



MORALI

Multi-Objective Robust Assessment of the Helicopter Improvements Final Report 1.1.2011 – 31.12.2014

IAG, University of Stuttgart
– Manuel Keßler –

Call ID: **SP1-JTI-CLEAN SKY-2010-01**
CfP topic number: **JTI-CS-2010-1-GRC-01-004**
Proposal number: **270629**

Coordinating person: Dr. Manuel Keßler
Universität Stuttgart
Institut für Aerodynamik und Gasdynamik
E-mail: kessler@iag.uni-stuttgart.de
Tel.: +49 711 685 63419
Fax: +49 711 685 53419

List of participants:

Participant no.	Participant organisation name	Country
1 (Coordinator)	Institut für Aerodynamik und Gasdynamik der Universität Stuttgart (USTUTT)	Germany
2	Macros Solutions Ltd. (MACROS)	Bulgaria



Table of Contents

1. Introduction.....	3
2. Recall of project.....	3
2.1 Objectives.....	3
2.2 Structure and schedule.....	3
3. Work package 1 – Management.....	5
3.1 Purpose.....	5
3.2 Progress.....	5
4. Work package 2 – Method development.....	5
4.1 Purpose.....	5
4.2 Progress.....	5
4.2.1 Task 2.1 – Blade element theory.....	5
4.2.2 Task 2.2 – Coupled CFD/CSD simulations.....	6
4.2.3 Task 2.3 – Optimisation.....	7
4.2.4 Task 2.3.1 – Rotor optimisation problem definition.....	7
5. Work package 3 – Performance assessment of innovative rotors.....	8
5.1 Purpose.....	8
5.2 Progress.....	8
5.2.1 Task 3.1 – Performance.....	9
5.2.2 Task 3.2 – Acoustics.....	9
5.2.3 Task 3.3 – Reliability.....	10
6. Work package 4 – Optimisation of innovative rotors.....	10
7. Summary.....	11
7.1 Deliverables.....	11
7.2 Milestones.....	11
7.3 Resources.....	12



1. Executive Summary

The MORALI project took place from 2011-2014 and was embedded in the Cleansky Joint Technology Initiative. It answered a call for proposals resulting from requirements defined by Eurocopter Germany, now Airbus Helicopters Germany. The proposal consortium consisted of two parties: University of Stuttgart from Germany, represented by the helicopter and aeroacoustics group of the Institute for Aerodynamics and Gasdynamics, as the coordinator and taking responsibility for aerodynamic simulations on different fidelity levels, and as the second partner MACROS Solutions from Bulgaria, taking care of advanced optimisation technology in a general framework.

The goal of the project was to establish helicopter rotor design capabilities at significantly advanced levels. Although other aspects as structural stability, mass and construction cost are finally certainly of great importance as well, the rotor as the main means of generating lift, propulsion and control is mostly aerodynamics driven, so the project concentrated on this aspect only, considering other factors only by appropriate boundary conditions. Of course, later on those other factors can be included as well, as soon as models become available and are accessible to numerical treatment.

Even on the aerodynamic side along the complexity of the problem is tremendous. About 20-30 design variables at least allow a fine-tuning of a given rotor to get optimal performance. However, "optimal performance" is difficult to define, as several flight states (hover, forward flight at different speeds, start and approach) have to be taken in to account, on terms as power requirements, loads and acoustics. Furthermore, a reliable evaluation at all degrees of freedom necessitates extremely demanding computational fluid dynamics, allowing only for a literally handful of variants to be considered in detail. Other approaches, as blade element theory or free wake simulations, are several orders of magnitude faster, but may mispredict the performance dependence on some parameters due to missing physical phenomena.

As stated, the University of Stuttgart took responsibility for the aerodynamic simulation on all modelling levels, enhancing the tools in order to improve the physical modelling (dynamic stall for blade element theory, transition prediction for CFD), to enable another method (free-wake model) or to boost performance (trim acceleration). Another task was the appropriate definition of performance – at least in view of power input – for a reliable differentiation between variants, where a new optimisation goal could be defined to drive the process. However, the exact balance between power at different speeds and acoustics is a strategic decision of industry, specific for a certain product with defined missions, and thus delegated to them, in this case Airbus Helicopters.

The optimisation procedure itself was taken care of and targeted by MACROS Solutions, who adapted their optimiser to the specific needs of MORALI. They regard the problem as high-dimensional, with various function evaluations at different reliability levels, from low (depending on the parameter, but very cheap) to high (and extremely costly). The tool can generate a new parameter set, which is then to be evaluated aerodynamically, and the result fed back to the optimiser to drive the process further. The goal here is to blend and integrate the different reliability and cost levels to create the best design with minimal computational effort.

In summary, the project was a huge success. In spite of some intermediate technical problems and political difficulties, all deliverables were deployed and all mile stones reached. Indeed, the technical program was more than fulfilled.. The improved tool chain is now in active use at Airbus Helicopters and supports the development of upcoming new products.



2. Summary of project

The MORALI project started in January 2011 and lasted for a duration of 48 months until December 2014. The strategic goal of MORALI was to improve the helicopter rotor design capability, including comprehensive analysis and evaluation skills of different designs, simulation competence at various modelling levels, and automated optimisation support.

2.1 Context

The development of a new helicopter rotor is – as any aerospace design task – clearly a multi-disciplinary and also multi-objective problem. While aerodynamic efficiency is obviously very important, other factors like structural stability, dynamic compatibility to the rest of the system, manufacturing issues and, last but not least, cost, do also influence the outcome. All these issues have to be taken into account before successful industrialisation of a specific configuration, consolidating the combined expertise of engineers experienced in different disciplines. An automated process chain delivering reliable results helps those engineers to explore substantially more effectively the huge design space of a rotor blade geometry. Such tools significantly support the development of new products, in order to generate better results in a shorter time frame.

As in probably any engineering process, no analytical solution to the design problem is available, and the task is to balance the differing and sometimes contradicting driving forces. Even taking into account only aerodynamic issues, the requirements on different parts of the rotor in various flight conditions are sufficiently diverse to necessitate such delicate balancing.

In order to predict the performance of new rotor designs, different simulation technologies at various confidence levels are available. Momentum theory, taking only mass and radius into account, gives only first impression for preliminary rotor sizing and is not used further in this project, while blade element theory at least uses geometric information of the rotor blades to give a more detailed view into the aerodynamic behaviour in a time frame of seconds. Obviously, some major phenomena are not represented in such a simplified model, and thus ask for more elaborate technology to successfully generate a detailed design.

Especially the wake and thus effective rotor inflow is coarsely modelled as more or less constant, which is the reason to use a free-wake approach. The variable circulation along the blade radius and during rotation generates vorticity, shed from the trailing edge to build up the wake. This vorticity is now convected freely along the inflow as well as its own induction, and thus gives a detailed local representation of the wake structure. For example, blade-vortex interaction phenomena and their accompanying noise emission can be obtained quite successfully with such a free-wake analysis, which consumes computing time on the order of hours.

However, the fully non-linear, non-stationary and three-dimensional flow structure can be only captured by solving the Navier-Stokes equations (in the Reynolds-averaged sense only due to the high Reynolds numbers, of course) in computational fluid mechanics. Of course, the aerodynamics have to be coupled to a structure dynamics simulation in order to take blade elasticity and deformations into account (which also improves blade element theory and free-wake significantly), and the collective and cyclic control angles have to be adjusted to produce the forces and moments targeted. In the end, such simulations are able to reproduce experimental values as power requirements, deformations and noise in sufficient quality to guarantee working designs before any hardware is built. The price for this data quality are substantial computational resources, which reach several ten thousand core hours for a single rotor in a single state of flight. More elaborate configurations, including the full helicopter with fuselage, empennage and tail rotor can even run into several millions of core hours for a trimmed free flight.



