

MASTER THESIS

FOR THE DEGREE OF MASTER OF SCIENCE IN AEROSPACE ENGINEERING
AT DELFT UNIVERSITY OF TECHNOLOGY

Multi-fidelity Method for Aerodynamic Shape Optimisation of Axial Compressor Blades

BY LENNART VERHEIJ

Supervisors: Dr. Robert Grewe, Siemens AG (Division Power and Gas)
Dr.ir. Matteo Pini, TU Delft (Faculty of Aerospace Engineering)

Completed in: Siemens AG (Division Power and Gas)
Mülheim an der Ruhr

Exam committee: Prof.dr.ir. P. Colonna di Paliano, TU Delft (Faculty of AE)
Dr.ir. M. Pini, TU Delft (Faculty of AE)
Dr. R. P. Dwight, TU Delft (Faculty of AE)
Dr.ir. R. Pecnik, TU Delft (Faculty of 3ME)

Thesis registration number: 166#17#MT#FPP

Matriculation number: 4142306

Date of submission: October 2017



Statement of originality

I hereby confirm that the content of this thesis is my own work, without contributions from any sources other than those cited in the text and acknowledgements.

This also applies to graphics, drawings and images included.

Lennart Verheij

*Mülheim an der Ruhr,
27th of October, 2017*

Abstract

In turbomachinery optimisation problems, run time is often a critical factor due to high dimensionality of the design search space. This work explores the use of the multi-fidelity method to speed up an aerodynamic optimisation algorithm applied to axial compressor blades. A rotor blade is optimised in a two-stage blade geometry for maximum isentropic efficiency. The single-fidelity reference optimisation uses a high-fidelity evaluation process employing Menter SST turbulence equations and a mesh of 903,000 cells. Five multi-fidelity optimisation setups are tested, which employ the same high-fidelity process, but distinct low-fidelity processes ranging from a fine mesh and RANS turbulence equations to a coarse mesh and inviscid Euler equations. It is found that multi-fidelity optimisation could cause a delay in run time of up to 39.9%, equivalent to almost three days and a loss in optimum efficiency of 0.11%. The best result is a speed-up of 14.1%, equivalent to 1 day of time savings and an improvement in efficiency of 0.02%. The speed-up of 50% demonstrated in literature could not be achieved since the high-fidelity model in this work is much cheaper. The best cost ratio achieved in this work is comparable with 0.14, but the correlation coefficient of 0.46 is insufficient. It is shown that at their optimum efficiencies, two selected single-fidelity and multi-fidelity optima have different geometries and aerodynamic behaviour. For improving performance using the multi-fidelity method, it is recommended to increase the fidelity gap between the low-fidelity and the high-fidelity processes. The cost ratio of a new low-fidelity process can be estimated with an error of at most 5%, by using 10 member designs. Furthermore, the correlation coefficient can be estimated with an error of at worst 35%, using 20 blade designs. From the results in this thesis, it is recommended to employ a cheaper low-fidelity process using for instance through-flow calculations or to make the high-fidelity process more expensive by adding more design features.

Keywords: multi-fidelity method, aerodynamic optimisation, axial compressor blade, fidelity reduction, Kriging, Co-Kriging, RANS turbulence modelling, Euler equations

Preface

This thesis considers the implementation of a new method to improve aerodynamic shape optimisation of compressor blades. The method is applied to an industrial four-row compressor design and measured against test results from reference works. This research project was issued by the R&D department of Materials and Technology for large industrial gas turbines of Siemens A/G. This report is written to fulfil the graduation requirements of the Master in Aerospace Engineering at Delft University of Technology in the Netherlands. I was engaged in writing this thesis from February 2017 to October 2017 at Siemens Power and Gas in Mülheim an der Ruhr.

Over a period of nine months, I have been given the chance to manage a unique collaborative project between my university and a frontrunner in the energy industry. Not only it has been an incredible experience to work with knowledgeable colleagues in a high-tech environment, I have also learned much about the German culture, customs and the inspiring work mentality. With deep respect for the people I have got to know within and outside of my department, I feel fortunate that I was given this chance.

I would like to express my gratitude to a number of people that have helped me to push my boundaries. First and foremost, I would like to thank my supervisor Robert Grewe, who has been a major support during the entire project. It was a great pleasure unravelling the mysteries of this method together and achieving major breakthroughs, even after periods of seemingly small progress. Furthermore, I like to thank my university supervisor Matteo Pini for his enthusiasm and for taking the responsibility to evaluate my work from an academic perspective. I say thanks to my colleagues Matthias Hüls and Gregor Schmid, for sharing their knowledge and advice on setting up the surrogate model and the aerodynamic simulation process. Furthermore, I appreciate how my colleagues and interns have made my time at Siemens pleasant and worthwhile. Finally, I say thanks to my family and close friends for their enjoyable visits to Germany, their vast support and boundless interest.

Lennart Verheij
Mülheim an der Ruhr, October 2017

Contents

Abstract	IV
Preface	V
Nomenclature	VIII
1 Introduction	1
1.1 Industrial context	1
1.2 Aerodynamic optimisation of compressors	2
1.3 Thesis objective	3
1.4 Report structure	4
2 Fundamentals of compressor design and modelling	5
2.1 Compressor performance and operating principles	5
2.2 Design philosophy	8
2.3 Basics of CFD simulation for compressors	9
3 Theory on aerodynamic shape optimisation	12
3.1 AutoOpti and optimisation terminology	12
3.2 Optimisation architectures	13
3.3 Evolutionary algorithm	14
3.4 Surrogate modelling	15
3.4-1 Kriging	16
3.4-2 Co-Kriging	19
3.5 Decision function	20
3.6 Optimisation progress quantification	24
3.7 Potential metrics	26
3.8 Literature review on multi-fidelity compressor optimisations	28
4 Methodology for shape optimisation	32
4.1 Problem specification	32
4.2 Optimisation settings	34
4.3 Setup of optimisations	36
4.3-1 Single-fidelity optimisation	36
4.3-2 Multi-fidelity optimisation	40
4.4 Low-fidelity model selection	40
4.4-1 Identification of fidelity reduction methods	40
4.4-2 Fidelity reduction study	43
4.4-3 Down-selection of models	47
4.5 Convergence criterion	49
4.6 Performance metrics	51

5	Results of the multi-fidelity method	54
5.1	Single-fidelity reference optimisation	54
5.2	Multi-fidelity results using variance decision function	55
5.2-1	Performance assessment	56
5.2-2	Potential assessment	60
5.2-3	Investigation of information transfer	63
5.3	Multi-fidelity results using constant decision function	65
5.3-1	Performance assessment	65
5.3-2	Potential assessment	66
5.4	Conclusions on multi-fidelity optimisation results	68
5.5	Multi-fidelity performance forecasting	69
5.5-1	Forecasting using one member	69
5.5-2	Forecasting using set of members	70
6	Aerodynamic analysis of optimum members	77
6.1	Three-dimensional blade shape	77
6.2	Radial distributions	78
6.3	Distributions in the blade-to-blade plane	80
7	Conclusions	83
7.1	Answers to research questions	83
7.2	Future work	85
	Appendix	87

Nomenclature

Abbreviations

B2B	Blade-to-blade
CFD	Computational Fluid Dynamics
DLR	Deutsches Zentrum für Luft- und Raumfahrt
HF	High-fidelity
GMC	General Mesh Connector
LF	Low-fidelity
MTU	Motoren Turbinen Union
RSL	Radial streamlines
RANS	Reynolds Averaged Navier-Stokes
SR	Sampling resolution
TRACE	Turbomachinery Research Aerodynamic Computational Environment

Greek symbols

ϵ	Turbulence dissipation rate	m^2/s^3
η	Efficiency	%
κ	Turbulence kinetic energy	m^2/s^2
μ	Dynamic viscosity	$kg/(ms)$
Π	Total pressure ratio	-
ρ_f	Flow density	kg/m^3
ρ	Correlation	-
σ	Standard deviation	-
σ^2	Variance	-
Θ	Hyperparameter	-
ω	Specific turbulence dissipation rate	$1/s$

Latin symbols

C	Cost	s
C_r	Cost ratio	-
f_r	Replacement ratio	-
G	Gain in CVG	%
\mathbf{K}	Covariance matrix	-
\dot{m}	Mass flow	kg/s
N	Rotational speed	rpm
p	pressure	N/m^2
p_t	Total pressure ratio	-
r^2	Correlation	-
\mathbf{R}	Correlation matrix	-
R	Radial span	%
S	Speed-up	%
T	Temperature	K
t	Time	s
U	Velocity	m/s
w	Weight coefficient	-
Δx	Mesh cell size	m
y	Wall distance	m
y^+	Dimensionless wall distance	-

1 Introduction

In the introduction chapter, the context of this master thesis is explained and the research question is motivated.

This chapter starts by explaining the context of this work in Section 1.1. Afterwards, a short summary of the current state-of-the-art in compressor optimisation is given by Section 1.2. The objective of this thesis is explained in Section 1.3. At last, in Section 1.4, the structure of this report is outlined.

1.1 Industrial context

Heavy-duty gas turbines have become an established means for power generation. The average power output of gas turbines has multiplied over a hundred-fold because of their rapid development in the last century. Primarily driven by advancements in the aerospace sector, the performance of stationary land-based gas turbines has seen a drastic increase, through increasing the overall pressure ratio and the firing temperature for combustion. In a combined cycle configuration, where waste heat from the gas turbine exhaust gasses is used to make steam and drive an additional steam turbine, an efficiency can be attained in the order of 50%. Recent demonstration has seen a combined efficiency of 60%. While gas turbines were initially used for emergency back-up in case of blackouts or peaking power in the USA and Europe, from 1980 there has been a trend to replace steam powered plants with gas turbines [2].

Gas turbines are now considered amongst the cleaner means to convert chemical energy from fossil fuel to electrical energy. In a forecast on the outlook of the world's energy consumption by the International Energy Agency [14], the demand for natural gas is expected to rise by 50% in the year 2040. This rise will be mainly due to rapidly developing countries that share a large interest in power provision using industrial gas turbines. However, there is an increasing share of renewable energy sources expected to be introduced into the system. Moreover, energy production is expected to take place in decentralised systems, instead of in the centralised system effective today. These future challenges drive innovation with a focus to improve powerplant flexibility for scaling up or down the power output and to accommodate various types of fuels and easy adaptation in configuration. In order to stay competitive with an increasing share of clean energy sources, the manufacturing and operating cost needs to be reduced.

Major gas turbine engine manufacturers, such as Siemens are challenged to stay competitive in a global market. Over a broad portfolio of turbines ranging from 4 to 450 MW, Siemens has set goals to provide high reliability, load flexibility, low life cycle cost and environmental compatibility. Besides searching for innovative new concepts, the main components of the gas turbine are individually improved for delivering optimal efficiency. Those components are generally separately designed: the compressor, combustion chamber and turbine. Achieving small improvements in efficiency can result in huge savings in operating cost spent on fuel. Those improvements also contribute to the objective of lowering emissions and reducing environmental impact. Within this context, in this thesis project, the optimisation of one blade row in an axial compressor is considered.

1.2 Aerodynamic optimisation of compressors

This section provides a short description on the current status of design and optimisation methods of an axial compressor. Many research efforts are made to improve the performance of the axial compressor. The main function of the axial compressor is to raise the total pressure of the air before it enters the combustion chamber. It thereby employs alternating stator and rotor blade rows, which can contain up to 20 blade stages in total. A fully assembled gas turbine shaft with 13 compressor stages is shown in Figure 1.1.



Figure 1.1: Siemens gas turbine SGT6-5000F in a manufacturing facility [26]

The development process of the compressor has seen a drastic improvement through the application of powerful simulation and optimisation codes. The aerodynamic design of compressor blades is established by employing various methods with increasing complexity, starting at mean line and through-flow calculations and followed by Computational Fluid Dynamics (CFD) simulations. Three-dimensional aerodynamic optimisation is performed to increase the isentropic efficiency of the machine. In aerodynamic shape optimisation, a variety of blade designs is explored, that improve one or more objectives as shown in a typical Pareto diagram in Figure 1.2. This figure gives an indication of how an optimum design can be selected on the basis of the objective functions.

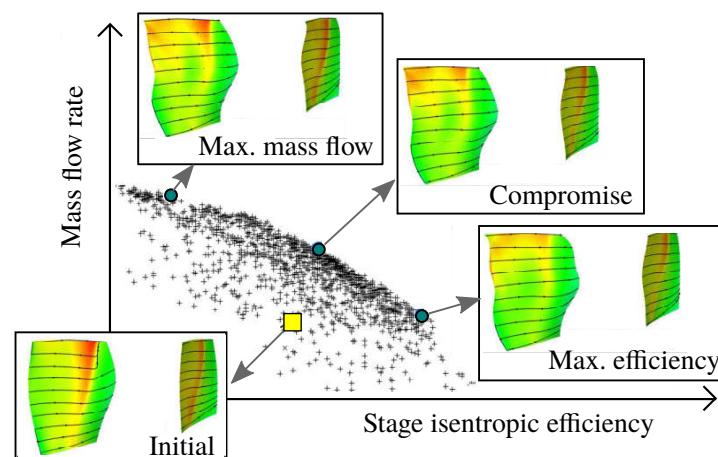


Figure 1.2: Pareto diagram with the initial blade design and optima per optimisation objective [12].

Generally, blade rows are separately optimised by applying inlet and outlet boundary conditions to ensure matching with neighbouring blade rows. The downside of single-stage optimisation, is that the full design is fixed intermediately, which restrains the design solution space. Instead, a more integrated approach like multi-stage optimisation is preferable. A schematic of this approach is shown in Figure 1.3, which shows how more blade rows are included in the multi-stage optimisation setup. Optimising multiple blade stages simultaneously increases the optimisation duration to several weeks, using traditional optimisation algorithm. Hence, multi-stage optimisations are hardly viable in industry due to limited computational resources.

Several methods are proposed in literature to speed-up the optimisation process. Among these is a method called “multi-fidelity optimisation” or “variable fidelity modelling”, which is often employed in areas of fluid mechanics and solid mechanics [8]. A multi-fidelity optimisation employs an additional “low-fidelity” simulation process, to evaluate the objective function values of a candidate design at a much lower cost, with slightly reduced accuracy. When correlating low-fidelity simulation results with results from the original high-fidelity process, the overall optimisation can be improved. In this thesis, the multi-fidelity technique is investigated to speed-up the aerodynamic optimisation of an axial compressor rotor.




Figure 1.3: Meridional diagram of the compressor, showing domains for single-row (left) and multi-row optimisation (right).

1.3 Thesis objective

The multi-fidelity technique is demonstrated successfully in several compressor optimisation problems, described in the works of Reimer [23] and Brooks et al. [3]. In those works, multi-fidelity optimisation provides a speed-up or an improvement of the objectives in the optimum, compared to a single-fidelity optimisation with the same computational budget. However, any enhancements are highly problem dependent and greatly rely on the cost and quality of the low-fidelity model [10]. Previous research generally does not substantiate the level of fidelity reduction required for the low-fidelity model in a multi-fidelity framework. Therefore this thesis investigates the potential of multi-fidelity optimisation with distinct low-fidelity processes.

The main research question answered in this work is the following:

1. *How much fidelity reduction should be applied to the low-fidelity process chain, in order to achieve the highest speed-up or improvement of the optimum by multi-fidelity optimisation with respect to single-fidelity optimisation?*

A second research question is posed which considers the differences in design optima, found by two optimisation frameworks. This research question cannot be answered fully within the framework of this thesis and is therefore of secondary importance. It is formulated as follows:

2. *Are optimum designs from multi-fidelity and single-fidelity optimisations similar or different in terms of shape and aerodynamic performance?*

These two research questions are answered in the closing chapter of this report. In this thesis, the multi-fidelity technique is employed in an optimisation problem consisting of one parametrised rotor cascade within a two-stage compressor. Five multi-fidelity optimisation setups are considered, with distinct low-fidelity processes. Optimisation performance is quantified through a difference in optimisation time and by a possible improvement of the objectives upon optimisation convergence. This research investigates the multi-fidelity approach itself, rather than seeking a physical optimum design. The rotor blade is optimised for its aerodynamic performance only, hence structural characteristics are not under consideration. Fidelity reduction applied in this work is limited to the domain of 3D CFD simulations, although also through-flow or meanline simulations would also be possible.

In order to choose a suitable fidelity level of the low-fidelity process, a designer is required to estimate the differences in cost and quality between the two processes, employed in a multi-fidelity architecture. This estimation can be based on the evaluation of various blade designs using the high-fidelity and low-fidelity processes. In Figure 1.4, it is shown that in order to improve the accuracy of the estimation, a designer needs to invest more effort by evaluating more designs. In the conclusions of this work, guidelines are provided on the recommended number of design evaluations, required to forecast the suitability of a low-fidelity process.




Figure 1.4: Schematic of the accuracy for estimating suitability of a low-fidelity process for multi-fidelity optimisation, by increasing the sample set of design evaluations.

1.4 Report structure

This thesis is structured as follows. In Chapter 2, the approach for establishing a compressor blade design and the involved disciplines are discussed. Chapter 3 introduces major elements of aerodynamic shape optimisation of compressors as well as parameters to quantify and estimate the performance of multi-fidelity optimisations. This chapter also provides a short literature review of works relevant for this research. Thereafter, Chapter 4 describes the optimisation problem, the setup of single-fidelity and multi-fidelity optimisations considered in this work, as well as the convergence criteria for termination and performance metrics for comparing the optimisation results. Subsequently, Chapter 5 presents results of five multi-fidelity optimisations with respect to performance of a single-fidelity reference optimisation. In an attempt to improve the potential of the multi-fidelity method, a slightly adjusted optimisation setup is also introduced and discussed in this chapter, followed by guidelines to forecast performance of a particular multi-fidelity optimisation. In Chapter 6, the aerodynamic performance of two selected design optima of the best multi-fidelity optimisation and single-fidelity optimisation is compared. Conclusions of this work are summarised in Chapter 7. This chapter also provides recommendations to improve the potential of a multi-fidelity optimisation, besides giving examples of applications that could benefit from the multi-fidelity method.

2 Fundamentals of compressor design and modelling

In this thesis, a compressor is optimised with the target to increase efficiency while maintaining a minimum mass flow and pressure ratio. This chapter discusses the general approach for establishing a compressor blade design and the disciplines involved. The resulting blade design, acts as a starting point for aerodynamic optimisation.

This chapter is structured as follows. Section 2.1 describes the major function of a compressor and provides a definition of pressure ratio and efficiency. The initial design is readily available in this work, and hence is only touched upon shortly in Section 2.2. In this work, a compressor blade design is optimised for maximum aerodynamic performance. Section 2.3 therefore provides an introduction into aerodynamic simulation of compressor blades.

2.1 Compressor performance and operating principles

The first section of this chapter starts by explaining the main function of a compressor, after which two parameters are defined to describe compressor performance: mass flow and pressure ratio. The operating conditions of the compressor are then schematically shown by means of a compressor map. Thereafter, typical features of the flow field in the compressor channel are described, which are important for aerodynamic simulation.

The main function of a compressor is to increase the total pressure of the inlet air from ambient pressure to an elevated pressure before it enters the combustion chamber. The ratio of channel inlet and exit pressures is calculated through Equation 2.1 and is called the “pressure ratio”. In this work, an axial compressor stage is optimised with the aim to increase its efficiency, while optimum candidate designs need to attain a minimum pressure ratio. A blade stage consists of a rotor and a stator blade row. As the rotor adds energy to the flow by means of rotation, which effectively creates a moment of momentum, the total pressure and total temperature of the flow are increased. The stator in its turn removes this rotational energy and converts kinetic energy from the circumferential velocity component to an increase in the static pressure of the flow. Since the flow is very sensitive to aerodynamic separations during the deceleration process, a large number of stages is required to reach the desired total pressure ratio. For industrial gas turbines, this number can go up to 10 to 20 stages, depending on the type of application.

$$\Pi_t = \frac{p_{t,2}}{p_{t,1}} \quad (2.1)$$

The operational limits of the compressor are often shown on a compressor map, which shows its operating performance as combinations of the total pressure ratio Π_t and mass flow, at different rotational velocities. A schematic of this map is shown in Figure 2.1 and will later be used to indicate the operating points for aerodynamic simulation. The mass flow and rotational velocity are reduced by Equations 2.2 and 2.3 respectively. These parameters are reduced to ensure comparable flow conditions for different inflow conditions $p_{t,1}$ and $T_{t,1}$. For normalisation, p_{ref} and T_{ref} consider the respective reference pressure and temperature at International Standard Atmosphere (ISA).

$$\dot{m}_{corrected} = \dot{m} \cdot \frac{p_{ref}}{p_{t,1}} \cdot \sqrt{\frac{T_{t,1}}{T_{ref}}} \quad (2.2)$$

$$N_{corrected} = N \cdot \sqrt{\frac{T_{ref}}{T_{t,1}}} \quad (2.3)$$

The surge line indicates the upper limit of the total pressure for a certain reduced mass flow. This line marks an area of unstable operating conditions where severe stall or surge occurs. The steady-state operating points of the compressor are on the working line.

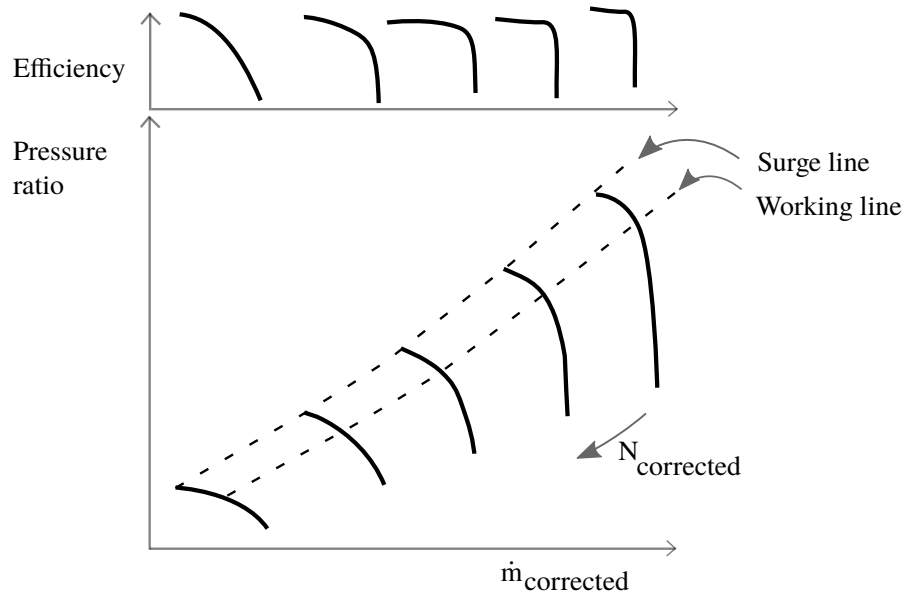


Figure 2.1: Compressor map

At a stationary operating point, ideally the compressor operates at a maximum isentropic efficiency, to minimise losses and maximise the amount of useful work to generate power. Hence, optimisations in this thesis aim to improve the isentropic efficiency in two operating points. In the top segment of the compressor map in Figure 2.1, the isentropic efficiency variation with mass flow for every speed line is shown. The isentropic efficiency expresses the ratio of actual compressor work over isentropic “ideal” work output of the compressor. This parameter can be calculated using air inflow and outflow conditions of the compressor, as shown by Equation 2.4. The heat capacity ratio, denoted by γ , is assumed constant over the entire considered range of temperatures.

$$\eta_{is} = \frac{\left(\frac{p_{t,2}}{p_{t,1}}\right)^{\frac{\gamma-1}{\gamma}} - 1}{\frac{T_{t,2}}{T_{t,1}} - 1} \quad (2.4)$$

Flow features and loss mechanisms

A reduction in compressor efficiency is caused by flow losses in the compressor channel. This paragraph identifies and describes those losses.

The flow field in an axial compressor is highly complex and features three-dimensional flow which is turbulent and viscous. Flow features that impose a loss in compressor efficiency, can be ascribed to classes of loss mechanisms. Losses are often attributed to a certain plane in which they are predominant. In this work, a system of so-called stream surfaces is used for describing those planes, as shown in Figure 2.2. A traditional description of flow losses is provided by Denton [5]. In that work, three categories are suggested, each approximately having an equal share towards all channel losses. Those three categories are summarised in the following.

The first category, profile loss, is caused by the viscous boundary layers originating from the blade surface, where flow can be assumed to be two-dimensional. The second loss mechanism, tip leakage loss, considers the vortices generated at the rotor blade tips and the hub clearance for stator blades. Whether a rotor blade is shrouded or un-shrouded plays a large role in the severity of tip leakage losses. The third category of losses considers end-wall losses, includes all flow losses from secondary flows as shown in Figure 2.3. Due to viscous friction at the hub and shroud casings, inlet flow has a non-uniform velocity profile and causes vortical flow structures to develop in the secondary plane of the channel. After passing through the passage, these vortices mix with the bulk flow, leading to mixing losses.

An important feature of the flow field in the transonic compressor is its shock structure. As the circumferential velocity increases along the blade radius, close to the tip region, local areas of supersonic flow are present. Because of the transonic nature of the flow, a complex shock structure is present in the channel, which interacts with the boundary layer at the blade surfaces.

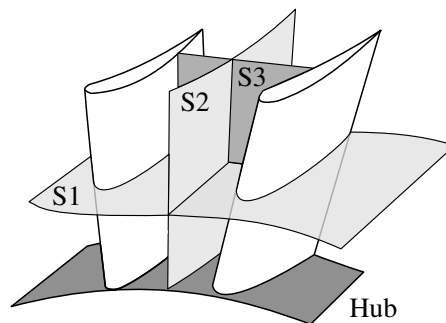


Figure 2.2: Definition of three principle stream surfaces within blade passage, showing the blade-to-blade surface (S1), meridional surface (S2) and secondary surface (S3) [22]




Figure 2.3: Two types of vortex structures in the compressor channel: (a) end-wall vortices predominant in the S2 surface and (b) secondary flows in the S3 surface [4].

Entropy generation

A convenient way to quantify loss and its effect on the compressor efficiency is through the notion of entropy. Since the value of entropy is independent of the frame of reference in which it is regarded (i.e. rotating or stationary), its application suits both rotor and stator rows. Moreover, there is a strong relation between the loss in compressor efficiency and the amount of entropy created in the compressor channel¹.

Entropy is generated in the compressor through several irreversible fluid dynamic processes that are characterised by the losses mentioned above. For compressors, shear forces and non-equilibrium processes are the most predominant, while heat transfer is less stringent compared to for example cooling flows of turbine blades. Shear forces, or viscous friction, causes losses elaborated on above. Shock waves can be attributed to the class of non-equilibrium processes. By imposing rapid compression, a shock wave involves heat conduction and high viscous normal stresses [5].

When optimising the isentropic efficiency, the aim is to resolve entropy generating flow features as accurately as possible. From above descriptions of irreversible processes, locations in the channel that have special interest are blade surfaces and end-walls for simulating the boundary layers, blade tip where a tip vortex develops, blade trailing edge in order to capture strong shear layers and the blade leading edge where the shock wave has its onset.

2.2 Design philosophy

The compressor design phase aims to deliver the initial design, which is the starting point for further optimisation. This section shortly touches upon the design process to illustrate the importance of aerodynamic simulation.

In the design process of an axial compressor, various disciplines such as aerodynamics, structures and mechanics need to be balanced with operating requirements related to noise, emissions and surge characteristics. Moreover, it is desired to achieve optimal performance at multiple operating points which often requires a trade-off. From these disciplines, the aerodynamic performance is generally the most influential. Therefore a common design approach starts at establishing a blade geometry to satisfy aerodynamic requirements, after which the structural integrity is checked to satisfy mechanical failure constraints [30].

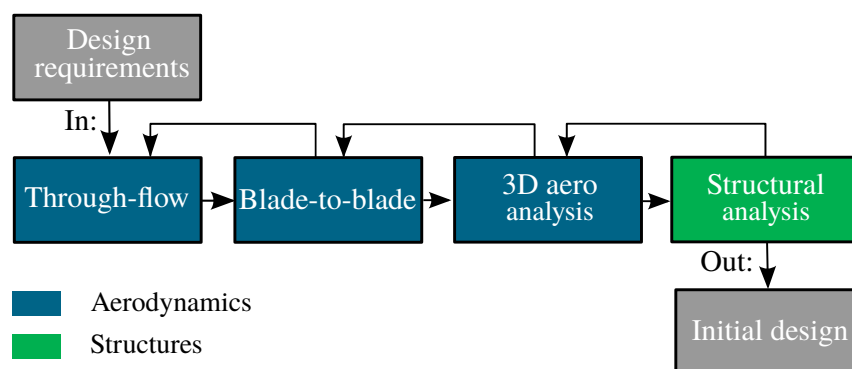


Figure 2.4: Flowchart of the compressor blade design process

¹Efficiency loss is directly proportional to the increase in specific entropy (total entropy over mass flow rate), assuming no external heat transfer and neglecting differences between static and stagnation conditions [5].

A typical design process is illustrated in Figure 2.4 which shows the process is highly iterative and involves multiple levels of modelling complexity. In the preliminary design phase, the multi-stage compressor is modelled using a one-dimensional through-flow code, that calculates inlet and outlet flow parameters for every stage using axi-symmetric meridional analysis and loss models. In this step, top-level requirements are specified like the desired pressure ratio and mass flow. The latter is imposed since it restricts the output power of the entire gas turbine. Next the blade rows are modelled in the blade-to-blade plane in a two-dimensional manner to allow for stacking in radial direction. Intermediately the design is already fixed to a certain degree [11]. In the following step, the flow can be accurately resolved using three-dimensional Computational Fluid Dynamics (CFD) simulations. Using CFD simulations, flow losses can be calculated with higher accuracy than lower-dimensional methods, against a higher computational cost. An example of a potential outcome from CFD simulation, the distribution of the Mach number in the S1 surface (blade-to-blade), is shown in Figure 2.5. Next, structural characteristics are simulated using Finite Element Analysis (FEA). An example of a potential result from FEA, the distribution of the von Mises stress over the blade, is shown in Figure 2.6. The *initial blade design* has now reached a high maturity level and serves as the starting point for optimisation. An overview of the optimisation process steps is given in Section 3.1.

Although the previously described design process is divided into three steps, it is not necessarily true that all of them are fully executed in practice. Due to considerable investments in research and development in the past, industry generally improves upon the latest technology standards, rather than starting a new design from scratch.




Figure 2.5: Illustration of the variation in Mach number throughout the flow field of a compressor channel, using aerodynamic flow simulation (CFD)




Figure 2.6: Illustration of the variation in von Mises stresses acting on a compressor blade, using structural analysis (FEA) [27].

2.3 Basics of CFD simulation for compressors

In this thesis, a compressor blade row is optimised for the maximum isentropic efficiency in certain operating conditions. For approximating the efficiency, Computational Fluid Dynamics (CFD) simulation is employed, using a RANS turbulence model or inviscid Euler equations for flow modelling. This section explains the basics of CFD simulation for compressors.

CFD simulation uses algebraic equations to solve and analyse fluid flows. These equations are solved at a grid of locations, which covers the geometry of the fluid problem. A CFD solution is characterised by two main elements. First of all the type of flow equations, which specify characteristics of the flow field. The second element is the mesh, which defines the locations at which flow equations are solved. Choosing the mesh structure and solver settings requires a trade-off between solution accuracy and simulation time. Both are tailored to capture the physical flow features in the channel, especially those mainly responsible for entropy generation and affecting the isentropic efficiency.

