Stretched Rotor:

Stretching Wind Turbine Blades using Optimization

Forlængelse af vindmøllerotorer ved brug af optimering

Final Report

EUDP journal nr.: 64015-0067

Frederik Zahle (frza@dtu.dk)†,
David Verelst†, Geoerg Pirrung†,
Albert Meseguer Urban†, Michael K. McWilliam†,
Witold Robert Skrzypinski†,
Casper Skovby‡

† Department of Wind Energy, Technical University of Denmark
‡ LM Wind Power

July 13, 2018

Contents

1 Project Details

Project Title	Stretched Rotor: Stretching wind turbine rotors using optimization
Project identification (program abbrev. and file)	EUDP Stretched Rotor, journalnr.: 64015-0067
Name of the programme which has funded the project	EUDP
Project managing company/institution (name and address)	DTU Wind Energy, Frederiksborgvej 399, 4000 Roskilde
Project partners	LM Wind Power
CVR (central business register)	30060946
Date for submission	01.07.2018

2 Short description of project objective and results

The objective of the project was to demonstrate the potential of using state-of-the-art multidisciplinary design tools to develop new passive load alleviating wind turbine blades. The aim was to design a 3 MW rotor with longer blades with a 3% higher annual energy production without increasing loads on the turbine platform. The design tool HAWTOpt2 was developed to meet the needs in the project, and used to design a series of rotors, a subset of which were evaluated in detail according to the design standards. Two final design candidates were selected, the SRB2 and SRB6 blades. SRB2 achieved a significant increase in AEP of 3.2%, but was found to exhibit increases in blade loads for some design load cases. SRB6 was a result of improvements to the design tool, and was more conservatively designed, and only achieved 2.4% increase in AEP, but [pending evaluation] did not exceed design loads of the platform. A detailed design of a manufacturable version of SRB2 was carried out by LM, which achieved an AEP increase of 2.8%, but with slight increase in loads due to higher stiffnesses and mass.

Formålet med projektet var at demonstrere potentialet for at bruge avancerede multi-disciplinære designværktøjer til at udvikle nye passivt aflastede vinger. Målsætning var at opnå 3% forøgelse i energiproduktionen for en 3MW platform gennem forlængelse af vingerne uden forøgelse af platformlasterne. Designværktøjet HAWTOpt2 blev udviklet til at imødekomme behovene i projektet, og anvendt til at designe en serie af rotorer, hvor nogle af disse gennemgik en detaljeret evaluering i overensstemmelse med de-sign standarderne. Two endelige kandidater blev valgt, SRB2 og SRB6. SRB2 opnåede en anseelig forøgelse i AEP på 3.2%, men udviste noget forøgede vingelaster. SRB6 var et resultat af forbedringer af designværktøjerne, og mere konservativt designet, og opnåede 2.4% forøgelse i AEP, men oversteg ikke designlasterne for platformen. Et detaljeret design med fremstillingsbegrænsninger for øje blev udført af LM, som opnåede 2.8% forøgelse af AEP, dog med nogen forøgelse af lasterne på møllen grundet øget stivhed af vingen samt øget vingemasse.

3 Executive Summary

The objective of the project was to demonstrate the potential of using state-of-the-art multidisciplinary design tools to develop new blade technologies for improved passively load alleviating wind turbine blades, which would allow greater energy production, without increasing loads, and therefore reduce cost of energy of the wind turbine. This was to be achieved through collaboration between DTU Wind Energy and LM Wind Power to develop a multidisciplinary blade design framework, and application of this framework to design a series of industrial grade blades aiming to increase AEP of 3% compared to a baseline blade defined by LM. The possibility to bring the designed blade or a specific blade technology to full scale demonstration was built into the project in a commercial milestone where LM would potentially partner with a turbine OEM to realize a full scale demonstrator.

The project was organised as follows:

- WP1: Administration and project management
- WP2: Software Implementation and Validation: Dedicated to development and implementation of the multidisciplinary design tool HAWTOpt2, and alignment with LM design practices.
- WP3: Preliminary Technology Evaluation: Establishment of the turbine design load basis and evaluation of possible load alleviating blade technologies.
- WP4: Concept Development: Blade design studies exploring different load alleviating technologies.
- WP5: Business Case Development: Business case and associated risks will be evaluated for blade designs that incorporate the new technologies.
- WP6: Detailed Design: Detailed blade design that incorporates one or more of the technologies explored in WP3.

Below is a summary of the main results achieved in the project:

- The multidisciplinary blade design framework HAWTOpt2 was developed to meet the requirements of LM. The framework was installed and is operational at LM. HAWTOpt2 enables us to explore advanced blade designs that incorporate passive load alleviation, and includes modelling of key design driving constraints. Compared to other available design tools, HAWTOpt2 is amongst the most sophisticated. This is due to the advanced aeroelastic and structural codes HAWC2, HAWCStab2, and BECAS.
- Implementation of new geometric parameterisation suitable for commercial blades: The parameterisation allows to quickly develop a blade structural geometry according to simple geometric rules. This was used for all designs developed in the project.
- Development of a new material fatigue model: Modern flexible wind turbine blades are often limited by fatigue in the composite materials used in the blades. A numerically efficient model for evaluating fatigue load capacity in the integrated design framework was therefore implemented.
- Implementation of a frequency domain based fatigue loads model: This model allows us to evaluate fatigue loads on the wind turbine in a numerically efficient manner. The model had already been implemented, but was refined and extended in this project to estimate extreme loads in normal operation with extreme turbulence.
- Development of inverse design method for blades: To establish the reference blade used in the project, a method was developed to create a detailed internal structural design based on an aeroelastic definition of the blade. This method solves a feasibility problem to match pre-defined structural characteristics by sizing material thicknesses and internal structural geometry.

- Development of a cross-sectional mesher: The project faced challenges to accurately model the geometric complexity of an industrial grade blade. It was therefore decided to pursue the development of a new cross-sectional meshing tool that could accurately model geometric details not possible with the meshing used. The tool was demonstrated to achieve this goal, but has not yet been incorporated in the design framework.
- Development of a multi-fidelity methodology to increase accuracy of loads constraints: A key challenge in integrated design is to capture the design driving loads correctly, while at the same time keep down computational time. This development aimed at addressing this through a multi-fidelity approach, where detailed load calculations were used in an outer loop to correct a lower fidelity load estimator inside the optimization loop. The methodology was demonstrated on a sample structural design problem and showed promising results but faced challenges in the application to loads correction. The multi-fidelity approach failed for several reasons, first was the challenge of calculating full DLBs within the framework and second errors in the gradients made the correction unfeasible.
- Development of a method to include multiple airfoil families and add-on devices in the integrated design tool: A multidimensional interpolation scheme was developed for including multiple airfoil families, airfoils with add-ons or varying degree of soiling in the optimization. It was demonstrated on a simple example, but not used for the final SRB designs due to time constraints.
- Development of a series of aeroelastically tailored blade designs: HAWTOpt2 was applied to design a number of blades aiming to maximise annual energy production subject to constraints to not exceed the loads of the baseline design.
  - The SRB2 design achieved a 3.2% increase in AEP while also reducing tower bottom extreme loads by up to 10%. However, blade root edgewise and torsional loads increased, and the materials were allowed to be unrealistically thin. Also, the design did not include constraints on material fatigue.
  - The SRB6 design was designed with realistic material thickness constraints, and achieved a 2.4% increase in AEP, including constraints on material fatigue and blade root edge and torsional loads that were now not exceeded.
- LM Blades Interface: LM developed an interface between HAWTOpt2 and their internal design tool LMBlades which was used to develop blade designs based on SRB2 and SRB6 according to LM design standards, materials and airfoil geometries, which was not shared with DTU due to confidentiality.
- Detailed Design: LM carried out a series of detailed designs of blades based on SRB2 taking into account LMs detailed manufacturing constraints and design rules. It was found that the SRB2_lmstruct_rev1 blade achieved 2.8% increase in AEP, but due to slight increases in stiffnesses and mass, minor increases in platform loads were observed. By loosening up some manufacturing constraints to lower the torsional stiffness an AEP increase of 3.4% was achieved. Work needs to be done in order to find a way to realize this.
- Commercial milestone on a demonstration project: LM decided not to pursue a demonstration project directly based on the SRB designs. However, an EUDP project proposal was submitted by DTU Wind Energy and LM Wind Power to design rotors targeting refurbishment of older platforms in the 250 kW to 2 MW range using the design tools and blade technologies developed in this project. The project was unfortunately not granted.

4 Project Objectives

The overall goal of the project was to demonstrate that aeroelastically tailored blades designed using a multidisciplinary design tool could lead to at least 3% increase in AEP compared to a commercial grade blade. This was to be achieved through implementation of the HAWTOpt2 blade design tool that incorporates the state-of-the-art aeroelastic and structural codes HAWC2, HAWCStab2 and BECAS, developed at DTU. The HAWTOpt2 design tool is based on an underlying open-source framework, and therefore allows for users to extend the tool with specific modules or interfaces to in-house tools, making it attractive for commercial users, compared to other closed source alternatives. The idea of the project was thus to leverage the core competences of the two partners into developing and applying a design tool. DTU carries many years of experience developing tool for aerodynamics, structural dynamics and optimization tools. Whereas LM, has years of experience with industrial blade design.

The project was divided into a number of work packages that would allow us to organise the work and track its progress. The project was organised as follows:

- WP1: Administration and project management: Daily management of activities, point of contact between DTU and LM, and organisation of project meetings. The project had a steering committee consisting of managers from DTU and LM.
- WP2: Software Implementation and Validation: Dedicated to development and implementation of the multidisciplinary design tool HAWTOpt2, and alignment with LM design practices.
- WP3: Preliminary Technology Evaluation: Establishment of the turbine design load basis and evaluation of possible load alleviating blade technologies.
- WP4: Concept Development: Blade design studies exploring different load alleviating technologies.
- WP5: Business Case Development: Business case and associated risks will be evaluated for blade designs that incorporate the new technologies.
- WP6: Detailed Design: Detailed blade design that incorporates one or more of the technologies explored in WP3.

Besides the work packages listed above, the project plan included a decision gate for which LM listed a series of requirements for the project to continue.

Overall, the project was a success. The project succeeded in implementing a state-of-the-art aerostructural blade design tool, enabling exploration of advanced blade design technologies. The project almost achieved the overall goal of designing a manufacturable rotor with a 2.8% increase in AEP according to LM detailed design rules, with only minor increases in platform loads. There was no direct commercial exploitation of the blade designs made in the project, but the potential for using passive load alleviation to increase AEP was clearly demonstrated, and can now be matured industrially by LM in further cooperation with DTU and other relevant partners.

The project followed the original plan, but technical and external challenges caused the finalization of the project to be delayed by six months. Some of the challenges are listed below:

Technical Challenges and Unforeseen Tasks:

- The development of the reference blade was planned to be an LM deliverable, based on an already existing blade. It was, however, decided that the blade should be based on openly available material characteristics, and the reference blade structural sizing was therefore carried out by DTU. This incurred several man months of extra work, since a tool for creating an inverse design had to be developed. Also, there were several design details that turned out to be important that was initially not supplied to DTU, which caused more delays.

- Requirements to include material fatigue in the design process: The HAWTOpt2 design tool did not include a model to predict material fatigue loads at the beginning. LM required this to be part of the design constraints, since modern flexible blades are often driven by fatigue constraints. The development of a material fatigue model was therefore prioritized, which incurred additional work not part of the original project plan.
- Fatigue in frequency domain model had to be improved: The method to compute fatigue in frequency domain was planned to be used from the beginning of the project. It, turned out that key constraints in the design of blades was not captured accurately. The model had to be improved and is now significantly more accurate. Additionally, the model is now able to predict extreme loads based on quasi-steady extreme load cases, such as DLC1.3 and DLC6.
- Multifidelity loads model: The multifidelity approach to improving extreme and fatigue loads estimation in an optimization context showed very promising initial results and resources were therefore prioritized to incorporate this model in HAWTOpt2. This, however, turned out to be a too ambitious goal to achieve within the present project, since the development of the model faced several implementational challenges. High-fidelity evaluations of loads were too large for the framework and gradients had too much error to build effect corrections.
- Development of a cross-sectional meshing tool: From discussion with LM it became clear that the finite element meshing tool used in HAWTOpt2 was not sufficient for modelling the geometric complexities in a modern wind turbine blade. It was therefore decided to pursue the development of a new meshing tool that could address these deficiencies. Due to time constraints the meshing tool has, however, not yet been implemented in HAWTOpt2, but this final step is expected to be completed in other projects with the calendar year of 2018 and made available to LM.
- Blade vibrations in the optimized design: The SRB2 blade design went through a detailed design loads evaluation, which revealed edgewise and torsional vibration, causing loads to exceed the design criteria of the platform. Several man months of work went into analysis of this issue which turned out to be a numerical artifact from using large time steps.
- Challenges in prediction of full design loads: During the project, DTU had to invest significant resources into streamlining the workflow around loads computations according to the IEC design standard. This development effort has now resulted in a significantly faster turn-around time including all platform loads and material strength limits.
- Controller for aeroelastically tailored blades: The current controller used at DTU Wind Energy and the workflow to tune controller parameters showed not to be sufficiently robust. This introduced additional challenges in the workflow surrounding the design load basis calculations. Based on this, it can be concluded that on one side there is a need for a higher fidelity controller but also an improved automatic controller tuning. Additionally, future work should also consider including the controller tuning parameters into the optimisation problem formulation.

Externally Generated Challenges:

- General Electic (GE) acquisition of LM Wind Power: During the project LM went through a change of ownership, which required significant resources, which affected this project adversely. This circumstance was one of the reasons for extending the project by six months.
- Change of employment and leave for key persons: Several persons part of the initial project group changed employment during the course of the project. The new members joining the project required training before they could make contributions.

5 Project results and dissemination of results

5.1 List Of Publications

- Zahle, Frederik; Tibaldi, Carlo; Pavese, Christian; McWilliam, Michael; Blasques, José Pedro Albergaria Amaral; Hansen, Morten Hartvig / Design of an aeroelastically tailored 10 MW wind turbine rotor. In: Journal of Physics: Conference Series (Online), Vol. 753, 062008, 2016.
- McWilliam M, Zahle F, Pavese C. Multi-fidelity optimization of horizontal axis wind turbines. In 35th Wind Energy Symposium. American Institute of Aeronautics and Astronautics. 2017.
- Skrzypiński W, McWilliam M, Zahle F. Airfoil Blender for Blade Optimizations. Journal of Physics: Conference Series. 2018;1037(4). 042010. Available from, DOI: 10.1088/1742-6596/1037/4/042010

5.2 The HAWTOpt2 Framework

The development of the HAWTOpt2 wind turbine design framework was initiated in the EUDP financed project "Light Rotor" (proj.nr. 64010-0107), with continued development in other projects, such as the EU-funded INNWIND project. Within the present project the HAWTOpt2 code base was to a great extent re-factored and improved to meet the requirements of the project, with some of the key improvements being driven by requirements made by the industrial partner LM Wind Power. This section provides an overview of the framework.

Figure 5.1 shows an overview of the HAWTOpt2 workflow, with the analysis codes framed in green boxes along the diagonal, connected by their data dependencies with names framed in in grey boxes.

One of the key design choices of the framework has been to carry out the aerostructural design of a blade in a fully coupled monolithic manner. The sensitivities for design objectives and constraints are computed for all design variables within the same scope. The advantage of this approach is that cross-couplings between disciplines are captured naturally, which leads the optimization towards bend-twist coupled blades naturally. The disadvantage is the high computational cost because all models of the work-flow must be executed for each function and gradient evaluation.

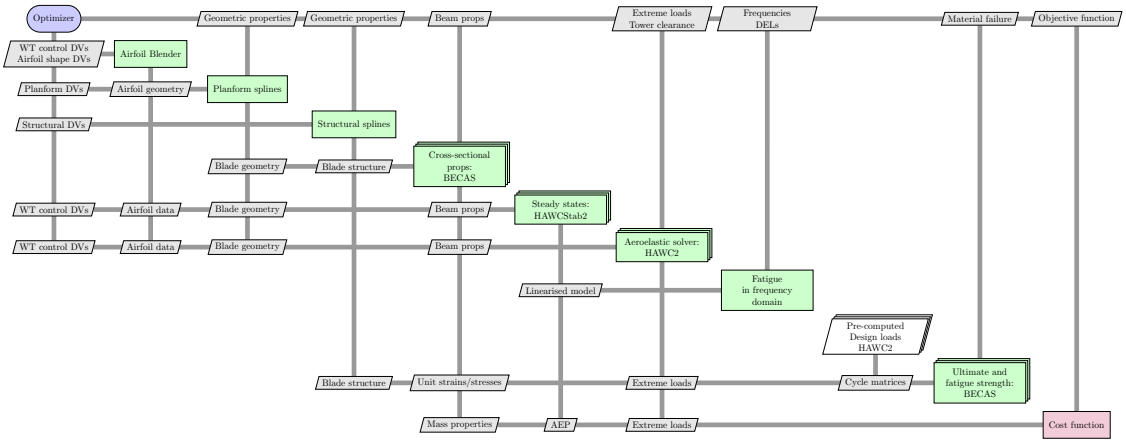

Figure 5.1: Extended Design Structure Matrix diagram of the workflow of HawtOpt2.

New Backbone Framework: OpenMDAO 1.x

HAWTOpt2 is based on the open-source systems engineering framework OpenMDAO[4] developed by NASA. It is used for defining the overall workflow and data transfer between models. It also interfaces to various optimization codes. OpenMDAO underwent a complete rewrite of its code and API prior to the present project, forcing HAWTOpt2 to be adapted to these changes. This has allowed us to take advantage of the MPI based parallelization made available in the updated OpenMDAO. This is used within a single execution of the workflow for the multiple cases (*e.g.* load cases, cross sections, *etc.*) and also parallel evaluation of finite-difference gradients (*e.g.* which requires a full execution per design variable). Being able to efficiently utilize high-performance computers has allowed us to reduce the computational time of optimizations by orders of magnitude.

OpenMDAO has an interface to the Python based framework PyOptSparse [3], which is used to interface to the two optimization algorithms SNOPT and IPOPT used in HAWTOpt2. The former is only available through a paid license, which made it unsuitable for the present project. IPOPT was therefore used exclusively in Stretched Rotor and has since also been adapted for most optimization work involving HAWTOpt2.

Geometric Parameterisation

The blade planform is described in terms of distributions of chord, twist, relative thickness and pitch axis aft leading edge, the latter being the distance between the leading edge and the blade axis. The lofted shape of the blade is generated based on interpolation of a family of airfoils with different relative thicknesses. Alternatively the shape can be defined by airfoils assigned to different spanwise positions.

The internal structure is parameterised according to the schematics in Figures 5.2 and 5.3. To define the structural reference plane, the blade is firstly rotated by the so-called structural mold angle, defined positive according to the right hand rule around an axis normal to and centered around the blade root center. The vertical reference plane is then formed that passes through the root center and the blade tip. Note that the structural mold rotation is applied on the lofted blade that includes both aerodynamic twist and possible prebend and sweep. The spar caps are then placed on the suction side (SS) and pressure side (PS) with a given horizontal offset with respect to the reference plane, and with a given width. Likewise, the shear webs are placed vertically with a given offset with respect to the reference plane. For blades with sweep, another method for placing the spar caps and shear webs is possible, in which these are positioned relative to the swept twist axis of the blade. The width of the trailing edge reinforcement is defined relative to the suction and pressure side trailing edge points, respectively, and the leading edge reinforcement is placed relative to the leading edge. The leading edge is defined as the point along the surface with maximum distance from the trailing edge.

The materials are placed in regions bounded by division points (DPs), as shown in Figure 5.2, within which a given stacking sequence of materials is used. Along the span the individual material thicknesses in regions can then vary and are allowed to have zero thickness. Although regions by definition extend from root to tip of the blade, a region can be removed by collapsing its associated DPs into a single point.

Computation of stiffness properties

The cross-sectional structural properties are computed using BECAS (BEam Cross-section Analysis Software), developed at DTU [1, 7, 6]. BECAS can output fully populated stiffness matrices, including the effect of off-axis material layups, but can also output stiffness properties directly in the standard HAWC2 format. The cross-sectional mesh used in BECAS is generated using Shellexpander, also developed at DTU. A pre-processor for Shellexpander, CS2DtoBECAS, translates the parameterised blade structure into Shellexpander inputs, and checks its consistency, e.g. removing zero-thickness materials and zero-width regions. The computations of all the cross-sections in the blade are parallelised using OpenMDAO's ParallelGroup. BECAS is written in Matlab and a file interface to this code has therefore been made. Since a typical optimization run issues hundreds of simultaneous calls to BECAS across the parallel set of compute nodes, Octave is used rather than Matlab.

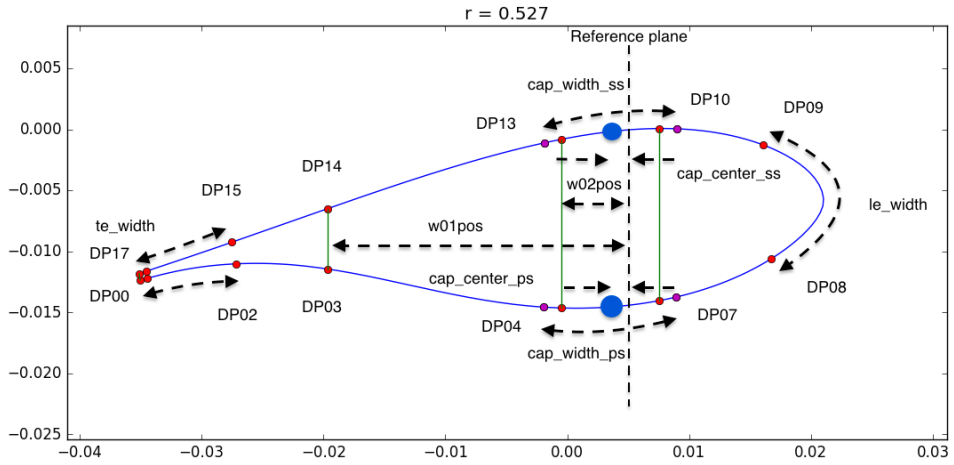

Figure 5.2: Schematic showing the geometric parameterisation of the blade structure used in the HAWTOpt2 blade parameterisation.