All in all, we have at our disposal three simulation models of increasing accuracy and fidelity, from blade element theory to the free wake approach until fully coupled CFD, each adding about four orders of magnitude of computational effort to the previous one.

Furthermore, the geometric design of a rotor blade offers many options, from different airfoil sections along the radius to planforms with sweep and taper, not to mention anhedral or dihedral and finally twist. Even though older blades seldomly were more complicated than an extruded airfoil with some linear twist (and adding some nice tip), better efficiency and acoustic behaviour can only be obtained by more elaborate geometries. Taking just a forward/backward-swept blade as an example, we find easily some 20 parameters to describe such a shape. Obviously, a parameter space consisting of 20 degrees of freedom or dimensions is hardly to be searched exhaustively, especially if a single evaluation takes thousands of core hours, as in the CFD case. Consequently, instead of a blind search in this huge space, an optimisation tool chain helps the qualified engineer to take a well-defined way of successive improvements.

This includes blade element theory, which represents the twist distribution quite well, for example, while leaving other parameters to a – quite limited – number of CFD simulations. Even then, successive searches in 2-3 dimensional subspaces proved more efficient than tackling the full space at once, as general trends and conflicting goals can such be much more easily visualised and understood. Acoustics, as another optimisation goal, may be included by means of free-wake, to look for BVI conditions, for example, that are especially significant in this regard, but not *a-priori* known for a specific rotor.

To sum it up, even only the aerodynamic design of a rotor blade geometry is a very delicate and demanding process. The high dimensionality of the parameter space in conjunction with the very costly evaluation necessitates a very clever combination of different optimisation approaches. Several fidelity levels, with their reliability depending on the parameter under investigation, ask for their integration into a common tool set, in order to help a competent engineer accomplishing its task. For the time being, this task remains difficult, but the tools developed in MORALI give a decisive impetus to bring about better performing rotor blades in less time and with smaller effort.

2.2 Objectives

The overall strategic goal of the MORALI project was – as already stated – to improve the rotor design capability in industry. Itemised into sub-objectives progress on several work areas was expected:

1. Improvement of fast design methods to include more physics and thus to enhance prediction quality.
2. Acceleration of high-fidelity methods to enable earlier adoption in the design process for a better understanding of detailed phenomena.
3. Enhancement of the optimisation tool chain to make a more automated work flow feasible.
4. Assessment of the potential of innovative rotor and control concepts.
5. Application of all the tools to an industrially relevant case.
6. Implementation in an industrial context.

2.3 Structure

There were four work packages in the project. Nearly all of them (besides WP1: Management) were split into several tasks and partly subtasks.

The first work package, WP1, dealt with the management of the project and included all related tasks, as negotiation and communication with the JU, industry and the partner, knowledge dissemination and results exploitation, documentation and financial issues.

The following work package, WP2, brought forward Method Development in three different fields, namely BEM and CFD, respectively, on the simulation side and additionally the optimisation procedure. Each of these tasks consisted of different subtasks, tackling a specific area identified as in need of improvement, respectively. On the boundary element side, dynamic stall was the most important effect unconsidered yet, in addition to the local wake development by vortex tracking. For CFD it was transition modelling and further refined control of the trim procedure. The multi-objective property of the multi-dimensional optimisation process asked for new concepts regarding effective and efficient handling of the information generated.

In contrast, WP3 was concerned with the assessment of simulation results, as produced by the tools improved in WP2. It has proven insufficient to look only at singular point values as, for example, a figure of merit, especially in an automated optimisation framework. Instead, the entire information provided by a detailed simulation needed to be taken into account, giving the very vague concept of “optimal performance” a quantifiable form. Acoustics come into play here, adding another direction besides fuel consumption. Either way this includes an appropriate consideration of errors and strategies to bound them.

WP No.	WP Title	2011				2012				2013				2014			
		Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
		Phase 1				Phase 1				Phase 3							
1	Management	2/1															
2	Method development	33/14															
2.1	Blade element theory	17															
2.1.1	Dynamic stall modelling			4													
2.1.2	Free wake modelling				10										3		
2.2	Coupled CFD/CSD simulations	13															
2.2.1	Transition modelling in CFD	8															
2.2.2	Trim robustness and efficiency											5					
2.3	Optimization	3/14															
2.3.1	Rotor optimization problem definition					1/1											
2.3.2	Existing tool chain review	0/1															
2.3.3	Surrogate modelling	0/6															
2.3.3.1	Proof of concept integration			0/2													
2.3.3.2	Surrogate modeling for industrial problems								0/3								
2.3.3.3	Parameter tuning													0/1			
2.3.4	Optimization coupling	2/6															
2.3.4.1	Proof of concept integration	0/2															
2.3.4.2	Implementation of industrial work flow							2/3									
2.3.4.3	Parameter tuning													0/1			
3	Performance assessment of innovative rotors	8															
3.1	Performance	3															
3.2	Acoustics			3													
3.3	Reliability							2									
4	Optimization of innovative rotors													7			
4.1	Rotor optimization													3			
4.2	High fidelity validation															4	

WP4 finally applied the full tool set to a rotor optimisation problem specified by the ITDL. As for the optimisation process several simplifications are necessary at different modelling levels, the outcome was to be assessed in terms of validity and accuracy by proven and trusted high-fidelity simulations.

3. Generated foreground and results

3.1 Work package 2 – Method development

In the second work package the tools available previously for the design of helicopter rotors were sharpened and further enhanced to include more flow physics and reduce the manual effort. Several levels of simulation quality were to be included in the design process, and consequently three different fields were tackled here: fast, low order methods for preliminary parameter studies, as the blade element theory, high fidelity computational fluid dynamics coupled to structural dynamics, and finally the automated optimisation loop including the different levels of physical representation quality and response time.

3.1.1 Blade element theory

Task 2.1 was concerned with blade element theory improvements, namely the inclusion of dynamic stall effects and a better representation of the rotor wake, taking local induction effects into account.

As blade element theory is based on tables of static two-dimensional airfoil data, instationary stall effects are not included more or less by definition. As those are very important for a proper representation of the rotor flow – for example on the retreating side of the blade – as a whole and especially the trim state (control angles), a heuristic model for lift hysteresis and moment overshoot was implemented and tested. Unfortunately, such models include model parameters, for which no universal data set exists, but has to be adopted to each airfoil used. While this reduces the generality of such an approach, the number of airfoils used for rotors is in fact quite limited, and the parameter set can be deduced by – albeit computationally expensive – detached eddy simulations. Once created, geometric variants of a single rotor as well as airfoil comparisons are possible, and the simulation quality is – depending on the flight state – considerably improved.

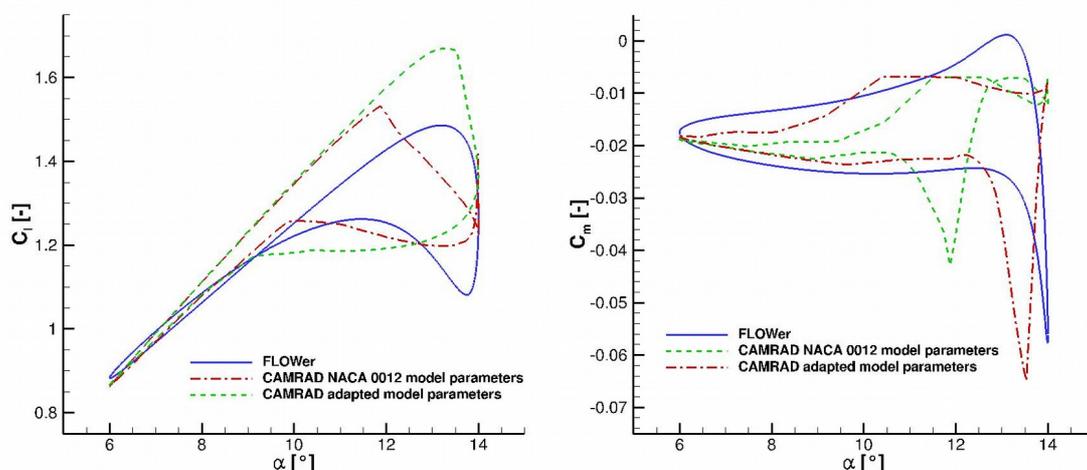


Figure 1: Advanced airfoil lift and pitching moment coefficients for different dynamic stall parameter sets (Leishman-Beddoes dynamic stall model); $Ma=0.3$, $Re=1.8e6$, $\alpha_{mean}=10.0^\circ$, $\Delta\alpha=4.0^\circ$, $f=6.6Hz$

This subtask has been worked on and finally completed in the first year of the project. Deliverable 2.1 documents the activities in this regard, a best practice guide for usage with the BEM tool CAMRAD II was written and handed over to the ITDL.