The following two sections provide a general description on selection of the mesh and solver equations for CFD simulation. This information is later used to decide on suitable settings for simulation processes employed in this work.

Mesh configuration

The accuracy of the flow simulation can be maintained by choosing a mesh that highly resolves the areas of interest mentioned above. Consequently, the error of the isentropic efficiency can be reduced. Whereas unstructured meshes are beneficial for fast generation and can easily be adapted to large variations in blade shapes seen throughout an optimisation, structured meshes allow for better control and better alignment of cells with flow streamlines. Using a structured mesh, a locally higher density of grid nodes can be implemented in regions where large velocity gradients are expected. Because of these advantages, in this work a structured mesh is applied.

In order to resolve the near-wall sublayer of the boundary layer, it is common to use one of the following approaches. The first approach is to use low-Reynolds models, which are relatively accurate. For resolving the flow subject to viscous friction, the cell thickness adjacent to the walls needs to be very fine. The width is determined by the dimensionless wall distance y^+ , the dynamic viscosity μ , flow density ρ and friction velocity U^* , through Equation 2.5. Using low-Re models, a dimensionless wall distance of $y^+ \approx 1$ is desired. A second approach is to apply wall-functions for calculating the velocity profile near to the wall using empirical relations. This is less expensive compared to using low-Reynolds models. As a result, the boundary layer behaviour is simplified and flow close to the walls is resolved with less accuracy. The thickness of the first cell can be increased to only roughly resolve flow in the boundary layer region, which reduces the mesh size. In this work, wall-functions are employed to save on computational time in the aerodynamic simulation of a blade design. This is necessary since thousands of member designs are evaluated.

$$y = \frac{y^+ \mu}{\rho U^*} \quad (2.5)$$

At the blade tip, the cell density should be high for capturing tip leakage flow, whose loss mechanism has been described previously. Using an increased number of radial streamlines in close vicinity to each other at the blade tip, leakage flows can be resolved. The number of closely spaced streamlines is adjusted to the tip gap size. In this work, including tip refinement is described as “enabling the tip gap”. In order to resolve the leading edge and the trailing edge, generally the mesh resolution is also increased locally at these locations. Usually it is hard to estimate the shock locations in advance, which prevents from increasing the cell density in advance.

Turbulence modelling

In a typical simulation problem, the mesh setup and type of flow equations are adapted to each other. For solving viscous flow, not only the mesh is to be refined in areas close to the wall, also a set of turbulence equations needs to be selected. Two common turbulence models for aerodynamic simulation of compressors are introduced in this section, from which one is selected in all viscous CFD simulations.

Solving the Reynolds Stress tensor

The introduction of a turbulence model is necessary when seeking an approximate solution of the Navier-Stokes equations for fluid flow. Several approaches are available to obtain a solution for these equations, which differ in accuracy. The most accurate approach would be to use a Direct Numerical Simulation (DNS). This simulation type solves for all length and time scales of turbulent flow and is therefore computationally expensive. A more practical approach for turbomachinery applications is

to average the smaller turbulence length scales and only solve the mean flow and largest turbulence structures, which is done using a Large-Eddy Simulation (LES). The last option, a Reynolds Averaged Navier-Stokes (RANS) simulation only solves the time-averaged solution of the Navier-Stokes equations, resulting in an approximation of the mean of the flow. With this averaging step, a term called the Reynolds stress tensor remains that needs to be approximated. In order to close the system of equations a turbulence model is used. The turbulence model considered in this thesis, belongs to the class of Eddy Viscosity Models (EVM). In this class of models, the Reynolds stress is assumed to be a function of the mean shear rate and the eddy viscosity ν . This parameter in its turn, can be expressed using simple algebraic 0-equation models or relatively complicated one-equation or two-equation models [13].

Two-equation models

A traditional example of a two-equation model is the $k - \epsilon$ model [17], that solves for the turbulence kinetic energy k and turbulence dissipation rate ϵ . Although it is relatively cheap in computational terms, it is known to poorly solve for complex flows with strong pressure gradients. A very popular two-equation RANS model Wilcox $k - \omega$ [31] solves the transport equation for turbulence kinetic energy k and the specific turbulence dissipation rate ω . Because it changes from wall-functions to a low Reynolds formulation, it works good to solve wall-bounded boundary layers, although it overpredicts the onset of separation. The later proposed Shear Stress Transport (SST) model, Menter SST $k - \omega$ [18], combines the previous two models by applying $k - \omega$ close to the walls and $k - \epsilon$ outside of the boundary layer region. In a study by Simoes et al. [28] on a rotor blade row of an axial compressor, simulation results from the three aforementioned turbulence models are compared. Using the Menter SST turbulence model has shown to best match the experimental data.

Inviscid flow simulation

Opposed to RANS models which solve viscous flow, Euler equations implement an additional simplification to the Navier Stokes equations, by neglecting the effect of viscosity. An inviscid Euler simulation is suitable for obtaining preliminary blade loadings and pressure distributions, on the assumption that inertial forces are much more prominent than viscous forces, which is the case at high Reynolds numbers. By neglecting shear stresses, boundary layers usually developing at solid boundaries are no longer resolved during simulation. This rigorous assumption discards a number of flow features including flow separation, viscous losses and secondary flows. Due to the strong connection between the entropy generation and associated reductions in isentropic efficiency, a rather large deviation in compressor efficiency is expected when solving using Euler equations, compared to RANS-based models. This deviation can either be positive or negative and is for instance, very dependent on positioning of mesh cells within the boundary layer. Regardless of this difficulty, successful use of the Euler equations is presented by Padr n and Alonso [20]. That work employs the multi-fidelity method to a 2D airfoil geometry optimisation, using Euler equations for the low-fidelity process and a RANS turbulence model for the high-fidelity process. More details on this work are provided in Section 3.8.

Concluding statement

Concluding from the basics of CFD simulation explained in the previous section, the mesh and solver equations affect the solution accuracy and simulation time for CFD analysis. This information is used in Section 4.4 to choose suitable simulation processes for evaluation of aerodynamic performance of compressor blades. This chapter has described the process to establish the initial blade design. The initial design is the starting point for optimisation, whose process is considered in detail in the next chapter.

3 Theory on aerodynamic shape optimisation

In this chapter, the reader is introduced to fundamental concepts of aerodynamic shape optimisation, applied to axial compressors.

First of all, Section 3.1 introduces the optimiser used in this thesis “AutoOpti”. Section 3.2 consequently presents the single-fidelity and multi-fidelity optimisation architectures. Respective Sections 3.3 and 3.4 provide an in-depth discussion on two essential elements of optimisations used in this work: evolutionary algorithms and surrogate models. In Section 3.6, the “Pareto diagram” and “Cumulative Volume Gain” are introduced, which are generally used to monitor and quantify optimisation progress. Section 3.5 then gives a description of the decision function types employed in AutoOpti, which control the budget allocated to low-fidelity members. Section 3.7 explains three metrics to estimate the potential of multi-fidelity optimisations. At last, Section 3.8 gives a review of reference literature on multi-fidelity optimisations applied to compressor blades.

3.1 AutoOpti and optimisation terminology

The first section of this chapter introduces the reader to the optimiser used in this work, called AutoOpti, and basic optimisation terminology.

In the current work, optimisations are performed using *AutoOpti*. This program offers multi-disciplinary, multi-objective optimisation using an evolutionary algorithm in combination with various surrogate models. Besides using single-fidelity optimisation, AutoOpti can also employ the multi-fidelity optimisation method. Development of AutoOpti started in 2006 by the Institute of Propulsion Technology of DLR (Deutsches Zentrum für Luft- und Raumfahrt). Its intended use was aerodynamic shape optimisation of compressor stages. With this application, attention was paid to allow for evaluation of expensive objective functions, which do not contain gradient information and are depending on large numbers of design variables [30]. Moreover, versatility of the program is demonstrated by the opportunity to parallelise internal calculation processes within as well as by possibilities for the user to adjust the region of interest during optimisation. Finally its suitability is evident from the freedom to setup a process chain of choice. AutoOpti has therefore also found application in various other disciplines such as fan optimisation and turbine stage optimisation.

The diagram in Figure 3.1 on the next page shows the basic elements of an optimisation scheme: 1) the optimiser performing and coordinating the search activities, 2) the parametrisation of the blade geometry and 3) an evaluation process which evaluates aerodynamic performance of a blade design. The optimiser tries to maximise certain parameters, called optimisation objectives, by making adaptations to the design variables $[d_1, d_2, d_3, \dots]$ in the design vector. The design vector is sent to the evaluation process where the blade is simulated and the objective function value is returned. The optimisations considered in this thesis are of the type “multi-objective”, since two objectives are optimised simultaneously. The starting point of an optimisation is the initial blade, whose design method was explained in Section 2.2. For a candidate design to be satisfactory, several constraints and restrictions need to be met, which can concern various disciplinary fields such as aerodynamics and structural mechanics.

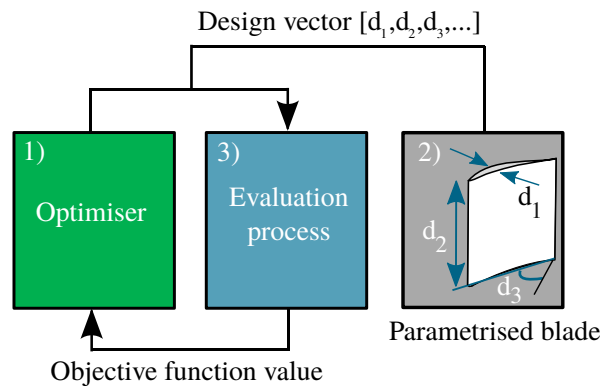


Figure 3.1: Diagram of the optimisation architecture for compressor blade shape optimisation

Automatic optimisation has allowed for design compromises between disciplines and in reduced design time and cost [27]. Currently, optimisation algorithms in turbomachinery design could apply simulation processes with a duration of over an hour. Such optimisations require several thousands of function evaluations to converge. Therefore the total runtime of an optimisation could be in the order of weeks.

3.2 Optimisation architectures

This section introduces the single-fidelity and multi-fidelity architectures and explains their differences. The single-fidelity technique represents the current state-of-the-art optimisation method employed within Siemens. Its architecture is discussed next, by means of a diagram. Slightly adapting this framework results in the multi-fidelity architecture that is explained consequently.

The single-fidelity optimisation architecture is presented in Figure 3.2. The central element in this configuration is the *master* that coordinates the optimisation and implements the search strategy. Besides, the *slave* executes a simulation process to evaluate performance of a blade design. The master sends each candidate design, called a *member*, to the slave which calculates and returns performance parameters of interest at simulated operating conditions. AutoOpti facilitates running multiple slaves in parallel, which allows simultaneous evaluation of multiple member designs. After evaluating a member, the slave sends member performance data to the master, which saves the results in a database.

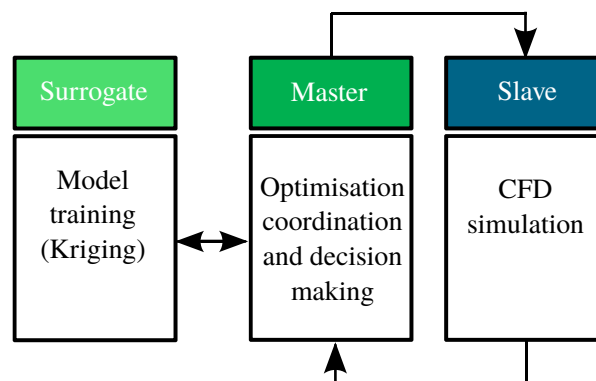


Figure 3.2: Diagram of the single-fidelity optimisation architecture

In AutoOpti, the search strategy for choosing new members is a type of *evolutionary algorithm*. This algorithm is explained in detail in Section 3.3. Since the evolutionary algorithm on its own requires many design evaluations to find the optimum, additionally a *surrogate model* is incorporated to generate many members at lower computational cost. As shown in 3.2, the surrogate model is trained during optimisation. Model training and model application are explained in Section 3.4.

Multi-fidelity optimisation implements an architecture that has a number of similarities with the single-fidelity framework. The master process is responsible for coordination and member evaluations are executed by one or more slave processes running in parallel. One slave merely evaluates one member at a time and relays performance information back to the master. In Figure 3.3, the multi-fidelity architecture is illustrated. Comparing Figures 3.2 and 3.3 it can be seen that the multi-fidelity framework employs an additional low-fidelity process for cheap evaluation of member designs as an alternative to the expensive high-fidelity process. The use of a much faster low-fidelity process enables a multi-fidelity optimisation to reduce the computational budget, compared to a single-fidelity optimisation.

When applying the multi-fidelity method, an optimiser can choose one of the two process chains to evaluate new members. This decision is performed by a decision function, denoted by DF in Figure 3.3. A discussion on potential decision function types in AutoOpti is considered in Section 3.5. Low-fidelity members can assist the optimiser in quickly identifying regions offering large improvement in the objective functions, while high-fidelity members are used to refine these regions and obtain performance information more accurately. A certain level of correlation between the process chains of different fidelity is required to fully benefit from the multi-fidelity technique. Successful applications of the multi-fidelity method in compressor optimisation are described in Chapter 3.8.

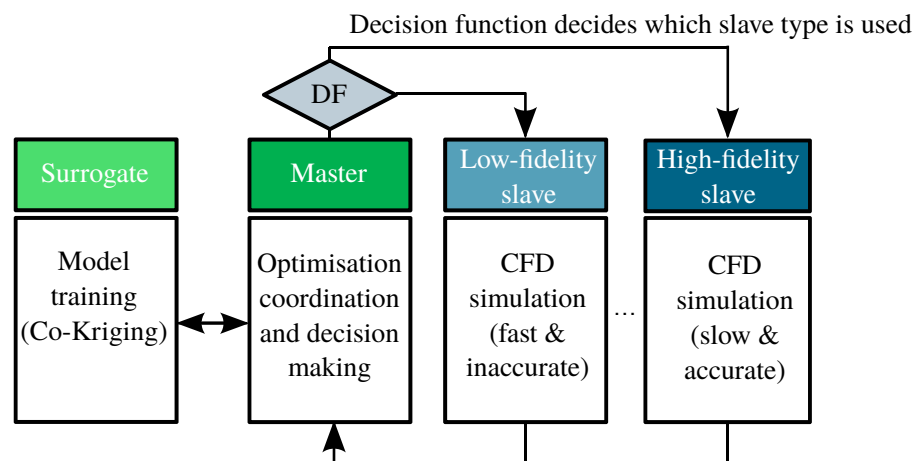


Figure 3.3: Diagram of the multi-fidelity optimisation architecture

3.3 Evolutionary algorithm

In order to find an optimal blade design, the optimisation code used in this thesis employs both a surrogate model and an evolutionary algorithm (EA). This section explains the theoretical background of the evolutionary algorithm.

Evolutionary algorithms are classified as *non-deterministic search methods* and saw their first development in the 1990's [15]. Non-deterministic search methods are employed in this work for two main reasons. First of all, in the current optimisation problem, the objective functions are not continuous. Since non-deterministic search methods do not require gradient information, like deterministic search methods,

they are more convenient to use in this case. Secondly, non-deterministic search algorithms allow for a trade-off between Pareto optimal members *after* finishing the optimisation. This is a benefit over deterministic search methods, which force automatic selection of a final design, when reducing the optimisation problem to a single-objective type [19].

Evolutionary algorithms consider a population of solutions and pick the most promising members to improve the objective functions. In this process, which is analogous to biological evolution, evolutionary algorithms use operations such as mutation and cross-over to propose new members showing improvement in the objectives [30]. Applying these generation methods, the optimiser can explore unknown parts of the search space and simultaneously prevent from getting stuck in local minima.

The optimiser *AutoOpti* implements various evolutionary operations interchangeably. Such operations manipulate design vectors of existing members showing superior performance, in order to establish a new member. First of all, the *mutation* operation scales all design parameters of an old member by a random percentage, chosen within a narrow bandwidth. Secondly, *cross-over* merges design parameters of two old members to establish a new member. The third operation, *differential evolution* is comparable to cross-over, but establishes the merger on the basis of three old members [12].

3.4 Surrogate modelling

This section explains the advantages of using a surrogate model to reduce the number of high-fidelity function evaluations. The following paragraphs provide a general motivation for using surrogate models as well as a short comparison of popular models. In Section 3.4-1, the Kriging model is explained, which is employed in this work by single-fidelity optimisations. In Section 3.4-2, the Co-Kriging model is discussed, which is used by multi-fidelity optimisations.

Motivation for surrogate models

The main drawback of using Evolutionary Algorithms is that establishing an accurate response surface requires many function evaluations, due to probabilistic nature of evolutionary operations [27]. Performing these evaluations repeatedly using a high resolution evaluation process, could be very expensive. The member selection process could be sped up considerably by applying an approximation model, also called a surrogate model, meta-model or response surface model.

An historical outlook of surrogate-assisted evolutionary calculation is provided by Jin [15], where it is suggested that this framework was first introduced in the mid 1990s. This work also mentions benefits and challenges of this framework. Due to its interpolation step, optima can be introduced in the surrogate model, that are not present in the original data set. One major challenge is to prevent the optimiser from pursuing these false optima. In Siller and Voss [27] it states that the main goal of using a surrogate model is to speed-up calculations by approximating the function of interest by means of a “functional relationship of acceptable fidelity”. Taking this point of view, the surrogate actually provides a lower fidelity approximation model alongside true function evaluations. *AutoOpti* allows parallel use of a surrogate model with evolutionary assisted generation, providing two options for selecting new member offspring. The surrogate model can only be used for member generation after the initialisation phase. During this phase an initial set of members is acquired using evolutionary operations, which serve as a training set for the surrogate model. The duration of the initialisation phase is decided by the engineer, by specifying a minimum number of converged members. In this work, the number is set to 80, which is equal to the number of design parameters used to generate a blade geometry.

Surrogate types

Popular surrogate types for optimisation problems in turbomachinery are Kriging (KG), Neural Networks (NN) and Response Surface Models (RSM). A suitable surrogate model type is selected on the basis of different criteria such as its accuracy, robustness, efficiency and transparency [16]. For a description of the application of these models in compressor blade shape optimisation, the reader is invited to consult the works of Samad [24] and Siller and Voss [27]. These papers demonstrate the use of the aforementioned surrogate types and their performance on reducing the number of function evaluations to find the optimum design. In this thesis, the different model types are not extensively compared but are shortly explained in the following. Response surface models fit a polynomial function to a training dataset of previous function evaluations. This surrogate type performs very well for small data sets and is generally easy to set up [24]. The application of a neural network is more preferable for complex, highly dimensional problems [7]. This model mimics the human brain of interconnected pathways between neurons, whose connections are given weights. The structure of a neural network imposes drawbacks such as a larger effort for model establishment and limited transparency in its use. An alternative to the previous two model types is Kriging, which is described as a Gaussian statistical model since it is based on normally distributed random variables. This modelling technique has gained much interest in recent years, owing to its flexibility and robustness [16]. Also it is successfully employed in aerodynamic optimisation problems within turbomachinery [3, 21, 23]. Moreover, taking into account the experience within Siemens of using Kriging for optimisations, this model type is used in this thesis in a single-fidelity optimisation setup. The model is explained in detail in Section 3.4-1. In this work, an extension of the Kriging model, called Co-Kriging, is used for multi-fidelity optimisation. Currently Co-Kriging is the only surrogate type available in AutoOpti to apply the multi-fidelity method. This model is discussed in Section 3.4-2.

Optimisation progress variability

In the work of Reimer [23], the concern is raised that two identical optimisations may not take the same path to converge towards an optimum. The variability, which can be interpreted as the variation in improvement of the objectives over time, can be explained by the following. In AutoOpti, the evolutionary algorithm selects the evolutionary operation type on a random basis, applying the weighting as previously discussed in Section 3.3. Moreover, a second probability parameter is used to decide for member generation using the evolutionary algorithm or using the surrogate model, which also influences the amount of potential progress. Results from Reimer [23] show that variability in single-fidelity optimisation progress is larger compared to variability in multi-fidelity optimisations. In the current thesis work, the progress variability is investigated by repeating a single-fidelity optimisation two times.

3.4-1 Kriging

In this section, the surrogate model *Kriging* is explained, which is used in this work for single-fidelity optimisations. The first paragraph explains Kriging model theory. The second paragraph describes the Kriging approximation of a test function in order to demonstrate the model to the reader.

Model theory

The Kriging model was originally developed by D.G. Krige in the 1950's as a statistical interpolation model to determine promising locations for mining. In due time, it has been applied in a variety of disciplines other than geostatistics, among which are environmental sciences, weather prediction and engineering applications. The Kriging method generally allows to estimate the value of a parameter in unexplored space, by interpolating values at surrounding probe locations where the values are known. The *correlation* between values at probe locations is spatially dependent and determines a so-called *covariance* of the dataset.

The known locations for probing, from now on called “support locations”, are denoted $\vec{x}_1, \dots, \vec{x}_n$. The associated values of the parameter under consideration are described by $y_1(\vec{x}_0), \dots, y_n(\vec{x}_n)$. The value to be estimated at the new location \vec{x}_0 , is written as $y^*(\vec{x}_0)$ and is calculated through Equation 3.1. Due to the statistical nature of the Kriging method, the variables will be modelled as random variables $Z(\vec{x}_1), \dots, Z(\vec{x}_n)$, each having an expected value and a variance. The first major assumption is that the parameter to be estimated $Z^*(\vec{x}_0)$, can be calculated by a linear sum of the known variables $Z(\vec{x}_1), \dots, Z(\vec{x}_n)$ and their weights w_1, \dots, w_n , as presented in Equation 3.2. The goal is to determine these weights [23].

$$y^*(\vec{x}_0) = \sum_{i=1}^n w_i y(\vec{x}_i) \quad (3.1)$$

$$Z^*(\vec{x}_0) = \sum_{i=1}^n w_i Z(\vec{x}_i) \quad (3.2)$$

A modelling error exists between the estimated variable $Z^*(\vec{x}_0)$ and its true value $Z(\vec{x}_0)$. The error function is defined in Equation 3.3. The Kriging approximation is formulated as a Best Linear Unbiased Estimator, which can be explained as follows. The fact that it is *best* implies that the variance of the error function is to be minimised, as described by Equation 3.4. This approximation approach can be said to be *linear* from an analogy with linear regression, which fits data through interpolation. However, instead of only taking spatial information, now also correlation information is used. Lastly, the estimation is *unbiased* since the expected value of the error function is zero as shown in Equation 3.5 [25].

$$F^*(\vec{x}_0) = Z(\vec{x}_0) - Z^*(\vec{x}_0) = Z(\vec{x}_0) - \sum_{i=1}^n w_i Z(\vec{x}_i) \quad (3.3)$$

$$\text{var}[F^*(\vec{x}_0)] = \min_{w_1, \dots, w_n} \text{var}[Z(\vec{x}_0) - \sum_{i=1}^n w_i Z(\vec{x}_i)] \quad (3.4)$$

$$E[F^*(\vec{x}_0)] = 0 \quad (3.5)$$

For determining the Kriging weights, the covariances between support locations need to be obtained. These are collected in a covariance matrix \mathbf{K} , calculated by the variance σ^2 and a correlation matrix \mathbf{R} , as shown by Equation 3.6. Each entry of the correlation matrix contains the correlation between two support locations, expressed by $\rho(\vec{x}_i, \vec{x}_j)$. For each pair of support locations x_i and x_j , the correlation is to be calculated by application of Equation 3.7. Here k indicates the dimensionality of a location (i.e. in three-dimensional space, $k = 3$) and l is the index of summation. The correlation function can be considered similar to a Gaussian normal distribution, whereas it can take values between 0 and 1, where a value close to 1 indicates strong correlation [25]. To solve its function, a so-called hyperparameter $\bar{\Theta}$ needs to be determined. The actual correlation between two support locations is high when they are in close proximity. The value of the hyperparameter, in turn, influences how much the correlation decreases for more remote support locations. It is recommended to consult the work of Schmitz [25] for a visual impression of these two factors.

$$\mathbf{K} = \sigma^2 \mathbf{R} \quad (3.6)$$

$$\rho(\vec{x}_i, \vec{x}_j) = \exp\left(-\frac{1}{2} \sum_{l=1}^{l \leq k} (e^{\bar{\Theta}_l} |\vec{x}_{i,l} - \vec{x}_{j,l}|^2)\right) \quad (3.7)$$

The optimisation of the hyperparameter is called *training* of the Kriging model. The hyperparameter is found by applying the Maximum Likelihood method, which is an estimation method from statistics

based on maximisation of the likelihood function. For a detailed description of this optimisation step, it is recommended to consider the work of Schmitz [25]. Due to high computational cost for surrogate model training, it is cheaper to generate members using an evolutionary algorithm compared to using a surrogate model [23].

Model demonstration

In this paragraph, the application of the Kriging model is demonstrated to gain a better understanding of the theoretical concepts explained in the previous section. By using Kriging to approximate a three-dimensional testing function, in the following it is shown that model accuracy can be improved by increasing the number of function evaluations on the domain.

The test function considered in this demonstration is a scaled variant of the *Mishra's Bird* function often seen in optimisation testing. The surface $z = f(x, y)$ is shown in Figure 3.4 on a rectangular domain $[0, 10]$ for both x and y . The three-dimensional surface sees a couple of minima and optima of unequal heights over the selected domain. This testing problem can be considered analogous to approximation of a performance parameter subject to variation of two design variables. For higher dimensional optimisation problems, this surface would scale up in multiple dimensions.

Over the effective domain, a square grid is established with a resolution of 30 in both x and y directions. The Kriging surrogate model is initialised and trained over all grid locations, when supported by a sample set of support locations. Training takes place over the course of 10 iterations, during which the optimum value of Θ is determined. The sample points, indicated as green dots in Figure 3.5, are exact evaluations of the test function at equally spaced support locations. In the following, the improved precision of the Kriging is demonstrated by increasing the number of sample points. The sampling resolution is indicated by SR , which denotes the number of samples in x and y directions. Precision of the surrogate model can be assessed locally by inspecting the height of the error bars indicated directly underneath the surfaces, showing the absolute difference between the Kriging approximation and the test function at a every grid node (x, y) . Additionally, the Root Mean Squared Error (RMSE) is calculated to quantify the global approximation error. The performance of the Kriging model is assessed for sampling resolutions $SR = 3$, $SR = 6$ and $SR = 9$. The resulting approximation from the Kriging model is shown in Figures 3.5, 3.6 and 3.7 respectively.

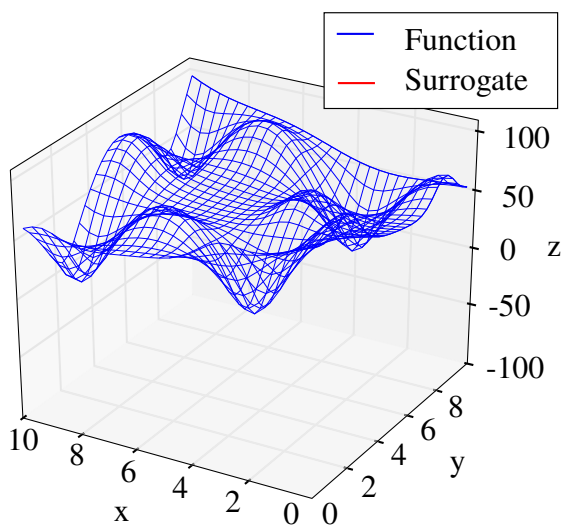


Figure 3.4: Fitting function

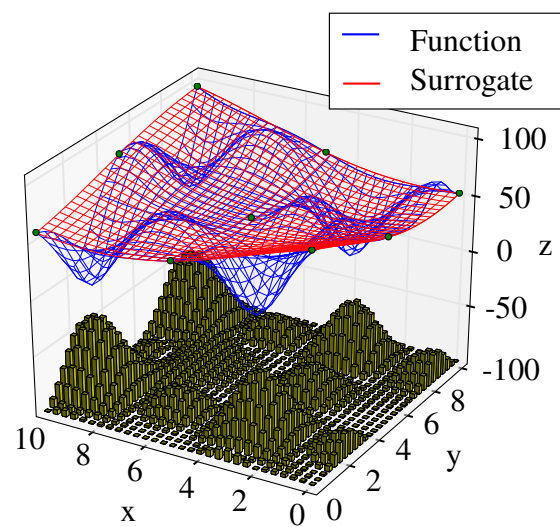


Figure 3.5: Kriging model for $SR = 3$

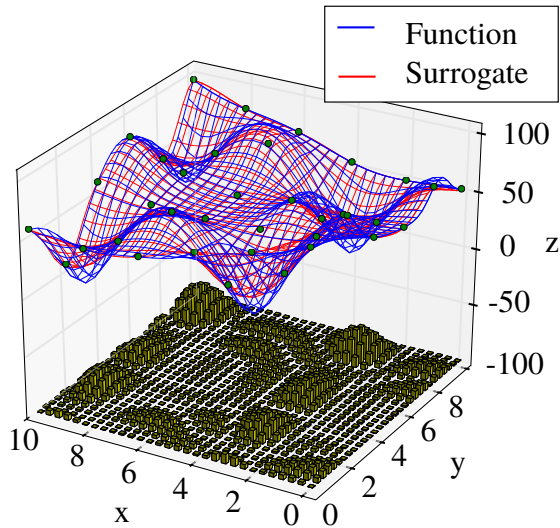

Figure 3.6: Kriging model for $SR = 6$ Figure 3.7: Kriging model for $SR = 9$

Table 3.1: RMSE of the surrogate approximation

Sampling resolution, SR	RMSE
3	19.0
6	5.9
9	1.9

By inspecting the error bars, it can be concluded that the Kriging model is able to approximate the region directly surrounding sample points with higher precision, compared to areas far away from those points. The error exactly at a support location reduces to zero. Considering the global error, which is quantified by the RMSE in Table 3.1, it is demonstrated that adding function evaluations enhances overall accuracy of the surrogate model. The process of calculating the hyperparameters during surrogate training, can be envisioned as matching of the valleys contained in the surrogate model with the minima in the true fitting function.

3.4-2 Co-Kriging

For implementation of the multi-fidelity method, in this thesis the Co-Kriging surrogate model is employed. This section considers the theoretical background of Co-Kriging and shortly explains the correlation assessment between two datasets. In the first paragraph, Co-Kriging theory is explained. Its demonstration is discussed in the second paragraph.

Model theory

Currently, the Co-Kriging model is the only model available in AutoOpti for multi-fidelity optimisation. In Co-Kriging, a second data set of lower-fidelity function evaluations is correlated with high-fidelity data. In this thesis, the evaluation process entails CFD simulation of two axial compressor stages. The low-fidelity process chain is established by reducing the mesh and solver settings of the high-fidelity process chain. Co-Kriging is, just like Kriging, also based on the BLUE estimation method, which implies a similar linear summation of random variables and weights is utilised. Now a summation is required over all fidelity levels. Since two fidelity levels are used in this thesis, $m = 2$ in Equation 3.8.

$$Z^*(\vec{x}_0) = \sum_{j=1}^m \sum_{i=1}^n w_{i,j} Z(\vec{x}_{i,j}) \quad (3.8)$$

The covariance matrix from Equation 3.6 is also adapted when including a data set of low-fidelity function evaluations. Elaborating on the subsequent calculation is refrained from due to its complexity. Instead, it is recommended to consider the work of Reimer [23] for a detailed approximation. In the following, only significant deviations from the Kriging model are summarised.