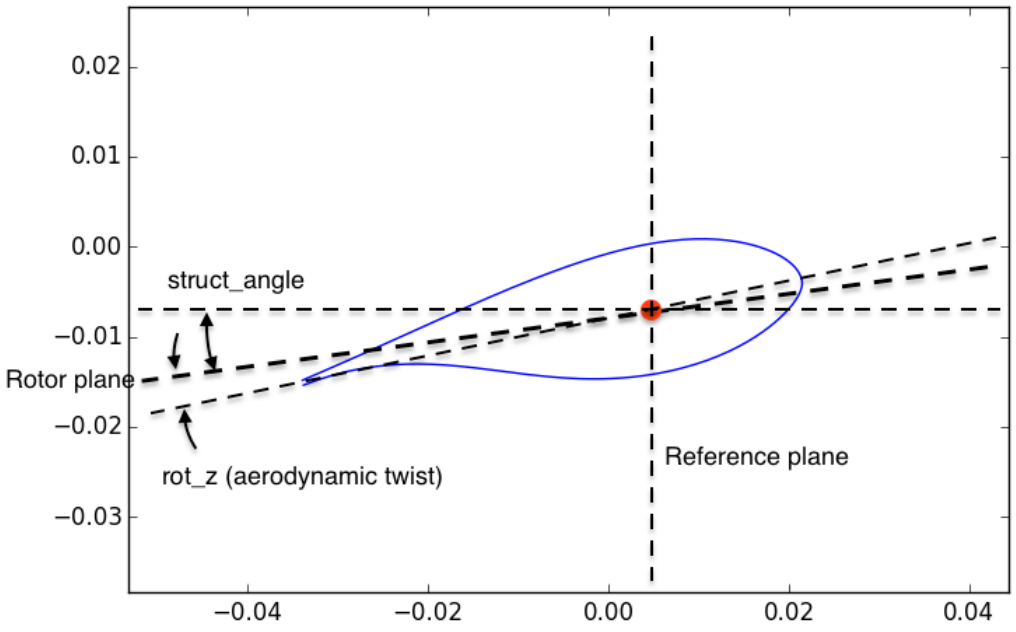

Figure 5.3: Schematic showing the definition of the structural angle used in the HAWTOpt2 blade parameterisation.

Computation of Aeroelastic Steady States

HAWCStab2 is an aeroelastic stability tool developed at DTU Wind Energy, and is in HAWTOpt2 used to compute the aeroelastic steady states of the turbine (loads, deflections and power production), aeroelastic frequencies and dampings. HAWCStab2 is developed to carry out stability analysis of wind turbines, and in an optimization context, the detailed outputs of HAWCStab2 can be used as design constraints on the

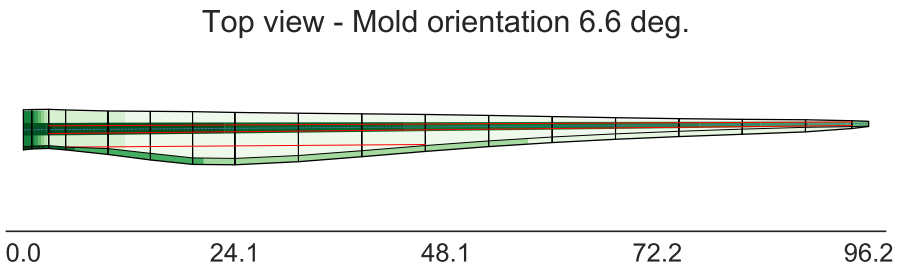

Figure 5.4: Topview of a blade showing internal structural layout.

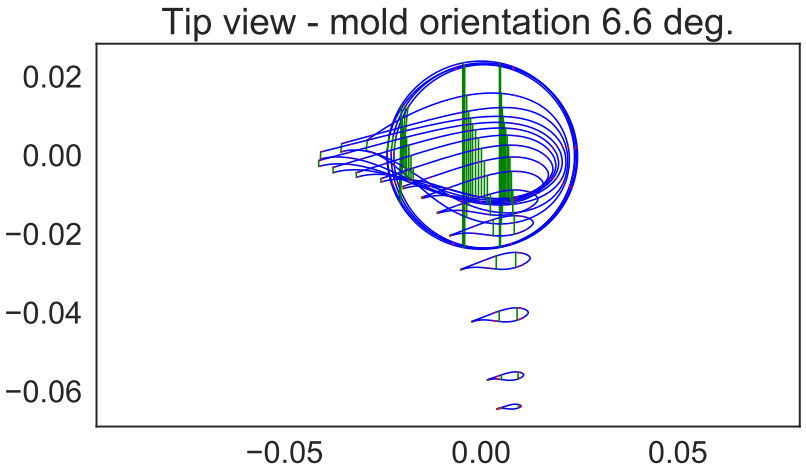

Figure 5.5: View from the tip of a blade showing blade sections with division points of regions and shear webs.

blade to e.g. constrain certain blade frequencies, ensure sufficient aeroelastic damping etc. Additionally, the linearised model computed in HAWCStab2 is used as input to the fatigue in frequency domain tool, described in Section 5.3. The rotor operational points are computed based on a simple single point optimization, tracking maximum C_P, subject to constraints on power and minimum pitch setting. Work was also done to in an efficient manner allow to include the turbine pitch and RPM schedule as design variables in the global optimization problem to allow for e.g. peak shaving strategies to reduce loads to be included in the full aerostrutural design problem.

Computation of Time Domain Aeroelastic Loads

One of the most critical aspects of an aerostructural design tool is the computation of ultimate loads, since these are used to size the internal structure of the blades. In HAWTOpt2 the loads are computed using the HAWC2 code, which is a time domain aeroelastic tool developed at DTU Wind Energy. Ideally, the turbine loads would be computed based on a full set of IEC-64100 design load cases, however, this is not computationally feasible. Instead, we use a very simplified approach with a set of load cases consisting of only gust cases without inflow turbulence. While the loads computed with this approach are not sufficient for a full design evaluation, they provide sufficient information about relative changes between the initial design and the optimized design. A key deficiency of having only approximate relative loads, is that the starting point of the design optimization has to be a well-defined, feasible turbine design. Another limitation of having only approximate loads is that additional safety factors have to be placed on hard

constraints such as blade tip tower clearance and material strength. As such the area of time domain load calculations for optimization purposes is still an area needing more research and improvements.

Computation of Fatigue in Frequency Domain

Since the current workflow only includes simplied time domain simulations using HAWC2 without inflow turbulence, it is not possible to compute fatigue loads based on these simulations. An alternative approach was implemented, which uses the linearised model that is an output from HAWCStab2 to compute fatigue loads in frequency domain. This methodology has the advantage that is is faster than HAWC2 but still retains sufficient fidelity for use in a blade design context. This model is described in more detail in Section 5.3.

Computation of Ultimate and Fatigue Loads Capacity of Materials

To evaluate the ultimate load capacity of the materials in the blade, two options are available in HAWTOpt2: Either the methods implemented in the BECAS tool or a more recently added pure Python version. The latter method is now used by default, since it is significantly faster than the BECAS method, due to the overhead of calling Octave. The Python stress recovery version uses the unit stresses and strains computed by BECAS as part of the computation of the stiffness properties, and is therefore only a simple matrix vector multiplication, that is numerically very efficient.

Another important constraint on blade designs is the material fatigue load damage, which at the initiation of the present project was not implemented in HAWTOpt2. Three methods were implemented which are described in detail in Section 5.3.

As to the implementation of the material level fatigue models in HAWTOpt2, the overall challenge was that time domain turbulent flow load cases were not included in the workflow due to prohibitive computational cost. The implemented workflow therefore involved pre-computed loads based on the initial design that were kept frozen throughout the optimisation. This approach is sub-optimal, since any improvement in the blade design leading to reduction of loads will not be reflected in the resulting constraint on material level fatigue damage. This leads to designs that are conservative with respect to fatigue damage.

Regarding the choice of material level fatigue damage method, the cycle matrix of strains method proved computationally too heavy to use, since it involves computing strains in the cross-section for every time step of each of the the load calculations. The best alternative was, as clearly shown in Section 5.3, the "Cycle Matrix of Loads" approach, which gives reasonably good correspondence to the cycle matrix of strains method, but at a significantly lower computational cost.

With the recent further development of the fatigue in frequency domain method (Section 5.3), it is now possible to compute the cycle matrix of loads based on the linearised model in HAWCStab2. While this will lower the fidelity of the load calculation itself, it has been shown that this method is sufficiently accurate to use in an optimisation context. Using this approach, the sensitivity to design changes leading to load alleviation would be directly reflected in the material level fatigue constraint, and thus allow to optimise blades where material level fatigue is a design driving constraint. Within the current project this development was not implemented due to time constraints, but will be an area of future work.

5.3 Computation of fatigue and extreme loads using a linear model

5.3.1 Introduction

Multidisciplinary design optimization of rotor blades requires fast methods for fatigue and extreme load evaluation that can be performed inside the optimization loops. In the context of the stretched rotor project, the results from these methods are necessary to evaluate if a change in design variables leads to an increase of loads beyond the permitted load envelope. The goal of the present work is to further develop a method for fatigue estimation based on a HAWCStab2 linear turbine model, [13]. The original method did not include the blade root torsional moments and had limited accuracy on the blade edgewise fatigue loads of the NREL 5 MW reference turbine [11], see Figure 5.6. The further developments presented in this report are intended to 1) improve the blade root edgewise fatigue estimation 2) add the blade root torsion sensor and 3) enable extreme load estimation based on a linear model.

Potentially these additions to the method can remove the HAWC2 runs in the optimization loop and lead to more accurate extreme loads and fatigue loads. Ultimately this results in less conservative blade designs and thus a reduced levelized cost of energy.

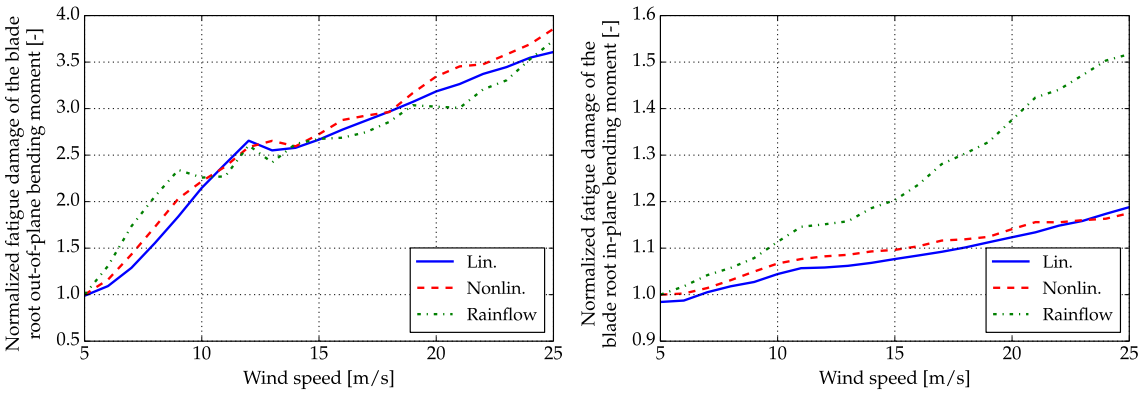

(a) Out-of-plane. *(b) In-plane.*

Figure 5.6: Trends of blade root out-of-plane and in-plane fatigue damage reported in [13] for the NREL 5 MW reference turbine [11]. The plots include results from HAWC2 using rainflow counting (denoted *Rainflow*), results from HAWC2 using a PSD load estimation method (denoted *Nonlin*) and results from the linear turbine model in HAWCStab2 (denoted *Lin*) using the same PSD load estimation method. The PSD method cannot predict the increasing loads with windspeed that are observed in the rainflow count for the in-plane loads.

5.3.2 Method

An overview of the method is shown as a block diagram in Figure 5.7. The Figure is an extended version of the diagram in [13]. The part developed in the last part of the stretched rotor project is indicated as 'new option'.

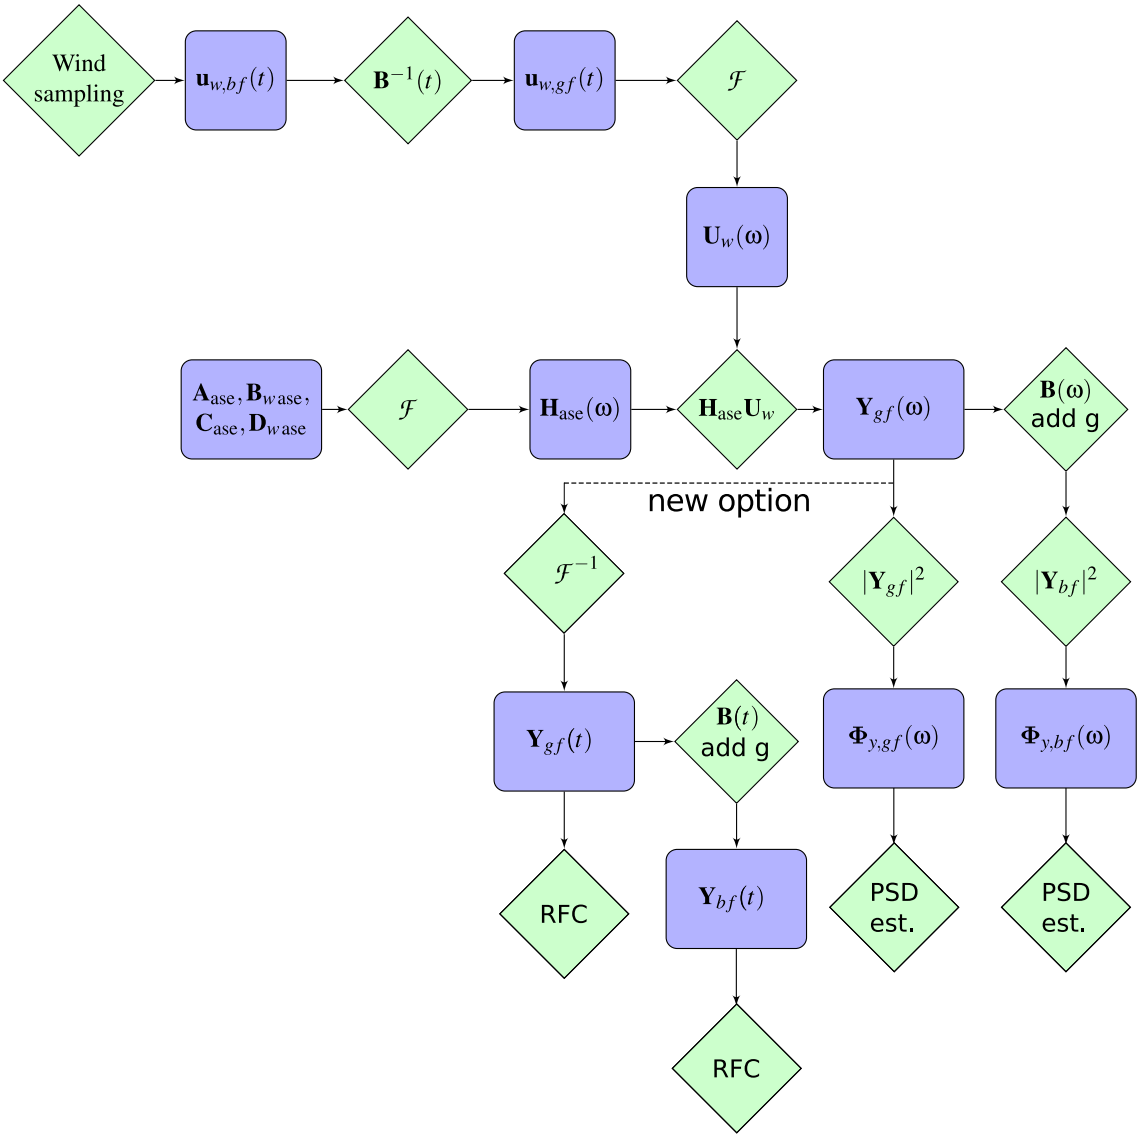

Figure 5.7: Diagram of the workflow in the frequency domain based fatigue estimation. The diagram is based on a Figure in [13], but extended by the *new option* that enables rainflow counting and determination of extreme loads from time series output.

Prior to the frequency domain fatigue estimation, a sampled turbulent wind in the blade fixed reference frame, denoted as $u_{w,bf}$, has to be obtained, see the top left of Figure 5.7. This sampled wind is computed by running time domain simulations of the Reference turbine. In these time domain simulations the turbine operates normally, which means that the rotor speed changes during the simulation. This has an influence on the sampled wind speeds along the radial stations on the three blades, such that the 1P peak due to shear for example is wider. The sampled wind is then inverse Coleman transformed from the blade fixed to the ground fixed frame and Fourier transformed into the frequency domain. The frequency domain wind input U_w multiplied with the aeroservoelastic transfer function H_{ase} yields the output in the ground fixed frame, Y_{gf}.

In the original implementation [13], this output Y_{gf} would be used to estimate the fatigue based on the corresponding PSD $\Phi_{y,gf}$. For the blade sensor outputs, a Coleman transformation has to be performed and the gravity contribution has to be added before estimating the fatigue loads. Both the Coleman transformation and the addition of the gravity contribution (only edgewise) are performed in frequency domain in [13], using an average rotor speed for each respective wind speed. Computing the PSD from the output removes the phase information, which can reduce the quality of the results, see [12].

An alternative method has been developed to avoid neglecting the phase information. The method is denoted as *'new option* in Figure 5.7 and uses rainflow counting instead of PSD estimation methods to compute the damage equivalent loads from the output sensors. Instead of computing the PSD, the output in ground fixed frame is inverse fourier transformed to obtain output time series $Y_{gf}(t)$. To obtain time series of the blade fixed sensors the coleman transformation is applied in the time domain. To ensure consistent phases of the rotor position with repect to the sampled wind input the time series of the azimuth angle is saved during the wind sampling and used in the coleman transformation. The computations of the gravity contributions on edge, flap and torsion also use these azimuth angles, along with the corresponding time series of the pitch angles. Especially the torsional gravity contribution is not straight forward. The approach to compute the contributions on flap, edge and torsion is described in [12]. Once the sensor time series are computed, a rainflow count can be performed to compute the DEL. An advantage of this method is also that the minima and maxima of the time series can be computed which can be used to estimate the extreme loads. For this purpose, sampled wind according to DLC1.3 can be used. Because the expensive part of the frequency domain fatigue method is computing the transfer functions, additional DLC1.3 wind input does not increase the computation time significantly.

5.3.3 Results

To validate the improved model, a detailed comparative study between the linearised model and the non-linear HAWC2 model was carried out comparing predictions of fatigue and extreme loads for Reference Blade and SRB2. The results are obtained using only rainflow counting methods on either HAWC2 or HAWCStab2 time series, because the load estimations from the PSD methods cannot be quantitatively compared with rainflow counting results. Further, the PSD methods do not yield extreme load results. The fatigue load results use DLC 1.2 input (normal turbulence), the extreme loads use DLC 1.3 input (extreme turbulence). The comparisons also include results obtained from HAWCStab2 using the sampled wind and controller tuning from the Reference Blade for the SRB2 computations. These results would be obtained in the optimization loop, where the sampled wind and controller tuning are not updated while the blade is optimized.

The results for DEL and extreme loads of the tower bottom fore-aft moment show that similar trends in the DEL are predicted by HAWC2 and HAWCStab2, but there is a offset by 15-20% at high winds. In this region HAWCStab2 under-predicts the loading. Both codes predict lower loading for the SRB2 blade compared to the Reference blade at high wind but the difference is exaggerated in the HAWCStab2 results. The impact of using Reference blade sampled wind and controller tuning for the SRB2 computations in HAWCStab2 is small. This observation holds for all sensors.

The minimum loads are under predicted at high wind speeds by HAWCStab2. The maximum loads at low wind speed are over predicted because the controller can't switch between regions and start pitching if the wind speed is temporarily increased. There is a peak in the minimal and maximal HAWC2 loading at 20 m/s that can't be captured in HAWCStab2. Otherwise both codes agree that the extreme loads are similar for Reference blade and SRB2.

The extreme loads for all 18 wind seeds shows a peak in the HAWC2 results at 20 m/s, present only in 2 seeds. Below rated the HAWCStab2 results are more spread than the HAWC2 results. This is partially due to the fixed controller region and in part due to the constant lift gradient in HAWCStab2: The effect of an extreme gust is not limited by stall as it would be in HAWC2.

The tower bottom side-side DEL and extreme loads are more difficult to predict than the fore-aft loads due to the low aeroelastic damping in the lateral direction. HAWCStab2 over predicts the damage equivalent loads significantly, and it also over predicts the difference in loads between Reference blade and SRB2. Both codes agree on lower loads for SRB2 above rated though. The extreme loads are greatly over predicted by HAWC2 except for wind speeds around rated and close to cut out.

The tower bottom torsion loads show good agreement on the DEL between HAWC2 and HAWCStab2, both in the level and in the trends. The extreme load reduction due to the SRB2 blade with passive load alleviation is also well predicted by HAWCStab2.