The second subtask in the blade element theory context was related to the interaction effects between the rotor blades by an improved representation of the wake structure. In simple cases, a global wake structure is prescribed and the induction effects are considered on a globally conservative basis. However, this neglects local effects, which are significant for acoustics as well as loads. Therefore a free wake code was developed, which sheds vortices and wake panels at the trailing edge, according to the spatial and temporal lift distribution on the rotor blade, and convects them downstream. Especially in descent flight, they interact with the following blades, causing local gradients with according changes in flow variables and forces.

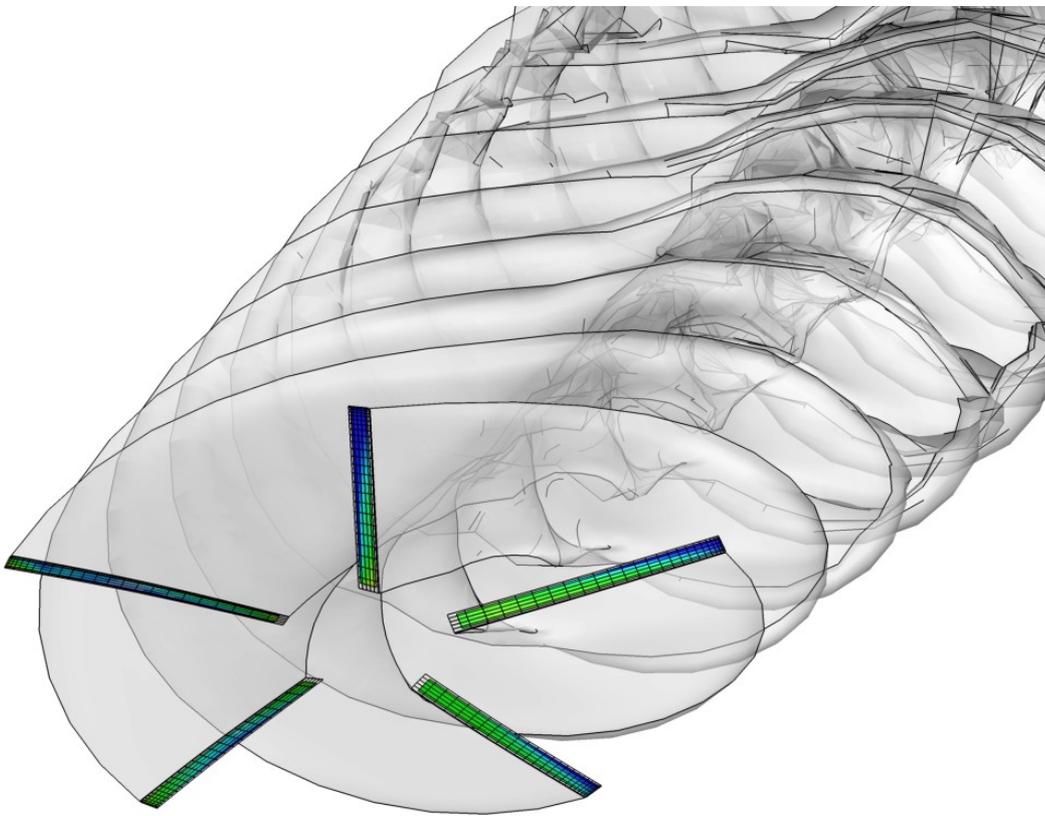
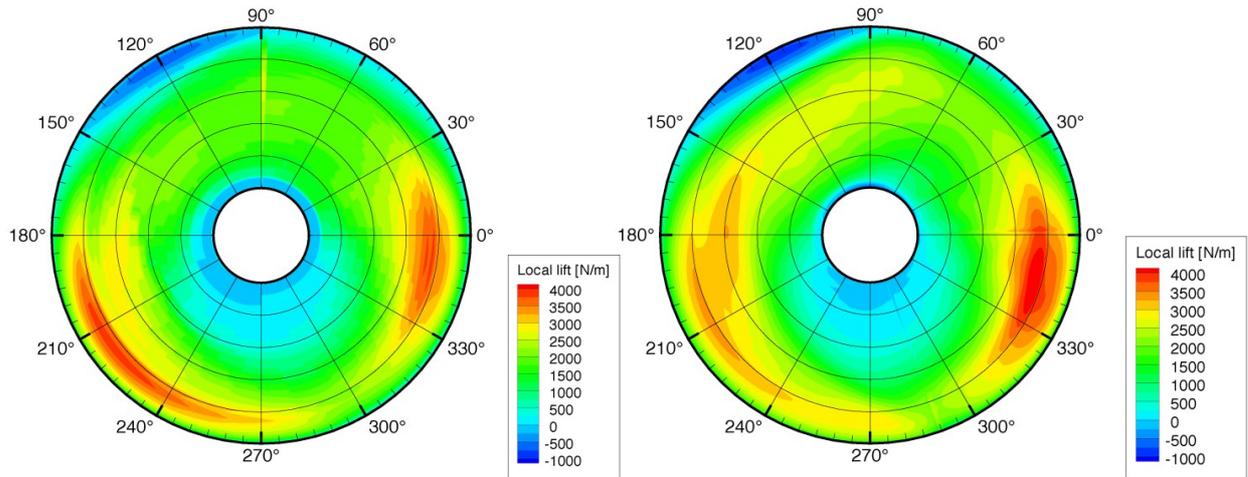


Figure 2: Pressure distribution on lifting surfaces and wake of 5 blade rotor at 125 kts

The code has been architected in a completely general way, enabling the simulation of an arbitrary number of rotors with possibly different blades and kinematics. Parallelisation was built in from the start, in order to make use of the increasing power of the computing architectures available nowadays. For the wake and vortex strengths the firstly implemented lifting line algorithm tailored towards rotating systems was complemented by a lifting surface and a fully flexible lifting body algorithm. Although the idea of coupling this free wake code to a near-blade CFD simulation is appealing, the interfacing proved sufficiently problematic to be handled in a consistent and conservative manner, such that the substantial projected effort did not relate favourably to the gains in computational efficiency. A large part of the CFD cells are located in the boundary layer anyway, and accordingly the reduction in computational time does not compensate for the loss of physicality in the free-wake part in addition to the interface difficulties.



(a) FIRST (b) FLOWer
Figure 3: Polar plots of local lift. FIRST in comparison with FLOWer

Stability issues occurring during the first validation period have finally been fixed. The performance could be massively increased by implementing a fast multi pole method to reduce the method complexity from n^2 per time step to $n \log n$. The free wake model has been tested and validated on simple cases. The software – including documentation, of course – has been delivered to AHD, which is about to take it into productive use. Preliminary performance measurements promise a speed up of at least an order of magnitude compared to other wake methods available in older codes.