First, it is assumed that every function result from the high-fidelity process y_{high} , can be decomposed in a function value from the low-fidelity process y_{low} and an error term y_{err} . Consequently, three types of covariance functions are required: one type considers covariance between function evaluations from two different data sets, and the other two types consider the covariance between two evaluations out of the same data set. Moreover, the error functions Z_{err} do not correlate with associated low-fidelity function values Z_{low} , as shown in equation 3.10.

$$y_{high} = y_{low} + y_{err} \quad (3.9)$$

$$Cov(Z_{low}, Z_{err}) = Cov(Z_{err}, Z_{low}) = 0 \quad (3.10)$$

Using the three covariance functions to correlate the results between two support locations, but also between two fidelity levels, the covariance matrix can be established as demonstrated by Fogel [9]. The Co-Kriging model uses twice as many hyperparameters compared to Kriging and therefore requires more time to train. Apart from the parametric background, also the nature of error between low-fidelity and high-fidelity results can influence training efforts. When, for instance, low-fidelity results are heavily fluctuating with a certain amount of noise, many support functions may be required to optimise the hyperparameters and calculate the covariances.

Model demonstration

In this work, the Co-Kriging model is not illustrated on the basis of an example. However, one can imagine its workings by referring back to the demonstration case of Kriging in Figures 3.5 to 3.7. In this example, the response surface of the Co-Kriging surrogate would be constructed from two sets of support locations. Function evaluations from the two datasets are not necessarily located identical x, y, z locations, but can be spread independently.

3.5 Decision function

In AutoOpti, the choice for selecting either the low-fidelity or high-fidelity process for evaluation is made by a *decision function*. The decision function has a large influence on how much computational budget is allocated to low-fidelity member evaluations, hence it has a considerable impact on the benefits of multi-fidelity optimisation.

In this section, the decision function types available in AutoOpti are explained in the first paragraph. In the second paragraph, the automatic decision function, called *variance*, is considered in more detail. This decision function takes into account the solver time of the two process chains and the quality of their results to choose a suitable low-fidelity process. In the closing paragraph of this section, the impact of this decision function on optimisation progress is considered.

Available decision functions in AutoOpti

Three decision function types can be employed in AutoOpti: *constant*, *linear* and *variance*. First, the type *constant* allows the user to take manual control over the decision function. The user chooses a constant ratio of low-fidelity over high-fidelity members for the entire optimisation. Second, the *linear* type employs a ramp that increases with the number of evaluated members. Applying one of the former two function types, the user needs to think of the desired ratio of low-fidelity and high-fidelity members at different phases of the optimisation. Third, an automated function type called *variance* is available in AutoOpti. This function bases a decision on the member results available at the moment a decision is to be made. The results from the work of Reimer [23] have shown that employing the variance decision function for multi-fidelity optimisation could attain the best speed-up and improvement in optimum efficiency, compared to an equivalent single-fidelity optimisation. Based on this proven result, in the current work the variance decision function is selected for most optimisations, unless stated otherwise.

Method of variance decision function

In this paragraph, the approach of the variance decision function is demonstrated by means of an application case of the Co-Kriging surrogate model. The variance decision function uses two types of information to decide for either low-fidelity or high-fidelity member evaluation. It considers the *time* for evaluating the processes and the benefit in *quality* each of them can offer.

The following explanation is a short summary of the description provided by Reimer [23]. Instead of revisiting the testing function of Section 3.4-1, where two design variables influenced a performance parameter, the example case now only considers one design variable, evaluated by two separate processes. The aim of this example is to demonstrate the influence of *two* data sets with different fidelity for the construction of the surrogate model.

The selected fitting function is a pure sinusoid. Both the fitting function and its Co-Kriging approximation are visible in the top part of Figure 3.8. The modelling inaccuracy with respect to the sinusoid function, expressed by the standard deviation of error, is presented in the lower part of the

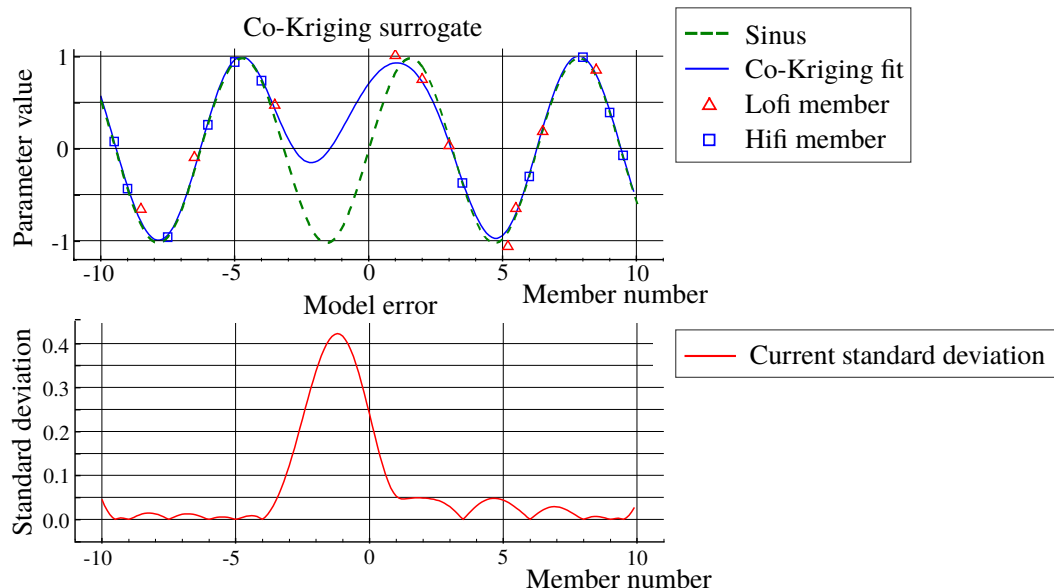


Figure 3.8: Co-Kriging surrogate model of a sinusoid function using several high-fidelity and low-fidelity members [23].

same figure. This plot shows that zero error is achieved at locations where high-fidelity members are included (blue rectangles), meaning the surrogate model exactly attains the value of the fitting function. At locations where only a low-fidelity member is included (red triangles), a non-zero deviation is still present. At $x = -1.5$ the standard deviation of the error is relatively large. In order to make a decision to either include a low-fidelity or high-fidelity member at this location, four steps are taken as explained in the following.

Step 1: Determine the effect on quality by high-fidelity member forecast

First a high-fidelity member is added at $x = -1.5$, indicated by the black rectangle in Figure 3.9. Due to inclusion of this member, the forecasted reduction in standard deviation of the surrogate at $x = -1.5$ reduces to zero. The reduction is denoted by $\sigma_{CK,HF}$, where index CK is the abbreviation for Co-Kriging. One should realise that, although the standard deviation reduces to zero, the surrogate will not perfectly approximate the sinusoid function due to small modelling error of the surrogate model.




Figure 3.9: Co-Kriging approximation, including a forecasted high-fidelity member at $x = -1.5$ [23]

Step 2: Determine the effect on quality by low-fidelity member forecast

In step 2, a low-fidelity member is added at $x = -1.5$ to assess the effect on standard deviation reduction of the Co-Kriging model, achieved when choosing the low-fidelity process. The reduction in standard deviation due to this anticipated member cannot be obtained directly. This can be accomplished by consequently adding a high-fidelity member and applying Equation 3.11. In this equation, the reduction after low-fidelity and high-fidelity member forecasts is denoted by $\sigma_{CK,(LF\&HF)}$, while the reduction in standard deviation from only a high-fidelity member (step 1) is written as $\sigma_{CK,HF}$. The result in anticipated Co-Kriging approximation and forecasted standard deviation can be observed in Figure 3.10.

$$\sigma_{CK,LF} = \sigma_{CK,HF} - \sigma_{CK,(LF\&HF)} \quad (3.11)$$

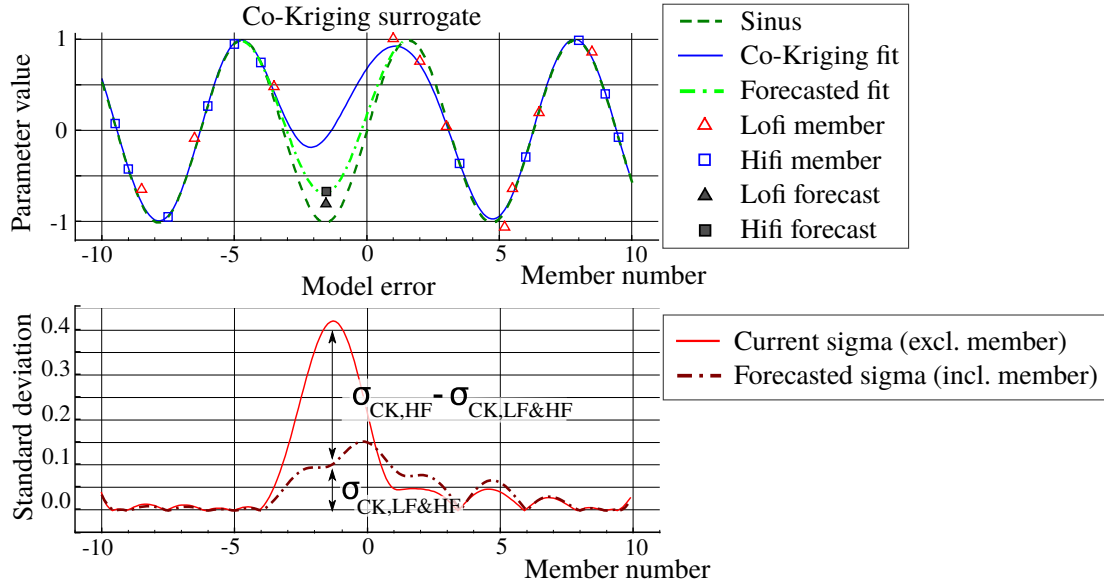


Figure 3.10: Co-Kriging approximation, including a forecasted low-fidelity member at $x = -1.5$ [23]

Step 3: Determine time effects

Besides the effects on quality $\sigma_{CK,HF}$ and $\sigma_{CK,LF}$, also the time influences the decision-making. The average process time is calculated for all high-fidelity members $t_{avg,HF}$ and low-fidelity members $t_{avg,LF}$. Consequently a ratio is calculated of the reduction in standard deviation of the model error and the average process time. The ratios are denoted by R_{LF} for the low-fidelity process and R_{HF} for the high-fidelity process and described by Equations 3.12 and 3.13 respectively. A performance factor $crit$ is obtained by dividing R_{LF} by R_{HF} , through Equation 3.14. This factor is used in the next step.

$$R_{HF} = \frac{\sigma_{CK,HF}}{t_{avg,HF}} \quad (3.12)$$

$$R_{LF} = \frac{\sigma_{CK,LF}}{t_{avg,LF}} \quad (3.13)$$

$$crit = \frac{R_{LF}}{R_{HF}} \quad (3.14)$$

Step 4: Make the final decision

Steps 1 to 3 are repeated for every available surrogate model, effectively applied during the optimisation. In this thesis, the total number of surrogate models is equal to $n_{CK} = 4$ and includes two objective functions and two restrictions. The parameter $critExp$ determines the final decision and is calculated by Equation 3.15. This parameter includes the arithmetic average off all contributions $crit_1$ to $crit_{n_{CK}}$. When $critExp$ reduces below 1, the high-fidelity process chain is selected. For its calculation, the parameter POI is introduced, called Probability of Improvement. In short, the POI allows to give higher priority to high-fidelity members compared to low-fidelity ones. This is necessary when the low-fidelity process is highly correlated and therefore constantly preferred over high-fidelity members. Since only high-fidelity members can contribute towards improvement in the optimum, they ought to be given priority. In another scenario where low-fidelity members show mediocre correlation with high-fidelity results, it is expected that the low-fidelity part of the surrogate model no longer improves its accuracy after evaluating many members. In this scenario, the reduction σ_{LF} becomes smaller than σ_{HF} , causing the final decision parameter $critExp$ to decrease and choose only high-fidelity member evaluations.

Implications on multi-fidelity optimisation

From the previous explanation of the variance decision function, the following can be concluded. The variance decision function in multi-fidelity optimisation, calculates the cost and the accuracy of the low-fidelity model with respect to the high-fidelity model and uses this information to decide which process chain is used for the next member evaluation. Thereby the variance decision function has a large influence on the number of low-fidelity members that is evaluated during an optimisation. This implication is reconsidered in Section 4.6, when estimation guidelines are introduced to forecast the potential of multi-fidelity optimisations. These guidelines take into account cost, quality and the number of low-fidelity member evaluations.

$$critExp = \frac{\sum_{i=1}^{n_{CK}} crit_i}{n_{CK}} (1 - POI) \quad (3.15)$$

3.6 Optimisation progress quantification

In this work, optimisation progress is described as the combined improvement of objectives over time. This section explains how progress is quantified in this work. The first paragraph of this section introduces the Pareto diagram used to visualise the improvement of the objectives. The second paragraph explains how improvement in both objectives is quantified by a single parameter, called the Cumulative Volume Gain.

Pareto diagram

The optimisation algorithm in this thesis aims to simultaneously improve two objectives: compressor isentropic efficiency in ISO and in HOT operating conditions. Objectives are enhanced until the point at which new members can only improve in one objective by degrading the other. These members are described as Pareto optimal or non-dominant solutions [1]. In AutoOpti, Pareto optimal members have a rank 1, while non-optimal members have higher Pareto ranks. In this thesis, improvement is visualised by means of a *Pareto diagram* that presents the two objectives attained by the entire population of members in one chart. A Pareto diagram of a single-fidelity optimisation is illustrated on the left of Figure 3.11.

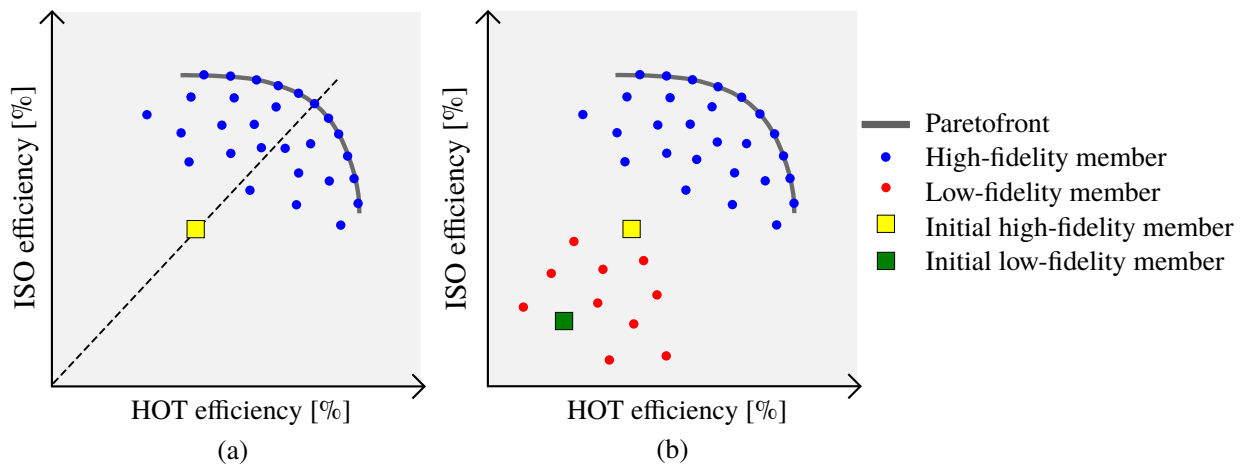


Figure 3.11: Pareto diagrams for single-fidelity optimisation (a) and multi-fidelity optimisation (b), visualised with Pareto fronts

The members achieving the highest improvement in both objectives are shown in the top right corner of Figure 3.11. The Pareto diagram of a multi-fidelity optimisation presents both member populations, as shown on the right in Figure 3.11. When the optimiser has identified a design space region providing high gain in both objectives, the Pareto optimal solutions line as a *Pareto front*. In multi-fidelity optimisations, only high-fidelity members form a Pareto front, since the optimisation algorithm uses the high-fidelity process chain to accurately resolve the design space region where the optimum is located. Due to an offset in accuracy, low-fidelity members are spread randomly as shown in the sketch of Figure 3.11. In the initial phase of the optimisation, improvement in objectives can be large over time causing the front to move relatively quickly in the direction of combined improvement. This movement slows down as achievable improvements become smaller. Movement of the Pareto front is an indication of optimisation progress. Therefore it is quantified by a new parameter, described in the next section.

Out of the final set of Pareto optimal solutions, a single optimum design is chosen to calculate the average efficiency gain with respect to the initial design. In this work, the isentropic efficiencies in ISO and HOT operating conditions are considered equally important, hence the optimum member is located on the diagonal of the Pareto diagram (dashed line in diagram (a) of Figure 3.11). The member closest to the point of intersection of the diagonal and the Pareto front is selected as the optimum design of the optimisation. The average efficiency of this member is calculated by applying Equation 3.16 and is used in Chapter 5 to compared optima from different optimisations.

$$\eta_{avg} = (\eta_{ISO} + \eta_{HOT})/2 \quad (3.16)$$

Cumulative Volume Gain

In this paragraph, the parameter Cumulative Volume Gain is introduced that quantifies combined improvement of the objectives. This parameter is used in this work to compare performance of multiple optimisations.

In Figure 3.12, a sketch of the Pareto diagram is shown for high-fidelity members at three time instances. At time instance 1, a front has been formed. At a later time instance 2, new members have been evaluated that shift the front and enclose an area with respect to the previous front location. This area is called the *Volume Gain* by AutoOpti.

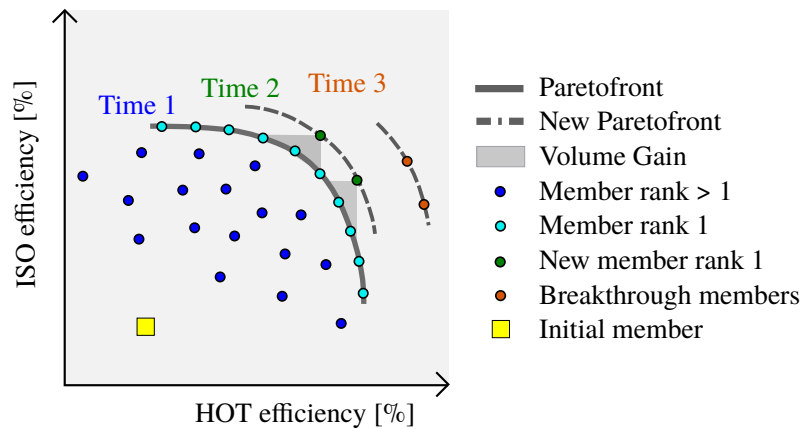


Figure 3.12: Pareto diagram visualising the Volume Gain at Pareto front shift

When summing all contributions to the volume gain, the *Cumulative Volume Gain* (CVG) is obtained. The CVG increases with optimisation time and is used in this work to quantify optimisation progress. An increase in CVG can give an indication of improvement in ISO and HOT isentropic efficiencies during an optimisation, as the CVG parameter captures the combined improvement by all Pareto optimal members. In Section 4.6, a difference in CVG will be used to compare performance of two optimisations.

Previously in Section 3.4 the occurrence of variability in optimisation progress was introduced, which is caused by two internal probability parameters in AutoOpti. In the remainder of this paragraph, the effect of variability on the CVG improvement over time is discussed. The concept presented here is referred to when defining the convergence criterion (Section 4.5) and when presenting the metrics for comparing multi-fidelity performance (Section 5.2). Optimisation progress, which is expressed in this work as the improvement in CVG over time, sees large variations. Improvements in terms of ΔCVG can be very small in the initialisation phase, when the optimiser has not yet identified the direction of largest improvement of the objectives. Conversely, periods of small improvement could be followed by a sudden steep increase in CVG. The optimisation then achieves a ‘breakthrough’, by identifying a new region in the design search space that can considerably improve the objectives. Figures 3.12 and 3.13 show that members evaluated at time 3 can drastically enhance the CVG. The moment these members are evaluated, could vary slightly through repetition of identical optimisations. When they are large, these unanticipated variations in optimisation progress complicate the assessment of convergence, as well as the comparison of two optimisations.

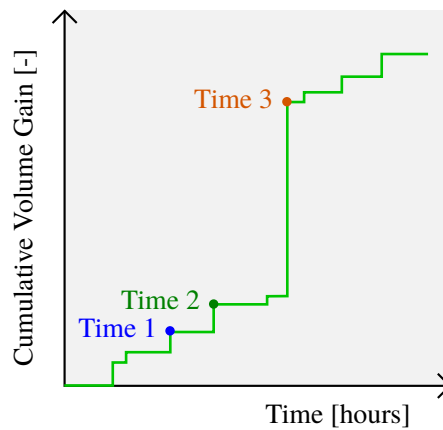


Figure 3.13: Demonstration of CVG improvement, with two time instances highlighted

3.7 Potential metrics

In this section, an estimation method is introduced which evaluates the potential benefit of a multi-fidelity optimisation. The estimation method is employed in this work to estimate potential of a multi-fidelity optimisation over a single-fidelity optimisation on the same computational budget. The method evaluates four guidelines developed in the work of Toal [29], which require the definition of three potential metrics.

The first three paragraphs of this section each look into one of the potential metrics: cost ratio, correlation coefficient and replacement ratio. The fourth paragraph explains the guidelines for potential estimation.

Potential metric 1: Cost ratio

The cost ratio C_r , indicates the ratio of computational cost for evaluation of the high-fidelity process compared to evaluation of the low-fidelity process, given by Equation 3.17. In this work, the cost of a

process is referred to as the time to run the process. The cost ratio is thus the speed-up factor of the low accuracy process chain with respect to high-fidelity process. The higher the speed-up, the lower is $t_{avg,LF}$ with respect to $t_{avg,HF}$, hence the lower is C_r .

$$C_r = \frac{C_{lofi}}{C_{hifi}} = \frac{t_{avg,LF}}{t_{avg,HF}} \quad (3.17)$$

Potential metric 2: Correlation coefficient

In this thesis, the correlation coefficient r^2 , quantifies the accuracy at which low-fidelity evaluations can approximate high-fidelity evaluations. Unless specified otherwise, the calculation of the correlation coefficient is performed *after the optimisation is finished*. This calculation is explained in the following.

First, every high-fidelity member evaluation is probed in the Co-Kriging surrogate model, to obtain its predicted low-fidelity response. Second, the so-called Pearson correlation coefficient r_{xy} is calculated using Equation 3.18, where x and y represent results from two different fidelity levels. The Pearson correlation r_{xy} can take values on the interval $[-1, 1]$. At last, Equation 3.19 is applied as recommended by Toal [29] to discard any negative correlation coefficients.

The approach described above allows to calculate the correlation coefficient only for parameters for which a surrogate model is established. The setup employed in this work uses four surrogate models, for two objectives and two restrictions. In this work, an average correlation coefficient is calculated for a multi-fidelity optimisation, by applying Equation 3.20.

Figure 3.14 presents an illustration of all member evaluations from well-correlated low-fidelity and high-fidelity processes. An illustration of evaluations from poorly correlating processes is shown in Figure 3.15. Comparing results in these two figures, one could observe that evaluations from poorly correlating processes are further away from the line of positive correlation.

$$r_{xy} = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^n (x_i - \bar{x})^2 \sum_{i=1}^n (y_i - \bar{y})^2}} \quad (3.18)$$

$$r^2 = r_{xy}^2 \quad (3.19)$$

$$r^2_{avg} = \left(\frac{\sum_{i=1}^4 r_i}{4} \right)^2 \quad (3.20)$$

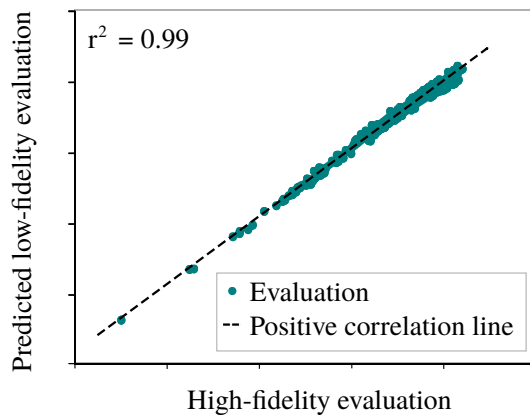


Figure 3.14: Sketch of member evaluations from well-correlated process chains

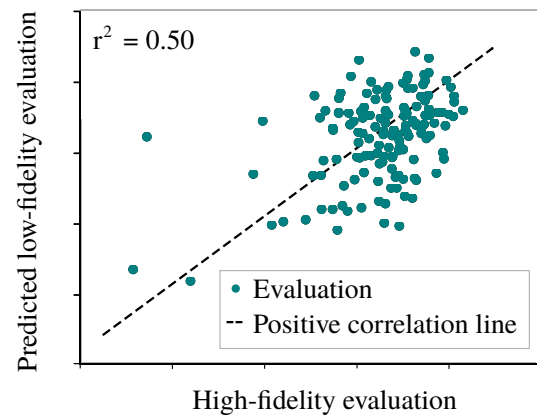


Figure 3.15: Sketch of member evaluations from poorly-correlated process chains

Potential metric 3: Replacement ratio

The last parameter to assess performance criteria is called the replacement ratio, f_r . This parameter describes the fraction of total budget allocated to low-fidelity members in multi-fidelity optimisation, compared to a single-fidelity optimisation with an identical computational budget. The replacement ratio is calculated by Equation 3.21, where $n_{HF,MF}$ and $n_{HF,SF}$ denote the number of high-fidelity members in multi- and single-fidelity optimisations respectively. Using optimisation diagnostics, $n_{HF,MF}$ can be obtained straight away. For obtaining $n_{HF,SF}$, first the total computational budget spent in a multi-fidelity setup needs to be calculated. This computation is shown in Equation 3.22. The budget for multi-fidelity optimisation equals the summation of total UDP times for high-fidelity processes ($n_{HF,MF} \cdot t_{avg,HF}$) and for low-fidelity processes ($n_{LF,MF} \cdot t_{avg,LF}$), where any training and waiting times are disregarded. In single-fidelity optimisations, all budget is spent on evaluating high-fidelity members, hence $t_{avg,HF}$ is in the denominator of Equation 3.22.

$$f_r = 1 - \frac{n_{HF,MF}}{n_{HF,SF}} \quad (3.21)$$

$$n_{HF,SF} = \frac{n_{HF,MF} \cdot t_{avg,HF} + n_{LF,MF} \cdot t_{avg,LF}}{t_{avg,F}} \quad (3.22)$$

Guidelines for potential estimation

This paragraph describes four guidelines recommended in the work of Toal [29] to estimate multi-fidelity potential. In this work, multi-fidelity optimisation performance is estimated by employing these guidelines. The results of the potential assessment are presented in Chapter 5. The four guidelines A,B,C and D are formulated as follows:

- A. The correlation between the low and high fidelity function should be reasonably high, $r^2 > 0.9$.
- B. More than 10% of the total evaluation budget should be converted to cheap evaluations, $f_r > 0.1$.
- C. No more than 80% of the total evaluation budget should be converted to cheap evaluations, $f_r < 0.8$.
- D. There should always be more cheap data points than expensive, according to $f_r > 1.75 / (1 + \frac{1}{C_r})$

3.8 Literature review on multi-fidelity compressor optimisations

In this section, literature is discussed which considers the multi-fidelity method in shape optimisation of compressor blades. From a short overview of previous work in this field, two major reference works are selected and described in detail. These two publications discuss the benefits of the multi-fidelity method, compared to single-fidelity optimisation of the same initial blade design.

A recent review by Fernández-Godino et al. [8] summarises the main areas for application of multi-fidelity models, the processes of different fidelity and the types of surrogate models frequently used. From information published in 18 papers, a diagram is established that shows both the cost savings of employing the multi-fidelity method and the cost ratio between high-fidelity and low-fidelity processes. This diagram is shown in Figure 3.16. In the work of Fernández-Godino et al. [8], it is concluded that there is no clear relationship between the cost of multi-fidelity (MUFI) versus single-fidelity (SIFI) optimisations and the cost ratio between low-fidelity and high-fidelity processes. That is because savings also depend on the correlation between low-fidelity and high-fidelity processes and are therefore highly problem dependent.

Figure 3.16 illustrates that a low cost ratio does not guarantee high savings. It also shows that a cost ratio of as low as $C_r = 0.10$ is required to achieve cost savings of up to 90%. An example case with high savings, is presented in the work of Padròn and Alonso [20], highlighted in Figure 3.16. In that work, the multi-fidelity method is used for optimising a 2D airfoil geometry. The setup employed inviscid Euler equations for the low-fidelity process and a RANS-based turbulence model for the high-fidelity process. Computational cost could be reduced by 64%, compared to cost from single-fidelity optimisation. Cost savings are interpreted as an optimisation speed-up of 64%. The observations presented in this paragraph are reconsidered in Chapter 5, to evaluate performance of optimisations presented in this thesis.

In the following paragraphs, two research works are discussed in more detail, which serve as references for the setup and findings presented in this thesis. To avoid confusion, in this section this thesis is called “the current work”.

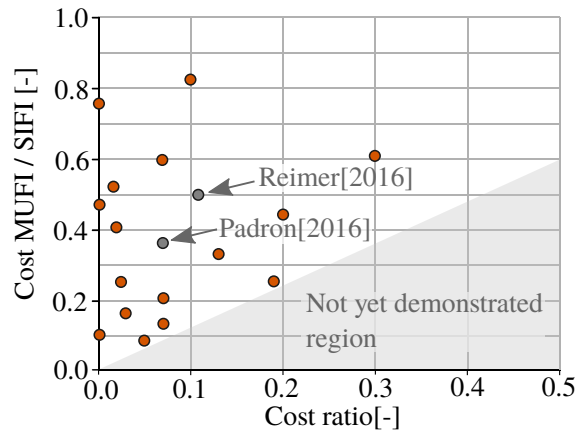


Figure 3.16: Cost of multi-fidelity versus single-fidelity optimisations described in 18 papers, plotted against the cost ratio of low-fidelity and high-fidelity processes [8]. Highlighted are the works from Reimer [23] and Padròn and Alonso [20].

Reimer - Multi-fidelity optimisation of a compressor fan

The first reference work is the thesis by Reimer [23], which employs the multi-fidelity method for aerodynamic optimisation of a single-stage compressor fan. The objective of this thesis is to test and evaluate various optimisation setups which differ in the type of decision function used. Hence the low-fidelity process is identical for all setups. The performance of these setups is compared to the performance of a single-fidelity reference optimisation. The optimisation objectives are the isentropic efficiency under nominal operating conditions and the increase in pressure ratio between two operating points, over a certain mass flow. All optimisations are repeated four times, in order to take into account progress variability. A total of 64 design parameters is used to parametrise the geometry of the rotor and stator. This number is 20% lower than the number used in the current work (see Chapter 4.1). Low-fidelity and high-fidelity process chains contain an aerodynamic simulation step using CFD and a structural analysis step using FEM. The FEM calculation is merely used to calculate blade strain, which is restricted during the optimisation. No information is provided on the convergence criterion employed in this thesis. Details about the models used in this thesis are presented in Table 3.2. In this table, the mesh ratio indicates the low-fidelity mesh size versus the high-fidelity mesh size. From this table it could be seen that the mesh size per row is 550,000 cells. This is 59% more than used in the high-fidelity model in this thesis. For the low-fidelity model of Reimer [23], the number of radial streamlines is reduced by 75%, which is 20% more than the reduction of the best low-fidelity

model in the current work. Furthermore, the tip vortex is not resolved by the low-fidelity model in the work of Reimer [23]. At last, both process chains employ the same RANS-based turbulence model Wilcox $k - \omega$. In the current work, also a two-equation turbulence model Menter SST is employed by both processes. The most important differences in optimisation setups are reconsidered in Section 4.3-1.