Both HAWC2 and HAWCStab2 predict an edgewise load increase for SRB2 at low wind speeds and a load reduction at high wind speeds. A part of the load increase at low wind speeds is only visible if the correct sampled wind and controller tuning is used in SRB2, which indicates that it is due to the increased tip speed ratio and thus the higher 1P frequency. At high wind the aerodynamic loading becomes more important and the load alleviation built into SRB2 can reduce the loads. The difference between the two blades at high wind is over predicted by HAWCStab2. The HAWCStab2 peak in DEL at 8-10 meters/s is also visible in the extreme loads. For Reference blade the edgewise loads are well captured when the rainflow counting method is used. The PSD estimation underestimates the edgewise load increase with wind speed, similar to what has been observed for the NREL 5 MW, see Figure 5.6. Results from the PSD estimation method are omitted here for brevity but can be found in [12].

HAWC2 and HAWCStab2 predict similar blade root flap loads. The main difference in the DEL is that HAWC2 predicts a very slight increase in the loading when changing the blade from Reference blade to SRB2 at 4-8 m/s where HAWCStab2 predicts a small decrease. The extreme loadings from HAWCStab2 are shifted towards positive loading and the difference between SRB2 and Reference blade is exaggerated slightly.

For the blade torsional DEL and extreme loads HAWCStab2 predicts exaggerated DEL and extreme loads at 8 and 10 m/s, especially for SRB2. The codes agree on lower DEL for SRB2 at high wind speed and higher DEL between 6 and 12 m/s. HAWC2 predicts larger absolute extreme loads for SRB2 at 24 and 26 m/s.

5.3.4 Conclusions

A frequency domain fatigue estimation method has been modified and extended to compute the blade root torsional fatigue load and improve the edgewise fatigue load prediction. The extended method can be used for extreme load computation in an optimization loop at very little added computational cost.

Because the method uses a linear wind turbine and controller model there are some inherent limitations regarding mainly the linearized controller behavior and the linearized airfoil polars. A comparison of time series generated from the linear turbine model showed generally good agreement with results from time simulation for most sensors, but there are deviations due to the controller behavior around rated and in the tower side-side and blade root torsion moments, [12]. The tower side-side loading is difficult to model due to the low aeroservoelastic damping of the tower side-side mode. The blade root torsion moments are nonlinear due to the influence of the blade deflection on the gravity contribution.

It is shown that previously observed deviations in the edgewise damage equivalent load have been due to limitations of the PSD fatigue estimation method, [12]. Performing rainflow counting on the linear ouput from the HAWCStab2 model improves the edgewise loads dramatically for moderately flexible blades. For more flexible blades the agreement still needs to be improved, for example by dynamically computing the edgewise gravity moment which was assumed constant in the present study.

The computation of the torsional gravity moment has to be improved. The edgewise, flapwise and torsional gravity moments could be added as an output to HAWCStab2 to improve the accuracy of the method. However such linearized gravity moments will still suffer from the linearized controller behaviour at and just below rated wind speed.

Overall the method is very promising for both fatigue and extreme loads but future work is needed to improve mainly the results on the tower side-side, blade edgewise and blade torsional sensors.

5.4 AMMF

One of the challenges of design optimization is the balance between modelling the problem with sufficient fidelity and running the calculation efficiently. Multi-fidelity methods try to achieve the best of both aspects by using the high-fidelity models occasionally as a correction, but performing a large portion of the calculation with faster low-fidelity models. Multi-fidelity design optimization methods were explored in this work to determine whether they could provide an advantage.

There are several approaches in literature, many of which rely on global surrogate methods. A problem with global surrogate methods is they require sampling within the design space to build surrogates. The number of samples required to construct a model of sufficient accuracy grows exponentially with increasing number of design variables. Typically these approaches are only practical for problems with up to 10 design variables. The problem in this work had 50 design variables, thus these methods are not practical.

Instead this project looked at the Approximation and Model Management Framework (AMMF) algorithm. This is basically a multifidelity version of standard gradient based algorithms. Its strength is that it corrects both the function and gradient values. The latter is required to properly converge optimization algorithms because the stationary conditions are defined by the gradients of the objective and constraints. The algorithm is described in greater detail in section 5.4.1.

In the development phase, the algorithm was tested on a simple structural optimization problem. This study is given in section 5.4.2. The early results showed promise and several attempts were made to apply the algorithm to the more difficult problem of evaluating aeroelastic loads in typical wind turbine Design Load Cases (DLCs). The results from this are given in section 5.4.3.

Finally, this section closes with giving recommendations on future work in section 5.4.4.

5.4.1 AMMF Algorithm

One attractive multi-fidelity approach is given by Alexandrov *et. al.* called Approximation and Model Management Framework (AMMF)[5]. It is a general approach that assumes there is an expensive high fidelity simulation and a fast low fidelity approximation. The goal is to find the optimal solution for the high fidelity response, yet save computational time by performing the search with the low fidelity analysis (see figure 5.8). This approach is similar to that of Fuglsang *et. al.* and Bottasso *et. al.*, but AMMF differs in 2 ways. First the low fidelity simulation (f_l) is corrected with either a multiplicative correction (equation (5.1a)) or an additive correction (equation (5.1b)), where $\tilde{f}$ is the corrected response of the low-fidelity solution, β is the correction, which are all functions of the design vector $\boldsymbol{x}$. The corrected response is made equivalent to the high-fidelity response $f_h(\boldsymbol{x})$ in both the value (*i.e.* $\tilde{f}(\boldsymbol{x}_0) = f_h(\boldsymbol{x}_0)$) and the gradient (*i.e.* $\nabla\tilde{f}(\boldsymbol{x_0}) = \nabla f_h(\boldsymbol{x_0})$) at the reference point $\boldsymbol{x_0}$ using the correction in equation (5.1c) for β_m or equation (5.1d) for β_a.

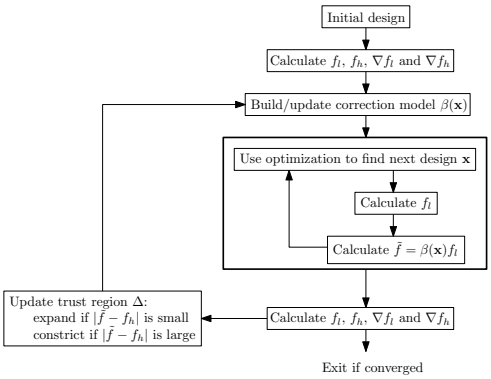

Figure 5.8: The AMMF algorithm

$$\tilde{f}(\boldsymbol{x}) = \beta_m(\boldsymbol{x}) f_l(\boldsymbol{x}) \qquad (5.1a) \qquad\qquad \tilde{f}(\boldsymbol{x}) = f_l(\boldsymbol{x}) + \beta_a(\boldsymbol{x}) \qquad (5.1b)$$

$$\beta_m(\boldsymbol{x}) = \left(\frac{f_h(\boldsymbol{x}_0)}{f_l(\boldsymbol{x}_0)}\right) + \left(\frac{\nabla f_h(\boldsymbol{x}_0)\cdot(\boldsymbol{x}-\boldsymbol{x}_0)}{f_l(\boldsymbol{x}_0)} - \frac{f_h(\boldsymbol{x}_0)\nabla f_l(\boldsymbol{x}_0)\cdot(\boldsymbol{x}-\boldsymbol{x}_0)}{f_l^2(\boldsymbol{x}_0)}\right) \tag{5.1c}$$

$$\beta_a(\boldsymbol{x}) = (f_h(\boldsymbol{x}_0) - f_l(\boldsymbol{x}_0)) + (\nabla f_h(\boldsymbol{x}_0)\cdot(\boldsymbol{x}-\boldsymbol{x}_0) - \nabla f_l(\boldsymbol{x}_0)\cdot(\boldsymbol{x}-\boldsymbol{x}_0)) \tag{5.1d}$$

Second, a trust region is employed to prevent errors in the low-fidelity model taking the optimization away from the high-fidelity minimum. Since the approximation is only based on local information, it will only give acceptable estimates close to the reference point. To prevent the optimization from moving into regions with poor approximations, the inner optimization is constrained to stay within a trust region (through boundary constraints). When the inner optimization finds a suitable solution, the trust region is re-centered on that solution and possibly resized. Then another approximation is constructed and the inner optimization proceeds in the new trust region. Figure 5.9 shows a hypothetical example of this where the colored squares are the trust regions for different iteration of the AMMF algorithm and the red line shows the optimization progress starting from the bottom left.

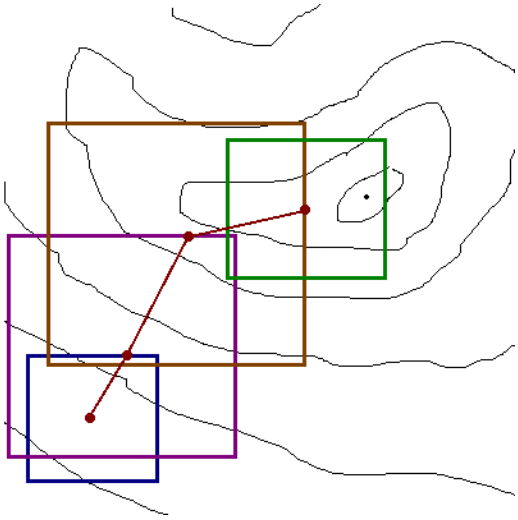

Figure 5.9: Example of the trust region algorithm

Extending the multi-fidelity approach to the constraints is an important aspect of wind turbine design engineering. In this research the constraints are corrected in the same way as the objective using equations (5.1). The constraints are incorporated into the AMMF algorithm by using the Lagrangian shown in equation (5.2), where $\boldsymbol{g}_e$ and $\boldsymbol{g}_i$ are the equality and inequality constraint respectively and $\tilde{\boldsymbol{\lambda}}_e$ and $\tilde{\boldsymbol{\lambda}}_i$ are the respective Lagrange multipliers. These Lagrange multipliers are estimated from the high fidelity gradient information from the most recent gradient information. Since, at the initial stages when when not all constraints are active, the Lagrange multipliers may be zero. To avoid this the minimum Lagrange multiplier is specified by the user. The merit of the inner optimization solution is based on improving Φ. Similarly Φ is used in equation (5.3) to calculate the merit M of the approximation. That information is used to either expand, contracts or maintain the trust region size, depending on how far it is from 1. In this work the the trust region was doubled in size if M was between 0.75 and 1.25, or contracted to half the original size if it was under 0.1.

$$\Phi = f_h + \tilde{\boldsymbol{\lambda}}_e^T|\boldsymbol{g}_e| + \tilde{\boldsymbol{\lambda}}_i^T \max(\boldsymbol{0}, \boldsymbol{g}_i) \tag{5.2}$$

$$M = \frac{\Phi_{i-1} - \Phi_i}{\Phi_{i-1} - \tilde{\Phi}_i} \tag{5.3}$$

5.4.2 Structural Optimization

A common application in wind turbine design optimization is in the design of the internal structure. One approach is to perform a set of aeroelastic calculations on a design. The results of this calculation define the loading on the blade. Then optimization is used to find the structure with the minimum weight that can withstand these loads. By assuming the aero-elastic loads are constant, the optimization only required Finite Element Method (FEM) analysis of the internal structure. Additional aero-elastic calculations would have to be performed as the design evolves to update the loads.

The AMMF algorithm was applied to this problem. The high fidelity model was based on BEam Cross section Analysis Software (BECAS)[8] for the cross section stiffness properties and linear beam model for the deflection. The low fidelity model is based on Classical Laminate Theory (CLT) for the cross section stiffness properties and the same linear beam model for the deflection.

Results from a full DLC evaluation were used to determine the deformed state of the blade when the maximum tip deflection occurred. Then an inverse problem was solved to determine the equivalent static loads to achieve the equivalent deformed state. The goal of the design optimization was to minimize the weight of blade by varying the spar-cap thickness without increasing the tip deflection when subjected to these loads. Gradient based optimization was used to solve this problem, finite difference gradients were used in all cases.

Figure 5.10 shows the optimization results. AMMF was compared with direct optimization with both CLT and BECAS. The results show that optimization based on CLT is not conservative and would lead to unsafe designs. Thus, one would have to apply larger safety factors which in-turn would lead to sub-optimal designs. AMMF managed to achieve identical results as direct optimization with BECAS. Thus, it was successful in incorporating the extra high-fidelity information to achieve a feasible and optimal design.

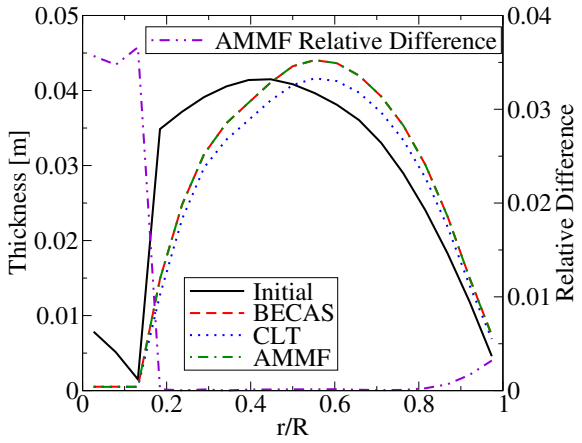

Figure 5.10: Structural Optimization Results

Figure 5.11 shows the progress of the AMMF algorithm in comparison with direct optimization with BECAS. It is clear that AMMF was able to converge the results much fast than with BECAS. In this example AMMF converged in 2 outer iterations, while direct BECAS required approximately 30 iterations. This shows that only a small amount of high-fidelity information is required to achieve equivalent results.

The first order correction in AMMF introduces increasing error moving away from the reference point. To compensate for this error the AMMF algorithm employs 2 features. First, a trust region is used so the optimization stays within a domain where the approximation is acceptable. Second, the Lagrange multipliers are used to penalize constraint violations when the AMMF algorithm evaluates whether an optimization solution should be accepted or not.

This problem was used to explore how these features improve the robustness of the optimization. Figure 5.12 shows how the optimization fails to converge when both of these features are disabled. Since CLT is not conservative the optimization makes a large jump into an infeasible region and cannot recover. Each of the robustness features was enabled and in both cases the optimization was able to converge. The results show that only one of these features needs to be enabled to achieve robust convergence. Thus, if the user

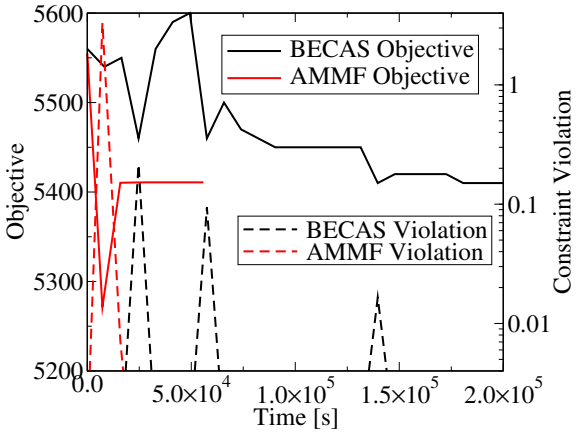

Figure 5.11: Structural Optimization Progress

does not know a suitable initial trust region they can use a large initial penalty. Otherwise if a good initial trust region is known, then the initial penalty is not important and can be set to a small value and subsequently calculated by the AMMF algorithm.

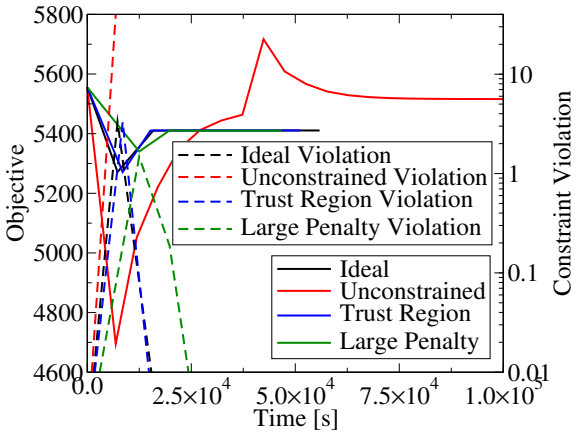

Figure 5.12: Robustness of the AMMF Algorithm

Both additive and multiplicative corrections were applied in this study. Both corrections gave similar results and show no significant advantage over eachother.

5.4.3 Application of AMMF for evaluating wind turbine loads

To apply the AMMF algorithm, one needs to choose both a low and high fidelity model. The high-fidelity model represents the desired fidelity for the optimization, then the low-fidelity needs to be chosen such that it runs significantly faster and still capture most of the trends. Many attempts were made to use the AMMF algorithm to correct the loads of the wind turbine with different low and high fidelity pairings.

In the first attempt, a simple load extrapolation scheme was chosen that simply scaled the steady state loads to get a dynamic equivalent. This model would run significantly faster than HAWC2, so it has a lot of potential for accelerating the optimization. However, it was found that it was so simple that it did not adequately capture the trends. So the optimization would still rely heavily on the high fidelity simulations. So this work was abandoned.

The results from the first study showed that the two models should be similar. Currently, only HAWC2 is incorporated into the framework, so only HAWC2 can be used for both fidelities. Instead of varying the fidelity on the simulation front, it was decided that it should be varied in the load cases used in the

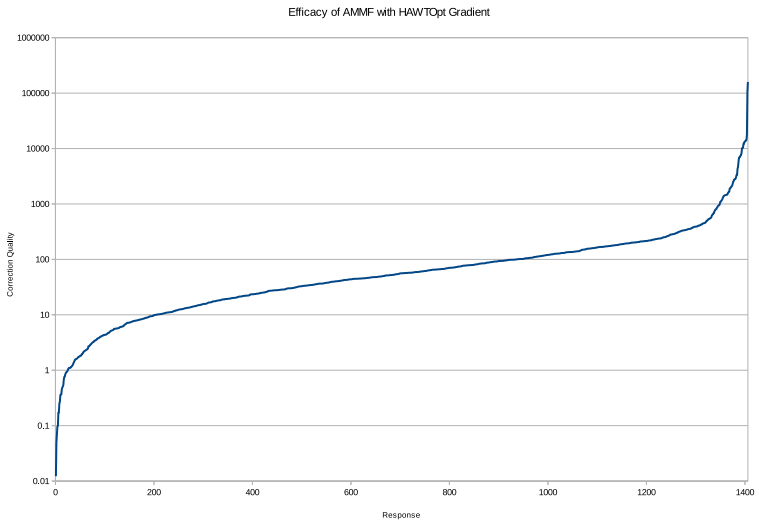

Figure 5.13: The quality of the AMMF correction with HAWTOpt2 gradients

optimization. To have a large difference in execution speed, the full DLB was chosen for the high fidelity simulation, then the simplified DLB used in the optimization is chosen for the low fidelity. This would give designs that satisfied the IEC design standards. However, this proved to be problematic because the full 61400 design load basis requires $\approx$ 1400 simulations. The tool chain could not execute such a large number of cases. This prompted additional development on the tools to enable these large calculations. The tools were improved, however the new code still failed to calculate a full DLB. Further development to solve this issue was beyond the scope of this project, so this work was abandoned.

In the final attempt, DLC 1.3 with turbulence was chosen for the high fidelity. While the simplified DLC 1.3 was chosen as the low fidelity. In the low fidelity model, there is no turbulence, instead a deterministic gust is used to generate extreme loads.

The first step was to evaluate how a corrected low fidelity model would match the high fidelity model. The study was based on the optimization of the DTU 10 MW RWT. The correction was generated for the initial point with the gradients as calculated by the HAWTOpt2 framework. Fortunately, this turbine has been optimized in this framework, thus several designs can be generated between the initial point and the optimum design. Both the corrected low-fidelity and the high-fidelity versions were executed for each of these points. The differences between these solutions determines the quality of the correction.

The quality of the correction was evaluated with equation (5.4), where f_h is the high-fidelity response, f_l is the low fidelity response and f_c is the low fidelity response corrected with the AMMF correction. This equation basically give the variance of the corrected response, compared to the uncorrected response. A perfect correction would give a quality of 0, and a quality of 1 means the corrected response is just as good as the original low-fidelity model, while values greater than 1 indicate that the correction is worse than the original low-fidelity model.

$$Q = \frac{\sum (f_h(\boldsymbol{x_i}) - f_c(\boldsymbol{x_i}))^2}{\sum (f_h(\boldsymbol{x_i}) - f_l(\boldsymbol{x_i}))^2} \tag{5.4}$$

There are several terms that must be corrected, so all the quality results are sorted and graphed in figure 5.13. It is clear that the corrected response is worse for most of the responses than the uncorrected model.

Figure 5.14 shows the signal that had the best AMMF correction. It is clear that this signal is quite simple to correct and that AMMF does an excellent job of correcting it. Thus, in ideal circumstances, AMMF could be an effective tool.

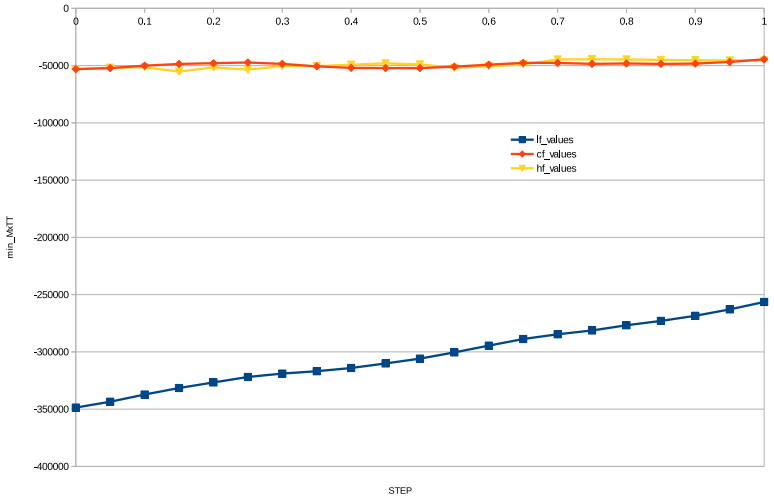

Figure 5.14: The best example of an AMMF correction

Figure 5.15 shows the signal that had the worst AMMF correction. It is clear that the gradient correction is dominating and the corrected signal bears little similarity to the original signal. This suggests that errors in the gradients are effecting this result.