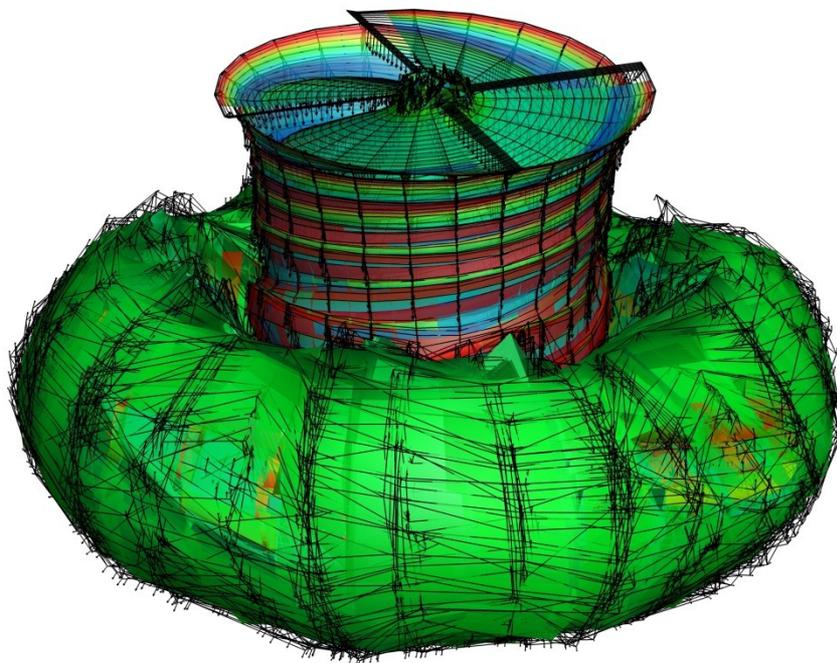


Figure 4: Four bladed rotor in hover including swimming off starting vortex.
Coloring shows chordwise vortex strength

3.1.2 Coupled CFD/CSD simulations

One evident shortcoming of current CFD simulations is the neglect of transition and thus of laminar flow at the leading edge. This accounts for a systematic overestimation of power requirements in the order of 10%. While this is not a problem for any optimisation result, if it stays approximately constant for all variants, this assumption is probably not valid for all possible cases. Therefore a transition model based on flow variables and development should switch between laminar and turbulent models in order to take laminar flow regions into account properly and consider their dependency on other degrees of freedom.

Therefore a transition prediction method has been implemented in FLOWer, based on a detection mechanism for the boundary layer edge, streamline tracking along the blade surface and finally an e^N -method for the transition onset. Depending on the parameters used, the results compare well to the (very few, unfortunately) experiments, and give a substantial improvement for global power predictions on the rotor.

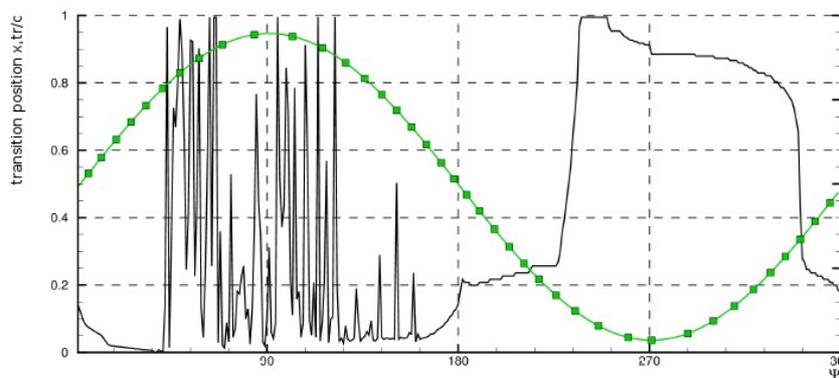


Figure 5: Transition positions (e^N envelope, TRACFREQ 10); OA209 airfoil, $Ma=0.31$, $Re=1.16e6$, $\alpha_{mean}=9.83^\circ$, $\alpha_{ref}=9.1^\circ$ and $\alpha=0.05$

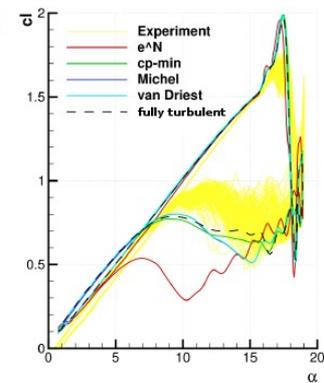


Figure 6: lift coefficient for different transition prediction criteria

In the other subtask – *trim robustness and efficiency* – firstly, some grid convergence studies of rotor blades with different base shapes in hover and forward flight attitudes have been undertaken to examine the dependence of absolute values on the grid resolution and to verify the grid independence of relative values between different blade shape designs.

As it turned out, differences between coarse, intermediate and fine grids can be easily observed, as shown in Figure 7, and add mostly high-frequency content with increasing resolution, going from 7 million cells for the coarse setup over 15 millions to 34 millions in the fine case. Overall, the global discrepancies in thrust or power, for example, account for some 2-3%. However, what is more important, the *trends* differentiating various configurations of blade geometries, are *consistent* even on the coarse meshes, as shown in Figure 8. That means, better configurations show improved performance of similar amount on all three resolutions. Consequently, an optimiser will chose the best configuration and thus the correct optimisation direction, regardless of the resolution – at least at these levels and for the criteria looked at (thrust, power) in this investigation. For different goals – acoustics under BVI conditions comes to mind – the required resolution may be quite different, as said high-frequency content is exactly the phenomenon to be minimised. Therefore the application of the tool chain obviously needs the correct specification of boundary conditions and driving parameters depending on the exact problem to be solved.

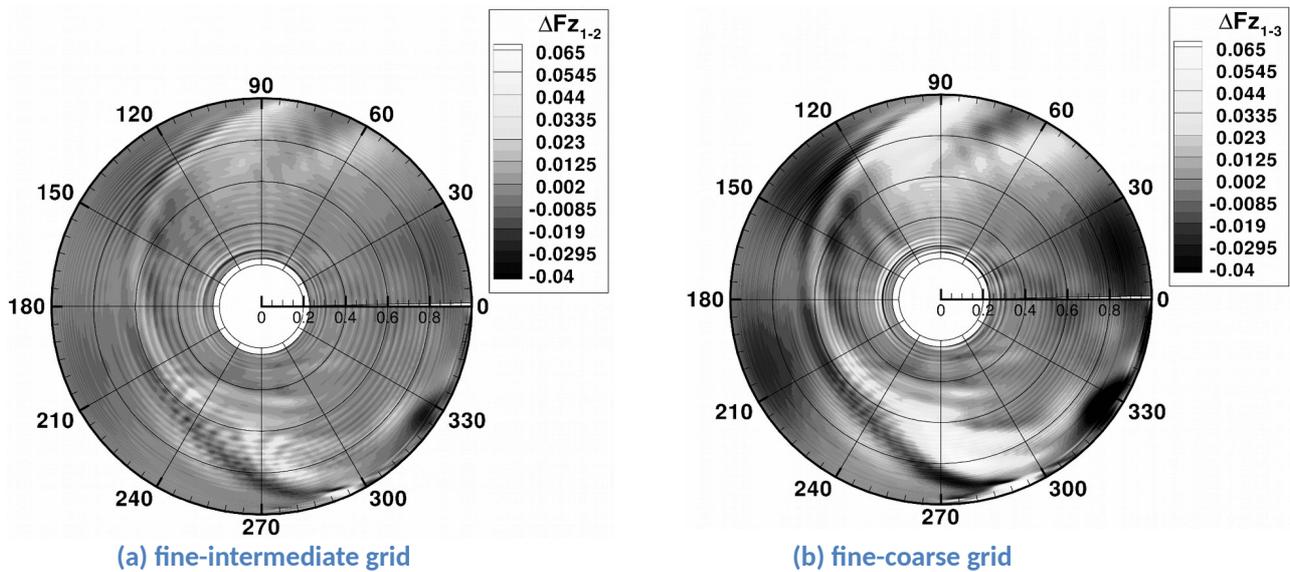


Figure 7: Thrust differences

It has to be stated, that the applicability of coarse meshes is probably highly dependent on the quality of the generated meshes, and especially of their comparable quality. Our automated mesh generation process is based on the geometric parameters defined by the optimiser and builds up a carefully chosen topology of block structured grids without the need for any manual intervention, guaranteeing the same high quality meshes even for quite different blade geometries. This process is necessary for mesh generation in the automated tool chain and enables the comparability of results on physical differences, not on mesh quality.