The results presented in the work from Reimer [23] are as follows. All optimisations have outperformed the single-fidelity reference optimisation. The best multi-fidelity optimisation outperforms the single-fidelity reference optimisation with a speed-up of 50% and the worst with a speed-up of 13%. The best multi-fidelity optimisation employs the variance decision function¹. This multi-fidelity optimisation has led to an efficiency improvement of 0.75% with respect to the initial design. This is an excess of 0.35% in efficiency on top of 0.40% of efficiency improvement using single-fidelity optimisation. The cost ratio between low-fidelity and high-fidelity processes is equal to $C_r = 0.11$. Information about the correlation coefficient is not available from this work. The replacement ratio is $f_r = 0.42$. In Section 5.3-2, these values are compared to optimisations considered in this thesis. For this work, aerodynamic performance of the optimum blade from multi-fidelity optimisation is only compared to performance of the initial blade design. Hence, no information is available on the comparison of multi-fidelity and single-fidelity optimum designs.

Table 3.2: Differences between high-fidelity and low-fidelity process in the work of Reimer [23]

Parameter	High-fidelity model	Low-fidelity model
Total mesh size [cells]	1, 100, 000	136, 000
Mesh per row [cells]	550, 000	68, 000
Mesh ratio [-]	1	0.12
Radial streamlines [#]	≈ 80	≈ 20
Tip gap resolution	(Enabled)	(Disabled)
Turbulence model	Wilcox $k - \omega$ (RANS)	Wilcox $k - \omega$ (RANS)

Comment: Although exact resolutions were not specifically stated, from figures in the thesis of Reimer [23] it could be observed that tip gaps were enabled in the high-fidelity model and disabled for the low-fidelity model.

Brooks et al. - Multi-fidelity optimisation of a transonic compressor rotor

The second reference work is a research paper by Brooks et al. [3], presenting the implementation of a multi-fidelity method for aerodynamic optimisation of a transonic compressor rotor. The aim of this paper is to demonstrate budget savings for optimisation, by comparing performance of the new setup with a single-fidelity optimisation. Both multi-fidelity and single-fidelity setups use the same computational budget. Many low-fidelity function evaluations on a coarse CFD mesh are combined with few high-fidelity function evaluations on a high resolution mesh, in order to improve accuracy of the Co-Kriging surrogate model. The optimisation objective is to maximise the isentropic efficiency of the rotor, while constrained by the pressure ratio and mass flow. A total of 28 design parameters is used to parametrise the blade geometry of the rotor. The process chain for member evaluation only contains CFD simulation. No information is provided about the convergence criterion employed to terminate optimisations. Additional data on the high-fidelity and low-fidelity models is presented below in Table 3.3. From this table, one could observe that a mesh ratio is used of 0.32. This is a 10% smaller reduction than applied to the best low-fidelity model in the current work. The number of radial streamlines is reduced by almost 70%, which is 16% more than the reduction applied to the best low-fidelity model of the current work.

¹The decision function type is called ‘variance without POI’. AutoOpti normally uses POI to favour high-fidelity members as explained in Section 3.5. However in the work of Reimer [23], 15% of all members were high-fidelity, and 85% were low-fidelity.

Furthermore, the tip vortex is resolved in both high-fidelity and low-fidelity model by a similar number of 6 streamlines. The tip vortex is resolved by 9 radial streamlines in the high-fidelity model (see Section 4.3-1) and 3 in the most suitable low-fidelity model. At last, both process chains employ the same RANS-based turbulence model from Spalart-Allmaras, while this work employed Menter SST turbulence equations. The impact of differences in settings on optimisation performance is discussed in Section 5.2-1.

The results presented in this paper are as follows. An efficiency improvement of 2.34% with respect to the initial design is gained using the multi-fidelity method. This is an excess of 0.55% in efficiency on top of 1.79% of efficiency improvement using single-fidelity optimisation. The cost ratio between low-fidelity and high-fidelity processes is equal to $C_r = 0.33$ and the correlation coefficient is equal to $r^2 = 0.90$. The replacement ratio is $f_r = 0.50$. In Section 5.3-2, these values are reconsidered.

The optimised blade designs from multi-fidelity and single-fidelity optimisations both show similarities in three-dimensional shape (blade leaning), and differences in shape (blade stagger). The optimum design from multi-fidelity optimisation improves the efficiency at a larger radial span (from 10% to 80% span), compared to the optimum design from single-fidelity optimisation (from 30% to 80% span). Hence it can be concluded that the geometrical shape and the aerodynamic performance of these optimum designs are different. From this conclusion in the work of Brooks et al. [3], no conclusions drawn on the search method multi-fidelity method. This fact is reconsidered in Chapter 6.

Table 3.3: Differences between high-fidelity and low-fidelity process in the work of Brooks et al. [3]

Parameter	High-fidelity model	Low-fidelity model
Mesh size [cells]	740,000	240,000
Mesh per row [cells]	740,000	240,000
Mesh ratio [-]	1	0.32
Radial streamlines [#]	≈ 70	≈ 22
Tip gap resolution	6 radial streamlines	6 radial streamlines
Turbulence model	Spalart-Allmaras (RANS)	Spalart-Allmaras (RANS)

Concluding statement

This chapter has provided an introduction to the fundamental concepts of shape optimisation, applied to axial compressors. Essential elements of single-fidelity and multi-fidelity optimisations have been presented. Most notably, the multi-fidelity architecture employs a second evaluation process with lower fidelity. The main conclusion of this chapter is as follows. Three parameters, so-called potential metrics, determine the benefit of a multi-fidelity optimisation: the cost ratio C_r , the correlation coefficient r^2 and the replacement ratio f_r . These parameters are strongly affected by the amount of fidelity reduction applied to the low-fidelity process. In the next chapter, it is described how suitable low-fidelity processes are selected for test optimisations.

4 Methodology for shape optimisation

The multi-fidelity technique is applied to an optimisation geometry, comprising of two blade stages at the front-end of an axial compressor. The benefits offered by the multi-fidelity method are assessed with respect to a single-fidelity reference optimisation that considers the same geometrical problem. This chapter describes the optimisation problem, the setup of both optimisation types as well as the convergence criteria for termination and performance metrics for comparison of optimisations.

In Section 4.1, the optimisation problem is specified. Consequently, Section 4.2 gives an overview of optimisation settings selected in AutoOpti. Setups of both single-fidelity and multi-fidelity optimisations are explained in Section 4.3. The selection of low-fidelity models for multi-fidelity optimisations is motivated in Section 4.4. Optimisations are terminated by employing a convergence criterion, explained in Section 4.5. This chapter is closed by an overview of performance metrics for comparing two optimisations, given in Section 4.6.

4.1 Problem specification

The single- and multi-fidelity optimisation techniques both aim to maximise the same optimisation objectives. A region of interest is defined around the initial design in the design solution space to guide the optimisation. The region of interest is specified by means of a range on the design variables at the input side and limits on operating parameters at the output side. In this section, the optimisation problem is explained by considering the objectives, restrictions, constraints, geometry and parametrisation.

Objectives

In search of an optimum, candidate blade designs are evaluated on the basis of two performance parameters, so-called optimisation objectives. The first objective is the isentropic efficiency in the design point at 100% operating speed, which is from now called the “ISO” operating point. The second objective is the isentropic efficiency in a point of operation where the ambient temperature is slightly higher. This operating point is denoted by “HOT”. As can be seen in the compressor map of Figure 4.1, the two operating points both lie on one single equilibrium running line, but are positioned along two different speed lines. That is because for the HOT operating point, the normalised rotational velocity of the machine is lower due to an increase in temperature. This relation was previously described by Equation 2.3.

Restrictions

In AutoOpti, a user can steer optimisation results and objectives towards a desired interval, through specification of certain intervals that define a *Region of Interest* (ROI). The ROI of optimisation objectives is determined by lower and upper limits. The ROI employed for objective functions in this work is given by Equation 4.1. This equation indicates an optimum design is desired which achieves an efficiency in the ISO and HOT operating points of at least 90%.

$$90\% < \eta_{is} < 100\% \quad (4.1)$$

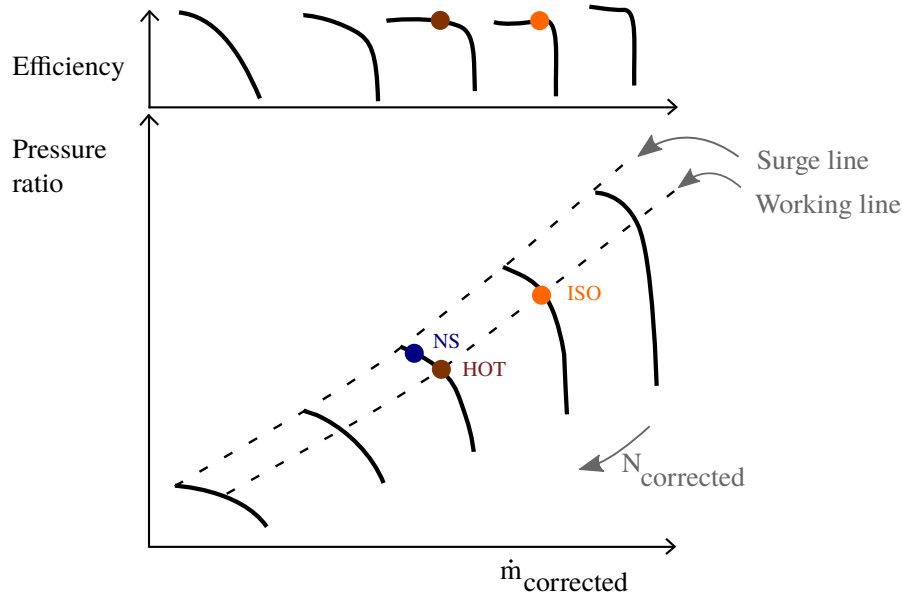


Figure 4.1: Compressor map with three operating points for analysis

Secondly, a Region of Interest can be specified on performance parameters by definition of *restrictions*. Restrictions constrain the design search space and serve as a secondary means to steer the optimisation and separate satisfactory members from inadequate designs. In this thesis, two restrictions are applied. The first restriction considers aerodynamic performance in the near stall operating point, that marks the onset of compressor surge. In Figure 4.1, the near stall operating point is denoted by “NS” and lies on the same speed line as the HOT operating point. The first restriction applies to the pressure ratio in NS (Π_{NS}) and is expressed as a range around the pressure ratio achieved by the initial blade design ($\Pi_{NS,init}$), as shown by Equation 4.2.

$$0.99 \cdot \Pi_{NS,init} < \Pi_{NS} < 1.01 \cdot \Pi_{NS,init} \quad (4.2)$$

The second restriction constrains the mass flow (\dot{m}) that passes through the channel. This restriction considers the true mass-flow instead of the non-dimensionalised mass flow. The restriction is defined as such, that the mass flow is equal to, or higher than the mass flow effective for the initial design (\dot{m}_{init}). This mass flow rate guarantees a minimal power output of the gas turbine. Equation 4.3 shows the effective limits.

$$\dot{m}_{init} < \dot{m} < \infty \quad (4.3)$$

As long as objective functions are outside the aforementioned intervals and restrictions are not satisfied, AutoOpti calculates the distance towards the restriction intervals. First it is attempted to minimise this distance, before the objective functions are maximised [12].

Geometry

The geometry considered by all optimisations described in this work, consists of two blade stages in the front-end of the axial compressor. This geometry is established using the Siemens Integrated Design program (SID). Taking advantage of rotational symmetry of the flow field, only one blade is included per row, as shown in Figure 4.2. Within this four-blade setup, only the second rotor (highlighted in Figure 4.2) is parametrised and optimised. Surrounding blade stages are included in CFD simulations to capture how a change in $R2$ leads to changes in the flow field for the other blades and vice versa.



Figure 4.2: Meridional view of geometry used for simulation

Parametrisation

In this research work an existing optimisation framework is used. Therefore the parametrisation scheme describing blade designs was readily available and not further adapted. The shape of the second rotor blade is described by 80 design variables, which is 16 design variables (25%) more than in the work of Reimer [23] and 62 variables (185%) more than in the work of Brooks et al. [3]. The blade consists of five separate cross-sectional slices which are stacked in radial direction. Within such a cross-sectional slice, the camberline is defined by three control point positions at chord-wise positions, dividing the blade into three camberline sections. Curvature is introduced by three coefficients within these camberline sections. Blade thickness varies smoothly from the leading edge towards the trailing edge, controlled by two coefficients as well as the maximum thickness. The latter is specified as a function of the blade chord length. The three-dimensional shape of the blade is obtained by stacking the slices in radial direction. This positioning requires centre of gravity (COG) coordinates x_{cog} and y_{cog} and stagger offsets, as shown in Figure 4.3. The design variables are varied over a range limited by a lower and upper value, as specified in AutoOpti input settings.




Figure 4.3: Parametrisation of a the rotor showing radial slices

4.2 Optimisation settings

Single- and multi-fidelity optimisations are controlled in AutoOpti through settings distributed across several input files. This section describes those settings, which have the most impact on optimisation progress. A complete overview of settings cannot be disclosed due to confidentiality reasons.

Optimisation problem

First of all, the objectives, regions of interest and restrictions formulating the optimisation problem are specified in a straightforward manner. Those settings have been explained in Section 4.1.

Optimisation phases

Initialisation duration is defined as the total number of valid members to be evaluated by evolutionary algorithms before enabling the surrogate model. The duration is set to 80 members, which equals the number of design variables described in Section 4.1. Choosing this value allows to vary every design parameter at least once. Members created in the initialisation phase, are derivations of the initial design and will therefore show many similarities. For multi-fidelity optimisations, in the initialisation phase also the percentage of low-fidelity members with respect to the total number of members is specified. To obtain a large share of low-fidelity members during initialisation, the parameter is set to 0.8. The value of this parameter is very comparable to the value used in the work of Reimer [23], in which good multi-fidelity performance was demonstrated.

Furthermore, a termination criterion is specified, expressed as the number of evaluated members after which the optimisation should be stopped. The disadvantage of this criterion is that a user would need to estimate the convergence time. Therefore in this work, optimisations are terminated on the basis of a custom convergence criterion presented in Section 4.6. The setting in AutoOpti is set to a member number which is much higher than can be achieved during the total optimisation time¹.

Generation methods

Thirdly, a weighting distribution is employed to rank the optional generation methods. For optimisations in this work, all four evolutionary operations are given a weight of 1, while member generation using the surrogate model² is given a weight of 8. With this probability distribution, there is a probability of 66% the surrogate model is used for member generation. Therefore the emphasis is put on directional search using the surrogate model, instead of employing random exploration using evolutionary operations.

Multi-fidelity decision function

When implementing the multi-fidelity architecture, the decision function type needs to be specified. For the first round of test optimisations described in this work, the *variance* decision function is selected, which is motivated in Section 3.5.

Surrogate model training

The definition and training behaviour of the surrogate model can be controlled through various settings parameters. In this work, default settings from AutoOpti are applied in general as suggested by Goinis et al. [12], since adaptations to Kriging and Co-Kriging models are not in the scope of this research. However, only the surrogate model training frequency is controlled, as specified through *TrainMM*. The frequency parameter regulates the occurrence of surrogate training, which is expressed as an interval of high-fidelity members for all optimisations in this work. Normally a user would manually decrease the training frequency in later stages of an optimisation. That is done because training efforts increase over time due to a growing member population. However, an automatic ramp option is not yet available in AutoOpti. Hence, the average of such ramp is chosen, which is approximately equal to $TrainMM = 5$.

Process control settings

AutoOpti can perform simulation processes in parallel by distributing member evaluations over several slaves. A user can control the number of slaves during optimisation, using two parameters specified in the file *ProcessControl.input*. First of all, the user can specify the time limit in hours before a new slave is created. Secondly, the user can set the target number of slaves. The number of slaves can be changed dynamically without changing optimisation progress considered for comparison. This is taken into account

¹The total number of members is set to $MAX = 10,000$.

²Member generation using the surrogate model is denoted by *EXPIMPR* in the AutoOpti settings file.

by the definition of optimisation progress applied in this work, which is independent of parallelisation of slave processes. Optimisation progress is defined in Section 4.6.

4.3 Setup of optimisations

The multi-fidelity technique will be tested and compared to a single-fidelity optimisation. In this work, the single-fidelity setup thereby serves as the reference optimisation for benchmarking multi-fidelity performance.

4.3-1 Single-fidelity optimisation

The single-fidelity framework employs only one evaluation process to simulate the aerodynamic flow field at high-fidelity, as previously explained in Section 3.2. In this section, the high-fidelity evaluation process applied in single-fidelity optimisation is considered in detail. Elements that are touched upon in the following, are the high-fidelity process chain, CFD mesh setup and solver settings.

Process chain

In single-fidelity optimisation, every member design is evaluated using a high-fidelity evaluation process. Its process chain is represented in a flow chart in Figure 4.4. Subprocesses used for CFD simulation in TRACE are highlighted in blue, while the mesh creation step in GMC is green. Processes meant to support the aerodynamic simulation in preparation and evaluation are grey. An indication of the time for every subprocess is presented by Table A1 in the Appendix.

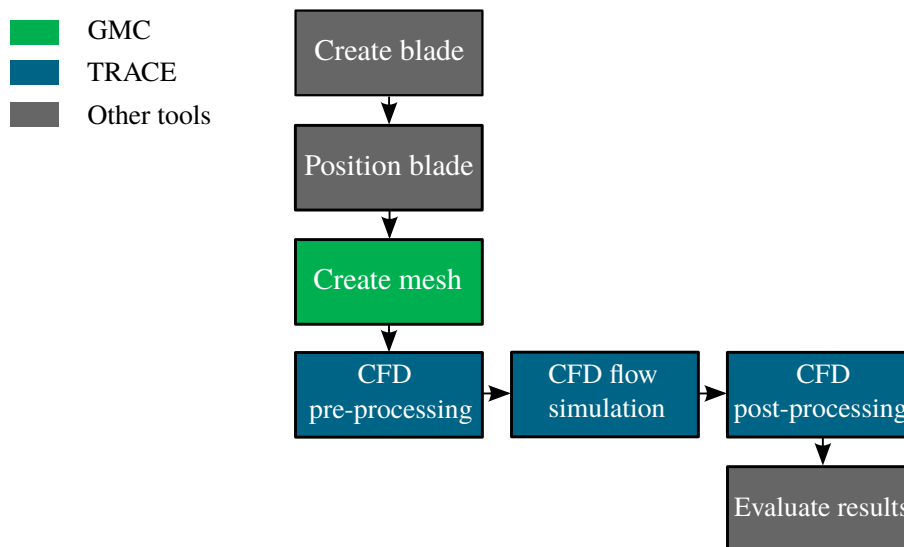


Figure 4.4: Process chain executed for high-fidelity member evaluation

The input for the first step is a design vector. This vector of geometrical design parameters is sent by the master to the slave, which then uses it to generate the blade geometry of the rotor. First the airfoil cross-sections are established after which stacking is carried out, as shown in Section 4.1. Consequently, the 3D blade is positioned in the next step, while subject to geometrical constraints. This assures that the blade fits within a fixed boundary box. Next, the geometry is extended by adding the other three blade rows and later sent for meshing. Boundary conditions and solver options for CFD simulation are specified in *GMC*, a pre-processor developed by *MTU Aero Engines* that is capable of connecting a multi-block

mesh. Pre-processing settings are explained later in this section. Simulation and post-processing steps are performed by flow solver *TRACE*. This aerodynamic solver was developed by DLR, specifically for development and optimisation of turbomachinery components. *TRACE* is particularly useful for multi-row setups, owed to its capability to split the mesh in blocks and solve processes in parallel. The mesh, pre-processing settings and solver setup for high-fidelity simulation are explained in the following subsections.

Mesh setup

A structured mesh is implemented in order to maintain control over mesh complexity and accuracy of resolving flow features, as motivated in Section 2.3. The resolution of the mesh for all four blades is defined by distributions in the S1 and S2 planes, previously shown in Figure 2.2. In the S1 plane, a dense blade-to-blade resolution is specified. Furthermore, in the S2 plane a parabolic distribution of radial streamlines is selected. For the sake of simplicity, these two resolutions are shown in Figure 4.5 for the second rotor blade only. A parabolic radial distribution of cells is chosen such that cell density is highest in blade root and tip regions, for better resolving the end-wall boundary layers. Furthermore, the green highlight at the blade tip indicates the mesh is highly refined by 9 streamlines in the tip region. This mesh refinement allows for resolving tip vortices.

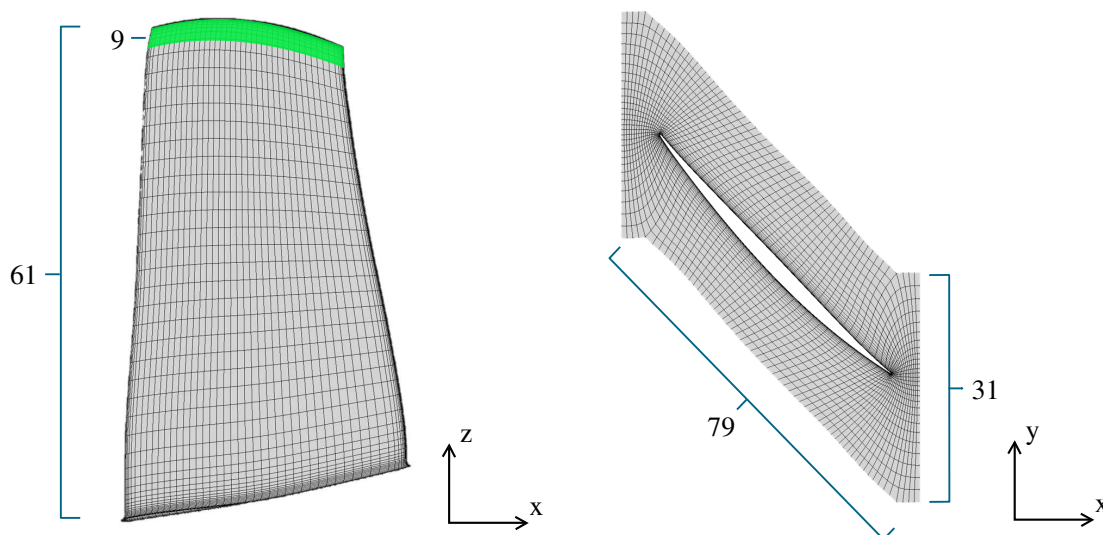


Figure 4.5: Radial and blade-to-blade mesh resolutions for high-fidelity CFD analysis for R2 only. (Note: geometry skewed for confidentiality reasons)

In this investigation, resolving viscous effects within the boundary layers is not of primary interest. Therefore wall-functions are applied to resolve near-wall regions instead of low Reynolds models. A discussion on both of these options was provided in Section 2.3. The dimensionless wall distance y^+ can be increased when using wall-functions, compared to using low-Reynolds models. Figure 4.6 shows the blocking structure applied in the blade-to-blade plane. This structure implements a O-H grid topology. The inner O-grid (indicated in blue) surrounds the blade surfaces with very high resolution for capturing large velocity gradients. The H-grid (indicated in green) aligns the mesh with the flow further away from blade surfaces. Details of the finalised mesh are summarised in Table 4.1.




Figure 4.6: Mesh blocking scheme of the rotor blade for high-fidelity CFD analysis (Note: geometry skewed for confidentiality reasons)

Table 4.1: Properties of high-fidelity CFD mesh

Parameter	Value
Total mesh size, four rows [#]	904,000
Mesh cells per blade row [#]	226,000
Dimensionless wall distance y^+ [-]	24.2

Pre-processing

In GMC, boundary conditions are selected for walls of the flow domain. These conditions determine the type of flow interactions with the blade geometry. The conditions are either inviscid, viscous or moving. Rotor blades are marked as “moving”, while the stator blade surface is labelled “viscous”, introducing a no-slip boundary condition. Moreover, stationary parts of the hub and shroud are given ‘inviscid’ boundary conditions. Between two consecutive blade cascades, mixing planes are introduced where flow is averaged circumferentially. This averaging step is enforced at the mixing plane such that a continuous flow field can be calculated, even though the flow domain switches from stationary to rotating when moving between stators and rotors.

Solver settings

Aerodynamic simulations for evaluating a blade design are run in steady-state. The convergence control settings specify limits to the relative changes in mass flow, efficiency and pressure ratio. If these physical quantities are all within their specified thresholds, an aerodynamic simulation is said to be converged and is automatically terminated in TRACE. For CFD simulations in this work, default settings from GMC are applied, in accordance with the thesis of Reimer [23]. All convergence settings are presented in Table 4.2. The time step is included at which convergence is checked for the first time. Furthermore, the convergence interval is indicated. This is the interval during which convergence criteria need to be fulfilled before starting another new measurement interval.

In TRACE an implicit time stepping scheme is employed to solve the turbulence equations. The turbulence model Menter SST is selected according to best practice in state-of-the-art optimisation cases and offers a good trade-off in accuracy and simulation time. The pseudo-time at which these calculations are performed, differs from the physical time step through the Courant-Friedrichs-Lewy number, as

Table 4.2: Convergence control settings

Convergence criterion	Value
Mass flow	0.01%
Efficiency	0.01%
Pressure ratio	0.01%
Time step at first convergence check [#]	200
Convergence interval for checking [#]	100

defined by Equation 4.4. This number indicates how fast the flow is moving with respect to the time step size Δt and the mesh cell size Δx . The CFL setting is shown alongside with other solver settings in Table 4.3. A ramp setting is used rather than a constant CFL value, because increasing the pseudo-time initially might destabilise the simulation when residual values are still large. Robust convergence is a solver feature provided by TRACE to neglect small scale fluctuations in the flow field and apply robust algorithms at solid walls. This helps to quickly obtain a stable flow solution in steady-state simulation [6]. However, this solver mode needs to be switched off before satisfying all convergence criteria, such that the original flow field is solved. Pre-processing, aerodynamic simulation and post-processing steps are performed for all three operating points.

The initialisation of the flow field is organised by means of a through-flow calculation, providing a two-dimensional preliminary solution which is homogeneous in circumferential direction. The three dimensional flow field that serves as a starting point for simulation is obtained by interpolating this flow domain. The initial flow field is described by velocity components U in axial, circumferential and radial directions, as well as the pressure p and temperature T variations in axial and radial directions.

$$CFL = U^* \frac{\Delta t}{\Delta x} \quad (4.4)$$

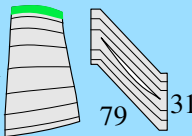
Table 4.3: High-fidelity solver settings

Parameter	Setting
Solver mode	Steady simulation
Turbulence model	Menter SST 2003 $k - \omega$
CFL ramp	1 to 50 in time steps 1 to 500
Robust convergence switch	Time step 600

Overview of high-fidelity process chain

A summary of settings of the high-fidelity process chain is presented in Table 4.4. This model is the only simulation process employed for member evaluation, in single-fidelity optimisations considered in this work. The model is given the name *HIFI* for a clear distinction with models of lower fidelity, which are described in Section 4.4. The radial number of streamlines and the blade-to-blade resolution are presented in a small diagram in column 4 of Table 4.4.

Table 4.4: Details of the high-fidelity process for single-fidelity optimisation

Name	Governing equations	Mesh size	Mesh resolutions (S2,S1)	Description
HIFI	Menter SST	903,000	 61 79 31	Highest accuracy evaluation process

4.3-2 Multi-fidelity optimisation

The multi-fidelity architecture used in this work employs a low-fidelity process besides a high-fidelity evaluation process. This short section outlines the setup of multi-fidelity optimisations which are tested in this thesis.

In this work, five multi-fidelity optimisations are tested in order to verify achievements of the multi-fidelity technique in comparison with single-fidelity optimisation. The high-fidelity process chain is identical to the process in single-fidelity optimisation, which was described in Section 4.3-1. Thereby the meshing, pre-processing and simulation settings are also identical. Whereas all five multi-fidelity setups apply the same high-fidelity process chain, the low-fidelity process chain is varied to study the effects of reducing fidelity on optimisation progress. By reducing the fidelity level of the low-fidelity process, both its accuracy and simulation time are affected. These two elements influence the capability of a low-fidelity process in supporting the multi-fidelity framework. The selection of suitable low-fidelity processes is explained in the next section. At the end of the next section an overview is provided summarising the employed process chains.

4.4 Low-fidelity model selection

In order to use the full potential of multi-fidelity optimisation, a low-fidelity process is required that is significantly faster to evaluate, but still sufficiently accurate compared to the high-fidelity process. In this section, suitable low-fidelity processes are identified and down selected for further investigation within the multi-fidelity framework. Multi-fidelity optimisations considered for testing are presented later in this section. The selection is based on an evaluation of the speed and prediction accuracy of a single CFD simulation with reduced fidelity with respect to a high-fidelity baseline case. In this thesis, this selection is called “single-member analysis”.

This section starts by introducing the scope of fidelity reduction and applicable reduction methods in Subsection 4.4-1. This introduction is followed by the single-member analysis, described in Subsection 4.4-2. In Subsection 4.4-3, five low-fidelity processes are selected for testing the multi-fidelity method.

4.4-1 Identification of fidelity reduction methods

In the following, a definition is provided of fidelity reduction. This definition helps to identify reduction methods suitable for establishing the low-fidelity process chain.

Fidelity reduction

is the act of neglecting certain flow phenomena within the aerodynamic field or resolving them with less accuracy in order to save process time, by which an error is introduced to the objective function.

Settings categories

The low-fidelity process of lower simulation time is obtained by applying fidelity reduction methods to the high-fidelity process. In this work, only fidelity reduction methods within 3D CFD simulation are studied, although normally it is also possible to consider one-dimensional and two-dimensional through-flow simulations to act as low-fidelity processes. The most expensive subprocesses in the high-fidelity process chain are mesh generation and flow simulations, which can be identified from the process times presented in the Appendix. Hence, only two categories of fidelity reduction methods are applied in this work: 1) mesh reduction and 2) changes in solver settings. The two categories are indicated in the high-fidelity process chain in Figure 4.7. Table 4.5 on the next page summarises possible mesh and solver adaptations to the process and their impact on time and accuracy as indicated by ✓ and ✗. An argumentation for the impact on time and accuracy by these adaptations is provided in the next two subsections. From the markings in Table 4.5, one can observe that four settings influence both time and accuracy. These settings are considered fidelity reduction methods because they are in agreement with the definition above. The fidelity reduction methods considered in this work are therefore: reduction in the number of radial streamlines, removing tip gaps, reduction in blade-to-blade mesh resolution and changing to a cheaper turbulence model. These four methods are illustrated in Figure 4.8.