Figure 5.16 shows another example where the AMMF correction did not provide any benefit. This example shows that the high fidelity response is simply too complicated to be corrected with such simple correction models.

This final result (figure 5.16) shows that AMMF may still fail even when accurate gradients are available. To evaluate the potential of AMMF, the solutions between the initial and optimal point were used to define an "apparent" gradient, which is basically the gradient that gives the best correction. The performance of this correction was evaluated by determining how far in the optimization could use this correction. This was evaluated based on two conditions, the first was based on the correction quality, it indicates the progress where the correction starts to degrade. The second metric looked at the overall quality, it shows the point in the optimization where the correction no longer provides useful information and must be updated. These results are shown in figure 5.17.

The results in figure 5.17 show that the correction starts to degrade early in the optimization. However, despite this degredation, the correction is effective for $\approx 60\%$ of the optimization. This shows that with ideal gradients, the AMMF algorithm would only need to be updated 2-4 times to achieve converged optimization results. This indicates that the algorithm may still be used, but the gradients need to be improved first.

5.4.4 Future work in developing AMMF further

Multi-fidelity methods could have the greatest impact in evaluating the loads within a design optimization loop. However, the results in section 5.4.3 show that the gradients given by the framework are not sufficiently accurate to build corrections that can be used in design optimization. This prompted a more detailed study into the gradient quality given in section 5.8 that has confirmed that the gradients are poor. Thus, this work cannot continue until the gradients in the tools are improved.

Another challenge in leveraging AMMF was finding an adequate pair of low and high fidelity models for the algorithm. Currently there is only one solver, thus another solver that has either a lower, or higher fidelity formulation needs to be incorporated into the framework. One possible example would be FLEX5.

Most of the attempts focused on using different load cases to represent the different fidelities. At the highest fidelity is the full IEC 61400 design load basis. However, the tools could not execute a full DLB without crashing. The tools were improved, but the final version still could not execute a full DLC automatically,

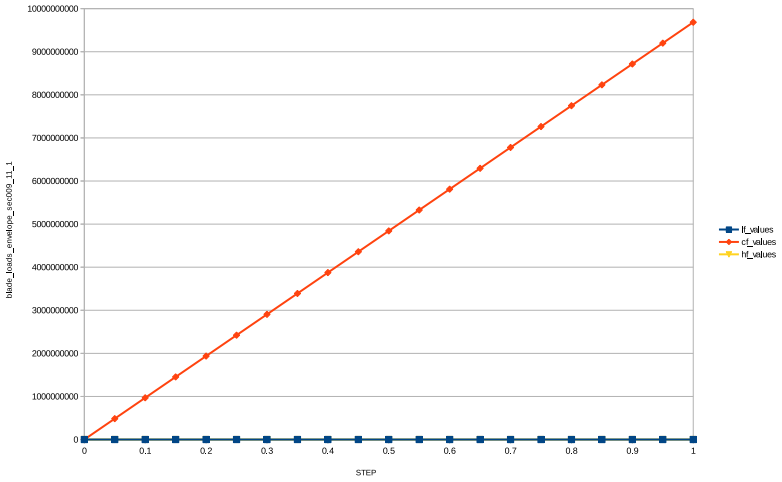

Figure 5.15: The worst example of an AMMF correction

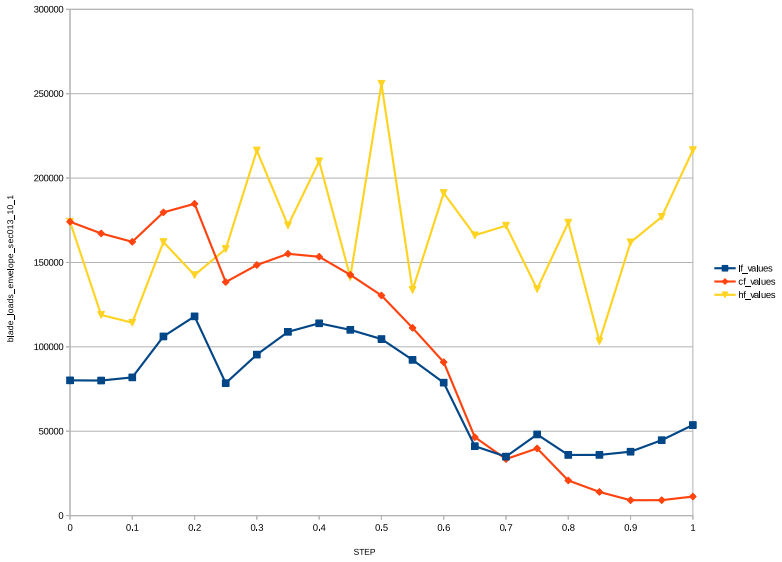

Figure 5.16: Moderately bad example of an AMMF correction

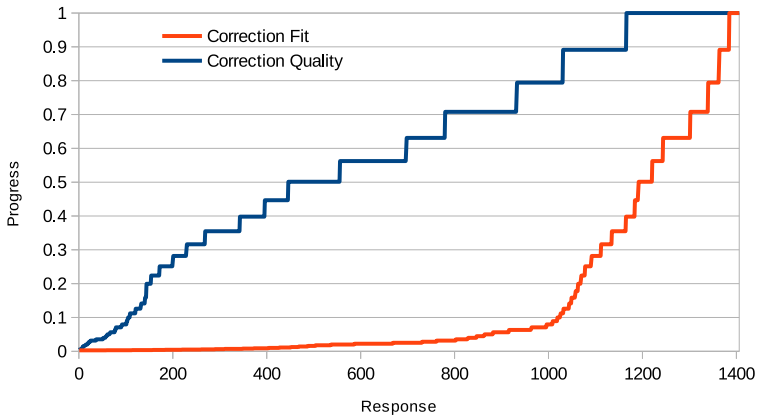

Figure 5.17: Potential optimization progress with an ideal AMMF correction

thus, cannot be used in its current state. The tolls need to be further improved to robustly execute full DLC calculations.

5.5 Fatigue Damage Calculation

5.5.1 Introduction

This chapter provides the information to calculate fatigue damage at material level using three different approaches. The methods are presented and their results discussed based on an example calculation on the DTU 10MW reference wind turbine.

The blade-cross sectional properties are obtained by means of BECAS (BEam Crooss section Analaysis Software), developed at DTU-Risø. The analysis of anisotropic and heterogeneous beam cross sections with arbitrary geometry can be taken into account,which make it suitable for wind turbine blades. The aero-elastic loading, shear and bending moments of/around the three axies are obtained by means of HAWC2. The general concepts used in all three methods are presented in this section.

Rainflow counting is used to reduce the spectrum of varying stress/strain in a simpler set. The output is used to build the cycle matrices which are two dimensional matrices where the axes represent the mean value and amplitudes of the cycles. The entries are the number of cycles with the specific mean and amplitude defined in the axis indices. This information is used, later on, to compute the damage.

Table 5.1: Cycle Matrix Example

Stress Range	**Mean Stress (MPa)**								
(MPa)	-400	-300	-200	-100	0	100	200	300	400
200	8	16	24	10	80	352	437	215	5
250	2	16	46	15	12	521	414	57	
300		8	22	35		437	174	94	
350	1			2	1	2	48	9	
400			3					7	2
450	3		7	6	2	7		2	
500		2			4			8	1

The number of cycles after which the material will fail, given the amplitude and mean, is computed using the Goodmann diagram as described in the BECAS manual [1, 7, 6]. The formulation relates the number of cycles to failure, N, with the ultimate tension and compression material properties for a given mean stress and amplitude.

$$N = \left[\frac{UTS+|UCS|-|2\sigma_m-UTS+|UCS||}{2\sigma_a}\right]^m \tag{5.5}$$

Once the number of cycles to failure for each mean stress and range is computed, it is possible to calculate the damage. The fatigue damage only takes into account the longitudinal component of the strains, perpendicular to the cross-section, disregarding the other elements of the vector. The damage, D, can be obtained by the linear Palmgreen-Miner rule equation as:

$$D = \sum_{c=1}^{cases} w_c \sum^{bins} \frac{n}{N} \tag{5.6}$$

where w_c is the Weibull scaling factor and n represents the number of cycles per bin.

5.5.2 Fatigue Damage Methods

Cycle Matrix of Strains

This method is considered as reference since it is the most accurate and all the available information is used. The full load vector is simultaneously applied to compute the strain time series at each element. Then, the time series of the strains are used to compute the cycle matrices for each element on the FEM model. Giving the amplitudes and the mean values of the cycles, the damage is computed based on Goodmann theory.

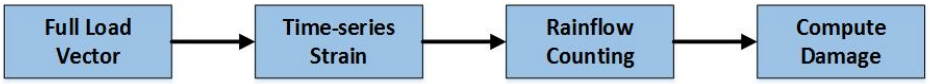

Figure 5.18: Workflow Cycle Matrix of Strains

Damage Equivalent Load

The second approach uses the Damage Equivalent Load (DEL) to compute fatigue damage where only the flapwise load component is taken into account. In order to encounter a possible flap-edge load combination, a series of load rotations are applied into the loading before computing the DEL. Then, the DEL is rotated back to its original reference and the strain values obtained in the same way as a static load would be applied. Finally, the damage is computed diving the number of cycles used to compute the DEL and the limit cycles.

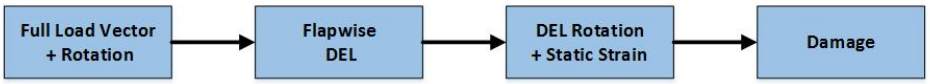

Figure 5.19: Workflow Damage Equivalent Load

Cycle Matrix of Loads

Cycle matrices are used to compute the fatigue damage using a similar approach as the damage equivalent load method. Rainflow counting is used to compute the cycle matrices on the flapwise load moment. Before computing the cycle matrices, the load is rotated to different angles and rotated back before computing the strains. Damage is obtained by Goodmann theory too.

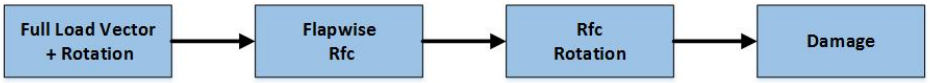

Figure 5.20: Workflow Cycle Matrix of Loads

5.5.3 Fatigue Damage Comparison

In order to show the differences between the three presented approaches, two cases are presented in this chapter. First, periodic loading is applied and its impact on the damage is discussed. Then, the full load vector retrieved from HAWC2 is used to compute the damage. All damage presented in this section corresponds to single 10 minute simulations where no probability has been applied.

Period Loading

The period loading case corresponds to a 1 Hz sinusoidal flapwise moment with a mean of 1MN and an amplitude of 1MNm. The three methods are compared in Figure 5.21, where damage for the outermost elements of the blade root cross section are plot.

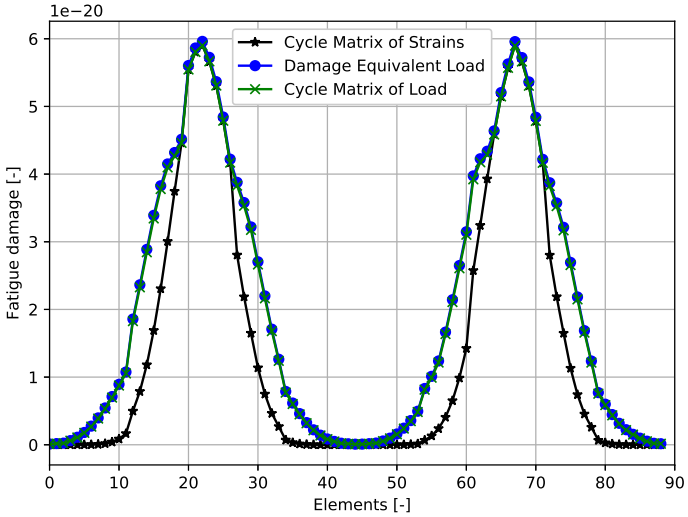

Figure 5.21: Comparison of Fatigue Damage Computed with the Different Methods

The three methods predict the same damage at peak positions, around elements 20 and 65. However, the DEL and cycle matrix of loads methods over-predict the damage at the rest of the elements, due to the load rotation where at regions away from the maximum, the damage obtained by rotating the load is higher than the actual damage.

Both methods, DEL and cycle matrix of loads, use load rotation to compute damage. An example of the load rotation impact on the damage is shown for the damage equivalent load case. From the given periodic load, the 1 Hz DEL of the 10 minute period corresponds to 2 MNm on pure flapwise loading. Figure 5.22 presents how the DEL distribution is translated into flapwise and edgewise when rotation is applied.

When simple loading is applied, the load rotation creates artificial combinations not defined in the input loading. It is possible to plot the normalized damage of an element over the load rotation angle. Two outer elements are plot in Figure 5.23. The black line represents an element where the maximum damage is found at rotation angle 0, pure flapwise loading and corresponding to the maximum damage region observed in Figure 5.21.

On the other hand, the blue line corresponds to an element where the maximum damage is found when the load is rotated 130 degrees and corresponds to the overestimation observed before the peak values in Figure 5.21.

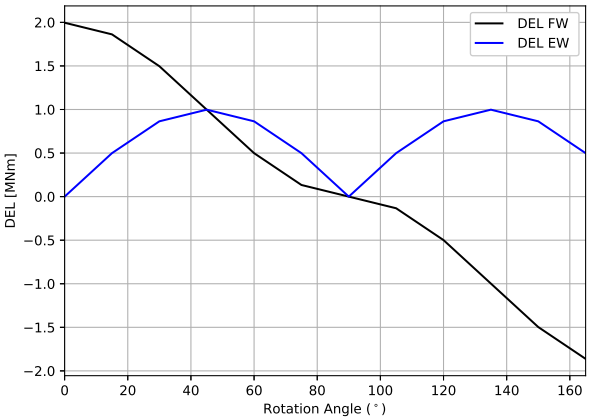

Figure 5.22: Flapwise and Edgewise DEL as Function of Rotation

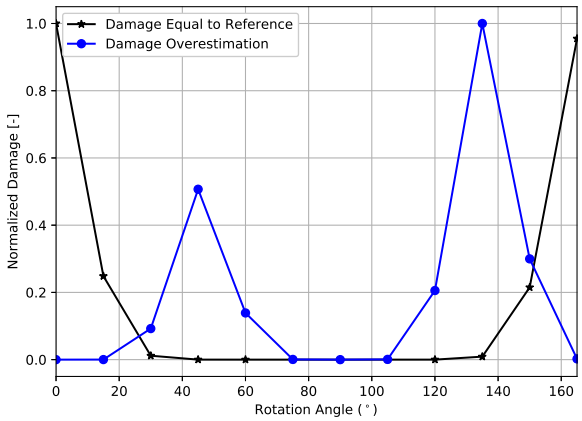

Figure 5.23: Normalized Damage as Function of Rotation Angle

Full Load Vector Loading

The full aero-elastic load vector is applied in this case by retrieving the data from HAWC2. The presented results are based on a DLC 1.2 using different mean wind speeds at the blade-root section. A zoom-in of the loading can be observed in Figure 5.24.

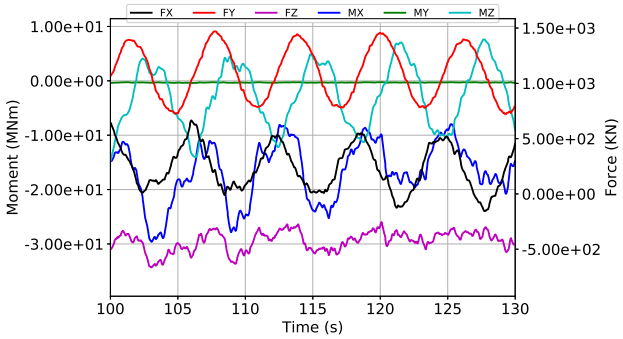

Figure 5.24: Zoom In: Load Time Series at Blade Root Moment Cross-section local coordinate system unsteady wind speed mean 12 m/s

The simplified method performance depends on the applied loading. For low wind speed where, the load variation is smaller, the simplified methods perform better than at rated wind speed, where the load variation is higher. Figure 5.25 presents the damage results for the outermost elements of the section for two different mean wind speeds.

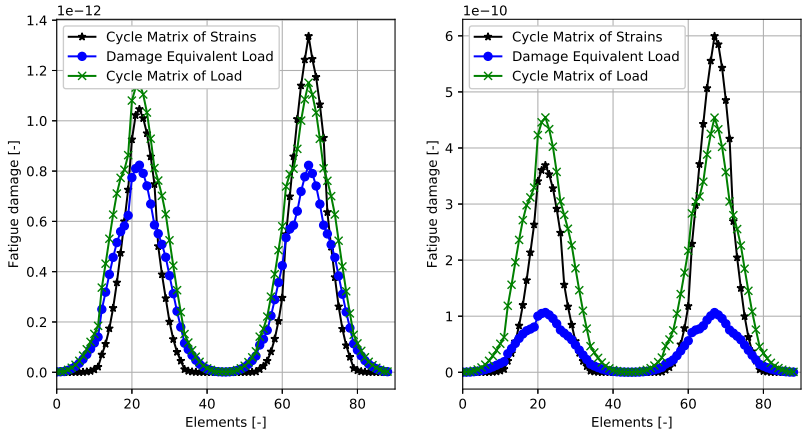

Figure 5.25: Left: Damage Unsteady Wind Speed mean 4m/s, Right: Damage Unsteady Wind Speed mean 12 m/s

In both cases, left and right in Figure 5.25, the cycle matrix of load method over-predict the first damage peak around the 25 element, while it under-predicts the maximum damage, the second peak, found around element 65. On the other hand, the DEL approach has significant differences for low wind speeds and its difference is increment as higher is the load variation.

In order to quantify that, the difference of the maximum damage of the section is plotted as a function of wind speed in Figure 5.26. The difference between the value predicted at maximum damage, around element 65, is obtained for the three methods.

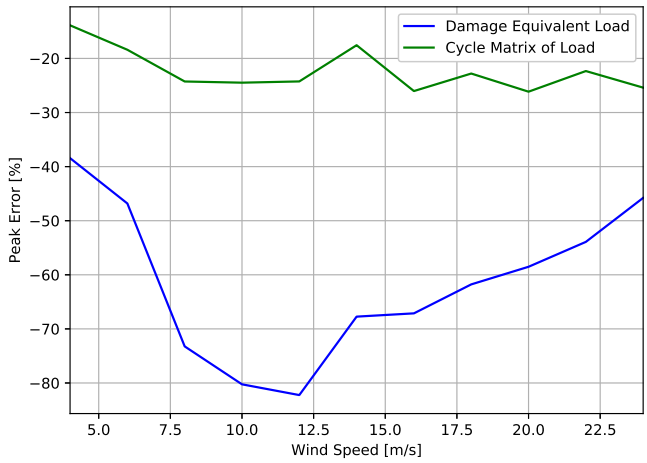

Figure 5.26: Normalised Maximum Error as Function of Mean Wind Speed

5.6 Airfoil Blender

5.6.1 Summary

Increasingly sophisticated and complex computational tools are necessary to design larger, more robust and effective wind turbine blades. In this context, it is natural to rely more on design optimization frameworks, HAWT Opt 2 being one of them. In such frameworks, it may be beneficial to increase the number of design variables and explore design spaces which have not really been explored before within optimizations. One of such spaces is the choice of airfoil family at each section of the blade. Nowadays it is still rather natural for designers to consider a relatively narrow range of airfoil families during blade design. In principle, allowing optimizers to do this job by selecting among a vast number of families and considering all design requirements could potentially point in design directions overlooked so far.

The same principle may be applied to aerodynamic add-ons where various add-on types, sizes and positions on the blades could be taken into account already in the design stage and included in the optimizations. Examples of such add-ons could be Gurney flaps and Vortex Generators. The latter may be used to mitigate the negative effect that blade surface erosion has on the aerodynamic performance.

In order to extend capabilities of HAWT Opt 2 by allowing for various airfoil families and aerodynamic add-ons to be used within optimizations, co called Airfoil Blender was developed in present work. It may be thought of as an airfoil database where lift, drag and moment coefficients are stored as a function of airfoil family, thickness, Reynolds number and angle of attack. Further, airfoil coordinates (shapes) are stored as a function of airfoil family and thickness. Note that airfoil family in this context refers to an aerodynamic data set where airfoils modified with any aerodynamic add-ons may be stored as separate families.

The essential feature of the tool is the ability to continuously interpolate airfoils across all aforementioned variables. Continuous interpolation is such that does not create discontinuities and hence will not misguide the optimizer upon computation of gradients. The interpolator works on an unstructured grid, hence allowing for different airfoils to be defined at different Re, thickness or angle of attack values.

The method relies on the robust SciPy's Radial Basis Function interpolation method being aided by a number of underlying parameters, allowing for some degree of tuning depending on the exact application.

Airfoil Blender also includes two different three-dimensional correction methods. Such methods are typically used to account for three-dimensional effects present on the rotating blade and not captured by typical two-dimensional wind tunnel measurements or any two-dimensional computations. Three dimensional effects are mostly pronounced on the inboard part of the blade. The closer to the tip of the blade, the more two-dimensional the flow becomes. Having those 3D correction methods within Airfoil Blender will allow for real-time correction of the data taking into account specific parameter values present during optimizations. This is in contrast to how this issue is handled in most cases where the lift and drag coefficients are corrected once and the correction itself remains insensitive towards parameters varying during optimizations e.g. radial position of the airfoil.