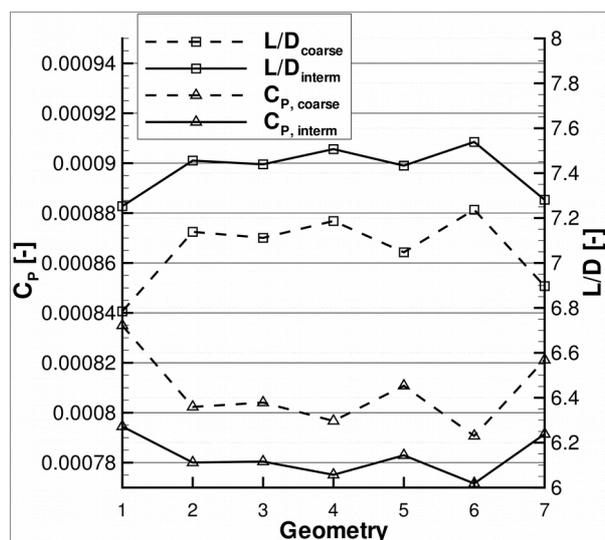


Figure 8: Power coefficient and L/D for different blades on coarse and intermediate meshes

Secondly, the weak coupling scheme requires periodic loads. Hence, CFD computations must be performed until periodicity of the flow field is obtained. A method was developed and implemented, which examines the current periodicity within each CFD time step. In doing so, less CFD time steps and consequently less computation time is required per trim iteration.

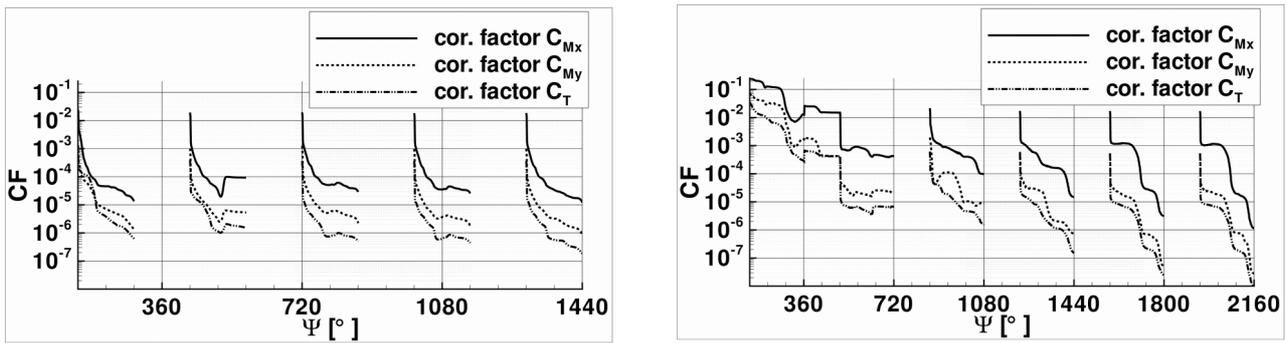


Figure 9: Trim auto-correlations for 360° and 288°-trims

Thirdly, a study was undertaken focussing on the reduction of necessary trim iterations. It was proven that the values of L/D obtained by a three component trim (with the objectives thrust, and pitching and rolling moment of the rotor) are nearly identical to values obtained by trimming rotors with a four component trim (with the objectives lift, propulsive force, rolling, and pitching moment). While power for a fully trimmed configuration is the ultimate criterion, L/D depends equivalently on the design parameters, resulting in a very high correlation. Thus it was shown, that the three component trim with L/D as quality criterion for the rotor in forward flight is adequate for optimisation purposes, whereby on average one trim iteration less is required before trim state converges.

The achieved reduction of computational costs for trimmed rotor simulations exceeded the original expectations of possible speed up by updating the trim Jacobian within each time step and using induced velocity information extracted from CFD results within the BEM model.

3.1.3 Rotor optimisation problem definition

This subtask consists of two parts. Firstly, the definition of optimisation objectives and secondly, the definition of geometrical parameters to be optimized during the rotor optimisation.

For performance, firstly it was considered to use the lift-over-drag ratio as a selective criterion for comparing different representations in the parameter space. In a later stage of development this ratio was improved to the product of propulsive force and flight speed divided by the power required. In a certain sense this would be an equivalent of the figure of merit for forward flight, meaning the fraction of power being converted into propulsive force to overcome fuselage drag. This lead to the definition of

$$\frac{L}{D} = \frac{L}{\frac{P}{v_\infty} - X}$$

with the propulsive force X of the rotor. If two rotor shapes yield different forces, the one with the stronger forward force is then preferred, as it may overcome more drag of the fuselage.

Based on these findings, an investigation of L/D values obtained by three component trims in comparison to the power coefficients obtained by four component trims was undertaken. As shown on representative configurations, L/D values of three component trims behave the same as power coefficients of four component trims. Hence L/D of three component trim was selected as performance objective, resulting in faster convergence than power of a four component trim and thus reducing resource requirements.

The function of power demand over rotor advance ratio is often used during preliminary design as it describes the performance of a rotor system for the most relevant flight conditions. A method was

developed to determine the empirical parameters of this function by only three CFD simulations per blade design. This allows easily the definition of complex, practical multi objective optimization functions for different mission scenarios.

Due to the high resource constraints of CFD rotor simulations the number of parameters which can be optimized is quite limited. Considering technical and geometrical limitations and aerodynamic relations, the design space used for rotor blade optimisations resulted to eight different geometrical parameters defining the base shape of the rotor blade tip section.

3.1.4 Surrogate Modeling and Optimisation Coupling

Due to IP problems the integration of the MACROS optimiser into the tool chain of Airbus Helicopters was severely delayed, even if at then end a solution satisfying for all partners could be found.

In the meantime, MACROS improved their optimisation algorithms and surrogate models in order to accelerate the optimisation process. As one parameter evaluation (namely, a trimmed and coupled CFD solution) is extremely expensive, it is of uttermost importance to keep their number as low as possible with respect to the parameter space.

Furthermore, a process was invented to combine evaluations at various confidence levels (CFD and blade element theory, for example), taking the different correlations of the results with the parameters into account. For example, the parameter twist is easily optimised with blade element theory, and the correlation of the respective results to the parameter high. The specific twist shape, on the other hand, results in highly three-dimensional flow, which is basicalle neglected with blade elements, or at most considered by some empirical corrections. Consequently, only CFD results provide a reliable optimisation basis for those planform parameters.

3.2 Work package 3 – Performance assessment of innovative rotors

While the evaluation of a single rotor design is difficult enough, as already extensively elaborated, the assessment of the outcome is possibly even more involved. Specific needs require different optimisation objectives, depending on the defined design driving features.

While this is clearly a multi-objective problem, the parallel handling of different optimisation directions is definitely out of question due to resource constraints. The multi-dimensional design space in conjunction with the heavily involved function evaluation makes a – more or less – dense exploration of it impossible, and even such a concept as a Pareto front becomes questionable at a small single-digit number of evaluations per design variable – at least for high fidelity methods.

The same argument holds for acoustics like noise emission during landing procedures. How to combine acoustic levels with performance measures – and on which level – to business values is at the end a strategic decision of the industrial partner, valid for a single rotorcraft and with respect to a representative mission profile. Thus it is very difficult to trade such optimisation directions without considerable and proper knowledge of the business model – and finally not the purpose of a research project like MORALI.