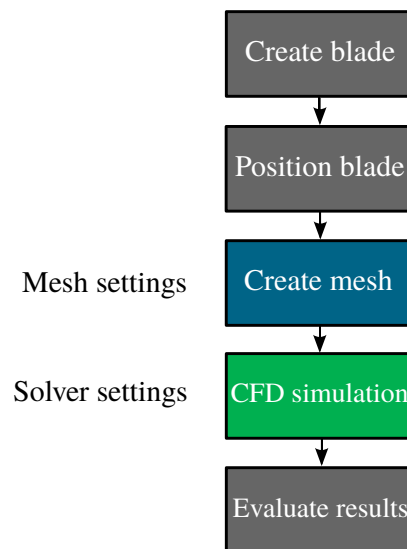


Figure 4.7: High-fidelity process chain with categorised settings




Figure 4.8: Fidelity reduction methods for establishing the low-fidelity process chain: (a) number of radial streamlines, (b) tip gap refinement, (c) blade-to-blade resolution, (d) Turbulence model

Table 4.5: Mesh and solver settings and their effects on simulation time and solution accuracy

Category	Method	Effect on time	Effect on accuracy
Mesh	Reduce b-to-b resolution	✓✓	✓✓
Mesh	Reduce radial streamlines	✓✓	✓✓
Mesh	Remove tip refinement	✓✓	✓✓
Solver	Switch turbulence model	✓✓	✓✓
Solver	Adapt CFL setting	✓	✗
Solver	Adapt convergence criteria	✓	✓

- ✓✓ Large effect
- ✓ Small effect
- ✗ Very little effect

Effects of mesh settings

In Table 4.5 it can be seen that all *mesh settings* influence simulation time and accuracy. Mesh settings determine the density of mesh cells in regions of the flow domain. Therefore the effect on simulation time is evident: reducing mesh resolution reduces the total number of cells and therefore the run time of a CFD simulation. Reducing the mesh resolution in certain areas also affects the local accuracy of the flow solution. In Section 2.3, it is described how the mesh contains local refinement regions to capture physical flow phenomena. By removing mesh refinements in the tip regions, tip leakage flows can no longer be resolved. Analogously, by reducing the radial number streamlines, secondary flows are resolved with less accuracy.

Effects of solver settings

In Table 4.5, one can observe that the first method in the category of *solver settings* affects both time and accuracy. By switching the turbulence model, potentially the simulation time can be reduced. This is expected in particular, when changing from a two-equation to a one-equation model. The accuracy is then influenced depending on how the model handles small-scale flow fluctuations or viscosity.

Other solver settings in Table 4.5 influence time, but do not affect accuracy directly. First of all, the CFL setting is strongly tailored to the Δx value of cells, whose relation was shown by Equation 4.4 in Section 4.3-1. In general, the CFL setting determines the time for simulation but has little influence on the simulation result itself. However, it can impact the numerical stability of the solver and hence lead to no converged solution at all. In the sketch of Figure 4.9, it is shown that using a baseline CFL ramp, the solver residual gradually decreases until the convergence point. Increasing the steepness of this ramp could prevent the simulation to converge. The CFL setting is hence not a fidelity reduction method and all low-fidelity models employ identical CFL settings.

The convergence criterion is the second setting that does not affect accuracy directly. For instance, reducing the mass flow convergence criterion improves accuracy of the solution. The sketch of Figure 4.10 shows two settings of the convergence threshold. While the simulation shown in this figure reaches convergence according to the baseline criterion, using a lower criterion threshold is unattainable for the depicted simulation. As a result, the simulation does not reach convergence. Concluding from this effect, adaptation of the convergence criteria is not a primary fidelity reduction method. This setting is kept constant for all low-fidelity models. The CFL ramp and convergence criteria specified for the low-fidelity processes are identical to those of the high-fidelity process chain, which was presented in Section 4.3-1.




Figure 4.9: Sketch of decreasing solver residual for two CFL ramp settings.




Figure 4.10: Sketch of decreasing solver residual for two convergence criterion settings.

4.4-2 Fidelity reduction study

This section assesses the impact of previously introduced adaptations in mesh and solver settings on simulation time and accuracy. This assessment motivates the choice of low-fidelity processes employed in multi-fidelity optimisations, whose results are described in the next chapter.

A low-fidelity process is established by applying various fidelity reduction methods to a high-fidelity baseline model. The high-fidelity baseline model is the CFD simulation described in Section 4.3-1, which evaluates the isentropic efficiency in the ISO operating point at 100% normalised speed. The performance assessment of a low-fidelity model is based on comparing two CFD simulations of a *single member*: one high-fidelity simulation and one low-fidelity simulation. The results of the single-member analysis are gathered in Table 4.6. In the following, first the parameters in each of the columns in Table 4.6 are explained. Thereafter, the data contained in each of the rows is discussed.

Columns 1 to 7 provide information about the model setup. In column 1, it is shown which fidelity reduction methods are applied to a given model. Here “B2B” means a reduction in blade-to-blade resolution and “RSL” means a reduction in the number of radial streamlines. Column 2 presents the model number. The high-fidelity model has been given the number 0, and every low-fidelity model takes an ascending number. The blade-to-blade resolution is presented in column 5 according to the numbering explained in Table 4.7 and Figure 4.11. Furthermore, column 6 indicates whether refinement in tip regions is enabled or disabled.

In Table 4.6, columns 8 to 11 contain the parameters that evaluate the performance of a low-fidelity model. These parameters are shortly explained next, in order to clarify the final scoring of each model. The time ratio t_{ratio} in column 8 is calculated by applying Equation 4.5. In this equation, subscript $hifi$ denotes the baseline high-fidelity model, and m represents a low-fidelity model index. After evaluating the time ratio for all models, the maximum time ratio $(t_{ratio})_{max}$ can be determined. The normalised time offset Δt_{norm} of column 9 is then determined for every model using Equation 4.6. Analogously, the efficiency ratio η_{ratio} of column 10 is calculated by applying Equation 4.7, while the normalised offset $\Delta \eta_{norm}$ of column 11 is found by Equation 4.8. Finally, the score of each low-fidelity process LF_{score} of column 12 is obtained by multiplying the normalised time and normalised accuracy offsets through Equation 4.9. This score can attain values within the interval $[0, 1]$, where a value close to 1 indicates a large reduction in simulation time is achieved while maintaining a high solution accuracy.

Table 4.6: Ranking low-fidelity models using single-member analysis. The high-fidelity baseline model is highlighted in blue. The most suitable low-fidelity process is highlighted in green.

	1	2	3	4	5	6	7	8	9	10	11	12
	#	Cells	RSL	B2B	Tip gap	Turb. model	t_{ratio}	Δt_{norm}	η_{ratio}	$\Delta \eta_{norm}$	LF_{score}	
	0	931,000	61	79 / 31	Enabled	Menter SST	0.00	0.00	0.00	0.00	0.00	
B2B	1	700,000	61	63 / 27	Enabled	Menter SST	0.38	0.39	0.07	0.96	0.37	
	2	486,000	61	47 / 19	Enabled	Menter SST	0.46	0.47	-0.04	0.98	0.46	
	3	310,000	61	35 / 15	Enabled	Menter SST	0.72	0.73	-0.53	0.70	0.51	
	4	236,000	61	27 / 15	Enabled	Menter SST	0.80	0.82	-1.27	0.27	0.22	
RSL	5	731,000	49	79 / 31	Enabled	Menter SST	0.22	0.22	-0.50	0.71	0.16	
	6	575,000	39	79 / 31	Enabled	Menter SST	0.39	0.40	-0.89	0.49	0.20	
	7	414,000	28	79 / 31	Enabled	Menter SST	0.51	0.52	-0.11	0.94	0.49	
RSL + B2B	8	401,000	28	79 / 31	Disabled	Menter SST	0.51	0.52	0.50	0.71	0.37	
	9	223,000	28	47 / 19	Enabled	Menter SST	0.75	0.77	-0.15	0.91	0.70	
	10	213,000	28	47 / 19	Disabled	Menter SST	0.72	0.73	0.45	0.74	0.54	
	11	205,000	28	35 / 15	Enabled	Menter SST	0.72	0.73	-0.60	0.66	0.48	
	12	132,000	17	47 / 19	Disabled	Menter SST	0.81	0.83	0.31	0.82	0.68	
	13	13,000	6	19 / 15	Disabled	Menter SST	0.98	1.00	-1.75	0.00	0.00	
All	14	931,000	61	79 / 31	Enabled	Wilcox $k - \omega$	0.11	0.11	-0.24	0.86	0.10	
	15	132,000	17	47 / 19	Disabled	Wilcox $k - \omega$	0.83	0.85	0.55	0.69	0.58	
	16	13,000	6	19 / 15	Disabled	Euler	0.98	1.00	-1.65	0.06	0.06	

- High-fidelity reference model
- Best low-fidelity model overall
- Best LF_{score} in category
- Worst LF_{score} overall

Table 4.7: Indicators of blade-to-blade (S1) resolution settings

Resolution indicator [X / Y]	Applied resolution [$x_1 - x_2 - x_3 / y_1 - y_2 - y_3$]
79 / 31	13-53-13 / 13-5-13
63 / 27	9-45-9 / 9-9-9
47 / 19	5-37-5 / 5-5-9
35 / 15	5-25-5 / 5-5-5
27 / 15	5-17-5 / 5-5-5




Figure 4.11: Diagram demonstrating conversion of blade-to-blade resolution

$$t_{ratio} = \frac{t_{hifi} - t_m}{t_{hifi}} \quad (4.5)$$

$$\Delta t_{norm} = \frac{t_{hifi} - t_m}{t_{hifi}} \bigg/ \left(\frac{t_{hifi} - t_m}{t_{hifi}} \right)_{max} = t_{ratio} / (t_{ratio})_{max} \quad (4.6)$$

$$\eta_{ratio} = \frac{\eta_{hifi} - \eta_m}{\eta_{hifi}} \quad (4.7)$$

$$\Delta \eta_{norm} = 1 - \left| \frac{\eta_{hifi} - \eta_m}{\eta_{hifi}} \right| \bigg/ \left| \frac{\eta_{hifi} - \eta}{\eta_{hifi}} \right|_{max} = 1 - \frac{|\eta_{ratio}|}{|(\eta_{ratio})_{max}|} \quad (4.8)$$

$$LF_{score} = \Delta t_{norm} \cdot \Delta \eta_{norm} \quad (4.9)$$

In the following, the scores of low-fidelity models are explained on the basis of the categories of column 1 in Table 4.6.

1. Reducing blade-to-blade resolution

Out of all four fidelity reduction methods described in Section 4.4-1, only the blade-to-blade resolution can be varied solitarily. Therefore models [1 – 4] see a descending blade-to-blade resolution, while all other settings are similar to those from the hifi baseline model. From the first category, model 3 shows the highest score ($LF_{score} = 0.51$), followed by model 2 ($LF_{score} = 0.46$).

2. Reducing radial streamlines

In the second model category, a reduction of the number of radial streamlines is applied in a range from $RSL = 49$ to $RSL = 28$. Further reduction in radial streamlines while using a 79/31 blade-to-blade resolution, results into formation of skewed cells, causing the CFD simulation to fail. Within the second model category, model 7 employing $RSL = 28$ obtains the highest score ($LF_{score} = 0.49$).

3. Reducing both blade-to-blade resolution and radial streamlines

Judging on high scores of models 3 and 7, a combined reduction in blade-to-blade resolution and radial resolution might give a high LF_{score} . Therefore models [8 – 13] implement different combinations of these two reduction methods. The following can be concluded from model scores attained within this category. Since model 8 scores lower than model 7, it is concluded that mesh refinement in tip regions should be kept for higher simulation accuracy of the isentropic efficiency. However, tip gaps cannot be maintained when the radial distribution reduces below $RSL = 28$. The scores of model 9 ($RSL = 28$) and model 12 ($RSL = 17$) are of a similar high values, being $LF_{score} = 0.70$ and $LF_{score} = 0.68$ respectively. These models achieve the best balance of reduced time while maintaining a low offset in efficiency with respect to the high-fidelity baseline model.

It could be observed that high-fidelity baseline model 0 as well as model 13 employing the most severe fidelity reduction both scored $LF_{score} = 0.00$. The reason for this is that the baseline model obtained $\Delta t_{norm} = 0.00$ due to zero time savings, and model 13 attains the highest efficiency offset, automatically resulting in $\Delta \eta_{norm} = 0.00$.

4. Adapting turbulence model

Changing turbulence equations for low-fidelity models could help to decrease the CFD simulation time even further. Hence, the last category of low-fidelity models considers two alternative solver options. In this work, the Wilcox $k - \omega$ model is considered as a first alternative to the Menter SST turbulence model, since it is often used for this type of turbomachinery applications as discussed in Section 2.3. The second alternative is to employ inviscid Euler equations. It is considered because it is very cheap to evaluate and suits simulations using rough meshes.

The settings of low-fidelity models 14 and 15 are similar to the respective models 0 and 12, except for that they employ Wilcox $k - \omega$ turbulence equations. On the basis of their scores, it is concluded that it is beneficial to use Menter SST. The time savings from switching to Wilcox $k - \omega$ are too small to compensate for the increased accuracy offset. This agrees with the expectation of achieving little benefit, since Menter SST and Wilcox $k - \omega$ are both two-equation models.

Instead of employing a turbulence model, one could simulate the flow using Euler equations which assume zero viscosity. In Section 2.3, it is explained that Euler equations are only suitable for low resolution meshes with high y^+ values, since boundary layers are not resolved. Euler equations might therefore cause a large accuracy offset. However, the reduction in accuracy might be compensated by a larger reduction in simulation time. The settings of model 16 are similar to model 13, except for that model 16 employs the Euler equations. The time savings and offset in accuracy for model 16 are similar to those obtained by model 13. Hence the score is also equally low, with $LF_{score} = 0.06$.

One major conclusion is drawn from the previous performance assessment of low-fidelity models. High scores could be attained by models out of different fidelity reduction categories. Hence, results

from this approach show that models with two very distinct levels of fidelity can both achieve a large simulation time reduction, while maintaining sufficient solution accuracy. In the following section, suitable low-fidelity models are selected for multi-fidelity optimisation.

4.4-3 Down-selection of models

The assessment strategy of low-fidelity model performance presented in previous sections, provides preliminary insight to select models for supporting multi-fidelity optimisation. In this section, the down-selection of low-fidelity models for multi-fidelity optimisation is described and motivated. First the expectedly most suitable low-fidelity process is identified and thereafter the selection of low-fidelity processes is motivated.

Best performing low-fidelity process

Out of all models presented in Table 4.6, low-fidelity model 9 attains the highest score of $LF_{score} = 0.70$. This score is based on a large reduction in simulation time of 77% and decrease in efficiency of only 15% with respect to the high-fidelity baseline model.

Lofi processes for multi-fidelity optimisations

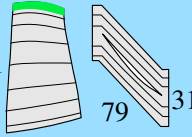
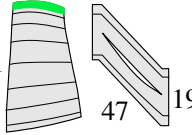
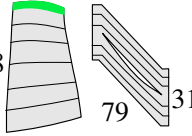
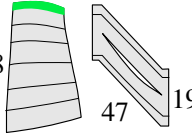
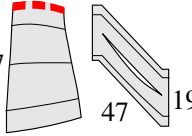
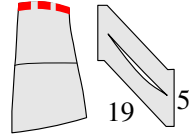
In this work, the influence of the fidelity level of the low-fidelity model is assessed by testing five multi-fidelity optimisation setups with distinct low-fidelity processes. The first low-fidelity process that is employed multi-fidelity optimisation is the best performing low-fidelity model 9. Furthermore, four extra low-fidelity models are selected, preferably from other fidelity reduction categories, previously presented in Table 4.6. This approach results in five contrasting multi-fidelity architectures, which help to answer the research question on the most suitable low-fidelity level for maximum multi-fidelity benefit. Also it assists to verify the applicability of single-member analysis. The verification of this assessment method is described in Section 5.5.

In Table 4.8, an overview is given of the process chains employed by multi-fidelity optimisations considered in this work. All multi-fidelity optimisation employ the same high-fidelity process chain, shown in the top row. The low-fidelity processes that distinguish multi-fidelity optimisations are shown from row 2 and onwards, and ordered by an increasing amount of fidelity reduction. In the following, first the columns are explained, after which the model selection is motivated.

In Table 4.8, column 1 shows the type of process chain. Column 2 presents the model number, previously shown in Table 4.6. Column 3 shows the turbulence model employed during the simulation. Column 4 shows the number of total mesh cells. In column 5, the rotor blade mesh is illustrated in radial and blade-to-blade planes by means of a schematic. Each figure presents three numbers (from left to right): the mesh resolutions in radial, axial and circumferential direction. Column 6 shows what type of fidelity reduction is applied with respect to the high-fidelity baseline model.

The first four low-fidelity models implement Menter SST, since models employing this turbulence model have attained higher scores than models employing Wilcox $k - \omega$. Besides model 9, also models 7 and 12 have obtained high scores in their respective categories. In the category of reducing the blade-to-blade resolution, model 2 is selected over model 3, although model 2 attains a slightly lower score. The reason is that model 2 has the same blade-to-blade resolution as models 9 and 12. The fifth low-fidelity model selected for multi-fidelity optimisation is model 16, which applies inviscid Euler equations. This model attains the lowest score in the entire performance assessment, and calculates efficiency with a large offset. However, it is employed to investigate if the multi-fidelity method can provide a benefit over single-fidelity optimisation, using low-fidelity process results with very little information content.

Table 4.8: High-fidelity process and selected low-fidelity process chains for multi fidelity optimisation

1	2	3	4	5	6
Process chain	Model number	Governing equations	Mesh size	Mesh resolutions (S2,S1)	Adaptation with respect to HIFI
HIFI	0	Menter SST	903,000		No adaptation.
LOFI 1	2	Menter SST	486,000		Reduced B2B.
LOFI 2	7	Menter SST	414,000		Reduced RSL.
LOFI 3	9	Menter SST	223,000		Reduced RSL and B2B.
LOFI 4	12	Menter SST	132,000		Reduced RSL, B2B, tip gap.
LOFI 5	16	Euler	12,800		RSL, B2B, tip gap and solver equations.

— Tip refinement enabled

- - - Tip refinement disabled

Conclusions of assessment method

This chapter has presented a simple approach for a designer to setup multi-fidelity optimisations. The models 9, 7 and 12 have been selected to support multi-fidelity test optimisations, because they attained high scores in their respective fidelity reduction categories. Moreover, model 2 is chosen to fit the fidelity reduction steps of other selected models. Lastly, model 16 is selected with lowest fidelity to test multi-fidelity potential providing low information content.

The method discussed in this chapter is a first attempt to forecast the behaviour of a thousand members simulated during the optimisation, on the basis of only a single member evaluation. At this moment, it is not certain whether a single member evaluation is sufficient to setup a successful multi-fidelity optimisation. In Section 5.5, the applicability of this method is reconsidered by evaluating optimisation results.

4.5 Convergence criterion

In this section, a convergence criterion is defined for optimisation termination. First, it is explained why a convergence criterion is required to terminate optimisations. Thereafter, the implementation of the convergence criterion is discussed.




Figure 4.12: Sketch of a possible speed-up achieved by multi-fidelity optimisations with respect to a single-fidelity optimisation and its dependence on the convergence criterion.

Motivation for convergence criterion application

The following motivates the use of a convergence criterion. Multi-fidelity and single-fidelity optimisations in this work employ the same high-fidelity process to resolve the design solution space. In Figure 4.12, a theoretical progress sketch of both optimisations is shown. In this idealised representation, both optimisations approach the same theoretical limit of improvement at infinite time, indicated by *Time C*. Provided that optimisations do not get stuck in local optima and that two identical global optimum designs do not exist, both optimisations converge towards the same *global* optimum. Obviously, using an infinite optimisation is not practical and optimisations need to be terminated before reaching the theoretical optimum. The duration of an optimisation is determined on the basis of two requirements. The duration needs to be 1) sufficiently long to allow for good optimisation convergence and 2) short enough to fit the optimisation phase within the compressor design schedule. Using the setup for single-fidelity optimisation described in Section 4.3, good convergence is achieved after a run time of 7 days using 10 parallel slave processes. Moreover, this run time is within the available time frame for single-row optimisations. Hence, an optimisation time of 7 days is taken as the basis for defining the convergence criterion in the next paragraph.

In Figure 4.12 it can be observed that at *Time C* there is no speed-up obtained using multi-fidelity optimisation. Considering two earlier instances *Time A* and *Time B*, one could observe that the speed-up at is larger at *Time B*. From this it could be concluded that the moment for comparing multi-fidelity and single-fidelity optimisations influences the measured speed-up by the multi-fidelity method. This dependency is later revisited in the conclusions of this work, to evaluate the selected convergence criterion. Concluding from what is stated above, the convergence criterion has two functions. First of all, it is used to terminate the single-fidelity optimisation at an optimisation time of approximately 7 days and thereby fixes the available computational budget. Secondly, using this criterion the convergence status of multi-fidelity optimisations is checked at the time instance of single-fidelity convergence. This allows to determine how well multi-fidelity optimisations have progressed for the same computational budget.

Criterion implementation

Next, the implementation of the convergence criterion is explained. Section 3.6 has illustrated the variation in CVG over time, caused by optimisation variability. Local variations in progress prevent one to identify a local trend in optimisation progress. The trend of the CVG curve is therefore approximated on a global scale by means of least-squares fitting. An inverse exponential function described by Equation 4.10 is selected for two fundamental characteristics: this function passes through the origin ($CVG_{fit}(0) = 0$) and has zero gradient for large values of x ($\lim_{x \rightarrow \infty} \frac{\delta CVG_{fit}}{\delta x} = 0$), regardless of coefficients a, b and c . For making a least-squares fit, the three shape coefficients a, b and c are optimised using an internal function in Python.

A fit is established for the CVG progress curve as a function of the number of valid high-fidelity members. This approach discards the time contributions by low-fidelity members in multi-fidelity setups and thus enables a proper comparison of multi-fidelity and single-fidelity convergence behaviour. The latter would not be possible when considering CVG improvement over time, since optimisation time also includes the time for low-fidelity member evaluations.

Now the application of the convergence criterion is discussed. The gradient of the fitting function is given by Equation 4.11. In Figure 4.13, it is shown that the member number at convergence, HF_{conv} , is found at the intersection of the gradient of the fit function and the convergence criterion threshold. This threshold is defined by Equation 4.12 and its value is selected as such that the single-fidelity optimisation is terminated at 7 days (as described above). Figure 4.14 demonstrates the CVG obtained at the convergence point, indicated as CVG_{conv} .

$$CVG_{fit}(x) = a \cdot \left(1 - e^{-b(x-c)}\right) \quad (4.10)$$

$$\frac{dCVG}{dx}(x) = a \cdot b \cdot e^{-b(x-c)} \quad (4.11)$$

$$\text{Threshold on CVG gradient: } \left(\frac{dCVG}{dx}\right)_{conv} = 0.25 \cdot 10^{-6} \text{ [1/member]} \quad (4.12)$$

As shown earlier in Figure 4.12, a small improvement in CVG is expected after reaching convergence. By testing two single-fidelity optimisations, it is determined that the CVG value could be increased by 6% when continuing the optimisation for a 20% longer time. This CVG gain amounts to an average efficiency improvement of approximately 0.07% with respect to the average efficiency attained in the convergence point. This efficiency gain is found to be sufficiently small. Hence in this thesis, the CVG gradient threshold of $0.25 \cdot 10^{-6}$ is applied as the convergence criterion.




Figure 4.13: Intersection of convergence limit with gradient of fitting curve, marking the high-fidelity member number at convergence.




Figure 4.14: Sketch of single-fidelity optimisation progress with gradient fit and convergence point.

4.6 Performance metrics

This section explains two performance metrics used in this work to compare actual performance of one optimisation with respect to another. The first metric is the speed-up measured in time. This metric is of primary interest when using the multi-fidelity method, because this method is employed to reduce the computational cost of optimisation and thereby provide savings in optimisation run time. The second metric is an improvement in Cumulative Volume Gain (CVG). The definition of this parameter was previously explained in Section 3.6.

In the first paragraph, the parameter “Theoretical Optimisation Time” (TOT) is introduced. In the second paragraph, the performance metrics are explained.

Optimisation time

In this work, optimisation progress is compared on the basis of the CVG development over time. The time scale considered in this thesis is *Theoretical Optimisation Time* (TOT), which essentially takes the sum of all process times and time spent on surrogate training. The key assumptions for defining this time scale are: 1) all members are evaluated by a single slave, 2) members are created *after* surrogate training is completed and 3) cluster waiting times are not taken into consideration. For this comparison, actual “wall-clock” time is not used, because wall-clock time highly depends on parallelisation of slave processes and any cluster waiting times.

The definition of TOT is given by Equation 4.13, where $t_{process}$ indicates the time to run the process chain for evaluation and $t_{training}$ represents time for surrogate training. In multi-fidelity optimisations, the contribution $t_{process}$ includes all time spent on both high-fidelity and low-fidelity member evaluations. Figure 4.15 on the next page shows a schematic of CVG improvement against TOT. It can be observed that the CVG can only be improved by high-fidelity member evaluations (H), while evaluation of low-fidelity members (L) and surrogate training time (T) merely add time and do not add CVG.

$$TOT = \sum t_{process} + \sum t_{training} \quad (4.13)$$




Figure 4.15: Schematic of CVG improvement over TOT, illustrating contributions of process times and training time.

Application of performance metrics

In this paragraph two modes of comparison are defined to determine the benefits of multi-fidelity optimisations with respect to the single-fidelity baseline optimisation.

Multi-fidelity performance is assessed at the time instant of single-fidelity convergence, according to the criterion described in Section 4.5. Figure 4.16 illustrates the two modes used in this work, for comparing a multi-fidelity optimisation to a single-fidelity optimisation.

The first mode concerns the optimisation speed-up, denoted by $t_{speedup}$. This speed-up considers the ratio of two TOT instances, as shown by Equation 4.14. In this equation, parameter t_{MF} describes the TOT at which a multi-fidelity optimisation reaches the CVG value obtained at single-fidelity convergence. The time instant of single-fidelity optimisation convergence is denoted by t_{SF} . The second mode of comparison considers the potential gain in CVG improvement at the time of single-fidelity convergence, expressed as CVG_{gain} . This parameter is calculated by applying Equatio 4.15.

$$t_{speedup} = (t_{SF} - t_{MF}) / t_{SF} \cdot 100\% \quad (4.14)$$

$$CVG_{gain} = (CVG_{MF} - CVG_{SF}) / CVG_{SF} \cdot 100\% \quad (4.15)$$




Figure 4.16: Modes for comparing multi-fidelity and single-fidelity optimisations

The convergence criterion has previously been motivated in Section 4.5. It is expected that the efficiency gain by the multi-fidelity method will not be significant, whereas the speed-up by multi-fidelity optimisation could be large, as shown at point B in Figure 4.12.

Concluding statement

This chapter has presented the optimisation problem, the setup of the single-fidelity reference optimisation and the setup of five multi-fidelity optimisations. Each multi-fidelity setup employs a low-fidelity process with a distinct fidelity level. Low-fidelity processes are selected using an assessment based on comparing high-fidelity and low-fidelity CFD simulations of one member design. This selection method is evaluated in the next chapter, by considering the performance of the five multi-fidelity test optimisations.

5 Results of the multi-fidelity method

In this chapter, performance results of six multi-fidelity optimisations are compared to the performance of a single-fidelity reference optimisation.

This chapter is structured as follows. First, the single-fidelity reference optimisation is presented in Section 5.1. Then in Section 5.2, an overview is provided of the first five multi-fidelity optimisation results, from which the best and worst multi-fidelity optimisations are identified. Besides evaluating previously introduced potential metrics, also the information transfer from the low-fidelity process is investigated. In order to identify if information transfer could be enhanced, results from a sixth multi-fidelity optimisation are presented in Section 5.3. Conclusions of all results are summarised in Section 5.4. Finally, Section 5.5 provides the reader with guidelines on how to select a suitable low-fidelity process for the multi-fidelity method *before* starting an optimisation.

5.1 Single-fidelity reference optimisation

This section presents the single-fidelity optimisation that is later used as a reference for comparison with multi-fidelity optimisations, presented in Section 5.2.

In this work, two single-fidelity optimisations of identical setup have been performed in order to assess the variability of optimisation progress. The variability is imposed by probability parameters employed in the optimisation algorithm of AutoOpti, whose influence was previously explained in Section 3.4. Due to time constraints, in the current work the number of optimisation repetitions could not be increased to more than two. The progress of the two single-fidelity optimisations referred to as SIFI 1 and SIFI 2 is presented in Figure 5.1. This figure shows the improvement in Cumulative Volume Gain (CVG) over Theoretical Optimisation Time (TOT), as well as the the optimisation convergence points. For more details on the two aforementioned parameters and the convergene criterion, the reader is referred to Section 4.5.




Figure 5.1: Progress diagram of SIFI 1 and SIFI 2 with CVG improvement over TOT.




Figure 5.2: Pareto diagram of the selected single-fidelity reference optimisation, SIFI 1.

Table 5.1: Convergence details and average efficiency gain of single-fidelity optimisations

Optimisation	Members [#]	Slaves [#]	Wall-clock run time [dd:hh]	TOT_{conv} [hrs]	CVG_{conv} ($\cdot 10^{-3}$)	$\Delta\eta_{avg}$ [%]
SIFI 1	3200	3	25 : 21	1666	1.23	1.15
SIFI 2	3220	5	18 : 18	1725	1.15	1.11

It can be observed in Figure 5.1 that the progress curves of SIFI 1 and SIFI 2 show a similar converging trend. In this figure, three regions can be identified of distinct behaviour. From optimisation start to $TOT = 250 \text{ hrs}$, marked as region A, SIFI 1 is faster than SIFI 2. Then from $TOT = 250 \text{ hrs}$ to $TOT = 600 \text{ hrs}$, optimisation progress is identical. In the third region B that spans from $TOT = 600 \text{ hrs}$ to approximately $TOT = 1700 \text{ hrs}$, again SIFI 1 outperforms SIFI 2. Convergence details of SIFI 1 and SIFI 2 are presented in Table 5.1. Concerning the number of evaluated valid members, in column 2, both optimisations have spent almost an identical computational budget. The SIFI 2 has more parallel slave processes, due to beneficial cluster availability compared to SIFI 1, as shown in column 3. The optimisation wall-clock run time of SIFI 2 was therefore 7 hours less, as shown in column 4. Comparing the CVG values at convergence in column 5, SIFI 1 attains $CVG_{conv} = 1.23 \cdot 10^{-6}$ which is $\Delta CVG = 6.5\%$ higher than $CVG_{conv} = 1.15 \cdot 10^{-6}$ attained by SIFI 2. This value is attained at an earlier TOT instance as shown in column 6 of Table 5.1. The average gain in isentropic efficiency by SIFI 1 is 1.15%, which is indicated in column 7.

The Pareto diagram of SIFI 1 is shown in Figure 5.2. From this figure, it can be observed that its front is resolved by a large number of Pareto optimal members. This generally indicates the optimisation has converged well. Although not shown, SIFI 2 has resolved an almost identical Pareto front. SIFI 2 achieves an average efficiency gain of $\Delta\eta_{avg} = 1.11\%$ which is 0.04% lower than SIFI 1. From the observations described above, it could be concluded that SIFI 1 and SIFI 2 progress is overall similar, but has slight differences in the amount of CVG improvement. This is due to progress variability¹, which was previously explained in Section 3.4.

Because of superior performance, SIFI 1 is selected as the single-fidelity reference optimisation for comparison with multi-fidelity optimisations, described next in Section 5.2. The CVG gain or speed-up of a multi-fidelity optimisation are hence measured with respect to SIFI 1 performance, obtaining $CVG_{conv} = 1.23 \cdot 10^{-3}$ and $TOT_{conv} = 1666 \text{ hrs}$ respectively, as shown in Table 5.1. From this point onwards, single-fidelity optimisation SIFI 1 is simply referred to as *SIFI*.