In principle, the interpolation method is not bound by any angle of attack regime. However, if the interpolations are meant to take relatively short time, especially when run on a single machine, they should be restricted to the operational angle of attack regime. This is why the Blender also includes two different 360-degree extrapolation methods meant to be applied on the data post interpolation in order to make it aeroelastic-code-friendly.

Airfoil Blender, including simple validations and an exemplary application is described in the attached work by Skrzypiński et al.[1]. The example comprises optimization of the radial layout of 1% Vortex Generators mounted on the suction side of the eroded DTU 10 MW Reference Wind Turbine blade by Bak et al.[2], aiming at maximizing energy production without any additional constraints.

[1] Skrzypiński, W., McWilliam, M., Zahle, F. (2018). Airfoil Blender for Blade Optimizations. Journal of Physics: Conference Series, 1037(4), [042010]. DOI: 10.1088/1742-6596/1037/4/042010

[2] Bak, C., Zahle, F., Bitsche, R., Kim, T., Yde, A., Henriksen, L. C., ... Natarajan, A. (2013). The DTU 10-MW Reference Wind Turbine [Sound/Visual production (digital)]. Danish Wind Power Research 2013, Fredericia, Denmark, 27/05/2013

5.6.2 Interpolation in multi-dimensional parameter space

In the present work, described in 1, airfoil families as such are characterized by a single parameter. This makes it most practical to consider not more than two families at once, e.g. by assigning them parameter values of zero and one, respectively. Then, it is left to an optimizer and the Blender to interpolate between both and find optimal blending factor. If in the actual design, a choice needs to be made between two families, then the one that the optimal blending factor fell closer to would be of choice. In principle, it is possible to add more airfoil families, e.g. three with parameter values of zero, half and one. The issue is that then, it is only possible to obtain a blend of two families at once, i.e. the first and the second if the optimal blending factor falls between the values of zero and half, or the second and third family if the optimal blending factor falls between the values of half and one.

In order to allow for blending of more than two families and at the same time for obtaining a blend of two without involving others, each family needs to be defined by n-1 independent parameters where n is the number of families in the interpolation. Then, the interpolation of families is carried out across n-1 independent variables and the optimum is defined by n-1 blending factors. In other words, optimum is defined as a point in n-1 - dimensional space. For example, to blend three families, their parameters should be defined in two-dimensional parameter space, preferably on three edges of an equilateral triangle, see 5.27 (left). Then, the optimization may be constrained to blending factors that define points falling within this triangle. If in the end, for a certain blade section, a choice needs to be made between the three families, then the one corresponding to the triangle edge that the point defined by the two optimal blending factors (being x and y coordinates of that point) falls the closest to should be of choice.

Elaborating on that abstraction, in the case of four families, each defined by three independent parameters, a preferred choice would be to place those families on the edges of an equilateral triangular pyramid, see 5.27 (right).

Note that theoretically there is no limit to the number of families provided that there is n-1 corresponding parameters, n being the number of families. However, addition of interpolation variables significantly increases CPU time of the interpolation, hence moderation is advised.

As for now, the Blender was extended to two-dimensional parameter space in terms of airfoil families which allows for unrestricted interpolation between three families at once. To present this abstraction, the example from 1 was extended and presented below. In the article, the Blender is used in a simple optimization to design the spanwise layout of 1% Vortex Generators on the model of an eroded DTU 10 MW RWT blade 2. There, two families are considered, one being the default airfoils with Leading Edge Roughness (erosion) without any Vortex Generators. The other family is the same set of airfoils with the same amount of erosion but with the addition of 1% Vortex Generators mounted at 20-25% chord length on the suction side. Parameter values of zero and one were assigned to the aforementioned families, respectively.

Now, an additional family was introduced, corresponding to 0.5% Vortex Generators installed at suction side, 55-60% chord length. In terms of parameter values, each of the three families was defined at the edge of an equilateral triangle, defined by independent parameters, x and y. Remember that x and y are no physical quantities in this context, see 5.27 (left).

Figure 5.28 presents how interpolation works in the present case on the example of the modelled lift and drag coefficients of the FFA-W3-241 airfoil at 15 deg Angle of Attack. The three edges of either of the triangles correspond to the three 'airfoil families', i.e. aerodynamic data sets. The coloured grids are a result of the interpolation between those sets.

Figure 5.29 presents results of previous modelling work, part of published by Skrzypiński et al.[3], being the reference for present results. The curves present the increase in the Annual Energy Production plotted as a function of the radial position of VGs. Two curves in the upper part of the figure relate to VGs mounted on clean blades and the curves in the lower part of the figure, on blades with the level of erosion corresponding to a 3.7% loss in the AEP compared to the clean blade. Solid lines represent 1% VGs whereas dashed, 0.5% VGs. The simplest way to interpret the curves is to say that the higher the slope of a curve at a radial location, the more gain in the AEP would be introduced by installing the corresponding VGs at that point. For example, installing 1% VGs on eroded blades between 15 and 25 meter radius would correspond to a 0.7% increase in the AEP (solid orange line). Installing 0.5% VGs between 50 and

[3]Skrzypiński, W. R., Gaunaa, M., Bak, C., Brønnum, N. F., & Brink Kruse Olsen, E. (2015). Increase in the Annual Energy Production due to a Retrofit of Vortex Generators on Blades

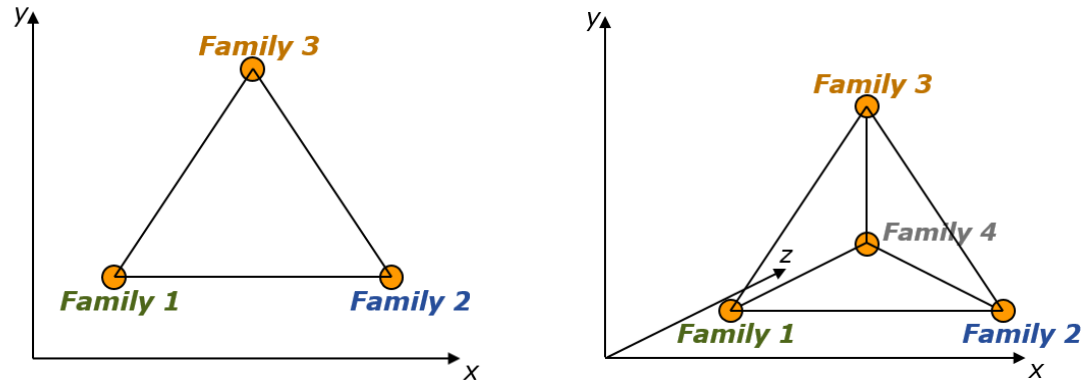

Figure 5.27: Definition of 'airfoil families' (aerodynamic coefficient sets) in two-dimensional (left) and three-dimensional (right) non-physical parameter spaces, for the purpose of interpolations

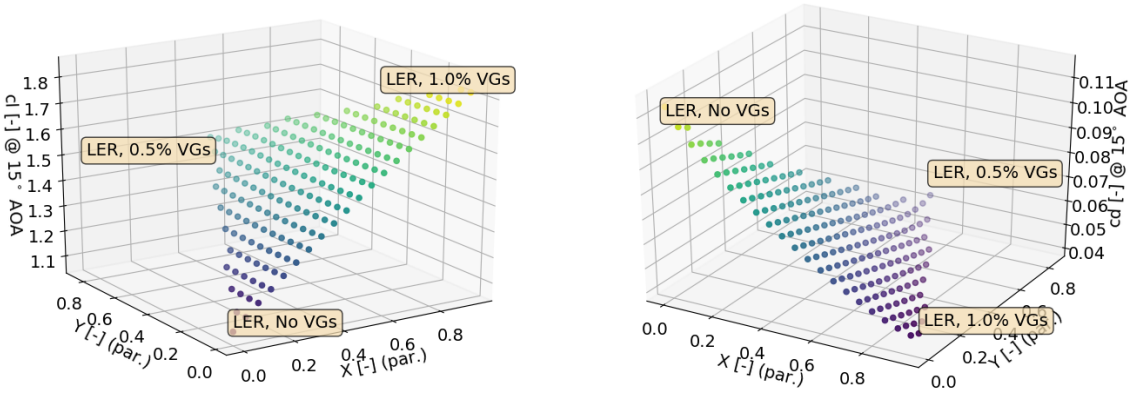

Figure 5.28: Exemplary lift coefficient (left) and drag coefficient (right) of the FFA-W3-241 airfoil at 15 deg AOA, corresponding to three different configurations, i.e. (1) modelled Leading Edge Roughness (LER) without Vortex Generators (VGs), (2) LER with 1% VGs modelled at 20-25% chord length (suction side), (3) LER with 0.5% VGs modelled at 55-60% chord length (suction side). Aforementioned configurations are positioned at the edges of an equilateral triangle in a non-physical two-dimensional parameter space. Colour grids represent interpolated values.

85 meter radius would gain additional 1.0% AEP. Note that the modelling assumes certain level of erosion whereas, according to the modelling, the same installation of 0.5% VGs on a clean blade would decrease the AEP.

Further, Figure 5.30 presents the results of the present design optimization with the solid orange line being the result. The markers present the radial positions for which the optimizations were carried out with the text boxes indicating the corresponding positions. An easy way to interpret the results is to say that for each of the considered radial positions, the edge of the triangle (airfoil family, i.e. aerodynamic data set) which the corresponding point on the orange curve fell the closest to should be of choice. However, there is actually more information contained in the result presented in the figure. That is, if the point falls approximately half-way between the [LER, No VGs] and [LER, 0.5% VGs] sets, it may actually indicate that the optimal solution at that location would be either to install smaller VGs than the 0.5% or install the 0.5% VGs farther from the leading edge. Either solution would likely result in the corresponding aerodynamic characteristics falling in-between those of the two data sets. It may be concluded that the present results are in agreement with the reference. The orange curve in Figure 5.30 approaching the [LER,

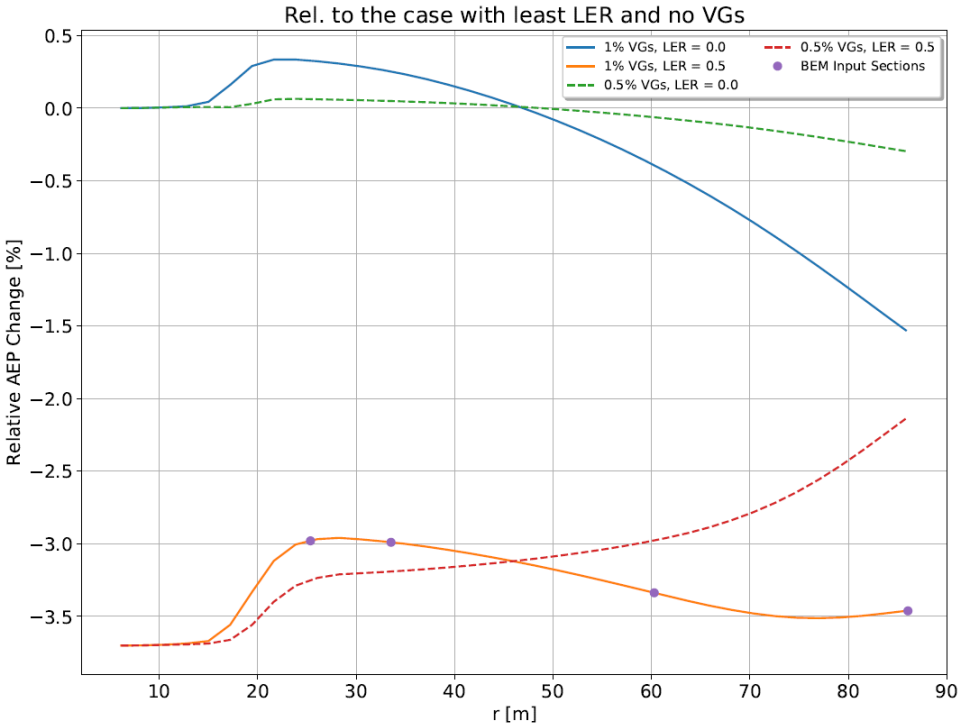

Figure 5.29: Reference modelled AEP curves indicating optimal radial layout of Vortex Generators (VGs) on the DTU 10MW Reference Wind Turbine blade.

0.5% VGs] point towards the tip of the blade corresponds to the red dashed line in Figure 5.29 increasing its slope towards the tip of the same blade.

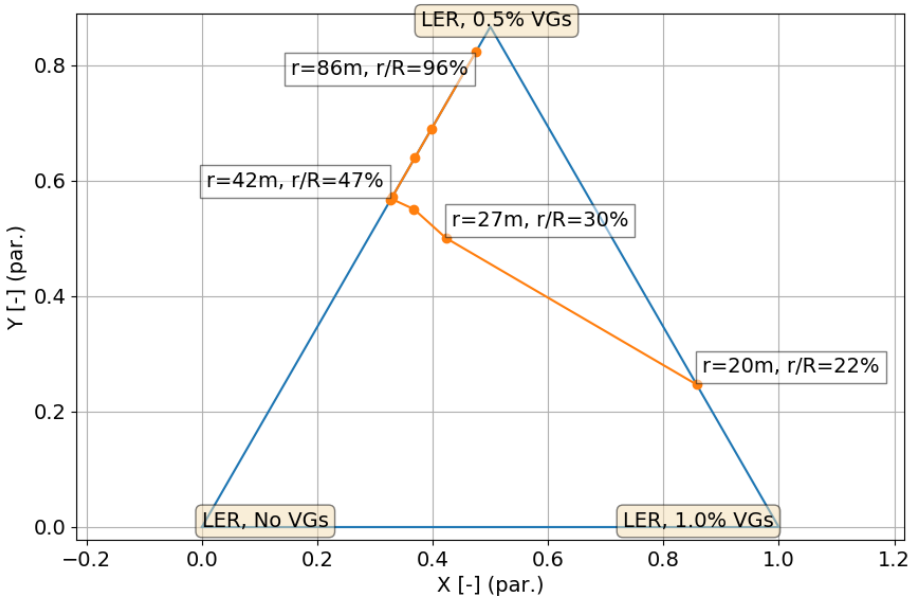

Figure 5.30: Result of the optimization with three aforementioned data sets

Figure 5.31 presents a VG layout that could be proposed based on the aforementioned results, with 1% VGs installed between 15 and 25 m radius, and 0.5% VGs installed between 50 and 85 m radius, increasing the modelled AEP by 1.7% relative to the turbine with eroded blades without VGs.

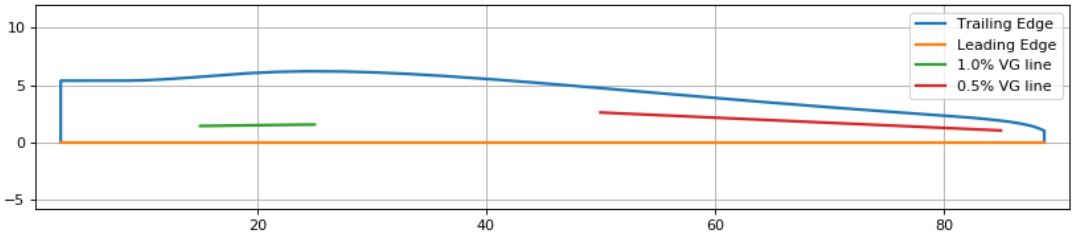

Figure 5.31: Resulting VG layout on the DTU 10MW RWT blade model with erosion

Note that the actual final layout of VGs is, in such a case, a result of an interpretation of the results by a human rather than directly output by the optimizer as factors such as cost of the installation and uncertainty in the estimation of the level of erosion must be taken into account prior to final decisions.

Finally, note that this section is primarily intended to show that the Blender may be conveniently used in a design optimization framework and that it works properly withing optimizations, here being utilized in a simple SLSQP optimization. The next step in this work is to plug the Blender into HAWT Opt 2 framework and to utilize and further develop its capabilities in complex blade design optimizations. Among other research ideas, one is to study the trade-offs between the use of aerodynamically superior airfoils (larger camber, larger L/D) vs structurally superior and thicker.

5.7 Cross Section Mesher

BECAS requires two dimensional cross section meshes of the internal structure as input. This can be quite challenging to generate robustly in design optimization. There are difficulties at the trailing edge and leading edge where there is high curvature. This is due to the fact that as you add material on the outer-mold-line it will start to interphere with the material on the opposite face. In the old meshing tool, this would create degernerate meshes or negative jacobians that would lead to BECAS failing to calculate accurate cross section properties and stresses. To overcome these problems a new mesh tool was created.

The objectives for this new mesh tool are as follows:

- Parametric geometry specification
- Robust geometry generation
- Analytic gradients

This tool was created in C++. It has three layers, a geometry parameterization layer, which takes design variables and constructs the geometry. A topology layer that describes how the mesh lines should travel through the geometry and a mesh generation layer that generates the actual mesh. Since the geometry is defined parametrically, gradients of mesh deformation with respect to design variables can be calculated analytically.

An example of a mesh is given in figure 5.32 of the DTU 10MW reference wind turbine. The mesh is coarse so individual cells can be seen.

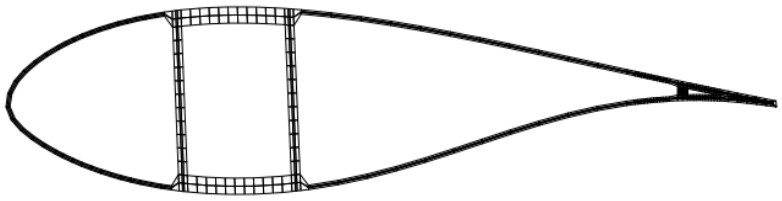

Figure 5.32: An example of the cross section mesher

Figure 5.33 shows the trailing edge region of the mesh. In this case, the material thicknesses for the wind turbine are too thick to fit within this region. In design optimization, a tool cannot simply fail in these circumstances because the optimization will easily generate such designs. Instead, the mesh tool will detect the situation and automatically scale the thickness of all materials so that there are no collisions. Since this scaled thickness differs from the design thickness, the mesh program will output these scaling parameters so they can be constrained in the optimization.

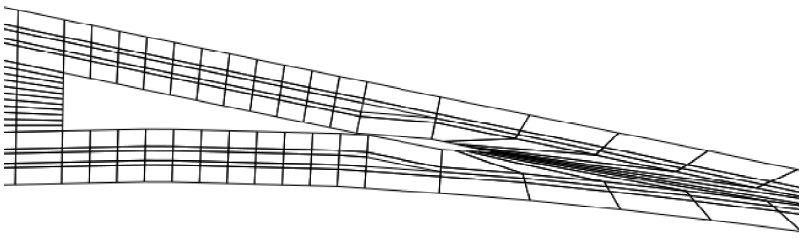

Figure 5.33: The trailing edge region of the mesh

Figure 5.34 shows the mesh at the web connection point. This is a complicated area where the shear web connects to the end of the spar-cap. The main laminate at the end of the spar-cap contains a plydrop and a tapering region. It is clear that the mesh program easily handles these difficult geometries.

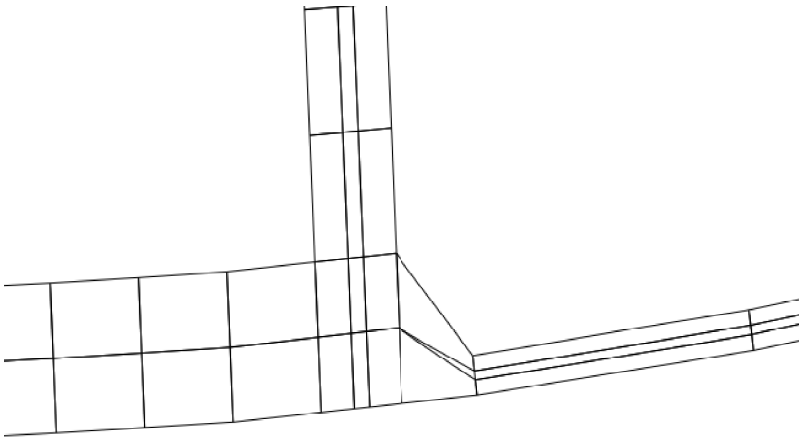

Figure 5.34: The mesh at the web connection point

Finally, figure 5.35 shows the gradient capabilities of this tool. The figure shows how the nodes would move if division point 4 was moved. This design variable controls the position of the lower right corner of the box-spar. This can be used to get very accurate gradients of stiffness and strength with respect to variations in geometry.

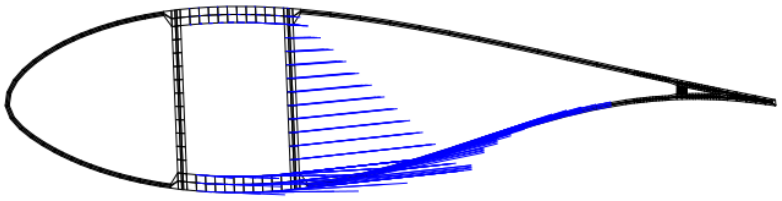

Figure 5.35: The analytic gradient of the mesh with respect to the web connection point

5.7.1 Future work

This new cross section mesh program is superior to the current version used in HAWTOpt2. This cross section mesher was not implemented in HAWTOpt2 for 2 reasons. BECAS has not been updated to use analytic gradients, theis means the gradient information from this tool could not be utilized. The main reason is the fact that the user interface is not mature enough to be used in an industrial setting. Currently the mesh is generated from a text file that describes all the entities in the cross-section. The input for the example shown here had over 2500 lines. Further improvements are needed in the user interface so new cross-sections can be generated quickly.