3.2.1 Performance

As coupled CFD/CSD simulations are computationally extremely expensive, it is very important to get the most possible information out of each single shot in the dark. For example in the trim procedure, control angles are adjusted to deliver prescribed values for thrust and mast moments.



While it is clearly pointless to compare power requirements for different thrust levels on different variants, it is not so clear how to include other values like drag into the global evaluation. A larger drag means more forward tilt of the mast to compensate, requiring more power. While this could in theory be handled by the trim of another control variable (mast angle), each trim variable considerably increases the computational effort to come to convergence. Therefore it was decided to define a lift-over-drag ratio similar to fixed wing aircraft for the rotor, which includes its forward force X. This force helps to overcome the fuselage drag and delivers as such usable power.

$$\frac{L}{D} = \frac{L}{\frac{P}{v_\infty} - X}$$

Of course, this is not fully equivalent to a full four component trim as described previously. However, experiments show, that this efficiency number is very closely coupled to the fully trimmed power requirements and thus leads to the same optimization landscape, in contrast to a simple power comparison ignoring differences in forward force. As a bonus, we get a trim acceleration by at least one iteration and thus a full revolution.

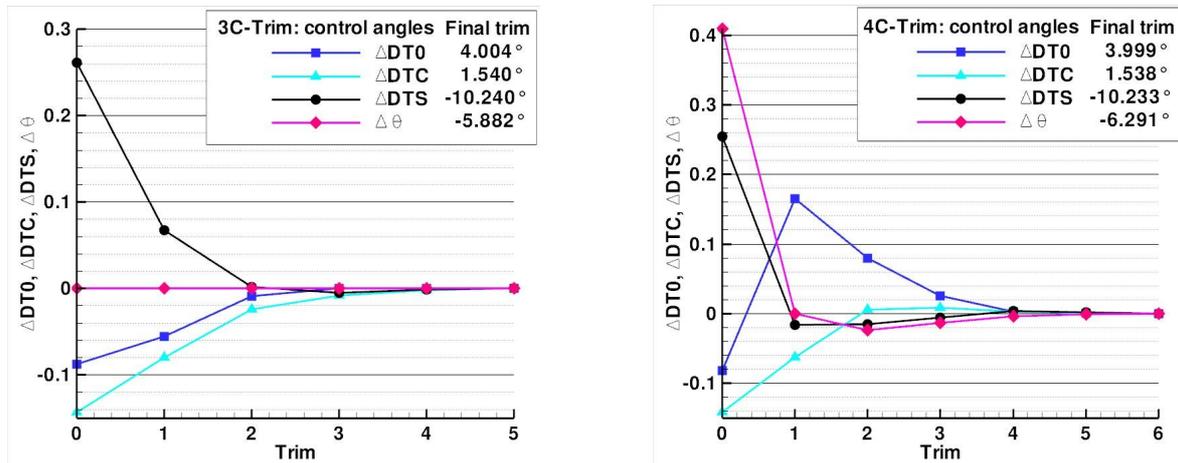


Figure 10: Convergence of three vs. four component trim

For the optimisations, fluid-structure coupled CFD simulations were carried out to examine the influence of an-/dihedral on the performance of an isolated rotor in forward flight and hover. While blades with dihedral performed well in forward flight and achieved an improvement of more than 11% compared to a rectangularly shaped reference blade with parabolic tip section, in hover those blades were less efficient. Blades with anhedral behaved in a contrary way. Yet the benefit of anhedral in hover was less intense than the performance loss in forward flight due to a downward deflection of the blade tip. Generally blades with a less intense deflection performed better both for anhedral and dihedral.

In another study two parameters were examined: radial position and size of an increased chord length. At first the parametric area was evaluated with a Finite State Unsteady Wake model, then CFD calculations were conducted. The necessary amount of CFD calculations was estimated to 10-15 data points to get a good idea of the behaviour of the two parameters. A broad optimum range was found in the parametric area yielding an improvement of up to 9% compared to the reference blade.

Furthermore, the influence of the rotor blade planform on the cruise flight performance of a five bladed isolated helicopter rotor were studied. Therefore, trimmed and fluid-structure coupled



Computational Fluid Dynamics simulations were performed. The rotor blade base shape geometry of a blade with increased chord length was described with eight parameters. The chord length at the blade tip was adapted to obtain blades with the same thrust weighted equivalent chord length.

The effect of the different parameters on the local thrust distribution was identified and related to the trim control angles collective pitch, longitudinal and lateral cyclic pitch. Improvements in L/D of up to 5 % compared to a rectangularly shaped reference blade with the same thrust weighted equivalent chord are achieved by different designs.

3.2.2 Acoustics

An acoustic evaluation of several rotor blades was performed. In a grid study the influence of grid refinement on the acoustic solution was found to be in an acceptable range for the intermediate compared to the fine mesh resolution concerning noise footprints. Larger differences occurred on the coarsest mesh resolution. Although the noise footprint of the four examined blade shapes did not show significant differences, the blade with the best performance (in terms of rotor L/D) from the present study of an increased chord length, had a reduced noise emission on the advancing side, while at the retreating side the noise level dropped only slightly. Transformation into the frequency domain showed that noise results were affected by non periodic disturbances, likely to be caused by numerical effects.

3.2.3 Reliability

As part of Task 2.2 systematic simulation inaccuracies like discretisation errors or errors caused by turbulence models have been proved to be mostly independent of the design of rotor blades. Thus, these modelling uncertainties are not hindering the shape optimisations of rotor blades when driving towards goal improvements.

But, to prevent wrong optimisation targets, the significance of a positive or negative effect of parameter variations must be examined in relation to the convergence error of a solution. It is shown, that for reliable shape optimisations a confidence level of 99.9% must be reached for the last two trim iterations. This confidence level is assured once the rotor trim angles converged to less than 0.01° .

3.3 Work package 4 – Optimisation of innovative rotors

The application of the improved tool chain to an industrial case was the final part of the project, in order to prove the practicability of the chosen approach and the full functionality of the developed and improved components.

3.3.1 Rotor optimisation

After the implementation of the optimisation tool chain at the ITDL it was run for a specific optimisation case as requested by Airbus Helicopters. This consisted of some 1000 runs of the blade element tool HOST to properly fill the design space for the blade geometry as given by the – in this preliminary test case – five parameters depicted in Figure 11. The blade has a double taper starting with a rectangular section up to radius.2, then the chord is linearly increased up to radius.6 to the value chord.6, finally the chord is decreased again using a Hermite spline function for the leading edge shape, which is specified by the two angles LEinnerdxdy.6 and LEouterdxdy.6.

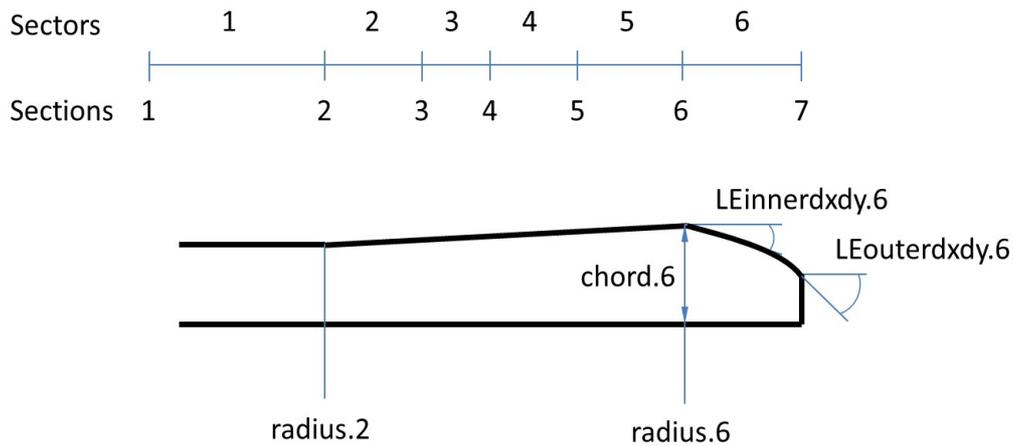


Figure 11: Blade parametrization with 5 design variables

Although such a highly dimensional space is difficult to visualise, Figure 12 gives an indication of some two-dimensional slices through it, rendering the improvements in L/D as colours. Trends are clearly visible, and by inspecting several parameter combinations in sequence it is possible to get an impression of the optimisation direction the process is going to move forward.