5.2 Multi-fidelity results using variance decision function

In this section, the performance of five multi-fidelity optimisations is presented in terms of CVG improvement and optimisation speed-up. These optimisations employ the same high-fidelity process (described in Section 4.3-1) and employ distinct low-fidelity processes (whose selection is described in Section 4.4-3). These low-fidelity processes differ in fidelity level, whereas the most expensive process uses a fine mesh and Menter SST turbulence equations and the cheapest uses a rough mesh and inviscid Euler equations. The main finding from the results shown in the following, is that benefits from the multi-fidelity method applied to this optimisation problem does not meet the expectations.

In the following, first optimisation results are presented, after which they are discussed and consequently recommendations for improvement are given.

¹Based only on SIFI 1 and SIFI 2 results, single-fidelity optimisations in this thesis are less affected by variability than optimisations in reference works. The difference in CVG at convergence of $\Delta CVG = 0.08$ is much smaller compared to deviations in single-fidelity optimisation shown in the work of Reimer [23], in which four optimisation repetitions were performed.

5.2-1 Performance assessment

An overview of the optimisation setups and results of the five considered multi-fidelity optimisations is given in Table 5.2-1. The following paragraph explains the structure and contents of this table.

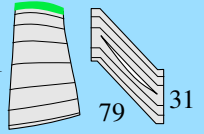
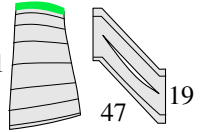
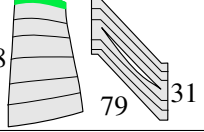

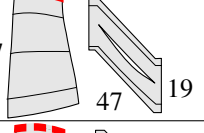
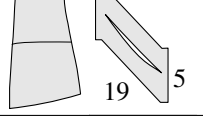
In the top row, the single-fidelity reference optimisation, denoted by SIFI, is included to allow for quick comparison of the employed evaluation processes. Column 1 indicates the optimisation name. For SIFI, columns 2 to 5 present the mesh size, mesh ratio, turbulence model and mesh resolution of the high-fidelity evaluation process. Multi-fidelity optimisations all employ this process for high-fidelity member evaluations.





Optimisations MUFI 1 to MUFI 5 distinguish themselves purely on the basis of the employed low-fidelity process. For these optimisations, columns 2 to 5 present the mesh size, mesh ratio, turbulence model and mesh resolution of the low-fidelity process chain. The motivation for selecting these models was previously presented in Section 4.4. The mesh resolution is presented by a small diagram illustrating the radial and blade-to-blade mesh resolutions. A green tip marker indicates the mesh is refined in tip regions, whereas a red tip marker indicates this refinement is disabled. Columns 6, 7 and 8 present the potential metrics: cost ratio, correlation and replacement ratio introduced previously in Section 3.7. Columns 9 and 10 indicate the initial member offset and the low-fidelity score described in Section 4.4-2. At last, columns 11 and 12 describe the two performance metrics CVG gain speed-up respectively, for which the reader is referred to Section 4.6. CVG gain is defined as the percentage gain in CVG by a multi-fidelity optimisation with respect to SIFI, measured at the TOT time of SIFI convergence. Speed-up is expressed as the percentage reduction in time by a multi-fidelity optimisation with respect to SIFI, measured at the CVG at SIFI convergence. The average efficiency gain, whose calculation was described in Section 3.6, is indicated within brackets in column 11.

In the following, the focus is on the performance metrics (CVG gain and speed-up), whereas values of potential metrics (C_r , r_{avg}^2 and f_r) are discussed later in this section. From column 11 in Table 5.2-1, it could be observed that the highest CVG gain is achieved by MUFI 5, equal to 2.57% on top of SIFI. This CVG gain translates to an efficiency gain of 0.02%, with respect to the efficiency gain of the SIFI optimum design. This gain is small relative to the gain obtained in previous research, as will be discussed in the next paragraph. The lowest CVG gain is achieved by MUFI 1 which underperforms SIFI by 6.03%, and thereby loses an average efficiency improvement of 0.11%. This loss is small relative to the average efficiency gain by SIFI equal to 1.15%. From column 12 in Table 5.2-1, one could observe that the highest speed-up is achieved by MUFI 3, equal to 14.1% measured from SIFI convergence. The lowest speed-up is attained by MUFI 1, which has slowed down by 39.9% with respect to SIFI. This speed-up is further discussed in the next paragraph.

Only based on the values of performance metrics (CVG gain and speed-up) of MUFI 1 to MUFI 5, there is no clear connection between the benefit of the multi-fidelity method and the fidelity level of the low-fidelity process. Because performance metrics alone cannot explain which fidelity level is most suitable, another assessment is required on the basis of potential metrics, as will be discussed in Section 5.2-2. In the next paragraph, the best and the worst multi-fidelity optimisations are examined more closely.

Table 5.2: Overview of multi-fidelity optimisation performance results

1	2	3	4	5	6	7	8	9	10	11	12
Opti. name	Mesh size	Mesh ratio	Turbulence model	Mesh resolution	C_r [-]	r_{avg}^2 [-]	f_r [-]	Initial offset [%]	LF_{score} [-]	CVG gain [%] (η_{avg} gain [%])	Speed-up [%]
SIFI	903,000	1.00	Menter SST		-	-	-	-	-	-	-
MUFI 1	486,000	0.54	Menter SST		0.53	0.90	0.06	-0.10	0.46	-6.03 (-0.11)	-39.9
MUFI 2	414,000	0.46	Menter SST		0.49	0.87	0.05	-0.05	0.49	-1.65 (+0.02)	-29.0
MUFI 3	223,000	0.25	Menter SST		0.29	0.98	0.07	-0.16	0.70	+1.12 (+0.01)	+14.1
MUFI 4	132,000	0.15	Menter SST		0.25	0.86	0.04	+0.55	0.68	-3.61 (-0.05)	-33.4
MUFI 5	12,800	0.01	Euler		0.14	0.46	0.01	-1.86	0.06	+2.57 (+0.02)	+12.0

-  Tip refinement enabled (Column 5)
-  Tip refinement disabled (Column 5)
-  Best result for the performance metric (Columns 11 and 12)
-  Worst result for the performance metric (Columns 11 and 12)

Progress of best and worst optimisations

This paragraph considers optimisation progress of the best and the worst performing multi-fidelity optimisations.

In Section 4.6, it was explained that the speed-up is of primary interest in this work. That is because the convergence criterion (discussed in Section 4.5) is chosen as such that the Pareto front is well-resolved upon optimisation termination. On the basis of the speed-up results in column 12 of Table 5.2-1, MUFI 3 is found to be the best optimisation out of all five multi-fidelity optimisations. Figure 5.3 shows that progress curves of MUFI 3 and SIFI look very similar. The progress curve of MUFI 3 has four clear intersections with SIFI, where it alternately underperforms and outperforms the single-fidelity optimisation. These intersections are at approximate TOT instances 200 hrs, 400 hrs, 800 hrs and 1400 hrs. After the last intersection at $TOT = 1400$ hrs, MUFI 3 outperforms SIFI and attains a speed-up of 14.1%, as indicated in Figure 5.3. In this figure several regions are indicated of high variability, where a sudden change in CVG gradient is present. The cause of progress variability was previously discussed in Section 3.6. Taking into account this variability, as well as the fact that SIFI and MUFI 3 progress curves intersect four times, one could conclude that there is little guarantee a similar speed-up can be achieved when using the setup of MUFI 3 repetitively. Furthermore this speed-up is very dependent on the convergence criterion, whose concern was previously raised in Section 4.5. The speed-up of 14.1% is equivalent to an absolute time savings of approximately 1 day, for an optimisation run time of 7 days. It is 36% lower compared to a speed-up of 50% described in the work of Reimer [23]. Concluding from the effect of variability described above, this time savings is a good result in absolute terms, but should be taken with care.

The CVG gain of MUFI 3 is equal to 1.12%, as shown in Figure 5.3. From this diagram, the multi-fidelity optimisation seemingly small with respect to the single-fidelity optimisation result. In the next paragraph, the gains and losses in CVG are put into perspective by considering the efficiency gain with respect to single-fidelity optimisation.

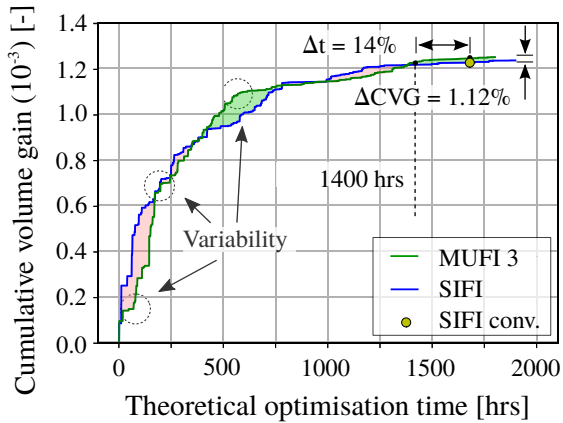


Figure 5.3: Progress diagram of the best multi-fidelity optimisation, MUFI 3, showing performance evaluated with respect to SIFI




Figure 5.4: Progress diagram of the worst multi-fidelity optimisation, MUFI 1, showing performance evaluated with respect to SIFI

In Figure 5.4, it is shown how the worst multi-fidelity case, MUFI 1, underperforms SIFI during two major parts of the optimisation. In region A, from optimisation start until a point of intersection at $TOT = 600$ hrs, MUFI 1 is slower than SIFI. In region B, from $TOT = 600$ hrs and onwards, MUFI 1 again underperforms SIFI. At SIFI convergence, MUFI 1 has a delay of 39.9%. This amounts to an absolute delay of 3 days over a run time of 7 days, which is hence very large. The progress curve in Figure 5.4 only shows one region of a large change in CVG gradient, located at $TOT = 200$ hrs.

Compared to progress of MUF1 3, multi-fidelity optimisation MUF1 1 therefore underperforms SIFI with a higher certainty. The reason for poor performance of MUF1 1 is explained in Section 5.2-2.

Lastly although not shown here, all multi-fidelity optimisations (MUF1 1 to MUF1 5) reach the convergence criterion faster than SIFI. Using the threshold defined in Section 4.5, this implies that multi-fidelity optimisations have approached their theoretical optimum using less high-fidelity members than SIFI. The implication of the number of high-fidelity members to reach the optimum is further explained in Section 5.2-3.

Evaluation of efficiency gains

This paragraph considers the efficiency gain of all optimisations and compares these values to achievements described in reference works.

In Figure 5.5, the Pareto fronts of all multi-fidelity test optimisations and the single-fidelity reference optimisation are presented. For every optimisation, the optimum member is identified closest to the intersections of the respective Pareto front with the diagonal in Figure 5.5.

It can be observed in Figure 5.5 that optimisations MUF1 1 and MUF1 4 underperform SIFI in terms of average efficiency achieved by the optimum member. Although these optimisations underperform SIFI using the current convergence criterion described in Section 4.5, they are expected to achieve the same efficiency gain for a longer optimisation duration. According to the convergence criterion applied in this work, the optimisations are terminated. Also it is shown in Figure 5.5 that Pareto fronts of MUF1 2, MUF1 3 and MUF1 5 overlap with the Pareto front of SIFI. Therefore, their optima have attained very similar gains in average efficiency, as shown by values in between brackets in column 11 of Table 5.2-1.




Figure 5.5: Pareto fronts of optimisations MUF1 1 to MUF1 5 and SIFI, at the time of SIFI convergence

The best multi-fidelity optimisation (MUF1 3) achieves an average efficiency gain of 0.01% with respect to SIFI performance. Optimisations in this work are affected by progress variability, a concept that was introduced in Section 3.4 and illustrated in Section 3.6. Due to variability, there is a considerable chance that repeating the best multi-fidelity optimisation does not result in a positive efficiency gain. Considering the height of the efficiency gain of 0.01%, this value falls within the accuracy limits of CFD simulations employed in this work. Those limits are defined by threshold criteria for convergence control, presented previously in Table 4.2 of Section 4.3-1. The value especially falls short compared to results from works of Reimer [23] and Brooks et al. [3], considered previously in Section 3.8. In these reference works,

the efficiency gain from employing the multi-fidelity method is in the order of 0.40% to 0.50%, which is approximately a factor of 25 higher than the highest efficiency gain of 0.02%, achieved in this work. An important factor that largely impacts the efficiency gain is the choice of convergence criterion for optimisation termination. Terminating the optimisations at a point in time close to *time A* in Figure 4.12, would give the multi-fidelity method a larger percentage gain over single-fidelity optimisation, but reduces the obtained efficiency increase in absolute terms. In Section 5.3-2, a more extensive explanation is given for the difference in benefit of the multi-fidelity method presented in this work and the performance achieved in previous research.

Because the improvement in average efficiency is comparable for the best multi-fidelity optimisation and single-fidelity reference optimisation, the question is raised whether the multi-fidelity algorithm has led to a similar or different blade design. This investigation is described in Chapter 6.

Observations from performance assessment

From the performance metrics presented in the previous paragraphs, below the three main observations are summarised:

1. In the worst case, employing the multi-fidelity method has caused an increase in run time of 39.9%, equivalent to a delay of almost 3 days. The highest obtained speed-up is equal to 14.1%, equivalent to a time savings of 1 day in run time. One should take the speed-up with care, due to uncertainty imposed by progress variability. It is not guaranteed a speed-up is achieved when running the best optimisation setup twice.
2. In the worst case, a decrease in CVG of 6.03 is obtained, equivalent to a decrease in average optimum efficiency of -0.11% . The largest gain in CVG is equal to $+2.57\%$, equivalent to an average efficiency improvement of 0.02%. This improvement is considered low relative to the accuracy of aerodynamic simulation employed in this work, as well as with respect to the achievements of the multi-fidelity method described in reference works.
3. Findings on performance metrics (speed-up and CVG gain) do not show a clear connection between the benefit of the multi-fidelity method and the fidelity level of the low-fidelity process. Therefore another assessment based on potential metrics is used, as discussed in the next section.

5.2-2 Potential assessment

In this section, performance of all five multi-fidelity test optimisations is evaluated on the basis of potential metrics, previously introduced in Section 3.7.

The first paragraph evaluates optimisation potential to explain why the multi-fidelity technique only gives relatively small benefits, when applied to the current optimisation problem. The second paragraph suggests a strategy to improve multi-fidelity potential. In the third paragraph, this strategy is assessed using an adapted optimisation test setup.

Evaluation of potential metrics

The potential of multi-fidelity optimisations is assessed by means of three potential metrics: cost ratio C_r , correlation coefficient r^2 and replacement ratio f_r . In Section 3.7, four guidelines developed by Toal [29] were explained which serve to assess the potential of a multi-fidelity optimisation setup. The guidelines are listed below as A,B,C and D. In Table 5.3, the guidelines are evaluated for all five multi-fidelity optimisations.

Guidelines for potential estimation (presented in Section 3.7)

- A. The correlation between the low and high fidelity function should be reasonably high, $r^2 > 0.9$.
- B. More than 10% of the total evaluation budget should be converted to cheap evaluations, $f_r > 0.1$.
- C. No more than 80% of the total evaluation budget should be converted to cheap evaluations, $f_r < 0.8$.
- D. There should always be more cheap data points than expensive, according to $f_r > 1.75/(1 + \frac{1}{C_r})$

Table 5.3: Evaluation of guidelines for multi-fidelity potential assessment

Opti. name	C_r	r^2_{avg}	f_r	A	B	C	D
MUFI 1	0.53	0.90	0.06	■	■	■	■
MUFI 2	0.49	0.87	0.05	■	■	■	■
MUFI 3	0.29	0.98	0.07	■	■	■	■
MUFI 4	0.25	0.86	0.04	■	■	■	■
MUFI 5	0.14	0.46	0.01	■	■	■	■

■ Satisfaction of guideline

■ Violation of guideline

In the following, the evaluation results in Table 5.3 are discussed. Optimisations MUFI 1 to MUFI 5 indicated in column 1, employ low-fidelity processes with a decreasing fidelity level of the low-fidelity process. As shown in column 2, the cost ratio C_r reduces along with the fidelity level, as intended. For optimisations MUFI 1, MUFI 3 and MUFI 5, the calculation of the cost ratio is presented per subprocess in the Appendix. The average correlation coefficient is given in column 3. With reducing fidelity of the low-fidelity process, the correlation does not decrease monotonically as would be expected². However, the correlation coefficient is generally similar for models MUFI 1 to MUFI 4, with approximately $r^2 = 0.90$ and much lower for MUFI 5, with $r^2 = 0.46$. The replacement ratio is presented in column 4. From all five multi-fidelity optimisations, MUFI 3 has the highest replacement ratio of $f_r = 0.07$ (fraction of budget allocated to low-fidelity member evaluations). MUFI 5 allocates the least amount of budget to low-fidelity members with $f_r = 0.01$.

When considering the evaluation of guidelines A to D in columns 5 to 8 respectively, the following is observed. Guideline A on the minimum correlation coefficient, $r^2 = 0.90$, is only satisfied by MUFI 1 and MUFI 3. All multi-fidelity optimisations violate guideline B on the recommended minimum value of the replacement ratio, $f_r = 0.10$. All multi-fidelity optimisations satisfy guideline C on the recommended maximum value of the replacement ratio, $f_r = 0.80$. At last, all optimisations violate guideline D, which considers the cost ratio and replacement ratio.

The main conclusions from these observations are as follows:

1. In case of MUFI 3, the variance decision function has allocated the largest share of computational budget to low-fidelity members, equal to $f_r = 0.07$. From this result it is concluded that the low-fidelity process employed by MUFI 3 is the most suitable for applying the multi-fidelity method to this optimisation problem, compared to other low-fidelity processes. This process employs a mesh of 223,000 cells, which is a mesh ratio of 4 compared to the high-fidelity process and employs the same Menter SST turbulence model.

²The correlation coefficient increases between MUFI 2 and MUFI 3, while fidelity of the low-fidelity process is reduced by only reducing the blade-to-blade mesh resolution. In Section 5.5, this increase in correlation coefficient is further discussed.

2. However, all five multi-fidelity optimisations (MUFI 1 to MUFI 5) evaluate too little low-fidelity members. The low-fidelity budget is at least 3% lower than the recommended bottom limit. For all five optimisations, the variance decision function has more often selected the high-fidelity process over the low-fidelity process. The cause of this problem is explained in the next paragraph.

Improving the fidelity gap

In this paragraph, the main shortcoming of the current optimisation setup is illustrated. In Figure 5.6, the correlation coefficient and cost ratio from Table 5.3 are plotted for all five multi-fidelity optimisations (MUFI 1 to MUFI 5). Cost ratio and correlation combinations leading to beneficial multi-fidelity performance are highlighted by region X. The horizontal upper limit to region X is $C_r = 0.11$. This cost ratio is derived from the work of Reimer [23], where it resulted in a high performing multi-fidelity setup (among which was a speed-up of 50%, mentioned earlier in Section 3.8). The vertical lower limit to region X is $r^2 = 0.90$. This correlation coefficient is recommended by Toal [29]. Figure 5.6 shows that all optimisations (MUFI 1 to MUFI 5) do not fall within region X of high potential. It is observed that when reducing fidelity of the low-fidelity process to reduce the cost ratio, the correlation coefficient sinks too fast. The main conclusion from this observation can be formulated in two ways: 1) the high-fidelity process currently used is of a too low fidelity or 2) the low-fidelity processes lose too much in correlation, when subject to the current fidelity reduction methods (mesh and solver adaptations in 3D CFD).

Figure 5.7 again shows a diagram with identical correlation coefficients for all optimisations (MUFI 1 to MUFI 5), which use the same low-fidelity processes as in the first scenario. However in this *hypothetical* scenario, the cost ratio is calculated with respect to a different high-fidelity process *HIFI B*. The cost of this high-fidelity process is theoretically increased by a factor of two compared to high-fidelity process *HIFI A*. This implies that the range for fidelity reduction, in this work called the “fidelity gap”, is increased between high-fidelity and low-fidelity processes. Because all multi-fidelity optimisations (MUFI 1 to MUFI 5) move towards lower cost ratios and presumably maintain similar correlation coefficients, the chances are higher they lie within region X. Hence, one can better exploit the benefit from multi-fidelity optimisation if the high-fidelity process is more expensive. Comparing diagrams shown in Figure 5.6 and Figure 5.7, it is concluded the fidelity gap needs to be increased to gain more benefit from the multi-fidelity method. A more suitable simulation problem, would for instance consider a compressor model with resolved shocks and real geometry features such as penny gaps. These require a very expensive mesh, hence a simulation model of higher fidelity.




Figure 5.6: Cost correlation diagram presenting *current* performance of multi-fidelity optimisations, with HIFI A as reference.




Figure 5.7: Cost correlation diagram presenting a fictitious *improved* performance of multi-fidelity optimisation, with HIFI B as reference.

The main conclusion from the previously explained potential assessment is as follows. The optimisation problem considered in this work can have *little benefit* from the multi-fidelity method, because the high-fidelity process is already at a low fidelity. The cost ratio cannot be sufficiently reduced using the current fidelity reduction methods in 3D CFD, without considerably reducing the correlation coefficient. The main recommendation is therefore, to apply multi-fidelity optimisation to a different optimisation problem with a greater fidelity gap. The new problem has either one of following attributes: 1) a more expensive high-fidelity process, or 2) a very cheap low-fidelity process. These processes have a large difference in cost ($C_r < 0.11$) and a high correlation ($r^2 > 0.90$), as previously derived from reference works.

The consequence of the lack in fidelity gap in the current problem, is that an optimiser employing the automatic decision function decides for too few low-fidelity members. In the next section, the effects on optimisation progress are examined from using a larger share of low-fidelity members. However, the last paragraph of this section first explains why a larger share of low-fidelity members could be useful.

5.2-3 Investigation of information transfer

Previously, optimisations MUFI 3 and MUFI 5 have shown similar benefit with respect to the single-fidelity baseline optimisation. This section aims to find differences between these optimisations in the implementation of the multi-fidelity method. In the first paragraph of this section, it is discussed how these optimisations differ in the amount of information transferred by the low-fidelity process. In the second paragraph, a recommendation is given to enhance information transfer and thereby improve the current multi-fidelity setup.

Information transfer in the current setup

The following investigates how the benefit of the multi-fidelity method differs between MUFI 3 and MUFI 5, by considering the number of low-fidelity members. In the end of the previous section, it was concluded that all multi-fidelity optimisations MUFI 1 to MUFI 5 have spent too little budget on low-fidelity member evaluations. For these optimisations, the total number of low-fidelity member evaluations is primarily determined by the variance decision function, whose method is explained in Section 3.5. Optimisation MUFI 3 has evaluated the highest fraction of low-fidelity members ($f_r = 0.07$) from all test optimisations, as was previously shown in Table 5.3. MUFI 5 on the contrary, has the lowest fraction of low-fidelity members ($f_r = 0.01$). Judging on the budget spendings by the variance decision function, the low-fidelity process of MUFI 3 seems more useful than the low-fidelity process of MUFI 5.

Next, on the basis of three observations it will be argued that the low-fidelity process of MUFI 3 indeed is more suitable for applying the multi-fidelity method, through the amount of useful information transferred to the high-fidelity process.

First of all, Figures 5.8 and 5.9 show the progress curves of respective optimisations MUFI 3 and MUFI 5 as the improvement in CVG over high-fidelity members. In Figure 5.8, it can be observed that MUFI 3 generally shows a similar converging trend as shown by SIFI. In Figure 5.8, there is one distinctive region indicated in green. In this region between 500 and 1500 members, MUFI 3 attains at most 17% higher CVG values compared to SIFI. This indicates that MUFI 3 has an advantage over SIFI, owing to the information intermediately provided by low-fidelity evaluations. In this thesis, the information supplied by low-fidelity members is called *information transfer*. In Figure 5.9, it can be observed that MUFI 5 also has a similar shape as the progress curve of SIFI. However, MUFI 5 attains lower CVG values between 1200 and 2800 members, by at most $\Delta CVG = -7\%$. This region is highlighted in red in Figure 5.9. It is concluded that comparing MUFI 3 and MUFI 5 progress in CVG, MUFI 3 has a large benefit from information transfer, while MUFI 5 has close to none.

Secondly, the idea that MUFI 3 has better information transfer compared to MUFI 5, is supported by considering the Pareto diagrams of these optimisations. The Pareto diagram of MUFI 3, showing high-fidelity and low-fidelity members, is shown in Figure 5.10. In this figure, one can observe that many low-fidelity members are evaluated and that the low-fidelity Pareto front has a shape which is similar to the high-fidelity front. It is located at lower objective function values, which is in agreement with the offset between initial members equal to $\Delta\eta = -0.16$, indicated in Figure 5.10 and mentioned in column 9 of Table 5.2-1. The Pareto diagram of MUFI 5 is presented in Figure 5.11. When considering the low-fidelity members in the Pareto diagram of MUFI 5, it can be seen that they do not resolve a Pareto front as clearly as low-fidelity members of MUFI 3. Also the initial member offset between high-fidelity and low-fidelity processes is equal to $\Delta\eta = -1.86$, which is almost 12 times larger compared to MUFI 3. This large offset is in agreement with expectations described in Section 2.3, which stated that employing Euler equations can introduce a large discrepancy in accuracy due to the assumption of inviscid flow.

Thirdly, the correlation can be considered. In line with this large accuracy offset, it was previously identified that MUFI 5 has the lowest correlation coefficient of all multi-fidelity optimisations considered so far, which is equal to $r^2 = 0.46$. Its low-fidelity model was selected in Section 4.4, already taking into account a low correlation.




Figure 5.8: Progress diagram of MUFI 3 showing CVG over high-fidelity members, with respect to SIFI.




Figure 5.9: Progress diagram of MUFI 5 showing CVG over high-fidelity members, with respect to SIFI.

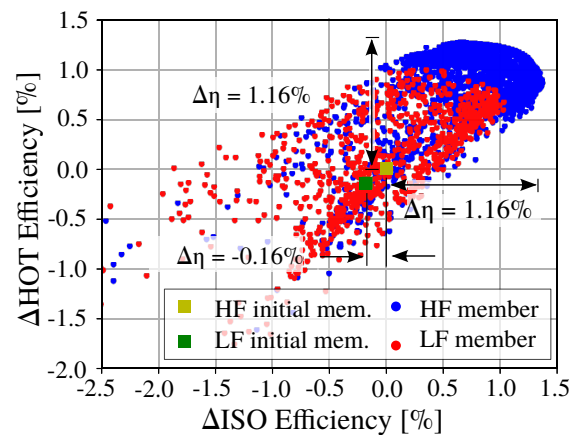


Figure 5.10: Pareto diagram with low-fidelity and high-fidelity members of MUFI 3.




Figure 5.11: Pareto diagram with low-fidelity and high-fidelity members of MUFI 5.

Based on these observations on the CVG progress diagrams, the Pareto diagrams and the correlations, it is concluded that low-fidelity members evaluated in MUF1 3 have informed high-fidelity member evaluations better in MUF1 5. In the next section, results from an adapted optimisation setup are discussed for which the amount of information transfer is enhanced.

5.3 Multi-fidelity results using constant decision function

In this section, results from an adapted optimisation (MUF1 6) are analysed, which employs a constant ratio of low-fidelity members over high-fidelity members after initialisation. Because in the previous section it was shown that all optimisation setups have spent too little budget on low-fidelity evaluations, optimisation MUF1 6 is meant to demonstrate if multi-fidelity potential can be enhanced by increasing the replacement ratio and thereby increasing the level of information transfer. Optimisation progress of MUF1 6 is analysed in Section 5.3-1. It will be shown that no speed-up is achieved with respect to an equivalent single-fidelity optimisation. Another performance assessment is performed and presented in Section 5.3-2 to determine if MUF1 6 fulfils the potential guidelines.

5.3-1 Performance assessment

In this section, performance of optimisation MUF1 6 is considered. The setup of optimisation MUF1 6 is identical to MUF1 3, except for the type of decision function employed. For MUF1 6, a probability is specified to select the low-fidelity process for member evaluations. This ratio is selected at a constant value of 85%, such that approximately half of the computational budget is spent on low-fidelity member evaluations. By choosing this decision function type for MUF1 6, the share of low-fidelity members approximately stays constant³ at 80% and attains a final value that is 60% higher than MUF1 3 as shown in Figure 5.12.

The progress diagrams of MUF1 6 is shown against valid hifi members in Figure 5.14. In this figure, it can be seen that MUF1 6 has evaluated a total of almost 800 members over its run time. Due to time constraints, this optimisation could not be run until the convergence point of SIFI. In Figure 5.14, the region is highlighted at which MUF1 6 outperforms SIFI and hence is well-informed. It is observed that high-fidelity members achieve at most 31% higher CVG values than SIFI, owing to information transfer from low-fidelity members. The enhancement in CVG by MUF1 6 is 14% larger than achieved by MUF1 3, whose progress was previously shown in Figure 5.8. This implies that the information from low-fidelity members is enhanced by increasing the fraction of low-fidelity members from $f_r = 0.07$ to $f_r = 0.51$, on the condition that processes are well-correlated ($r^2 = 0.98$).

Considering CVG progress over theoretical optimisation time in Figure 5.15, it is observed that MUF1 6 underforms SIFI over the entire optimisation run time, until the point manual termination at $TOT = 1000 \text{ hrs}$. Although MUF1 6 is not run up to SIFI convergence, from the previous observation it is expected that no speed-up will be achieved. The fact that MUF1 6 outperforms SIFI in terms of CVG gain over high-fidelity members, but underperforms in terms of CVG gain over TOT is explained by two competing effects. These are shown in Figure 5.13. First of all, additional low-fidelity evaluations *inform* high-fidelity member evaluations. Secondly, those low-fidelity members *delay the optimisation*, because they take too long to evaluate. Although the correlation of MUF1 6 results is high enough, the cost ratio is not sufficient. In the next section, previously introduced performance metrics are employed to investigate the influence of the fraction of low-fidelity members on expected performance.

³The share of low-fidelity members stays constant at 80% over time, while it was set to 85%. The small offset in low-fidelity member fraction is not investigated in this work, but is assumed to be of minor importance in this investigation.




Figure 5.12: Low-fidelity members over total evaluated members for MUFI 3 and MUFI 6.




Figure 5.13: Schematic of progress curves showing competing effects on mufi benefit.




Figure 5.14: Progress diagram of MUFI 6 showing CVG over high-fidelity members.




Figure 5.15: Progress diagram of MUFI 6 showing CVG over optimisation time.

5.3-2 Potential assessment

The performance assessment metrics of Toal [29] are employed in this section to evaluate the expected performance of multi-fidelity optimisation MUFI 6. These metrics support the comparison of MUFI 6 with optimisations from reference works extensively discussed in Section 3.8. Recommendations for further improvement of the multi-fidelity setup are given at the end of this section.

Guidelines A to D used in this assessment, have been introduced in Section 3.7. In Table 5.4, performance of MUFI 6 is evaluated as well as reference optimisations *Brooks* Brooks et al. [3] and *Reimer*⁴ For convenience sake, MUFI 3 is added to this list which attained best multi-fidelity performance using the variance decision function.

Results in Table 5.4 show that MUFI 6 satisfies *all criteria* and attains similar values of C_r , r^2 and f_r compared to *Brooks*. However, MUFI 6 does not achieve a speed-up nor it is expected to enhance optimisation improvement with respect to single-fidelity optimisation, as concluded in the previous section. Possible explanations for this contradiction are given on the next page.

