-skinny matrices in MapReduce architectures," in *IEEE International Conference on Big Data*, Santa Clara, 2013.
- [64] T. Sayadi, C. W. Hamman and P. J. Schmid, "Parallel QR algorithm for data-driven decompositions," in *Center for Turbulence Research Summer Program*, Stanford, 2014.

- [65] C. W. Rowley, "KITP talk: Discussion on Koopmanism II," 6 1 2017. [Online]. Available: <http://online.kitp.ucsb.edu/online/transturb17/rowley/rm/jwvideo.html>. [Accessed 28 12 2018].
- [66] A. Heft, T. Indinger and N. Adams, "Investigation of Unsteady Flow Structures in the Wake of a Realistic Generic Car Model," in *29th AIAA Applied Aerodynamics Conference*, Honolulu, Hawaii, 2011.
- [67] S. Mack, T. Indinger, N. Adams, S. Blume and P. Unterlechner, "The Interior Design of a 40% Scaled DrivAer Body and First Experimental Results," in *Proceedings of the ASME 2012 Fluids Engineering Summer Meeting*, Puerto Rico, USA, 2012.
- [68] L. Miao, S. Mack and T. Indinger, "Experimental and Numerical Investigation of Automotive Aerodynamics using DrivAer Model," in *Proceedings of the ASME 2015 International Design Engineering Technical Conferences & Computers and Information in Engineering Conference*, Boston, Massachusetts, USA, 2015.
- [69] S. Ahmed, G. Ramm and G. Faltin, "Some Salient Features Of The Time-Averaged Ground Vehicle Wake," in *International Congress & Exposition*, Detroit, Michigan, USA, 1984.
- [70] F. Wittmeier and T. Kuthada, "Open Grille DrivAer Model - First Results," *SAE International Journal of Passenger Cars - Mechanical Systems*, pp. 252-260, 2015.
- [71] A. I. Heft, T. Indinger and N. A. Adams, "Introduction of a New Realistic Generic Car Model for Aerodynamic Investigations," *SAE International*, 16 04 2012.
- [72] M. Grandemange, O. Cadot and M. Gohlke, "Reflectional symmetry breaking of the separated flow over three-dimensional bluff bodies," *Physical Review E : Statistical, Nonlinear, and Soft Matter Physics*, *American Physical Society*, 2012.
- [73] M. Grandemange, M. Gohlke and O. Cadot, "Turbulent wake past a three-dimensional blunt body. Part 1. Global modes and bi-stability," *Journal of Fluid Mechanics*, pp. 51-84, 2013.
- [74] G. Pavia and M. Passmore, "Influence of Short Rear End Tapers on the Unsteady Base Pressure of a Simplified Ground Vehicle," in *SAE International*, 2016.
- [75] G. Pavia, M. Passmore and C. Sardu, "Evolution of the bi-stable wake of a square-back automotive shape," *Experiments in Fluids*, 2018.
- [76] A. Heft, "Aerodynamic Investigation of the Cooling Requirements of Electric Vehicles," Ph.D. thesis, Technical University of Munich, Germany, 2013.

- [77] C. Collin, S. Mack, T. Indinger and J. Mueller, "A Numerical and Experimental Evaluation of Open Jet Wind Tunnel Interferences using the DrivAer Reference Model," in *SAE International*, 2016.
- [78] J. Reiß, "Numerische und experimentelle Untersuchungen zu instationären Strömungsphänomenen in der Fahrzeugaerodynamik," Masters's thesis, Technical University of Munich, Germany, 2018.
- [79] C. Strangfeld, D. Wieser, H.-J. Schmidt, R. Woszidlo, C. Nayeri and C. Paschereit, "Experimental Study of Baseline Flow Characteristics for the Realistic Car Model DrivAer," in *SAE World Congress & Exhibition*, 2013.
- [80] N. Simmonds, J. Pitman, P. Tsoutsanis, K. Jenkins, A. Gaylard and W. Jansen, "Complete Body Aerodynamic Study of three Vehicles," in *SAE World Congress*, 2017.
- [81] M. Rüttgers, J. Park and D. You, "Large-eddy simulation of turbulent flow over the DrivAer fastback vehicle model," *Journal of Wind Engineering & Industrial Aerodynamics*, pp. 123-138, 2019.
- [82] J. C. R. Hunt, A. A. Wray and P. Moin, "Eddies, streams, and convergence zones in turbulent flows," in *2. Proceedings of the 1988 Summer Program*, 1988.
- [83] M. Pastoor, L. Henning, B. R. Noack, R. King and G. Tadmor, "Feedback shear layer control for bluff body drag reduction," *Journal of Fluid Mechanics*, pp. 161-196, 2008.
- [84] C. Granger, "Investigating Causal Relations by Econometric Models and Cross-spectral Methods," *Econometrica*, pp. 424-438, 1969.
- [85] G. Tissot, A. Lozano-Duran, L. Cordier, J. Jimenez and B. R. Noack, "Granger causality in wall-bounded turbulence," in *1st Multiflow Summer Workshop, Journal of Physics: Conference Series*, 2014.
- [86] D. Wieser, H. Lang, C. Nayeri and C. Paschereit, "Manipulation of the Aerodynamic Behavior of the DrivAer Model with Fluidic Oscillators," *SAE International*, pp. 687-702, 2015.
- [87] K. Nisugi, T. Hayase and A. Shirai, "Fundamental Study of Aerodynamic Drag Reduction for Vehicle with Feedback Flow Control," *JSME International Journal*, pp. 584-592, 2004.
- [88] L. Haag, M. Kiewat, T. Indinger and T. Blacha, "Numerical and Experimental Investigations of Rotating Passenger Car Wheel Aerodynamics on the DrivAer Model with Engine Bay Flow," in *ASME Fluids Engineering Division Summer Meeting*, Waikoloa, HI, USA, 2017.