5.8 Gradient Study on HAWTOpt2

The earlier work on multi-fidelity optimization indicated that errors in the gradients could be contributing to AMMF failing to adequately correcting the model. This prompted a detailed study into the gradients. Section 5.8.1 gives the motivation for improving gradients in general. Then, section 5.8.2 gives the evaluation. Finally, some improvements in the gradients are shown in section 5.8.3.

5.8.1 Motivation for improving the gradients

To understand the effect of error in the gradient on optimization, a simple test case was constructed. The test case is based on minimizing the Rosenbrock function shown in figure 5.36 and given in equation (5.7). To test the effect of error, random values were added to the gradients, then the function was optimized.

$$z = (1.0 - x)^2 + 100(y - x^2)^2 \tag{5.7}$$

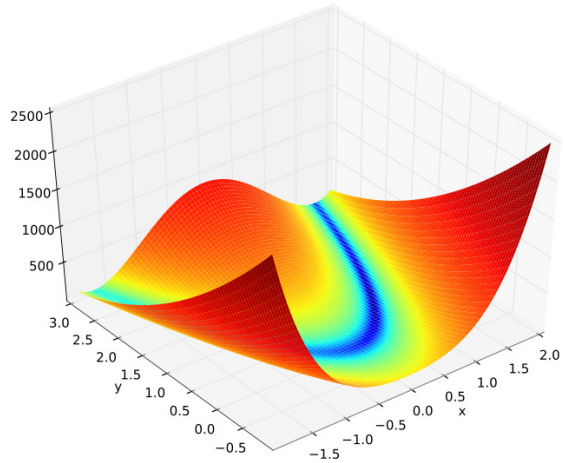

Figure 5.36: Rosenbrock Function

Figure 5.37 shows the effect of error in the gradients on various aspects of the optimization. Overall, the optimization proceeds as normal up until a point where the error in the gradients is similar to the Kuhn-Karush-Tucker norm (*i.e.* the degree of convergence in the optimization). At that point the optimization algorithm request many more function calls and fails to make significant progress.

The gradients in HAWTOpt2 are based on finite differencing, equations (5.8), (5.9) and (5.10), which shows the forward, backwards and central differencing formula respectively. An important parameter in each of these schemes is the step size h. When h is too small, the gradient will have round-off error, due to finite precision. When h is too large, the gradient will have truncation error due to nonlinearities.

$$\frac{df}{dx} \approx \frac{f(x+h) - f(x)}{h} \tag{5.8}$$

$$\frac{df}{dx} \approx \frac{f(x) - f(x-h)}{h} \tag{5.9}$$

$$\frac{df}{dx} \approx \frac{f(x+h) - f(x-h)}{2h} \tag{5.10}$$

To undertand the errors in HAWTOpt2, the gradient of DLC 1.3 of the DTU 10MW RWT with turbulence was evaluated over a large range of step-sizes. Figure 5.38 shows an example of these errors. Figure 5.38a shows how the normalized blade failure varies with different steps in spar-cap thickness. The round-off error is seen below step sizes of 10^{-5} and the nonlinearity is seen at step-sizes larger than 0.01. Figure 5.38b

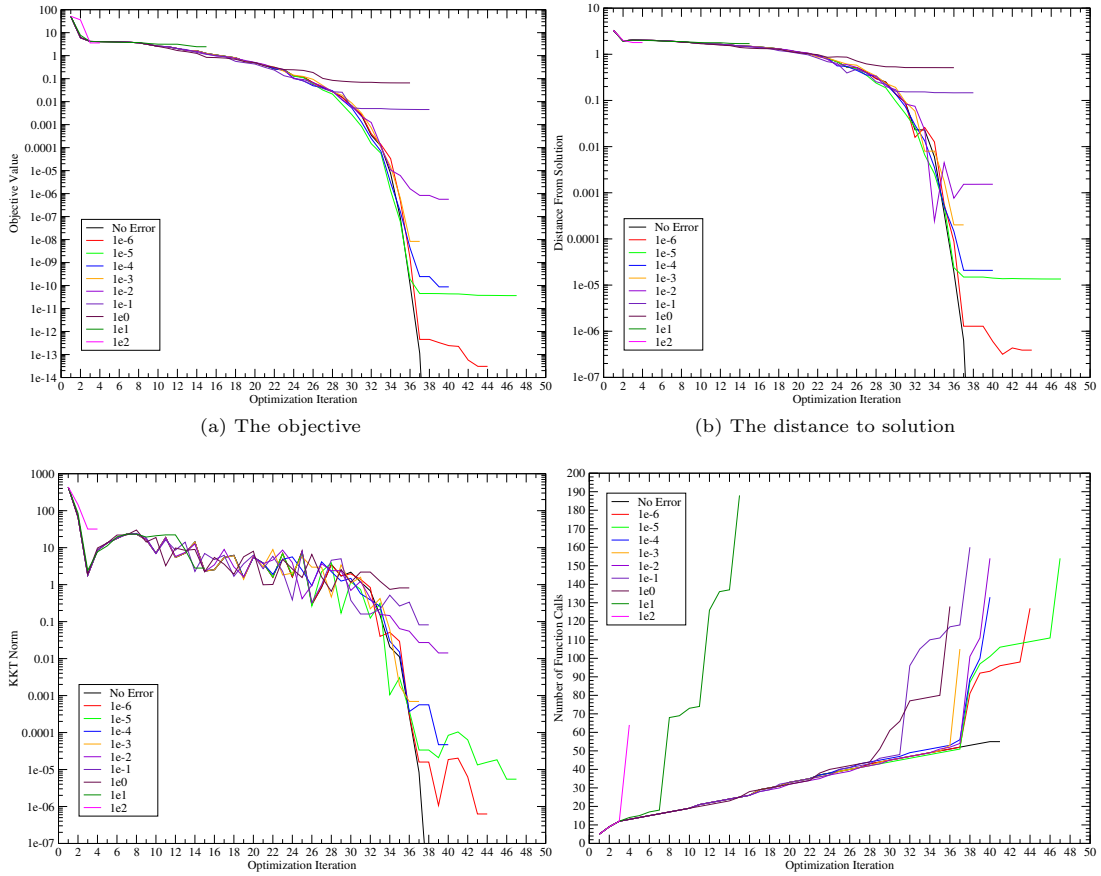

(a) The objective

(b) The distance to solution

(c) The Karush Kuhn Tucker norm

(d) The number of function calls

Figure 5.37: The effect of gradient error on optimization

shows the gradients from this trend, the thin brown lines are the gradients predicted from polynomial regression. The gradients are considered most accurate when all the gradient curves predict the same gradient. Yet, figure 5.38b shows that there are large regions where the gradients are not in agreement. This shows that finite differencing tends to amplify errors.

Minimizing the error in finite differencing is based on finding the optimal step size where both sources of error are minimised. However this is quite challenging for a number of reasons. Figure 5.39 gives one example where there are two regions of agreement, but with different gradient values. Furthermore, there are over $70,000$ terms that must be evaluated in HAWTOpt2 to properly tune the gradients. Such a large number could not be tuned practically by plots alone.

5.8.2 Evaluation of HAWTOpt2 Gradients

Due to the number of terms in the gradients, a statistical technique is needed to evaluate the gradients. Several statistical measures were explored and the best predictor was the coefficient of determination that can be calculated from polynomial regression. Figure 5.40 compared the coefficient with the gradients. The plot shows that the coefficient needs to be greater than 0.99 to have reliable gradients, and coefficients less than 0.9 are effectively random numbers.

This statistical measure can be used to understand the error in the gradients and the best step size to minimize this error. All the gradient terms were evaluated to find the optimal step size. The cummulative distribution of optimal step size is shown in figure 5.41. The figure shows that the current step size used

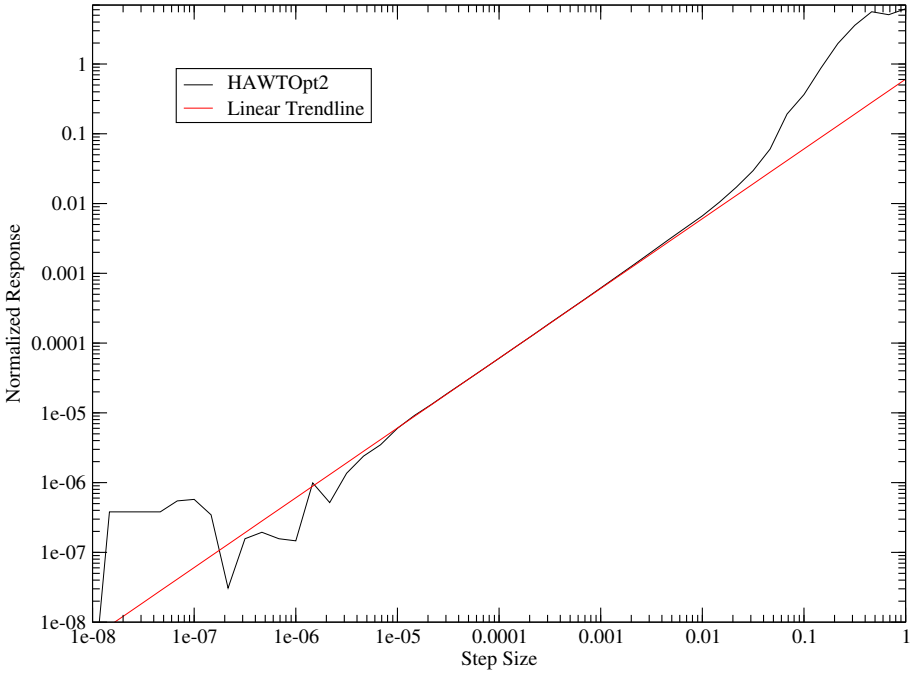

(a) Blade failure vs chord

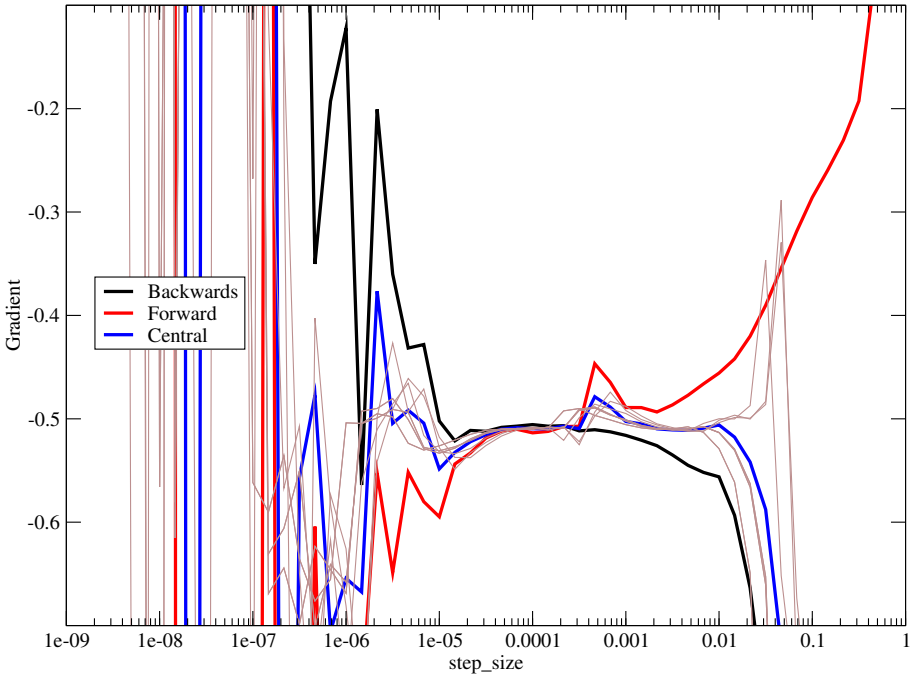

(b) The gradient according to finite differencing

Figure 5.38: Example of the different sources of error in finite differencing gradients

in HAWTOpt2 is optimal for most gradients, however there are still many that can be improved with different step-sizes.

The gradients can be improved by having a different step-size for each design variable. Figure 5.42 shows the error of the gradient with respect to different steps in the blade size. It is clear that with the typical

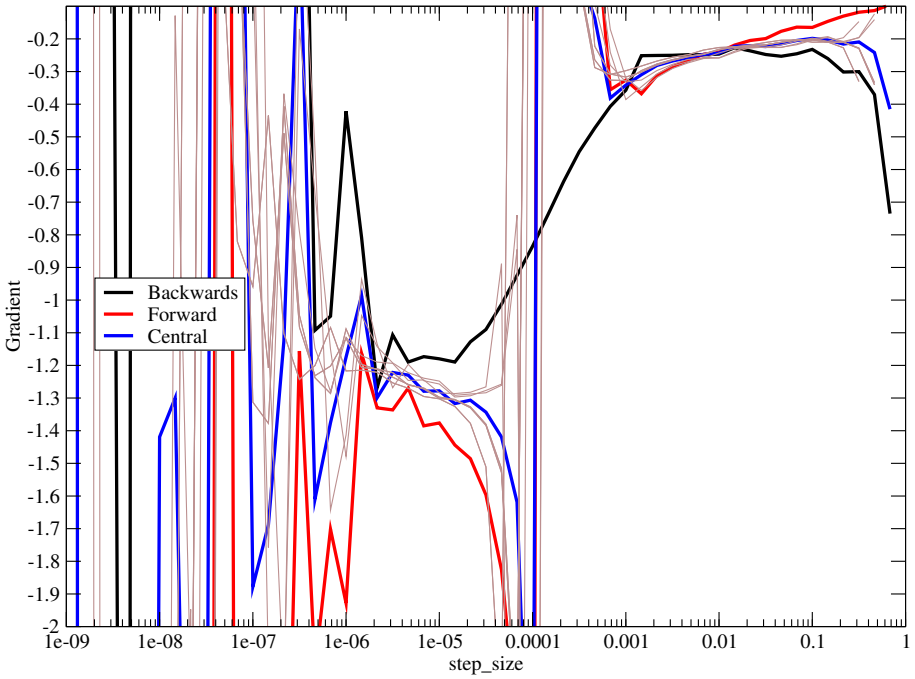

Figure 5.39: Blade failure vs. Chord gradients

step size of 0.01, the forward differencing error is $100\% \pm 200\%$. However, selecting a step size of 0.001 and central differencing, the error reduces to $40\% \pm 50\%$.

Figure 5.43 shows a cummulative distribution of the coefficient of determination statistics for HAWTOpt2. The figure shows that 30% of the gradients had coefficients less than 0.9 and only 40% of the gradients can be deemed reliable.

5.8.3 Improvements in the gradients

Section 5.8.2 shows that the gradients in HAWTOpt2 are quite poor. One source of error is the underlying tools. Within HAWC2 HAWC2 the output data is stored as 16 bit numbers. This increases the round-off error in the gradients. To eliminate this error a new version of HAWC2 was created that uses an output DLL to store the results in full precision. This is shown in figure 5.44.

Figure 5.45 shows that this new version of HAWC2 can dramatically improve gradients quality. However, looking at the coefficient of determination statistics in figure 5.46 shows that there still remains $\approx 40\%$ of the gradient terms have a coefficient less than 0.99 and $\approx 15\%$ of the terms have a coefficient less than 0.9. Thus, there is still more work required to improve gradient quality.

Figure 5.47 gives a summary of the different improvements one can make on the gradients. The current HAWTOpt2 configuration has a median error of $\approx 30\%$. Using central differencing brings the error down to $\approx 25\%$. Then tuning the step size for each variable brings the error down to $\approx 12\%$. Finally, improving the underlying tools brings the median error down to $\approx 2\%$. Improving the underlying tools will have the greatest impact on improving the gradients.

5.8.4 Future work in improving gradients

This study shows how the gradients of HAWTOpt2 can be improved by using different algorithms, different step-sizes and different tools. Each of these developments should be used in future optimizations. However the study could not eliminate all sources of error. Further improvements in the gradients will require further work on the tools. Analytic gradients can greatly improve the quality.

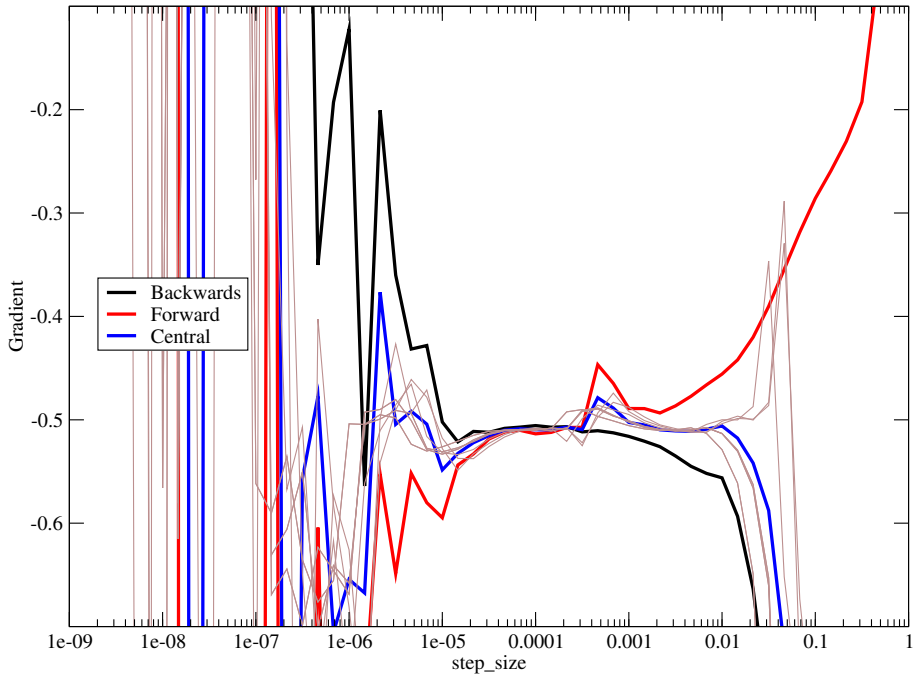

(a) The gradient of blade failure vs chord

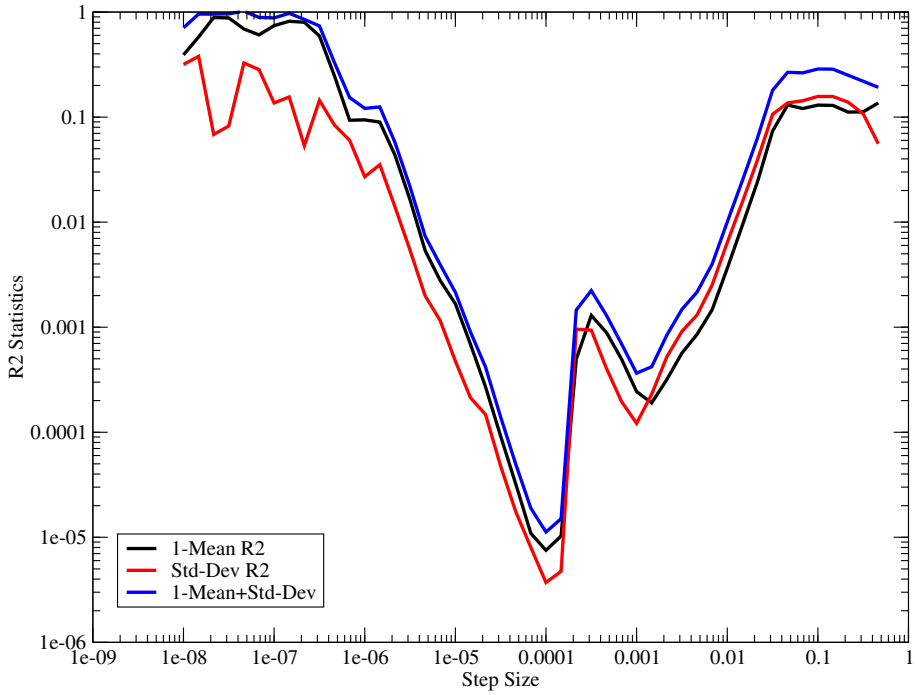

(b) The coefficient of determination

Figure 5.40: Example of the Coefficient of Determination in predicting gradient quality