⁴Optimisation *Reimer* employs the variance decision function and attains the highest speed-up of 50%.

Table 5.4: Evaluation of guidelines for assessment of multi-fidelity potential

Source	Opti. name	r^2_{avg}	C_r	f_r	A	B	C	D
Current thesis	MUFI 3	0.98	0.29	0.07	0.98	0.07	0.07	-0.33
	MUFI 6	0.98	0.24	0.51	0.98	0.51	0.51	+0.15
Brooks et al. [3]	Brooks	0.90	0.33	0.50	0.90	0.50	0.50	+0.06
Reimer [23]	Reimer	N/A	0.11	0.42	N/A	0.42	0.42	+0.25

Comment: In this table, it could be seen that no correlation information was provided by Reimer [23]. It is assumed the correlation of the optimisation in this work was above $r^2 = 0.90$, because the multi-fidelity optimisation achieved a speed-up of 50%.

■ Satisfaction of guideline

■ Violation of guideline

The contradiction that MUFI 6 satisfies all guidelines for potential estimation, without gaining from the multi-fidelity method could be explained by the following:

1. Potential estimation guidelines B and C on f_r define an interval between 10% and 80%. Regarding that these guidelines were developed on the basis of four analytical testing functions, this guideline might not be applicable or too large for the current physical problem. It could be for instance, that the replacement ratio should be at least $f_r = 50\%$ for this optimisation problem, as is the case in the work of Brooks et al. [3].
2. Potential metrics of Toal [29] consider the status of an optimisation at its very end, while the timing of low-fidelity member evaluations is important. For instance, if first only high-fidelity members are evaluated and then only low-fidelity members, benefit of multi-fidelity is low while f_r may seem sufficient.
3. Potential metrics do not take into account the quality of the Co-Kriging surrogate model. In multi-fidelity optimisation, the duration of initialisation and the frequency of training are two important factors of influence. The dataset for initialisation for Reimer [23] is approximately 1.5 times the number of design variables (100 members for initialisation). The dataset in this work is 1.0 times the number of design parameters (80 members for initialisation).
4. Potential metrics do not consider the employed methods for member generation. In this work after initialisation, members are generated using the surrogate model according to a ratio of 2/3. The remaining 1/3 of all members is generated using evolutionary operations. It is not certain whether the random search strategy of EA has provided any benefit to multi-fidelity optimisations in this thesis. In the work of Reimer [23], all members are generated using the Co-Kriging model. By only using the Co-Kriging model for member generation, benefits from the multi-fidelity method can be fully exploited.

The main conclusion of this section is that using a manual ratio of many low-fidelity members for a well-correlated optimisation MUFI 6, a positive effect has been demonstrated on the amount of information transfer. However, a negative effect on optimisation time has been observed. The benefit of the multi-fidelity method could therefore not be enhanced by increasing the replacement ratio.

5.4 Conclusions on multi-fidelity optimisation results

In this section, the main conclusions are summarised on performance of multi-fidelity optimisations considered in this work.

1. *One single-fidelity optimisation (SIFI) is chosen from two test cases of identical setup.* The best single-fidelity optimisation is selected as a reference optimisation for comparison with multi-fidelity optimisations. The selected optimum has an increased average isentropic efficiency of $\Delta\eta = 1.15\%$ with respect to the initial blade design and the run time until convergence was 1666 hours. This time is equivalent to an absolute duration of approximately 7 days.
2. *Five multi-fidelity optimisation setups are tested (MUFI 1 to MUFI 5) and their performance is compared to the single-fidelity reference optimisation.* The worst multi-fidelity optimisation (MUFI 1) has a delay in optimisation run time of 39.9%, equivalent to almost three days of time loss. The selected optimum design found by this optimisation has a loss in average efficiency of 0.11%. The best multi-fidelity optimisation (MUFI 3) has a speed-up of 14.1%, equivalent to 1 day of time savings. The selected optimum design has an improvement in average efficiency of 0.02%. This result is 36% lower compared to the result of 50% in time savings described in reference literature. Although a 14.1% speed-up is useful in absolute time savings, it is low when taking into account progress variability and dependency on the convergence criterion. There is little guarantee this speed-up can be achieved when repeating the same optimisation. If a designer is not careful choosing the fidelity level of the low-fidelity process, an optimisation might be delayed by almost 40%.
3. *The multi-fidelity method is not suitable for improving the current optimisation problem.* The fidelity level of the high-fidelity process is too low, since it uses only 226,000 mesh cells per blade row. The current fidelity reduction methods, consisting of mesh reduction and adaptation of solver settings in 3D CFD, do not allow a sufficient reduction in the cost ratio without a considerable reduction in the correlation coefficient. The best cost ratio is achieved by MUFI 5, equal to $C_r = 0.14$. However, MUFI 5 only attains a very low correlation coefficient of $r^2 = 0.46$. The best correlation coefficient is achieved by MUFI 3, equal to $r^2 = 0.98$. However, MUFI 3 attains a mediocre cost ratio of $C_r = 0.29$. Using the (automatic) variance decision function, at most a fraction of only $f_r = 7\%$ of computational budget is allocated to low-fidelity member evaluations. This fraction is manually increased in the setup of MUFI 6, to a value of $f_r = 51\%$ which is very comparable to the fractions used in reference works. By MUFI 6, the maximum improvement in CVG with respect to SIFI is increased by 14%, owing to enhanced information transfer from low-fidelity members. However, low-fidelity members are too expensive and delay the optimisation in terms of its run time.
4. *Multi-fidelity optimisation would be more suitable to an optimisation problem with a larger fidelity gap.* A more suitable setup either has a very expensive high-fidelity process or it has a very cheap, but well-correlated low-fidelity process. As an example, in Section 3.8 it was shown that the high-fidelity model presented in the work of Reimer [23] has a 59% larger mesh size per row, allowing for a mesh reduction of 0.12 to establish the low-fidelity model. The most successful low-fidelity model in this thesis only achieves a mesh ratio of 0.25, which illustrates a too low fidelity gap. Examples of recommended optimisation setups are given below.

Based on the previous conclusions, it is recommended to investigate the benefits of multi-fidelity optimisation for another problem. First of all, the setup should have a large fidelity gap. One example is to use an expensive high-fidelity process with detailed design features such as penny gaps, while keeping similar low-fidelity processes as considered in this work. Alternatively one could consider the same high-fidelity process, but use through-flow simulation as a low-fidelity model. The cost of the low-fidelity process should be reduced to at most 11% of the high-fidelity process, while the correlation coefficient

should be above 90%. These values were previously derived from the works of Reimer [23] and Toal [29].

Secondly, it is expected using the settings mentioned in the following, the performance of the multi-fidelity method can be increased. An initialisation duration of 1.5 times the number of design variables could be used. Furthermore, only the surrogate model should be selected for member generation after initialisation. Lastly, it is recommended to employ the variance decision function, such that the optimiser can automatically decide on the optimum ratio of low-fidelity members to total members and the timing of low-fidelity member evaluations.

In the next section, a forecasting method is presented which a designer could use to determine the suitability of a low-fidelity model for a new multi-fidelity optimisation setup.

5.5 Multi-fidelity performance forecasting

This is the last section that considers the multi-fidelity optimisation algorithm. The previous three sections of this chapter have considered results of multi-fidelity optimisations tested in this thesis project. On the basis of these results, the current section provides the reader with guidelines to forecast multi-fidelity performance with a certain level of confidence, *before* starting a multi-fidelity optimisation.

In subsection 5.5-1, the method is evaluated, which was used to select the low-fidelity processes for multi-fidelity optimisations. In subsection 5.5-2, an approach is presented on the number of members required to approximate the cost ratio and correlation coefficient applicable to a particular low-fidelity process.

5.5-1 Forecasting using one member

This section verifies the suitability of single-member analysis, which was used to choose low-fidelity processes for multi-fidelity optimisation.

Suitability of a low-fidelity process chain for multi-fidelity optimisation is governed by the cost ratio between the two employed processes and the correlation of their results. In this thesis, low-fidelity processes for multi-fidelity optimisation are selected using a single-member analysis, which was discussed in Section 4.4. This approach evaluates the performance of a low-fidelity process in terms of simulation time and solution accuracy merely on the basis of one member simulation. The cost ratio is estimated by the time spent on one high-fidelity and one low-fidelity CFD simulation. The correlation coefficient is estimated as an offset in efficiency, calculated by the two simulations. The combined effect of time and quality is expressed in a single parameter, called LF_{score} . In the following, it is discussed whether accuracy and time effects can be captured in a single parameter.

Table 5.5: Optimisation performance of multi-fidelity optimisations for verification of single-member analysis.

Opti. name	LF_{score} [-]	C_r [-]	r^2_{avg} [-]
MUFI 3	0.70	0.29	0.98
MUFI 4	0.68	0.25	0.86
MUFI 2	0.49	0.49	0.87
MUFI 1	0.46	0.53	0.90
MUFI 5	0.06	0.14	0.46

- Best forecast score or best performance value
- Worst forecast score or worst performance value

The results of five multi-fidelity optimisation setups (MUFI 1 to MUFI 5) are shown in Table 5.5 on the previous page, ordered by the value of the LF_{score} . These results only include the potential metrics which were previously discussed in Section 5.2-2. Performance metrics (speed-up and CVG gain) are not considered in this discussion, because it was concluded in Section 5.2-1 that no clear connection was observed between the values of performance metrics and the level of fidelity reduction employed.

First, the scores are compared to the cost ratio, shown in column 3 of Table 5.5. MUFI 5 has the best cost ratio, but its low-fidelity process has obtained the lowest LF_{score} . MUFI 1 has the worst cost ratio, while its low-fidelity process has the second worst LF_{score} . It is concluded the cost ratio cannot be estimated easily using the scoring system.

Second, the scores are compared to the correlation coefficient, shown in column 4 of the table. MUFI 3 has the highest correlation coefficient of $r^2 = 0.98$ and its low-fidelity model has obtained the highest LF_{score} . MUFI 5 has the lowest correlation coefficient of $r^2 = 0.46$ and its low-fidelity model has obtained the lowest LF_{score} . Although the scoring system has given insight into a large contrast in correlation, subtle differences in the correlation coefficient for MUFI 1 to MUFI 4, could hardly be estimated.

Concluding from these observations, the scoring system could not capture both the time and quality of the low-fidelity process using one parameter, the LF_{score} . In the next subsection, an alternative method is presented which uses a larger sample set of members, in order to better forecast low-fidelity suitability.

5.5-2 Forecasting using set of members

In the previous subsection it was concluded that single-member analysis is inadequate to analyse small differences in correlation and cost ratio values, for a particular low-fidelity process. Hence in this subsection, a forecasting method is presented which forecasts suitability of a low-fidelity model on the basis of a larger sample set of member evaluations. A designer can use this estimation method to decide on which low-fidelity model to employ for using the multi-fidelity method.

In the following, the average correlation coefficient and cost ratio are calculated from a sample set of members, with a size in the range of [2, 3, 5, 10, 20, 50, 100, ..., 1500, 2000] members. This analysis is performed for three multi-fidelity optimisations specifically because of their distinct correlation values: MUFI 3 (best correlation $r^2 = 0.98$, average cost ratio 0.29), MUFI 5 (worst correlation $r^2 = 0.46$, best cost ratio $C_r = 0.14$) and MUFI 1 (average correlation $r^2 = 0.90$, worst cost ratio $C_r = 0.53$). Performance results of these optimisations were presented in Section 5.2-1. Members are randomly selected from all members evaluated during these optimisations. In order to calculate the correlation coefficient, a random sample set of high-fidelity members is probed into the Co-Kriging surrogate model to obtain their low-fidelity approximations. Estimating the low-fidelity response is possible with high accuracy, because the Co-Kriging model is based on the entire finalised database consisting of > 3000 evaluated members. The cost ratio is calculated on the basis of randomly selected high-fidelity members and randomly selected low-fidelity members. Two repetitions are performed for each sample size, except if indicated otherwise. Results from these repetitions are consequently averaged. In the following two paragraphs, the results of this analysis are discussed.

Forecasting cost ratio

In Figure 5.16 on the next page, the cost ratio error for the three optimisations is shown, calculated with respect to the average cost ratio of the full set of evaluated members.

Two observations are made from the forecasting error of the cost ratio in Figure 5.16:

- The error in cost ratio generally decreases when increasing the sample size of member evaluations.
- The reduction in cost ratio error of optimisations MUFI 3 and MUFI 5 is similar, with a difference of only 0.05%. The cost ratio error of MUFI 1 is 2% higher than for MUFI 3 and MUFI 5. This result is unexpected because the variation in simulation duration (cost) is not expected to be related to the fidelity level of the low-fidelity process. However, a 2% difference in error is considered very small.




Figure 5.16: Forecasting error of cost ratio of three multi-fidelity setups

The main conclusion is that the cost ratio can be estimated with a good accuracy, regardless of the average cost ratio of the considered processes. An estimation can be performed on the basis of a small sample set of 5 members, with an accuracy of 5% or less. This accuracy is sufficient for estimating multi-fidelity potential using guidelines from Toal [29], presented in Section 3.7.

Forecasting correlation coefficient

In Figure 5.17, the correlation error is shown for the three optimisations. In total, five repetitions are performed for sample sets [5, 10, 20, 50] because large deviations were observed for small sample sets of members. For these sample sizes, an uncertainty range around the average error is shown in Figure 5.17, indicating the minimum and maximum error obtained from five repetitions.

Three observations are made from the forecasting error of the correlation coefficient in Figure 5.17:

- The error in correlation coefficient generally decreases when increasing the sample size of member evaluations.
- Forecasting the correlation has a larger error than the cost ratio. While the maximum error calculated for the cost ratio is 7.8%, the maximum error of the correlation coefficient is 98%.
- The higher the correlation coefficient r^2 of an optimisation, the smaller is the error in forecasting its correlation. MUFI 3 ($r^2 = 0.98$) has only 2% correlation error after 50 members, MUFI 1 ($r^2 = 0.90$) has 15% error and MUFI 5 ($r^2 = 0.46$) has 35% error.




Figure 5.17: Forecasting error of correlation coefficient of three multi-fidelity setups with uncertainty ranges

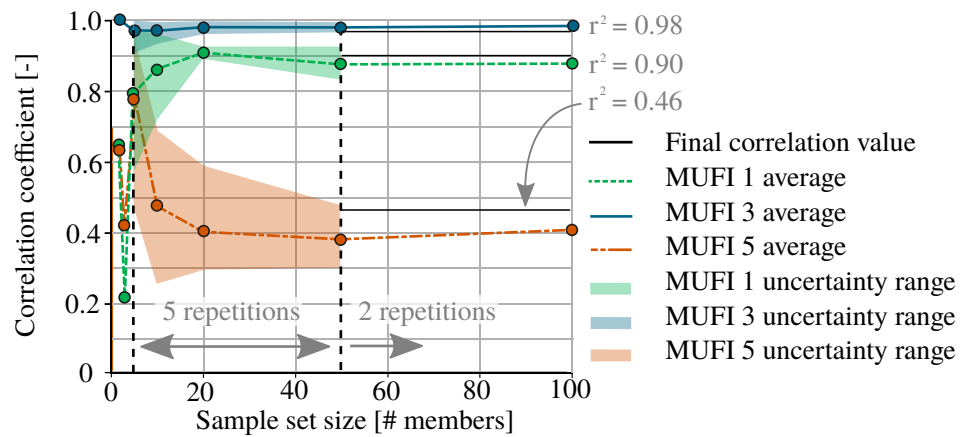


Figure 5.18: Absolute correlation coefficient convergence of three multi-fidelity setups with uncertainty ranges

The main conclusion is that the accuracy at which the correlation can be estimated, depends on the actual correlation of the process results. When the actual correlation is relatively high ($r^2 = 0.98$) an estimation with an accuracy of 5% can be performed on the basis of a sample set of 5 members. When the actual correlation is slightly lower ($r^2 = 0.90$) an estimation with similar accuracy can be performed using a slightly larger sample set of 10 members. When the actual correlation is very low ($r^2 = 0.46$) an estimation accuracy of 5% can only be achieved using a much larger sample set of 200 members.

Because the main goal is to estimate if correlation is above the limit of sufficient correlation $r^2 = 0.90$ for good multi-fidelity performance (according to guideline A defined by Toal [29] in Section 3.7), the convergence trends towards *absolute* correlation coefficients need to be investigated. In Figure 5.18, the absolute convergence trend of the correlation coefficient is shown for the three optimisations (MUFI 1, MUFI 3 and MUFI 5). Again a similar number of repetitions is performed at indicated sample sets to show the uncertainty range of the correlation coefficient. Three observations are made:

- The absolute correlation coefficient approaches but does not reach the final correlation coefficient for all three multi-fidelity optimisations, as indicated in Figure 5.18.

- Using a sample size of 50 members, the uncertainty range of a well-correlated optimisation (MUFI 1) is $\Delta r^2 = 0.03$. This is smaller than the ranges 0.09 and 0.18 of optimisations MUFI 1 and MUFI 5 respectively, having lower correlation values. This corresponds to the percentage error presented in Figure 5.17.
- If the actual correlation value is very high (MUFI 3, $r^2 = 0.98$), the average correlation coefficient is likely to be above the limit of sufficient correlation for all sample sizes. Additionally, if the actual correlation coefficient is low (MUFI 5, $r^2 = 0.46$), the uncertainty range of the estimation is found below the limit of sufficient correlation after 10 member evaluations. Lastly, if the actual correlation value is around the limit of sufficient correlation (MUFI 1, $r^2 = 0.90$), the observed uncertainty range contains the limit itself. For such an optimisation, more members are required to estimate if the correlation value is above the limit of sufficient correlation.

The main conclusion from observations of Figure 5.18, is that in order to estimate the correlation for performing a potential assessment by Toal [29], it is often not necessary to simulate up to 200 high-fidelity members when a model under consideration is of poor correlation. First one should check if the upper limit of the uncertainty range of r^2 is crossing the limit of sufficient correlation. Using this insight, in the next paragraph an approach is presented that guides a designer through the process of selecting a suitable low-fidelity model.

Estimation approach of low-fidelity model suitability

In this paragraph, the steps are discussed for estimating low-fidelity model suitability for multi-fidelity optimisation. In Figure 5.19, the approach is shown by means of a flow chart. As shown in this figure, the first step is to generate a low-fidelity process by applying fidelity reduction methods to a high-fidelity process. Next, 10 members need to be simulated using both high-fidelity and low-fidelity process chains. It should be noted that a Co-Kriging model cannot be used for this analysis, since this has not been established before starting an optimisation. After this step, the cost ratio can be calculated with an error of at most 5%. Consequently, a check needs to be performed concerning the uncertainty range of the correlation coefficient, bounded by the minimum and maximum obtained correlations. If the range overlaps with the limit of sufficient correlation $r^2 = 0.90$, then 10 additional members are simulated and the aforementioned steps are repeated. If there is no overlap of the range with the limit, one needs to check if the average correlation is above the limit of sufficient correlation. If this is not the case, this low-fidelity process is not suitable for multi-fidelity optimisation. However, if a correlation coefficient above 0.90 has been obtained, the low-fidelity process chain under consideration can be used for multi-fidelity optimisation. The error at which the correlation is estimated has been reduced to at most 10%.

The following is concluded from this section. The cost ratio of a new low-fidelity process can be estimated with at most 5% error using only 2 random member evaluations. For the same process, an initial estimate of the correlation coefficient can be made by evaluating 10 random member designs. Depending on the obtained average correlation and uncertainty range, it might be necessary to evaluate 10 additional members. On the basis of the multi-fidelity optimisation with worst correlation $r^2 = 0.46$, it has been demonstrated that 20 members are sufficient to check if the correlation coefficient is higher than the limit of sufficient correlation. This number of members however, is merely an indication and could be higher depending on the simulation process under consideration.

Previously it was explained that the average correlation is calculated by contributions from two restrictions and two objectives for each of which a Co-Kriging model is established. When performing the aforementioned calculation steps, a designer may see large differences in these contributions. Therefore, in the next paragraph the individual Co-Kriging model correlations are considered to investigate their impact on the average correlation.

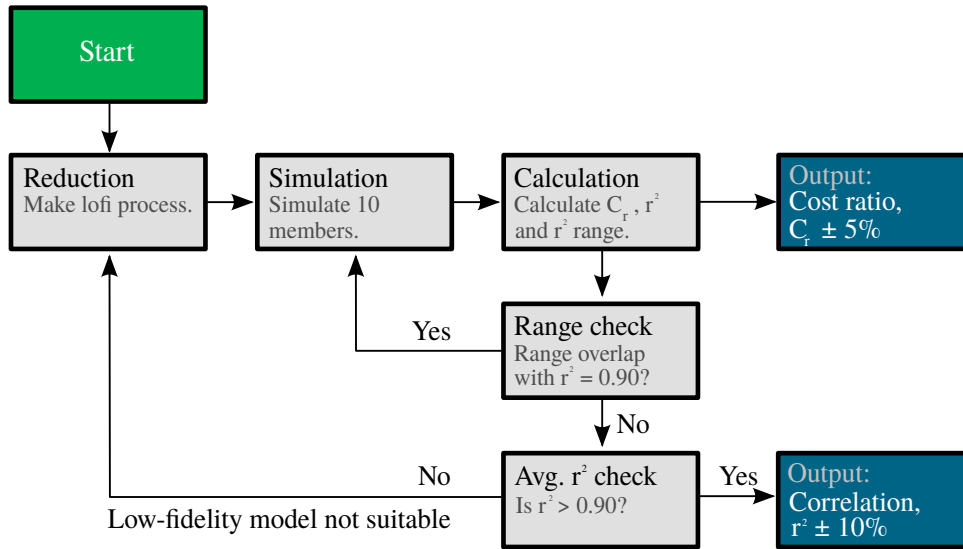


Figure 5.19: Flow chart to estimate the cost ratio and correlation coefficient and estimate low-fidelity process suitability.

Investigating Co-Kriging model correlations

In this work, the average correlation coefficient r_{avg}^2 is calculated by averaging the individual Co-Kriging model correlations, as explained in Section 3.7. Because contributions might be differently affected by a certain fidelity reduction method, in the following the effects of fidelity reduction on individual model correlations is investigated.

The optimisation problem in this thesis had two objectives and two restrictions. For each of these four parameters, a Co-Kriging model is established. The individual model correlations are $r_{\eta_{ISO}}$, $r_{\eta_{HOT}}$, r_{PR} and $r_{\dot{m}}$. By taking the average of these correlations, for each multi-fidelity optimisation an average correlation coefficient r_{avg}^2 is calculated. In Table 5.6 below, the individual correlations are shown in column 2 to column 5, for all five tested multi-fidelity optimisations. The average correlation coefficients are shown in column 6 of this table. Next, observations are made on the basis of these values.

Table 5.6: Individual correlations of objectives and restrictions and the average correlation coefficients for multi-fidelity optimisations tested in this thesis.

Opti. name	$r_{\eta_{ISO}}$	$r_{\eta_{HOT}}$	r_{PR}	$r_{\dot{m}}$	r_{avg}^2
MUFI 1	0.98	0.99	0.85	0.98	0.90
MUFI 2	0.98	0.97	0.80	0.98	0.87
MUFI 3	0.99	0.98	0.99	0.99	0.98
MUFI 4	0.96	0.93	0.90	0.92	0.86
MUFI 5	0.99	0.34	0.84	0.56	0.46

By inspecting the correlations in column 2 to column 5 in Table 5.6, one could notice that there might be a large difference between individual coefficients and the average correlation coefficient. The largest variation can be seen for optimisation MUFI 5, where $r_{\eta_{ISO}} = 0.99$ is 115% higher than the average correlation $r_{avg}^2 = 0.46$ and almost 200% higher than the individual correlation of the second objective $r_{\eta_{HOT}} = 0.36$. From this observation, it is concluded that two Co-Kriging models of the *same* optimisation could have very different correlations.

Previously in Section 5.2-2, it has been observed that MUF1 3 obtains a higher average correlation than MUF1 2. This is unexpected, since the low-fidelity process of MUF1 3 is subject to more fidelity reduction which usually results in a reduction in correlation. By inspecting individual correlations of these optimisations in Table 5.6, one can observe that the average correlation of MUF1 2 is highly reduced by the correlation of the Co-Kriging model approximating the pressure ratio in near stall operating conditions. The correlation of the pressure ratio $r_{PR} = 0.80$ for MUF1 2 is almost 20% lower than the correlation $r_{PR} = 0.99$ for MUF1 3.

In Figure 5.20 and Figure 5.21, the complete set of member evaluations of the pressure ratio in the near stall point are shown for MUF1 2 and MUF1 3 respectively. The interpretation of this type of diagram was previously explained in Section 3.7. When comparing the spread in member evaluations, it can be observed that the evaluations of MUF1 2 are indeed located further from the line of positive correlation compared to MUF1 3, as is suggested by difference in value of the correlation coefficient. In Figure 5.20 for every member also an error bar is included, which shows the absolute error to the low-fidelity approximation of the Co-Kriging model of MUF1 2. On the contrary for MUF1 3, error bars of member evaluations in Figure 5.21 are so small that they are not visible. Considering diagrams of other Co-Kriging models (not shown here), it is found that the error bars increase in size for models with low correlation coefficients. Regardless of the error bars, due to the spread of evaluations for MUF1 2, the correlation is certainly lower compared to the pressure ratio evaluations of MUF1 3.




Figure 5.20: Diagram of pressure ratio evaluations of MUF1 2 by low-fidelity and high-fidelity processes




Figure 5.21: Diagram of pressure ratio evaluations of MUF1 3 by low-fidelity and high-fidelity processes

An explanation for a change in model correlation of the PR in NS operating could be a sensitivity of the pressure ratio to the employed fidelity reduction methods. Compared to MUF1 2, the low-fidelity process of MUF1 3 is subject to a larger mesh reduction in the S1 plane. The number of axial cells is reduced from 79 to 47 and the tangential resolution is reduced from 31 to 19. However in this thesis, the dependency of the pressure ratio on this type of mesh reductions could not be investigated.

As shown by correlation values in the bottom row of Table 5.6, the mass flow (\dot{m}) and efficiency in HOT operating conditions (η_{HOT}) are mostly affected by severe fidelity reduction. For future research, one is advised to take into account the effects of fidelity reduction on individual model correlations.

Concluding statement

This chapter has presented and discussed the performance results of six multi-fidelity optimisations. By identifying the best and the worst multi-fidelity performance with respect to single-fidelity optimisation in Section 5.2-1 and evaluating these optimisations using potential metrics in Section 5.2-2, it has been concluded that the multi-fidelity method is not suitable for improving the current optimisation problem. The method should be applied to a problem with a higher fidelity gap. When an appropriate problem is identified and several low-fidelity processes are established, it is possible to calculate the correlation coefficient and cost ratio with an accuracy explained in Section 5.5. A designer can then estimate suitability of a low-fidelity process for multi-fidelity optimisation. Because the best multi-fidelity optimisation and the single-fidelity reference optimisation have found very comparable design optima as shown in Section 5.2-1, in the next chapter the geometry and aerodynamic performance of these blade designs are investigated.

6 Aerodynamic analysis of optimum members

Optimisation results from Chapter 5 have shown that selected optimum members from the best performing multi-fidelity optimisation MUFI 3 and single-fidelity reference optimisation SIFI attain almost the same improvement in the objectives. In this chapter, the geometrical shape and aerodynamic performance of optimum members from MUFI 3 and SIFI are investigated, to verify if those optima either show similarities or differences.

This chapter starts by visually inspecting the 3D blade shape in Section 6.1. A discussion of the radial distribution of isentropic efficiency and total pressure is described in Section 6.2. Finally in Section 6.3, the Mach distribution is presented at two radial span locations. This chapter is closed by explaining physical differences and their implications on aerodynamic performance.

6.1 Three-dimensional blade shape

Optimum members of MUFI 3 and SIFI are selected using the approach described in Section 3.6. In this section, the these optimum blade designs are investigated on the basis of their geometrical shapes.

The position of the member designs in the Pareto diagram, as well as the three-dimensional blade shapes are shown in Figures 6.1 and 6.2. It is observed that the Pareto fronts have a very similar shape. The optimum members of SIFI and MUFI 3 are located at the intersection point of the diagonal line (dashed) and the Pareto front. The motivation for this selection approach can be found in Section 3.6. The optimum of SIFI attains an average efficiency gain of $\Delta\eta = 1.15\%$. The gain by the MUFI 3 optimum is $\Delta\eta = 1.16\%$ and is hence only 0.01% higher. The blade shapes of the selected optimum blade design from SIFI (blue) the initial rotor design (grey) are shown in Figure 6.1.




Figure 6.1: Pareto diagram of SIFI with isometric view of selected SIFI optimum blade design and initial design. (Note: blade shapes skewed for confidentiality reasons).

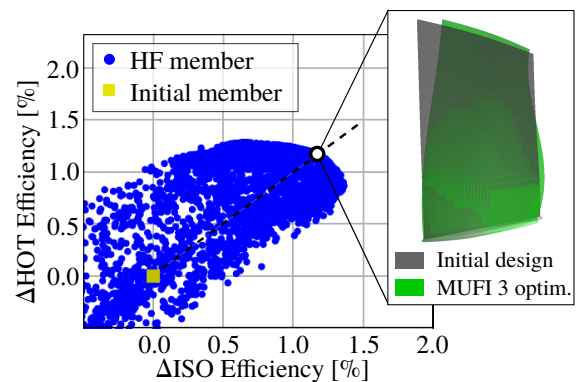


Figure 6.2: Pareto diagram of MUFI 3 with isometric view of selected MUFI 3 optimum blade design and initial design. (Note: blade shapes skewed for confidentiality reasons).

Purely from a visual perspective¹, it is observed that at the blade root, the chord length of the SIFI optimum is smaller than the chord length of the initial blade. By inspecting Figure 6.2 showing the selected optimum from MUFI 3 (green) and the initial design (grey), it is observed that close to the blade root, the chord of MUFI 3 is larger than the initial design. From these observations, it is concluded that MUFI 3 and SIFI optimum designs are geometrically different. In Section 3.8 it was explained that differences between optimum designs have been identified in previous research. In the work of Brooks et al. [3], single-fidelity and multi-fidelity optimisations spending the same computational budget have found blades with different blade leaning and blade twist. In this thesis, study on the blade design parameters was not performed, hence geometrical differences are not quantified in this section. In the next section, differences in aerodynamic performance are investigated by analysing radial variations.

6.2 Radial distributions

In this section, the optimum members of single-fidelity and multi-fidelity optimisation are compared on the basis of radial distributions of isentropic efficiency and total pressure ratio.

The radial distributions of isentropic efficiency for the initial blade design, SIFI optimum and MUFI 3 optimum are presented in Figure 6.3. The isentropic efficiency applies to the second stage efficiency simulated in the ISO operating point is investigated, which is an optimisation objective in the current setup. In Figure 6.3, it is observed that both SIFI and MUFI 3 distributions attain higher efficiency values compared to the initial design, over a blade span ranging from $R = 8\%$ to $R = 91\%$. Taking the integral of these radial efficiency distributions should amount the same efficiency gain of approximately $\eta_{ISO} = 1.15\%$. However, it can be observed in Figure 6.3 that the efficiency gain is obtained at different radial locations. The MUFI 3 optimum has a higher efficiency gain than the SIFI optimum between $R = 8\%$ and $R = 60\%$, by at most $\Delta\eta = 0.3\%$. The SIFI optimum has a higher efficiency gain than the MUFI 3 optimum between $R = 60\%$ and $R = 91\%$, by at most $\Delta\eta = 0.5\%$. However, both attain their maximum gains at approximately $R = 10\%$. A similar analysis is performed of the radial variation in isentropic efficiency in the HOT operating point. Although results from η_{HOT} analysis are not shown here, they are in agreement with previous observations based on the ISO efficiency variation.