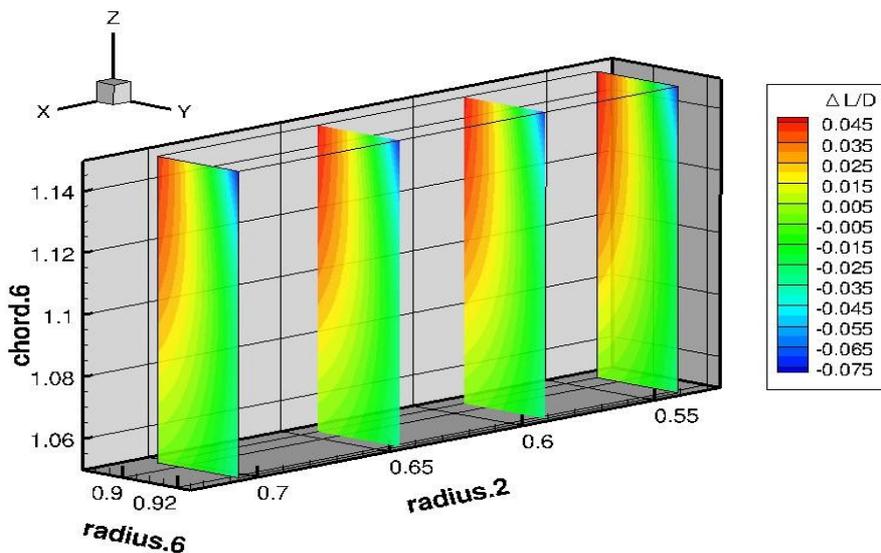


Figure 12: HOST computations: effect on L/D of parameters radius.2, radius.6 and chord.6

It has to be said, that even such a high number of evaluations leads to less than 20 singular points only on each of the slices shown. The smooth contours stem from the surrogate model, which was constructed from the blade element runs in order to constitute a continuous hypersurface, from which the optimisation process can be driven further. For CFD, even given substantial computing resources are available, a number of the order of 100 is more realistic, resulting in less than three points per dimension if using a full space coverage. Obviously, those computing power can be spent more sensible to only look in parts of improvements and neglect apparently less promising candidates.

3.3.2 High fidelity validation

In order to check the reliability of the blade element results from the previous task, a total number of 32 CFD fully coupled simulation runs have been run in the same five-dimensional design space, resulting in a single evaluation in each corner. Consequently, only linear dependencies can be tracked, which is acceptable for a first validation attempt.

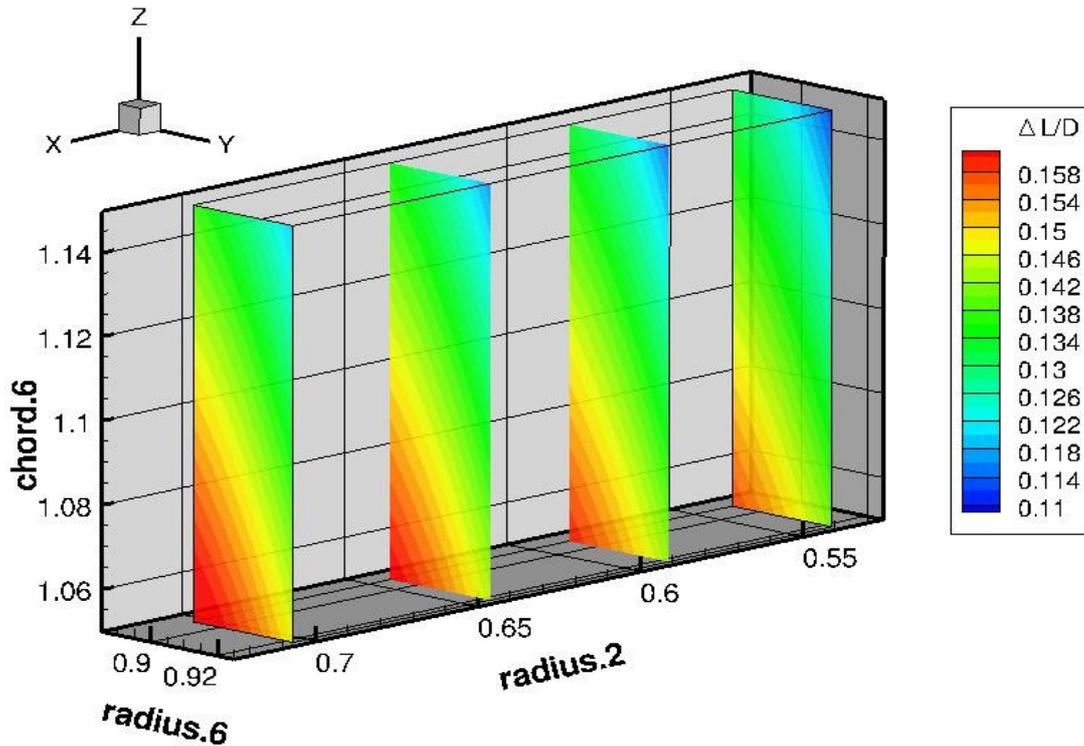


Figure 13: CFD computations: effect on L/D of parameters radius.2, radius.6 and chord.6

Comparing Figure 13 with the corresponding blade element results of Figure 12 it can be seen, that the general trend (and nothing more can be extracted from that few CFD results) on radius.2 and radius.6 is reproduced, while chord.6 differs significantly with a wrong sign of the derivative. While HOST prefers larger chord increases, the CFD simulations point to a less tapered design. Possibly the three-dimensional flow due to the chord (and thus circulation) change negatively influences the performance of the rotor. However, HOST is not able to capture this phenomenon, as it inherently relies on two-dimensional polars integrated along the radius.

This supports the earlier assumption, that some parameters are accessible to blade element optimisation, while others definitely need CFD simulations in order to model the physics happening in reality. Chord changes were expected at first to be well represented in blade elements (other than sweep, for example), but the results showed otherwise. Therefore, this experiment proves that a differentiation *a priori* is not straightforward, but needs substantiation for any individual parameter.

However, it would be a waste to throw away all blade element results generated already in favour of the scarce CFD results, so an attempt was made to blend both to a single surrogate model. This was accomplished by emphasising the CFD data and adding the less reliable blade element results as a minor correction. The goal was to fetch the global trends from CFD and “enrich” them with some local nonlinearities from HOST. One could even weigh the blade element input in various dimensions differently, depending on the reliability found in preliminary investigations. However, the available data did not substantiate a quantitative qualification of different dimensions of the

design space, so equal weight was applied to generate Figure 14, which reproduces the trends from the CFD results in Figure 13 and adds some small disturbances due to the merging of the blade element data from Figure 12.