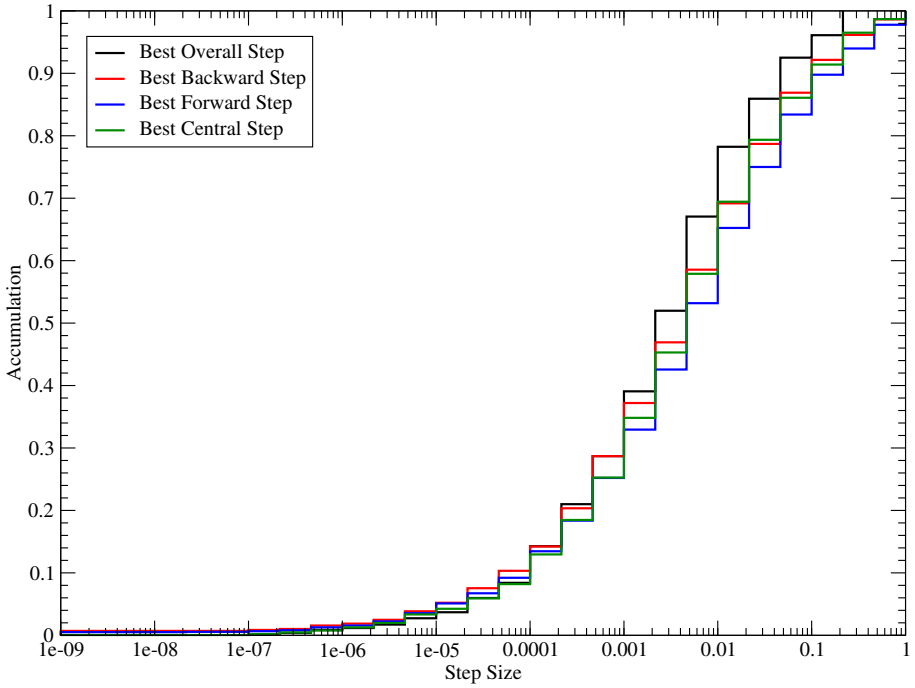

Figure 5.41: Cummulative distribution of optimal step size

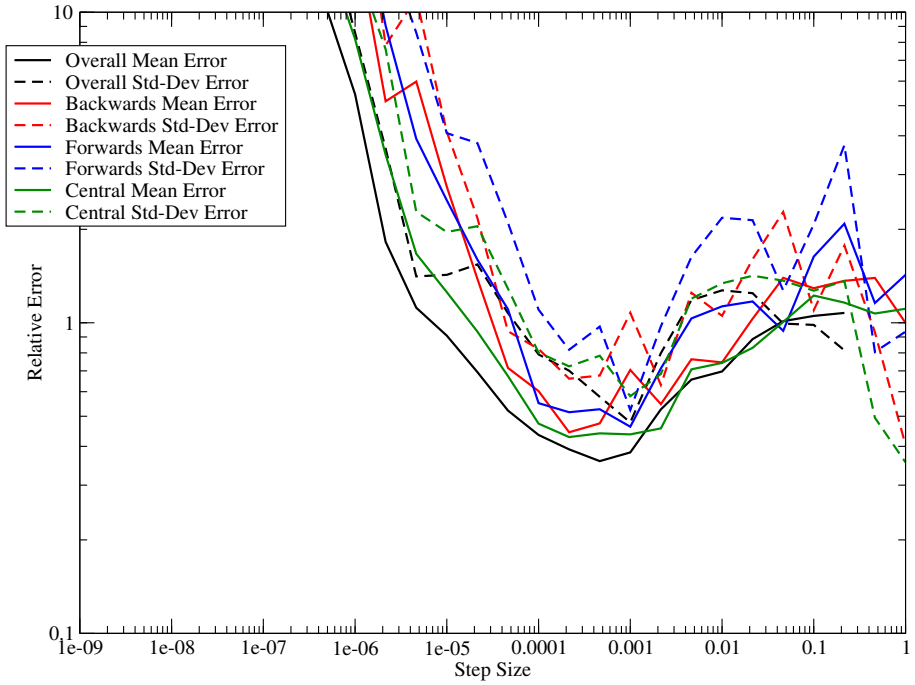

Figure 5.42: Relative error in blade-scaler gradient with respect to step-size

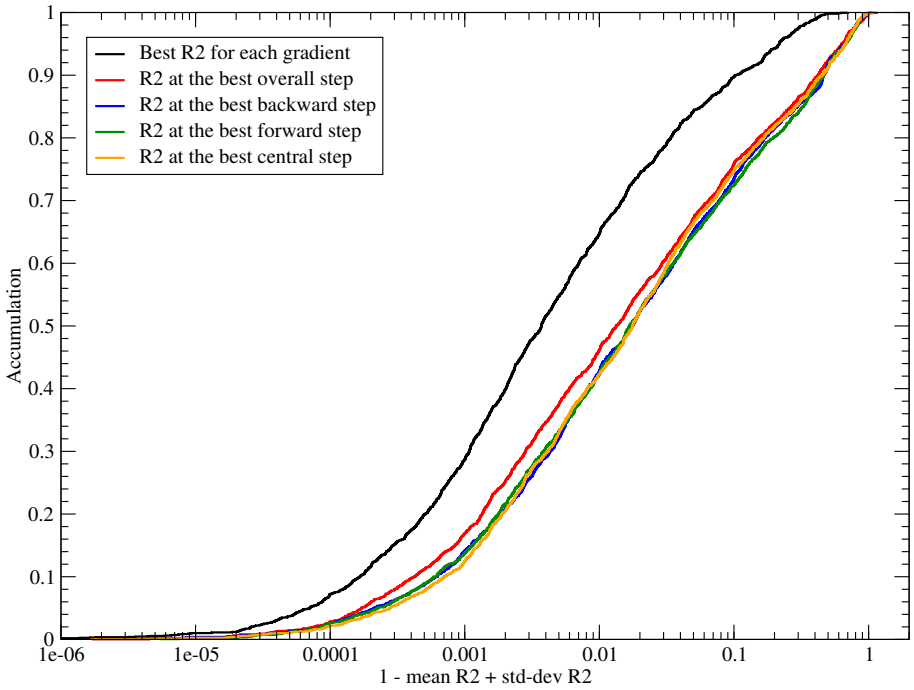

Figure 5.43: A summary of coefficient of determination statistics

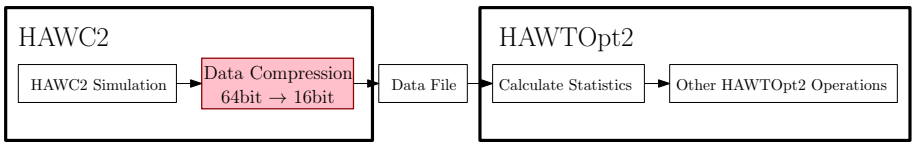

(a) Old HAWC2 Operations

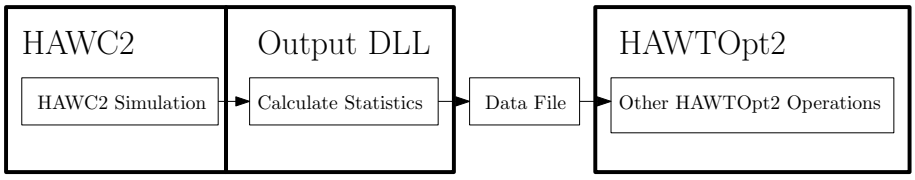

(b) Improved HAWC2 Operations

Figure 5.44: HAWC2 operations within HAWTOpt2

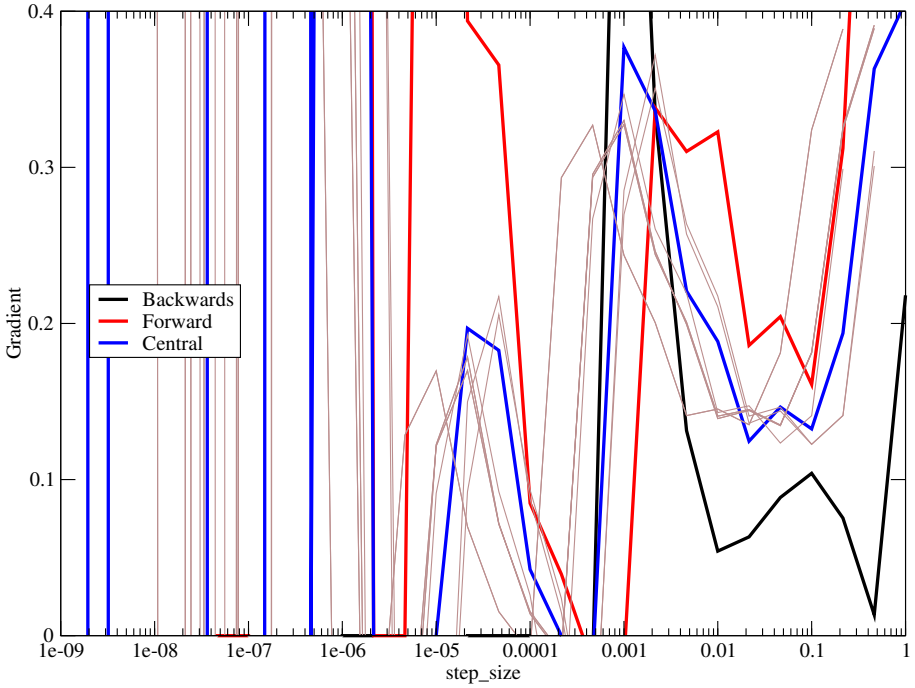

(a) Old HAWC2

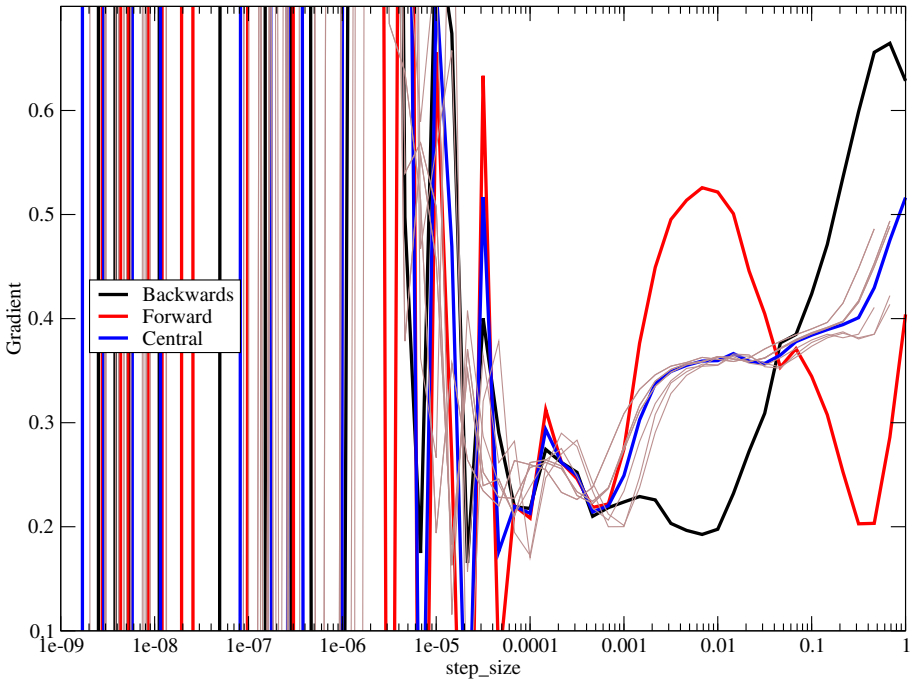

(b) Improved HAWC2

Figure 5.45: Effect of improving HAWC2 on the tip-deflection vs. blade scale gradient

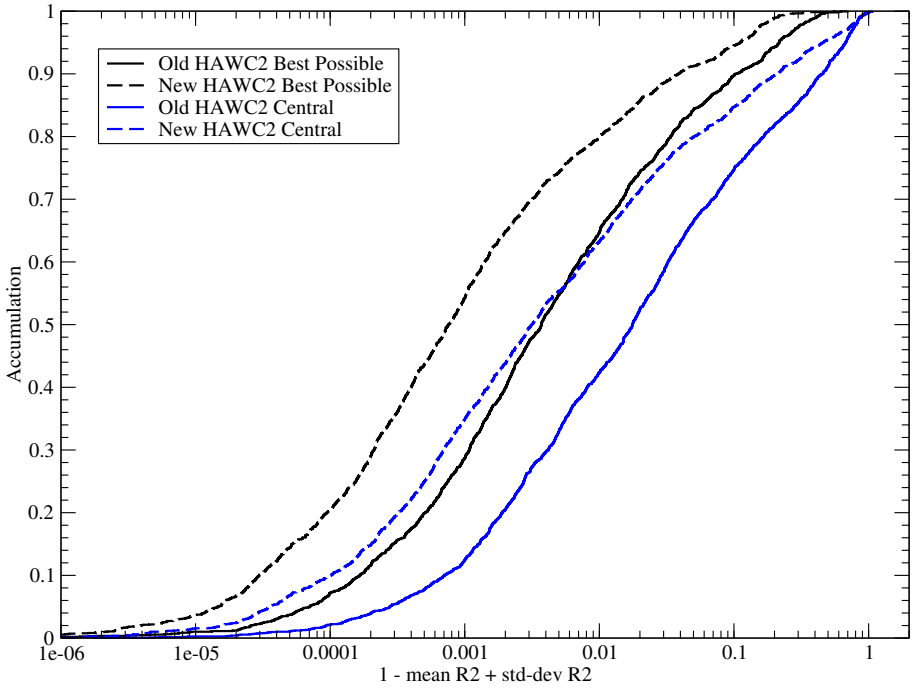

Figure 5.46: Coefficient of determination statistics with an improved HAWC2

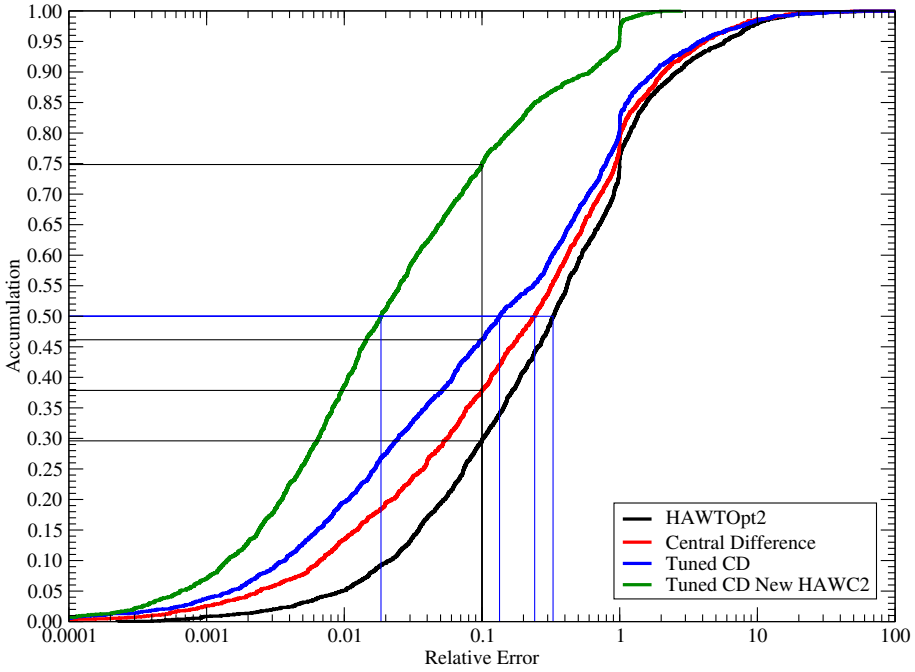

Figure 5.47: Summary of the different improvements in gradient quality

5.9 Stretched Rotor Blade Designs

The goal of the Stretched Rotor project was to apply new design tools and blade technologies to design an improved blade starting from a reference blade provided by LM. Since the original blade could not be shared fully in the project, the first task was thus to develop a blade matching its properties to the greatest extent possible using openly available materials and airfoil shapes. The detailed blade design resulting from this work is referred to as "Reference blade". Following this design, a number of optimized rotors were developed using HAWTOpt2, all of which are designated SRB¡X¿, where ¡X¿ is the incremented version number of the blade series. Below, a summary of the design iterations is provided:

- **Reference blade**: Based on LM supplied HAWC2 model and airfoil cross-sectional shapes, as well as basic structural geometric layout such as placement of shear webs and spar cap width.
- **SRB1**: First aerostructural design based on the Reference blade design. This design did not feature a minimum pitch setting for operation in Region 2, which lead to operation too close to C_{l-max}, and was therefore scrapped, and not documented in this report.
- **SRB2**: Same optimization setup as SRB1, but with the addition of a minimum pitch setting of -1.5 deg.
- **SRB3**: Based on the SRB2 optimization setup, with the addition of constraints on blade root edge and torsional extreme loads. The design was very similar to SRB2, and did not go through a detailed design loads evaluation, and is not documented in this report.
- **SRB4-5**: Experimentation with new constraints based on the material fatigue and fatigue in frequency domain models. Not documented in this report.
- **SRB6**: Blade design with identical design variables to SRB2, but with the addition of constraints based on the material fatigue and fatigue in frequency domain models. This design has gone through a detailed design loads evaluation and is reported in detail.
- **SRB7**: Experimental blade design with the addition of backward sweep as a degree of freedom. Constraints were identical to those used in SRB6. More work is needed to improve the optimization setup, and this blade did not go through a detailed design load comparison, and is not described in this report.
- **SRB8**: Blade design with a modified blade structural design featuring only one shear web. No optimization was carried out on this blade, but the simplified evaluation of the blade shows potential for achieving the necessary torsional flexibility in the blade for effective passive load alleviation. Details of the design are not reported here.

5.9.1 Reference Blade

The Reference blade was developed based on the aeroelastic and basic structural definition of the blade provided by LM. Since only basic structural geometric characteristics were provided, a number of choices were made, among them material characteristics, layup sequence of materials along the blade.

The sizing of the panel thicknesses was done using HAWTOpt2 where a feasibility problem was solved using SNOPT to match the beam properties supplied by LM. The quantities that were specified as constraints in the feasibility problem were: flapwise stiffness (EI_x), edgewise stiffness (EI_y), torsional stiffness (GJ), the chordwise position (relative to C/2) of the shear center (x_{sh}), and the chordwise position (relative to C/2) of the center of elasticity x_e. Since the Reference blade did not include added mass from e.g. the root solution, glue, gelcoat and other quantities not relevant for the stiffness properties, the mass distribution was not constrained, and any difference between the matched blade and the reference data was added manually to the matched blade as static added mass. The results show very good agreement within 1% along the majority of the blade, with only minor differences in one section at the root and at the tip.

5.9.2 SRB2

The SRB2 blade was optimized using HAWTOpt2 involving a full aerostructural design setup. Design variables included blade planform (chord, twist, and relative thickness), material thicknesses in key regions along the blade, as well as blade length and rotor tip speed ratio. The objective in the optimization was AEP, with constraints placed on the turbine platform ultimate loads, such as tower bottom moments, tower top forces, and blade root moments, as well as tower clearance and material limits. At the time of the work, the fatigue in frequency domain model and the material fatigue model was not yet implemented in HAWTOpt2 and constraints based on these models were therefore not taken into account.

The optimization problem is formally stated as follows:

$$\begin{aligned} \underset{\mathbf{x}_p,\, \mathbf{x}_s,\, \mathbf{x}_{oper}}{\text{minimize}} \quad & -\frac{AEP(\{\mathbf{x}_p,\, \mathbf{x}_s,\, \mathbf{x}_{oper}\}, \mathbf{p})}{AEP(\{\mathbf{0},\, \mathbf{0},\, \mathbf{0}\}, \mathbf{p})} \\ \text{subject to} \quad & \mathbf{g}(\mathbf{x}_p) \leq \mathbf{0}, \\ & \mathbf{h_g}(\mathbf{x}_s) \leq \mathbf{0}, \\ & \mathbf{h_s}(\mathbf{x}_s) \leq \mathbf{0}, \\ & \mathbf{k}(\{\mathbf{x}_p,\, \mathbf{x}_s\}) \leq \mathbf{0} \end{aligned} \tag{5.11}$$

where

- $\mathbf{x}_p$ is the planform variables,
- $\mathbf{x}_s$ is the structural variables,
- $\mathbf{x}_{oper}$ is the wind turbine control variables,
- p is turbine parameters kept constant.

Tables 5.2 and 5.3 provide a summary of design variables and constraints used in this study.

Parameter	# of DVs	**Comment**
Chord	5	Root chord fixed
Twist	5	Root twist fixed
Relative thickness	4	Root and tip relative thickness fixed
Blade prebend	4	-
Chordwise offset	4	-
Blade length	1	-
Trailing edge uniax	3	Symmetric pressure/suction side
Trailing edge triax	3	Symmetric pressure/suction side
Trailing panel triax	3	Symmetric pressure/suction side
Spar cap uniax	4	Symmetric pressure/suction side
Leading panel triax	3	Symmetric pressure/suction side
Leading edge uniax	3	Symmetric pressure/suction side
Spar cap center SS	2	Pressure side spar cap position/rear web attachment
Spar cap center SS	2	Pressure side spar cap position/rear web attachment
Shear web offset	2	Pressure side spar cap position/front web attachment
Mold angle	1	Suction side spar cap position/front web attachment
Tip-speed ratio	1	-
Total	50	

Table 5.2: Design variables used in the optimization of SRB2.

Figure 5.48 shows the optimization history of the SRB2 optimization. Within 60 major iterations, the optimization has improved the objective by over 4%, but to an unfeasible point with a maximum constraint violation of approximately 5%. As can be seen, rather large constraint violations are observed of up to 20%, before the optimization settles at an objective value of -1.42 at a feasible point.

Table 5.4 summarises the overall characteristics of the SRB2 blade design. The SRB2 design achieves a 3.7% increase in AEP computed based on steady state HAWCStab2 results, thus not including inflow turbulence. The increase in AEP is achieved partly through a 2% increase in blade length combined with an increase in aerodynamic efficiency.In this optimization, blade mass and blade mass moment were directly constrained, and we see that these two constraints were satisfied. The optimized blade is capable

Constraint	Value	**Comment**
max(chord)	< 4.0 m	Maximum chord limited for transport.
max(prebend)	< 4.0 m	Maximum prebend limited for transport.
min(relative thickness)	> 0.18	Same airfoil series as used on the reference design.
min(material thickness)	> 0.0	Ensure FFD splines do not produce negative thickness.
min(tip tower distance)	$>$ ref value	Tip deflection cannot exceed that of the reference design.
Blade root flapwise moments (MxBR)	$<$ ref value	Simple DLB loads cannot exceed starting point.
Blade root edgewise moments (MyBR)	$<$ ref value	Simple DLB loads cannot exceed starting point.
Tower top thrust (FyTT)	$<$ ref value	Simple DLB loads cannot exceed starting point.
Tower bottom fore-aft moments (MxTB)	$<$ ref value	Simple DLB loads cannot exceed starting point.
Rotor torque	$<$ ref value	Ensure that the rotational speed is high enough below rated to not exceed generator maximum torque.
Blade mass	< 1.05 * ref value	Limit increase in blade mass to maintain equivalent production costs.
Blade mass moment	< 1.01 * ref value	Limit increase in blade mass moment to minimise edgewise fatigue.
Lift coefficient @ $r/R = [0.5 - 1.]$	$<$ ref value	Limit operational lift coefficient to avoid stall for turbulent inflow conditions.
Ultimate strain criteria	$<$ ref value	Aggregated ultimate material failure in each section for 12 load cases.