The radial distributions of total pressure ratio for all three blade designs (initial, MUFI 3 and SIFI) are presented in Figure 6.4. From this figure it can be observed that the radial variation of at most $\Delta PR = 0.3$ in total pressure ratio of the initial blade design is smaller than the radial variations of the SIFI and MUFI 3 optima. Compared to the radial distribution of the initial blade design, the SIFI optimum design attains a higher pressure ratio between $R = 8\%$ and $R = 70\%$, with a maximum difference of $\Delta PR = 0.1$ at $R = 45\%$. The MUFI 3 optimum designs attains a higher pressure ratio compared to the initial design between $R = 0.30\%$ and $R = 0.85\%$, with a maximum difference of $\Delta PR = 0.75$ at $R = 60\%$.

From this analysis, it is concluded that SIFI and MUFI 3 optimum designs are not only geometrically different, but that they also differ in terms of aerodynamic performance. Whereas the SIFI optimum design obtains a higher efficiency gain in the tip region, the MUFI 3 optimum design obtains a higher gain in the hub region of the blade compared to the initial design. To further investigate aerodynamic differences, two cuts are generated at span locations $R = 0.40\%$ and $R = 0.75\%$. The reason for choosing these spanwise locations is that both the isentropic efficiency and the total pressure ratio show distinct values, as is shown in Figures 6.3 and 6.4. The cross-sectional blade shapes at $R = 0.40\%$ and $R = 0.75\%$ are shown in Figures 6.5 and 6.6 respectively. In both figures, one could observe the differences in the blade stagger and maximum thickness of the optimum blades with respect to the initial blade design. At $R = 0.40\%$, compared to the initial blade design, the maximum blade thickness of the

¹Although the 3D blade shapes are scaled, geometrical differences relative to the initial blade design can still be observed.




Figure 6.3: Radial distribution of stage 2 isentropic efficiency in ISO operating point




Figure 6.4: Radial distribution of stage 2 total pressure ratio ISO operating point




Figure 6.5: Blade cross-sections at $R = 0.40\%$ of SIFI optimum, MUFI 3 optimum and initial rotor design. Cross-sectional shapes are scaled for confidentiality reasons.




Figure 6.6: Blade cross-sections at $R = 0.75\%$ of SIFI optimum, MUFI 3 optimum and initial rotor design. Cross-sectional shapes are scaled for confidentiality reasons.

MUFI 3 optimum is 20% less and the maximum thickness of the SIFI optimum is almost 35% less. The stagger of the MUFI 3 optimum is 5% less and the stagger of the SIFI optimum is 5% more. At $R = 0.75$, compared to the initial blade design, the maximum blade thickness of the MUFI 3 optimum is 45% less and the maximum thickness of the SIFI optimum is 40% less. The stagger of the MUFI 3 optimum is 4% more and the stagger of the SIFI optimum is 2% less. Due to time constraints, the link between aforementioned differences in blade geometry and the radial distributions presented in this paragraph have not been investigated.

In the next section, the aerodynamic flow field in the blade-to-blade plane (S1) is analysed at the span locations indicated previously.

6.3 Distributions in the blade-to-blade plane

In this section, the aerodynamic performance in the blade-to-blade plane of the initial blade and the two optimum blade designs is considered. The Mach number distribution is presented for these blade designs at two radial positions, whose selection was motivated in the previous section.

The Mach number distributions of the initial member, SIFI optimum and MUFI 3 optimum at $r = 0.40\%$ and $r = 0.75\%$ can be inspected in Figure 6.7 to Figure 6.12. In these figures, only the second stage of the two-stage blade geometry is considered. In the following, the local regions of supersonic flow are compared, bounded by the shock wave at the rotor blades.

When comparing the Mach number distributions of the SIFI optimum and the initial blade at a radial span of $R = 0.40\%$ (Figure 6.7 and Figure 6.9), it can be observed that the region of supersonic flow at the SIFI rotor covers a 14% larger portion of the blade suction side (indicated by L). Comparing the distributions of the MUFI 3 optimum and the initial blade at the same radial span (Figures 6.7 and 6.11), the local supersonic region at the MUFI 3 rotor blade covers 11% more of the suction side. Making the same comparison at $R = 0.75\%$, comparing the Mach distributions of the SIFI optimum and the initial blade (Figures 6.8 and 6.10), the shock wave at the SIFI rotor covers 20% more of the suction side. From a comparison of the MUFI 3 optimum and the initial blade design (Figures 6.8 and 6.12), a similar increase in length of the supersonic region over the blade suction side is observed, equal to 20%.

The maximum shock strength of the MUFI 3 optimum at $R = 0.40\%$ is $M = 1.20$, which is 7% lower compared to a strength of $M = 1.29$ for the SIFI optimum and 6% lower compared to $M = 1.27$ for the initial design. At spanwise location $R = 0.75\%$, the maximum shock strength of the MUFI 3 optimum is $M = 1.25$, which is 2% higher compared to both a strength of $M = 1.23$ for the SIFI optimum and $M = 1.22$ for the initial design.

These observed differences in shock location and strength could determine the variations in isentropic efficiency at hub and tip, identified in the previous section and illustrated in Figure 6.3. However, in Section 2.3 it is explained that many flow phenomena have an impact on the isentropic efficiency, besides the shock structure. Hence, an extended aerodynamic study needs to be performed in order to explain radial variations in isentropic efficiency. That study could investigate the flow field at more than two radial locations. Furthermore, such a study should be targeting losses in the boundary layer and in the tip region due to tip clearance, amongst others.




Figure 6.7: Mach number distribution of the initial design at $R = 0.40\%$




Figure 6.8: Mach number distribution of the initial design at $R = 0.75\%$




Figure 6.9: Mach number distribution of SIFI optimum at $R = 0.40\%$




Figure 6.10: Mach number distribution of SIFI optimum at $R = 0.75\%$




Figure 6.11: Mach number distribution of MUFI 3 optimum at $R = 0.40\%$




Figure 6.12: Mach number distribution of MUFI 3 optimum at $R = 0.75\%$

Concluding statements on aerodynamic analysis

In the last paragraph of this chapter, conclusions are drawn from the comparison on geometrical features and aerodynamic performance of MUFI 3 and SIFI optimum designs and the initial blade.

This comparison study has considered two members out of all Pareto optimal solutions, obtained by the optimisations MUFI 3 and SIFI. From results discussed in this chapter, it is concluded that the optimum blades of MUFI 3 and SIFI optimisations are geometrically and aerodynamically different. Compared to the initial design, the MUFI 3 optimum design obtains a higher gain at the blade hub at a radial span from $R = 8\%$ to $R = 60\%$ by at most $\Delta\eta = 0.3\%$. The SIFI optimum design obtains a higher efficiency gain in the tip region at a radial span from $R = 60\%$ and $R = 91\%$ by at most $\Delta\eta = 0.5\%$. Furthermore, comparing the radial distribution of total pressure ratio for MUFI 3 with the initial blade design, a maximum difference in pressure ratio is obtained of $\Delta PR = 0.75$ at $R = 60\%$. The SIFI optimum attains a maximum difference of $\Delta PR = 0.1$ at $R = 45\%$. Furthermore the Mach number distributions of the three blade designs are considered at two spanwise locations. Concerning the region of supersonic flow as a fraction of the overall rotor suction side, for SIFI and MUFI 3 optima these regions are both 20% larger at $R = 0.75\%$ compared to the initial blade design. At $R = 0.40\%$, this region is also 14% and 11% larger, respectively. The highest flow velocity for the MUFI 3 is $M = 1.20$ at $R = 0.40\%$. This shock is 7% weaker than the shock at the rotor blade of the SIFI optimum design and 6% weaker than the shock of the initial design. Differences observed in shock location and strength could determine the variations in isentropic efficiency at hub and tip. However, an extended aerodynamic study needs to be performed on other flow mechanisms causing channel losses, in order to fully explain the radial variations in isentropic efficiency. Because distinct aerodynamic performance of two optimum blade designs is demonstrated, three hypothetical scenarios are formulated to describe the multi-fidelity algorithm:

1. Optimisations SIFI and MUFI 3 are not equally converged. The optimisations might eventually search the same design space, but only at a later point in time. In Section 4.5, it has been shown how the convergence criterion influences the amount of CVG improvement that is remaining after termination.
2. The blade geometries look different, but design parameters causing geometrical differences hardly affect the optimisation objectives. To verify this scenario, one is required to analyse the sensitivity of the objectives to all blade design parameters. An approach for analysing the sensitivities, is to consider the hyperparameter $\vec{\Theta}$ for every Co-Kriging model. In Section 3.4-1, it was explained that model hyperparameters are optimised during surrogate training. It was shown that every design parameter has its own hyperparameter $\vec{\Theta}_l$. Therefore impact of a design parameter on a Co-Kriging approximation can be interpreted through the value of its hyperparameter. If the objectives are insensitive to many blade parameters, one can say the optimum blade designs are located in the same design space.
3. Optimisation MUFI 3 did resolve a different part of the design space. The multi-fidelity optimisation has found a different optimum design, which obtained an almost identical improvement in compressor efficiency. To visualise the design space, one needs to reduce a multi-dimensional solution space to an intuitive graphical representation. A clustering algorithm can be used to visually group similar designs in a 2D or 3D diagram. Such a diagram could help to identify shared design features.

Because analysis of two optimum members does not allow to distinguish one of these scenarios, it can not be concluded whether the multi-fidelity optimisation found a different global optimum. In order to test the hypotheses above, the current research needs to be extended by an analysis of various Pareto optimal member designs and the development of an optimum blade design over the course of an optimisation.

7 Conclusions

In the final chapter of this thesis, the research questions are answered and supported by conclusions from the previous chapters.

In Section 7.1, the most important conclusions are summarised. Subsequently in Section 7.2, recommendations are given for future research on the multi-fidelity method, applied to aerodynamic shape optimisation of axial compressor blades.

7.1 Answers to research questions

This paragraph summarises the main conclusions of this thesis. In this work, the use of the multi-fidelity method is explored to speed up the aerodynamic shape optimisation process of axial compressor blades. A multi-fidelity optimisation differs from a single-fidelity optimisation by employing an additional low-fidelity process for cheap evaluation of member designs, as an alternative to the expensive high-fidelity process.

In Chapter 3, it is concluded from the work of Toal [29] that three metrics determine the benefit of a multi-fidelity optimisation: the cost ratio of process expenses, the correlation between process results and replacement ratio representing the fraction of low-fidelity members. Previous work in which the multi-fidelity method is employed to compressor optimisation problems has shown a 50% savings in overall optimisation run time can be achieved, besides an increase in the optimum efficiency of 0.35%. In the works of Reimer [23] and Brooks et al. [3] cost ratios of 0.11 and 0.33 respectively were achieved. The optimisation problem considered in this thesis consists of two compressor blade stages from which only the second rotor is parametrised. The optimiser “AutoOpti” is used for automatic optimisation of the rotor blade geometry. Aerodynamic performance is evaluated on the basis of the isentropic stage efficiency, using the flow solver “TRACE”. The optimisation objective is to increase the isentropic efficiency in two operating points, while subject to two constraints.

In Chapter 4, the high-fidelity process for CFD simulation is presented. This model employs a mesh with a total of 903,000 cells and solves using Menter SST turbulence equations. Five multi-fidelity optimisation setups are tested, which all employ this high-fidelity process, but distinct low-fidelity processes ranging from a fine mesh and Menter SST turbulence equations to a coarse mesh and inviscid Euler equations. The five low-fidelity models are selected based on the simulation of one member design using both high-fidelity and low-fidelity processes.

In Chapter 5, the results from the single-fidelity baseline optimisation and five test optimisations are presented and discussed. The single-fidelity reference optimisation has increased the average isentropic efficiency by 1.15% with respect to the initial blade design, in a theoretical optimisation time of 1666 hours. This time is equivalent to an absolute duration of approximately 7 days. The worst multi-fidelity optimisation has a delay in optimisation run time of 39.9%, equivalent to almost 3 days of time loss. The selected optimum design found by this optimisation has a loss in average efficiency of 0.11%. The best multi-fidelity optimisation has a speed-up of 14.1%, equivalent to 1 day of time savings. The selected optimum design has an improvement in average efficiency of 0.02%. The primary research question of this work is answered on the next page.

Research question 1

How much fidelity reduction should be applied to the low-fidelity process chain, in order to achieve the highest speed-up or improvement of the optimum by multi-fidelity optimisation with respect to single-fidelity optimisation?

Answer

The most suitable low-fidelity process employs a mesh of 223,000 cells in total (a mesh ratio of 4 compared to the high-fidelity process), with 28 radial streamlines and a blade-to-blade resolution of 47 axial and 19 tangential cells. The low-fidelity employs the same Menter SST turbulence model as the high-fidelity process. The multi-fidelity optimisation employing this low-fidelity process, MUFI 3, achieves a speed-up of 14.1%. This result is useful in terms of absolute time savings, but 36% lower compared to findings described in reference literature. Due to uncertainty imposed by progress variability, there is little guarantee this speed-up can be achieved when repeating the optimisation. The main conclusion of this work is that the fidelity level of the high-fidelity process in this work is too low to obtain a large benefit from the multi-fidelity method.

The current fidelity reduction methods, consisting of mesh reduction and adaptation of solver settings in 3D CFD, do not allow a sufficient reduction in the cost ratio without a considerable reduction in the correlation coefficient. The best cost ratio is achieved by optimisation MUFI 5, equal to $C_r = 0.14$. However, MUFI 5 only attains a very low correlation coefficient of $r^2 = 0.46$. The best correlation coefficient is achieved by MUFI 3, equal to $r^2 = 0.98$. However, MUFI 3 attains a mediocre cost ratio of $C_r = 0.29$. From all test optimisations using the (automatic) variance decision function, for MUFI 3 the highest fraction of computational budget was allocated to low-fidelity member evaluations, equal to $f_r = 7\%$. This fraction is 3% below the lower limit recommended by Toal [29] and 35% lower than used in the work of Reimer [23]. Repeating the optimisation as MUFI 6 with an increased computational budget equal to 51%, the maximum improvement in CVG with respect to SIFI is increased by 14%, owing to enhanced information transfer from low-fidelity members. However, low-fidelity members are still too expensive and delay the optimisation in terms of its run time. It is recommended to apply the multi-fidelity method to a different optimisation problem, whose setup is described in Section 7.2.

In Section 5.5, convergence trends of the cost ratio and correlation error are analysed by calculating these parameters for randomly sampled sets of member designs with increasing size. From this analysis it is found that a designer needs to evaluate 2 blade designs using a low-fidelity and high-fidelity simulation process, to estimate the cost ratio with an accuracy of 5%. Furthermore, a designer needs to evaluate at most 20 blade designs to estimate the correlation coefficient with an error of at most 35%. It is argued that this accuracy is sufficient to check the minimum required amount of correlation of $r^2 = 0.90$, specified by Toal [29]. In order to investigate the suitability of a low-fidelity process with respect to a selected high-fidelity baseline model, a designer can apply these guidelines.

In Chapter 6, the geometrical shape and aerodynamic performance of selected optimum members from the best multi-fidelity optimisation (MUFI 3) and the single-fidelity reference optimisation (SIFI) are compared, which attained almost the same average efficiency gain of $\Delta\eta = 1.16\%$ and $\Delta\eta = 1.15\%$ respectively. In the following, the second research question of this thesis is answered.

Research question 2

Are optimum designs from multi-fidelity and single-fidelity optimisations similar or different in terms of shape and aerodynamic performance?

Answer

The selected multi-fidelity and single-fidelity optima are geometrically different and perform aerodynamically different. The observed differences are summarised below. Three possible scenarios are presented which could explain why a different optimum blade design may have been found using the multi-fidelity method.

It is shown that compared to the initial design, the MUF1 3 optimum obtains a higher gain in the hub region of the blade at a radial span from $R = 8\%$ to $R = 60\%$ by at most $\Delta\eta = 0.3\%$. The SIFI optimum design obtains a higher efficiency gain in the tip region at a radial span from $R = 60\%$ and $R = 91\%$ by at most $\Delta\eta = 0.5\%$. At two radial locations the Mach number distribution is considered. Concerning the region of supersonic flow on the suction side as a fraction of the overall rotor chord length at $R = 0.75\%$, this region is 20% larger for both SIFI and MUF1 3 optima compared to the initial blade design. At $R = 0.40\%$, the supersonic regions are respectively 14% and 11% larger. The highest reduction in shock strength is at $R = 0.40\%$ by MUF1 3, with a shock strength of $M = 1.20$. The shock strength is 7% lower than the shock strength at the rotor blade of the SIFI optimum design and 6% lower than the shock strength of the initial design. Differences observed in shock location and strength could cause the variations in isentropic efficiency at hub and tip. However, the aerodynamic study needs to be extended to other flow mechanisms in order to fully explain the radial variations in isentropic efficiency. This discussion cannot ascertain that multi-fidelity and single-fidelity algorithms have led to different parts of the design solution space. Firstly, one of the optimisations could be non-converged and could later find a more similar design. Secondly, contradicting geometrical parameters could have a very weak influence on the current design objectives. The investigation should be extended, by monitoring the development of the optimum blade design over the course of an optimisation.

7.2 Future work

In this thesis it is found that the multi-fidelity method is not suitable for improving current optimisation performance. However, the method can be tested to a different optimisation problem with a higher potential. This section gives recommendations of such an optimisation setup and suggestions for future work in this field of research.

The first paragraph of this section describes a more suitable setup for multi-fidelity optimisation with higher potential. The second paragraph describes adaptations to optimisation settings to fine-tune the setup.

1. Adaptations to optimisation setup

For better performance of the multi-fidelity method, the fidelity gap between the low-fidelity and the high-fidelity process should be increased. In Figure 7.1, the fidelity gap is illustrated for the current setup and a recommended optimisation setup. The current high-fidelity process simulates the four-row geometry using 903,000 mesh cells and solves the flow using Menter SST turbulence equations. The process of lowest fidelity uses a mesh of 12,800 cells and solves the flow using Euler equations. The first option is to use a setup with a more expensive high-fidelity process that simulates detailed design features such as penny gaps by employing a more refined mesh. The second option is to employ a cheaper low-fidelity process for simulation such as through-flow calculations. Alternatively, additional simulation methods could be added to the high-fidelity process chain such as heavy structural simulations, which can then be greatly simplified in the low-fidelity process. One should consider if these methods provide additional value to compressor blade shape optimisation. Generally a larger fidelity gap is seen for other problems, such as optimisation of the rear compressor stages requiring a larger mesh to solve for stronger cross-flow. Another example is optimisation of turbine blades where heat transfer plays a predominant role besides flow aerodynamics. A new setup satisfies the guidelines defined by Toal [29] and attains

similar values in correlation, cost ratio and replacement ratio as demonstrated for the works of Reimer [23] and Brooks et al. [3]. The cost of the low-fidelity process should be at most 11% of the high-fidelity process ($C_r = 0.11$) and a high correlation coefficient should be maintained ($r^2 = 0.90$). A considerable part of the computational budget should be taken up by low-fidelity members ($f_r = 0.40$).

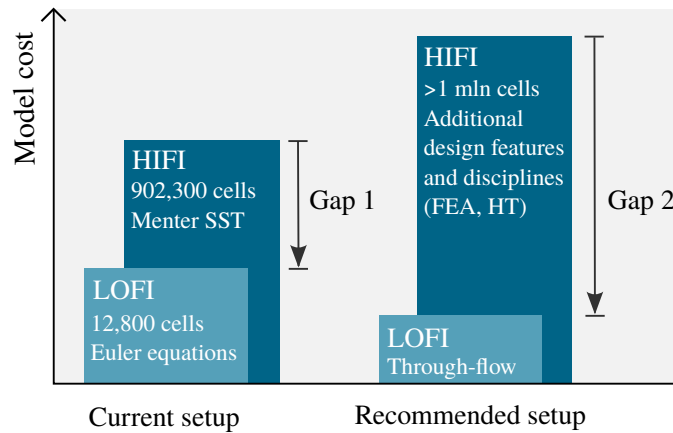


Figure 7.1: Fidelity gaps of the current and recommended optimisation setup. Increasing the cost ratio of processes in an improved setup, increases the chances of successful multi-fidelity optimisation, provided that process results are well-correlated.

2. Adaptations to optimisation settings

In this subsection, optimisation settings are recommended to possibly improve the benefit of multi-fidelity optimisation. These settings are based on reference literature and inspired by a discussion with AutoOpti developer C. Voss¹.

1. Based on successful multi-fidelity optimisations described in the work of Reimer [23], it is recommended to employ the *variance decision function* to let the optimiser decide on whether to evaluate high-fidelity or low-fidelity members. When a more suitable optimisation problem is available, the automatic decision function is better capable of estimating when to evaluate using the low-fidelity process, compared to manually setting the low-fidelity member ratio.
2. The *initialisation duration* could be increased to at least 1.5 times the number of design variables. Increasing the duration, improves quality of the Co-Kriging model and enhances information transfer from low-fidelity members, hence it could improve potential of the multi-fidelity method.
3. Third, for maximising the difference between single-fidelity and multi-fidelity optimisations, the *member generation method* should discard evolutionary operations. Only the surrogate model should be used for member generation after initialisation, which allows to take full advantage of the Co-Kriging model in multi-fidelity optimisation.
4. It is recommended to employ *asynchronous training*. The training time of the Co-Kriging is longer than the Kriging model used in single-fidelity optimisation. By using asynchronous training, the master no longer needs to wait for training to finalise before generating new member designs. Multi-fidelity optimisation then has less of a disadvantage compared to single-fidelity optimisation.

¹Personal communication, October, 2017

Appendix

Table A1: Process times of the hifi and lofi process chains of MUFI 1 (multi-fidelity setup with lofi process of highest accuracy and worst optimisation performance).

Process name	Hifi time	Lofi time	Cost ratio
CleanSlaveDir	00:00:06	00:00:03	0.50
MultistageAirfoilCreation	00:00:07	00:00:05	0.71
FEM_MeshBL_Ccx	00:00:02	00:00:02	1.00
GAT3003_1_Mesh	00:02:40	00:02:05	0.78
GAT3003_2_SetupTrace	00:00:03	00:00:02	0.67
GAT3003_2_SetupTrace	00:00:03	00:00:02	0.67
GAT3003_2_SetupTrace	00:00:03	00:00:02	0.67
GAT3003_3_RunTrace	00:11:21	00:05:45	0.51
GAT3003_4_PostProTrace	00:00:37	00:00:28	0.76
GAT3003_3_RunTrace	00:10:10	00:04:52	0.48
GAT3003_4_PostProTrace	00:00:36	00:00:28	0.78
GAT3003_3_RunTrace	00:09:41	00:04:26	0.46
GAT3003_4_PostProTrace	00:00:36	00:00:27	0.75
EvaluateMember	00:00:03	00:00:03	1.00
CleanUnimportant	00:00:01	00:00:01	1.00
Total	00:36:56	00:19:31	0.53

■ Expensive subprocess

Table A2: Process times of the hifi and lofi process chains of MUFI 3 (multi-fidelity with lofi process of moderate accuracy and best optimisation performance).

Process name	Hifi time	Lofi time	Cost ratio
CleanSlaveDir	00:00:05	00:00:05	1.00
MultistageAirfoilCreation	00:00:06	00:00:05	0.83
FEM_MeshBL_Ccx	00:00:01	00:00:01	1.00
GAT3003_1_Mesh	00:02:45	00:01:21	0.49
GAT3003_2_SetupTrace	00:00:03	00:00:02	0.66
GAT3003_2_SetupTrace	00:00:03	00:00:02	0.66
GAT3003_2_SetupTrace	00:00:03	00:00:02	0.66
GAT3003_3_RunTrace	00:08:03	00:01:55	0.18
GAT3003_4_PostProTrace	00:00:38	00:00:21	0.55
GAT3003_3_RunTrace	00:06:56	00:01:32	0.22
GAT3003_4_PostProTrace	00:00:38	00:00:21	0.55
GAT3003_3_RunTrace	00:06:40	00:01:30	0.23
GAT3003_4_PostProTrace	00:00:37	00:00:21	0.57
EvaluateMember	00:00:03	00:00:02	0.66
CleanUnimportant	00:00:02	00:00:01	0.50
Total	00:27:00	00:07:57	0.29

■ Expensive subprocess

Table A3: Process times of the hifi and lofi process chains of MUFI 5 (multi-fidelity setup with lofi process of lowest accuracy).

Process name	Hifi time	Lofi time	Cost ratio
CleanSlaveDir	00:00:05	00:00:04	0.80
MultistageAirfoilCreation	00:00:06	00:00:05	0.83
FEM_MeshBL_Ccx	00:00:01	00:00:01	1.00
GAT3003_1_Mesh	00:02:46	00:00:45	0.27
GAT3003_2_SetupTrace	00:00:03	00:00:02	0.66
GAT3003_2_SetupTrace	00:00:03	00:00:02	0.66
GAT3003_2_SetupTrace	00:00:03	00:00:02	0.66
GAT3003_3_RunTrace	00:07:53	00:00:36	0.08
GAT3003_4_PostProTrace	00:00:38	00:00:15	0.40
GAT3003_3_RunTrace	00:07:07	00:00:33	0.08
GAT3003_4_PostProTrace	00:00:38	00:00:15	0.40
GAT3003_3_RunTrace	00:06:44	00:00:34	0.08
GAT3003_4_PostProTrace	00:00:37	00:00:15	0.41
EvaluateMember	00:00:03	00:00:01	0.33
CleanUnimportant	00:00:01	00:00:02	2.00
Total	00:27:07	00:03:43	0.14

■ Expensive subprocess

Bibliography

- [1] Alonso, J. J. and Chung, H. (2004). *Multiobjective optimization using approximation model-based genetic algorithms*. Technical Report 4325, American Institute of Aeronautics and Astronautics.
- [2] Boyce, M. (2011). *Gas turbine engineering handbook*. Butterworth-Heinemann, Houston, Texas, 4 edition.
- [3] Brooks, C. J. et al. (2011). *Multi-fidelity design optimisation of a transonic compressor rotor*. Presented at: 9th European Conference for Turbomachinery, Fluid Dynamics and Thermodynamics, Istanbul, Turkey, Computational Engineering & Design Group.
- [4] Cohen, H. et al. (2009). *Gas turbine theory*. Pearson Education Limited, 6th edition.
- [5] Denton, J. (1993). Loss mechanisms in turbomachinery. *Journal of Turbomachinery*, 115(4):621–656.
- [6] DLR (2013). *TRACE PRE global settings*. Retrieved in June, 2017, from website: http://www.trace-portal.de/wiki/index.php/PRE_Global_settings.
- [7] Ellbrant, L., Eriksson, L. E., and Mårtensson, H. (2012). CFD optimization of a transonic compressor using multiobjective GA and metamodels. volume 4, pages 2698–2712. Volvo Aero Corporation, Chalmers University of Technology, International Council of the Aeronautical Sciences.
- [8] Fernández-Godino, M. G. et al. (2016). *Review of multi-fidelity models*. Presented at: ECCOMAS Congress 2016, Crete Island, Greece.
- [9] Fogel, D. B. (1994). An introduction to simulated evolutionary optimization. *IEEE Transactions on Neural Networks*, 5(1):3–14.
- [10] Gano, S. E. et al. (2005). *Update strategies for Kriging models for use in variable fidelity optimization*. Presented at: 1st AIAA Multidisciplinary Design Optimization Specialist Conference, Austin, Texas.
- [11] Ghisu, T. et al. (2011). An integrated system for the aerodynamic design of compression systems, part 1: development. *Journal of Turbomachinery*, 133:2.
- [12] Goinis, G. et al. (2017). *AutoOpti documentation*. DLR, Köln, Germany, february 2017 edition.

-
- [13] Hickel, S. (2015). *CFD for aerospace engineers*. Delft University of Technology, Delft, The Netherlands.
- [14] International Energy Agency (2016). *World energy outlook 2016*. Retrieved in July, 2017, from website: <https://www.iea.org/publications/>.
- [15] Jin, Y. (2011). Surrogate-assisted evolutionary computation: Recent advances and future challenges. *Swarm and Evolutionary Computation 1*, 61-70.
- [16] Jin, Y. et al. (2000). Comparative studies of metamodeling techniques under multiple modeling criteria. *Structural and Multidisciplinary Optimization*, 23(1):1–13.
- [17] Jones, W. and Launder, B. E. (1972). Prediction of laminarization with a two-equation model of turbulence. *International Journal of Heat and Mass Transfer*, 15(2):301–314.
- [18] Menter, F. R. et al. (2003). *Ten years of industrial experience with the SST turbulence model*. Presented at: 4th International Symposium on Turbulence, Heat and Mass Transfer, Antalya, Turkey.
- [19] Oyama, A. and Liou, M. (2002). *Multiobjective optimization of a multi-stage compressor using evolutionary algorithm*. Presented at 38th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit, Indiana, USA.
- [20] Padròn, A. S. and Alonso, J. J. (2016). *Multi-fidelity methods in aerodynamic robust optimization*. Presented at: 18th AIAA Non-Deterministic Approaches Conference, California, USA.
- [21] Peter, J. and Marcelet, M. (2008). Comparison of surrogate models for turbomachinery design. *WSEAS Transactions on Fluid Mechanics*, 3(1).
- [22] Pini, M. (2015). *Cascade aerodynamics*. Delft University of Technology, Delft, The Netherlands.
- [23] Reimer, E. (2016). *Vergleichende Optimierung eines Fans mit High- und Multi-Fidelity Verfahren*. Master's thesis, FH Aachen, Aachen, Germany.
- [24] Samad, A. (2008). Multiple surrogate modeling for axial compressor blade shape optimization. *Journal of Propulsion and Power*, 24(2):301–310.
- [25] Schmitz, A. (2013). *Entwicklung eines Objektorientierten und Parallelisierten Gradient Enhanced Kriging Ersatzmodells*. Master's thesis, FernUniversität in Hagen, Hagen, Germany.
- [26] Siemens A/G (2013). *First gas turbine manufactured in the new facility in Charlotte*. Website. Retrieved in September, 2017, from website: <https://www.siemens.com/>.
- [27] Siller, U. and Voss, C. (2009). *Automated multidisciplinary optimization of a transonic*

- axial compressor*. Presented at: 47th AIAA Aerospace Sciences Meeting, Florida, USA.
- [28] Simoes, M. R. et al. (2009). Validation of turbulence models for simulation of axial flow compressor. Proceedings of COBEM 2009 20th International Congress of Mechanical Engineering, Gramado, Brazil.
- [29] Toal, D. J. (2015). Some considerations and regarding the use of multi-fidelity and Kriging in the construction of surrogate and models. *Structural and Multidisciplinary Optimization*, 51(6):1223–1245.
- [30] Voss, C. and Nicke, E. (2008). *Automatische Optimierung von Verdichterstufen*. Technical Report 0327713 B, DLR, Institut für Antriebstechnik.
- [31] Wilcox, D. C. (1988). Reassessment of the scale-determining equation for advanced turbulence models. *AIAA Journal*, 26(11):1299–1310.