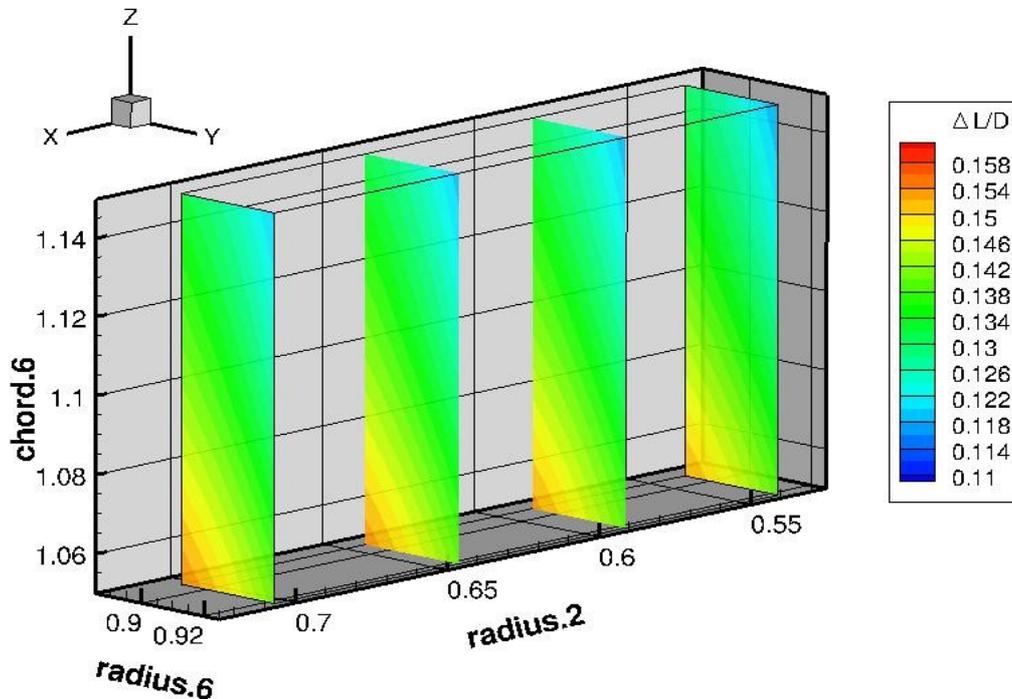


Figure 14: Multi-fidelity model: effect on L/D of parameters radius.2, radius.6 and chord.6

All in all, the developed tool chain is now in place and enables the automated optimisation process of rotor blade geometries, as demanded and proposed in the project. Further refinement on specific parameters for the application on next-generation industrial rotors will certainly improve the efficiency and efficiency of the process. However, the non-reliability of some blade element results and huge computational demands of a few accurate CFD simulations together with the high-dimensional design space still necessitate an experienced engineer to get working and significantly better rotor designs.

4. Impact and Dissemination

The general purpose of the CleanSky JTI is to support European aerospace industry in competitiveness and sustainable growth by enabling cleaner technologies for air transport. More specifically, the MORALI project was the response to a call for proposals created to support Airbus Helicopters (then Eurocopter) as the ITD leader relevant for MORALI.

So the main driver for MORALI was the supporting action for the optimisation tool chain for future helicopter rotor blades on demand of Airbus Helicopters. Consequently, the expected impact was to considerably improve the design capabilities of the ITDL for rotorcraft with respect to improved performance, efficiency and noise generation.

All goals described in the call for proposals have been reached, with respect to technology improvements (dynamic stall treatment, transition prediction), tool developments (free wake creation, design parameter driven automatic mesh generation, optimiser integration), result production (trim acceleration, periodicity detection, goal oriented evaluation) and finally application. Of course, no specific blade design for a real future helicopter product has been made



within MORALI, as this process involves many more relevant technologies and dependencies on other than aerodynamic properties, and furthermore contains considerable valuable intellectual property of the ITDL just by specifying goals and boundary conditions, but the process chain has been implemented and established at industry, in order to be applied for the next rotor generations. So MORALI's impact is the support in improved performance, reduced operation cost and noise emissions.

This industrialisation of technology and tools is the primary impact of MORALI, according to the definition in the call for proposals and the work plan negotiated at the beginning. However, the knowledge generated is also available now at the project partners for future research, furthering the capabilities in any forthcoming projects. For example, at IAG a completely new project, funded by the DFG (Deutsche Forschungsgemeinschaft), is dedicated to research on dynamic stall within the rotating system, which would not have been possible without the fundamental work carried out within MORALI. So, besides the industrial exploitation, even at our university a new job (another PhD candidate) was created in reaction to the MORALI effort. More projects are planned, hopefully leading to further improvement of the simulation capabilities at IAG and increased understanding of aerodynamic phenomena in the context of rotorcraft, for example in the context of the just started JTI CleanSky2.

Besides, the general ideas regarding the underlying technologies have been introduced at several conferences and published in international journals. Although some very specific details, especially at the application level, had to be spared to protect sensitive intellectual property rights of Airbus Helicopters, the concepts and their interaction and interfacing are publicly available for other interested parties as well. Thus the MORALI project has contributed to a significant increase in the general applicability of simulation technology for rotorcraft.

Specifically, the following conference contributions and journal papers directly resulted from work done within MORALI:

1. Martin Hollands and Manuel Keßler and Andree Altmikus and Ewald Krämer. Trade Study: Influence of Different Blade Shape Designs on Forward Flight and Hovering Performance of an Isolated Rotor (European Rotorcraft Forum, Gallarate, 2011)
2. Martin Hollands and Manuel Keßler and Ewald Krämer. Influence of An-/Dihedral and of Different Blade Shapes on Performance and Aeroacoustics of an Isolated Rotor (European Rotorcraft Forum, Amsterdam, 2012)
3. Patrick Kranzinger and Martin Hollands and Manuel Keßler and Sigfried Wagner and Ewald Krämer. Generation and Verification of Meshes Used in Automated Process Chains to Optimize Rotor Blades, (50th AIAA Aerospace Sciences Meeting, Nashville, 2012)
4. Hollands, Martin and Keßler, Manuel and Krämer, Ewald. Planform Design for a Five Bladed Isolated Helicopter Rotor Using Fluid-Structure Coupled CFD Simulations (31st AIAA Applied Aerodynamics Conference, San Diego, 2013)
5. Martin Hollands and Manuel Keßler and Ewald Krämer. Blade Shape Design: Interpolation Based on Fluid-Structure Coupled Simulations of an Isolated Rotor in Forward Flight (in: Dillmann, A., Heller, G., Krämer, E., Kreplin, H. P., Nitsche, W., Rist, U. (Hrsg.): *Notes on Numerical Fluid Mechanics and Multidisciplinary Design, Vol. 124, New Results in Numerical and Experimental Fluid Mechanics IX*. Springer, 2014)
6. A. Klein and Th. Lutz and E. Krämer and K. Richter and A. D. Gardner and A. R. M. Altmikus. Numerical Comparison of Dynamic Stall for Two-Dimensional Airfoils and an Airfoil Model in the DNW-TWG (Journal of the American Helicopter Society, Vol 57, No. 4, 2012)



More papers regarding MORALI technology or results are currently either in work or in the review process. Tools improved within MORALI (ACCO, the acoustic postprocessing tool), have been industrialised and several licenses sold to third parties. Further commercial exploitation of MORALI is within the responsibility of Airbus Helicopters.

5. Conclusion

In the aftermath the project MORALI has to be seen as a great success. Despite some problems (mostly due to political difficulties within the ITDL, although technical obstacles had to be tackled as well) all technical work as proposed and negotiated or adapted in accordance with Airbus Helicopters, respectively, was completed, in time and on budget. Due to careful planning and shifting of work assignments as necessary, even some additional features requested by the ITDL (as in the free wake tool) could be implemented.

The informal project management by (mostly) email and phone calls, and personal meetings only in case of necessity proved very effective and freed up more resources for technical instead of administrative work. This clearly shows, that a small and dedicated team of competent people can generate substantial progress on ambitious goals at high efficiency, instead of wasting a large part of available resources on coordination, meetings and controlling activities. Such a concentration on the technical effort certainly contributed significantly to the exceptionally successful completion of MORALI.

For future upgrades and new products, Airbus Helicopters will make use of the delivered and improved tools and the corresponding know-how laid down in the accompanying documentation and best practice guidelines. It is now up to the industry to turn the European tax money, invested into MORALI, into products, jobs, and finally economic success.