Table 5.3: Non-linear constraints used in the design process of SRB2.

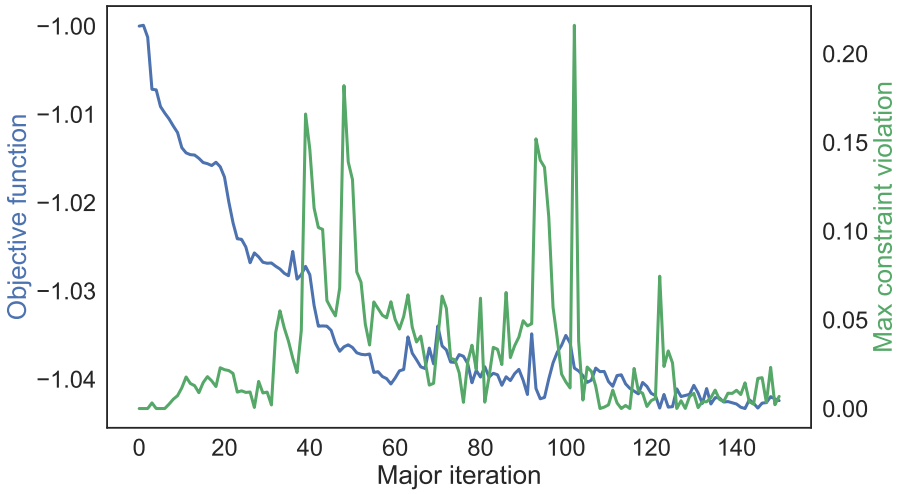

Figure 5.48: SRB2 optimization convergence history showing objective function and maximum constraint violation.

of respecting the loads constraints imposed in the optimization primarily due to the increase of torsional flexibility, which unloads the blade towards rated wind speed.

Quantity	**Difference** [%]
Static AEP (mean_wsp=7.5, k=2)	3.7
Blade length	2.0
Blade mass	-2.9
Blade mass moment	0.2

Table 5.4: Overall characteristics of the SRB2 design compared to the Reference blade.

5.9.3 SRB6

The SRB6 blade was designed at a later stage in the project, following the implementation of the fatigue in frequency domain model [12] and the material fatigue model, which enabled us to model constraints on platform fatigue loads as well as material fatigue strain limits, which for modern flexible blades can be a design driving constraint. The addition of fatigue loads in the optimization has allowed us to remove constraints on blade mass and mass moment, since increases in e.g. mass moment of the blade should now be reflected directly in increased edgewise fatigue loads. The optimization setup otherwise follows that of SRB2, with minor adjustments in design variables. Tables 5.5 and 5.6 provide a summary of design variables and constraints used in this study. In addition to the new model constraints, a constraint on minimum skin thickness of 2 mm was also added to ensure a fairly realistic outer skin.

Parameter	# of DVs	**Comment**
Chord	5	Root chord fixed
Twist	5	Root twist fixed
Relative thickness	4	Root and tip relative thickness fixed
Blade prebend	4	-
Chordwise offset	4	-
Blade length	1	-
Trailing edge uniax	3	Symmetric pressure/suction side
Trailing edge triax	3	Symmetric pressure/suction side
Trailing panel triax	4	Symmetric pressure/suction side
Spar cap uniax	4	Symmetric pressure/suction side
Leading panel triax	3	Symmetric pressure/suction side
Leading edge uniax	4	Symmetric pressure/suction side
Shear web biax	3	Identically sized front and rear shear webs
Spar cap center	2	-
Mold angle	1	-
Tip-speed ratio	1	-
Total	54	

Table 5.5: Design variables used in the optimization of SRB6.

Figure 5.49 shows the optimization history of the SRB6 blade. As observed for the SRB2 blade, the optimization converged almost linearly in the first phase of the optimization, but in the case of SRB6 only reaching an improvement of 2.5%, before a large constraint violation of 400% occurs. The optimization does not improve the objective further for the remainder of the optimization, but reaches a feasible point at iteration 68. A later analysis of the optimization showed that the large constraint violations stems from the fatigue in frequency domain that for some iterates predict very large tower bottom moment loads for specific wind speeds. This indicates that more work is needed to make this model more robust.

Table 5.7 summarises the overall characteristics of the SRB6 blade design. Compared to SRB2, this design does not achieve as high an increase in AEP, achieving a 2.5% increase. Notice that the blade length in this optimization has not increased compared to the Reference Blade, so the improvement comes purely from increase in rotor C_P below rated wind speed. The addition of blade root fatigue loads constraints can be seen to also limit increases in blade mass and mass moment. The torsional flexibility of SRB6 is reduced compared to SRB2, primarily due to more realistic material thickness constraints for the skin material.

Constraint	Value	**Comment**
max(chord)	< 4.0 m	Maximum chord limited for transport.
max(prebend)	< 4.0 m	Maximum prebend limited for transport.
min(relative thickness)	> 0.18	Same airfoil series as used on the reference design.
min(material thickness)	> 2.00 mm	Ensure realistic material thickness in skin.
min(tip tower distance)	$>$ ref value	Tip deflection cannot exceed that of the reference design.
Blade root extreme flapwise moments (MxBR)	$<$ ref value	Simple DLB loads cannot exceed starting point.
Blade root extreme edgewise moments (MyBR)	$<$ ref value	Simple DLB loads cannot exceed starting point.
Blade root extreme torsion moments (MzBR)	$<$ ref value	Simple DLB loads cannot exceed starting point.
Tower top extreme thrust (FyTT)	$<$ ref value	Simple DLB loads cannot exceed starting point.
Tower bottom extreme fore-aft moments (MxTB)	$<$ ref value	Fatigue in frequency domain loads cannot exceed starting point.
Blade root flapwise fatigue moments (MxBR)	$<$ ref value	Fatigue in frequency domain loads cannot exceed starting point.
Blade root edgewise fatigue moments (MyBR)	$<$ ref value	Fatigue in frequency domain loads cannot exceed starting point.
Blade root torsion fatigue moments (MzBR)	$<$ ref value	Fatigue in frequency domain loads cannot exceed starting point.
Tower bottom fore-aft fatigue moments (MxTB)	$<$ ref value	Fatigue in frequency domain loads cannot exceed starting point.
Rotor torque	$<$ ref value	Ensure that the rotational speed is high enough below rated to not exceed generator maximum torque.
Lift coefficient @ $r/R = [0.5 - 1.]$	$<$ ref value	Limit operational lift coefficient to avoid stall for turbulent inflow conditions.
Ultimate strain criteria	$<$ ref value	Aggregated ultimate material failure in each section for 12 load cases.
Fatigue strain criteria	$<$ ref value	Aggregated fatigue material failure based on a frozen loads approach.

Table 5.6: Non-linear constraints used in the design process of SRB6.

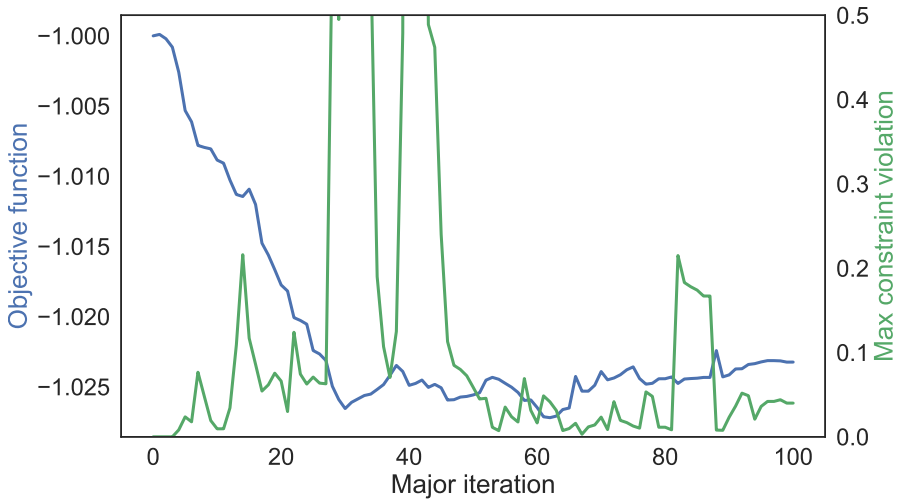

Figure 5.49: SRB6 optimization convergence history showing objective function and maximum constraint violation.

Quantity	**Difference** [%]
Static AEP (mean_wsp=7.5, k=2)	2.5
Blade length	0.3
Blade mass	0.4
Blade mass moment	-1.4

Table 5.7: Overall characteristics of the SRB6 design compared to the Reference blade.

5.10 Design Load Comparisons

The design load basis for the turbines evaluated in this report are based on the IEC-64100-ed3, and described in detail in "Design Load Basis for onshore turbines - Revision 00" by Hansen et al. [10]. The controller used in the project is the DTU Basic Wind Energy Controller [9].

The full design load basis was calculated using HAWC2 [2]. The design load cases included were: 1.2, 1.3, 1.4, 1.5, 2.1, 2.2b, 2.2y, 2.3, 2.4, 3.1, 3.2, 3.3, 4.1, 4.2, 5.1, 6.1, 6.2, 6.3, 6.4, and 8.1. All the cases included dynamic inflow except for 6.1, 6.2, 6.3, 6.4, and 8.1. The dynamic stall model was active in all the cases except for cases 2.2y for wind direction from 60 deg to 300 deg and 6.2 for wind direction from 60 deg to 300 deg.

5.10.1 Discussion

The AEP increases of SRB2 and SRB6 relative to the Reference blade were found to be 3.2% and 2.4%, respectively. The AEP was computed based on DLC12 turbulent inflow cases with a turbulence intensity of 10% for wind class 3B with a mean wind speed of 7.5 m/s and a Weibull scaling factor of 2. Compared to the static AEP increases we observe a slightly lower increase likely due to the increase in torsional flexibility of the blades.

The overall conclusion of the design load evaluation of the SRB2 blade is that the platform extreme loads are for the most part reduced compared to the Reference blade. Notably, the tower bottom fore-aft moment is reduced by 11%, and blade root flapwise moments reduced by 5%, despite the increase in blade length of 2%. However, blade root edgewise and torsional loads are increased by 12% and 20%, respectively, where the dominating load case is DLC22y, which is a case with extreme yaw error of up to 90 degrees. Looking closer at the increase in edgewise loads we see that for normal operation DLC12 and DLC13 cases, the SRB2 blade does not exhibit increases in edgewise loads compared to the Reference blade. On the contrary, the increase in torsional loads are observed under normal operation, and is a consequence of the larger edgewise and torsional flexibility of the blade. It should also be noted that constraints on edgewise and torsional loads were not included in the design optimization of SRB2. Turning to the lifetime equivalent loads, SRB2 overall does not increase platform loads, except for 6% increase in blade root edgewise and torsional moments. Lastly, the ultimate and fatigue strain levels in the materials were evaluated. As described in 5.9.2, constraints on material fatigue were not included in this design, and in fact we observe that the fatigue damage is significantly increased in SRB2 compared to the Reference blade. The ultimate strain levels in the SRB2 blade does not exceed those of the Reference blade, showing that the constraints enforced in the optimization setup captures the extremes sufficiently well. One issue to note is, however, that strain levels in individual materials are not taken into account, only the absolute max of all materials. This simplification results in the strain limits of the shear web biax to be somewhat increased in SRB2 compared to the Reference blade.

Turning to the SRB6 design, this blade was as discussed in Section 5.9.3 designed including new constraints based on the the fatigue in frequency domain model described in Section 5.3, as well as the model for material fatigue, see Section 5.5, with the aim of addressing the design constraint violations observed in SRB2. The load reductions on the platform observed for SRB2 are not realised for SRB6. In fact all platform loads are increased, although most only by small amounts. Turning to the blade root extremes, we observe reductions on all sensors. The flapwise moment is reduced by a small amount, but the edgewise and torsional extremes are reduced by 10% and 24%, respectively, which is significant. Turning to the lifetime equivalent loads, we observe a small increase in tower loads, and similarly small differences for the blade root loads. We can tentatively conclude that the addition of fatigue constraints on the platform loads, seems to limit the large increases in edgewise and torsional loads seen on SRB2. However, clear conclusions cannot be drawn before more detailed investigations are carried out on the individual DLC design load cases. Looking finally at the material damage from extreme and fatigue loading, we observe that the material damage level due to fatigue is only moderately increased for SRB6 compared to the Reference blade. Considering that the constraint is based on a frozen loads assumption, this result is promising. For the damage due to extreme loading, we can observe that the failure index on the outer part of the blade is somewhat higher than for the baseline blade. The explanation for this is most likely that the extreme flapwise case dominating on the outer part of the blade is DLC22y at 26 m/s with extreme yaw error of 90 degrees, which first of all is not a load case included in the simplied set of load cases considered in the optimization, and secondly, it could be questioned whether these special fault cases are

modelled sufficiently correctly to be used in preliminary blade design evaluations. But again, more detailed investigation of these load cases is needed before clear conclusions can be drawn.

5.11 LM Evaluation

In this section the optimized blades made in the project are evaluated using LM tools in order to judge how close the optimized blades are from an actual production ready blade. Because a lot of time went to unforeseen issues described in 4 it was chosen only to evaluate one design SRB2.

Due to IP reasons only certain elements in the LM blade design were shared between LM and DTU. Therefore it was agreed to use standard materials in the optimization process. The exact layup and other materials that go into a LM Blade are taken into consideration in this evaluation step done by LM. By doing this the overall blade properties and aero-elastic responses are close to what will be seen for a production ready blade.

5.11.1 LM's Standard Laminate Plan

LM has more than 30 years of experience in manufacturing blades. This experience is stored in LM's so called Standard Laminate Plan. When doing conceptual blade designs like in this project it ensures that all materials are put in the blade in a standardized way that respects manufacturing rules - rules that are based on massive testing and experience.

5.11.2 The Tool LM Blades

LM Blades is an in-house tool that is used for structural evaluation and optimization. It is based on beam theory coupled with a kind of finite-element approach. Given a set of loads with extreme loads and Markov matrices it can evaluate the strain safeties in a given structure both in fatigue and extreme. Furthermore it can optimize the structure such that the structure is strong enough to carry the loads and furthermore making sure that the deflection is not violating the maximum allowable deflection - doing this by making sure that only the minimum needed amount of material is put into the blade. It uses the Standard Laminate Plan to make sure that manufacturing rules are taken into account.

5.11.3 Evaluation Approach

In order to evaluate the SRB2 from a LM perspective a new LM structure was layed out in the blade using LMBlades where mass and static moment were optimized with strain and maximum allowable deflection used as constraints. This is in the following called **SRB2_lmstruc**. The loads used for this was calculated with HAWC2 using a setup provided by DTU.

The structure in the original reference blade was based on loads calculated by LM in the low fidelity aero-elastic tool LMFlex. In order make the LM structure in SRB2 comparable with the reference blade HAWC2 loads were also calculated on the reference blade. The loads were then passed to LMBlades and the strain safeties and the maximum deflection were computed. These strain safeties and deflection were then used for specifying the allowable minimum strain safeties and maximum allowable deflection for the new LM structure. Because SRB2 has a slightly larger prebend of 0.8 this should be accounted for in the allowable deflection (0.7 is used as safety factor).

With the new LM structure loads were calculated with the same setup as SRB2 and the structure was then updated for the new loads. In the following this is named **SRB2_lmstruc_rev1**. Because the LM structure resulted in responses that were not similar as the original SRB2 a modified version was also considered called **SRB2_lmstruc1**. Here an attempt was made to loosen up some of the manufacturing constraints in the Standard Laminate Plan by using same amount of skin material as in SRB2. It might not be very realistic but at least it might give an indication of what this can bring. Furthermore there might be other ways of achieving the same tendencies which maybe should be looked into in the near future. Like with **SRB2_lmstruc** new loads were calculated and the structure was updated and this is in the following called **SRB2_lmstruc1_rev1**. Below is a summary of the different variants.

- **SRB2_lmstruc**.
  - Standard LM structure in SRB2 using original loads from SRB2.
- **SRB2_lmstruc_rev1**.

  - Standard LM structure in SRB2 using loads from **SRB2_lmstruc**.

- **SRB2_lmstruc1**.
  - Standard LM structure in SRB2 with modified skin layers using original loads from SRB2.
- **SRB2_lmstruc1_rev1**.
  - Standard LM structure in SRB2 with modified skin layers using loads from **SRB2_rev1**.

It should here be mentioned that large loads were observed in some of the extreme load cases (dlc 2.2 pitch stuck and extreme yaw) which therefore for some of the optimized LM structures and due to lack of time it has not been investigated any further. Probably a re-tuning of the controller could have solved the problems.

5.11.4 Results

In this section the overall evaluation results are presented.

SRB2_lmstruc and SRB2_lmstruc1 have a bit larger static moment than SRB2. Despite of that the flap and edge frequencies are also larger meaning that they are stiffer in both flap and edge. They are at the limit regarding deflection but no extra material has been added for deflection. Assuming that strain safeties are not violated on SRB2 it indicates that damage calculation model and/or material properties are not 100% aligned in the two evaluations. Furthermore it can be observed that the torsional stiffness of SRB2_lmstruc is somewhat larger than for SRB2. For this reason SRB2_lmstruc1 was made and it is seen that it boosts the energy yield with about 0.6%. It also corresponds pretty well with the AEP calculated from the time series which is 1%. Having a more torsional soft blade makes it possible to achieve a better aerodynamic performance at a larger range of wind speeds. In addition it can have a positive effect on the flap wise and turbine loading which surprisingly is not very clear observed here since the tip deflection and the flap wise frequency are almost identical in the two variants. In addition the structural design driver is a mixture of fatigue and extreme so it seems not possible to pin point what causes this to e.g. a single load case event. Looking at some of the turbine loads it also looks mixed between increase and decrease in loads for the two variants SRB2_lmstruc_rev1 and SRB2_lmstruc1_rev1. One thing that is quite clear to see is that the pitch moment increases when softening the torsional stiffness, which of course is one drawback. It would be interesting to dig a bit more into the different load cases to find an explanation why a decrease is not clearly observed as it is the case for the SRB2.

5.11.5 Remarks

This evaluation of SRB2 shows that we end up with a blade that produces 3.4% more AEP with a mass increase of 6% and a static moment increase of 9% if the turbine loads remain the same. It is probably doubtful that this can be totally achieved but it clearly shows the potential of using an integrated aerostructural blade design tool like HAWTOpt2. For fairness it should be mentioned that the reference blade was designed years back without including torsion in the twist distribution. Furthermore the evaluation showed that when achieving more torsional response it increases energy yield by 0.6%-1% when using a constant pitch below rated. Computation shows 1.4% increase when applying a pitch ramp. Unresolved issue is to gain more knowledge about why no more load alleviation is not seen in the case when comparing SRB2_lmstruc_rev1 and SRB2_lmstruc1_rev1.

6 Utilization of project results

The two main results of the project are: Implementation of a state-of-the-art aerostructural wind turbine design tool HAWTOpt2 in an industrial context, and demonstration that aerostructurally tailored blade designs with torsional passive load alleviation can increase AEP by over 3%, while not increasing platform loads, thus allowing for considerable savings in levelized cost of energy. The project also showed that there is more work to be done both on the software side and blade design side before such blade designs are technologically mature to be commercialized. Both project partners will continue to pursue this field of research and development towards maturing further the technology for industrial use.

Key outcomes for the two project partners DTU Wind Energy has long been established as a leading institution for research in wind energy. This project has helped position DTU among the few research institutions with a strong competence within the field of multidisciplinary design of wind turbines. A direct outcome of this effort is that DTU Wind Energy has a leading role in the IEA Wind Task 37: Wind Energy Systems Engineering - Integrated RD&D, which is an international platform for research in integrated design of wind turbines and wind plants. Other collaborations on integrated blade design have also been established, both of commercial and research oriented nature. The collaboration with LM Wind Power has provided DTU with invaluable insight into design drivers for industrially designed wind turbine blades. This helps DTU to mature the software developed towards addressing the challenges that the industry faces. The use of HAWTOpt2 in LM Wind Power for blade design, provides an important basis for future collaboration between DTU and LM Wind Power. This linked to the fact that LM are also users of a series of other software developed at DTU.

LM Wind Power can utilize several aspects of the results of the project. The implementation of HAWTOpt2 in the company has as a start provided a source of second opinion on specific design challenges, but could potentially be integrated further in the future and be adopted as the primary design tool, given that the design framework is matured further. The open source nature of the framework surrounding the analysis tools used, also makes it possible for LM to customize the tool to their specific needs, and allows for integration of additional in-house software modules. As to the blade designs developed in the project, incorporation of new technologies in products is central for LM to remain industry leading on the global market. From current market trends, it is clear that large capacity factor rotors is key to success on many markets. Aeroelastic tailoring is a key blade technology to continue lowering LCOE, and the results of the project are therefore very relevant for LM. While the final SRB designs developed in the project only partially fulfilled the project goals, the designs shows promising results, which LM can use as basis for future technology development towards further optimization of wind turbine blades including passive load alleviation.

The project model used in Stretched Rotor, in which DTU and LM worked closely together to develop open source blade design software and apply these to develop advanced blade design technologies, is seen as a success, and very valuable for both organisations. On one hand, DTU gained valuable insight into the industrial requirements for such a tool, and on the other hand LM could directly influence the development and specifications of the tool. From DTU's perspective a follow-up project continuing to develop the design tools collaboratively towards designing a specific blade would be very attractive. There are, however, at present no confirmed financing for such a project, but this will be pursued.

7 Project conclusion and perspective

The objective of the project was to implement a new state-of-the-art multidisciplinary design tool in a collaboration between DTU Wind Energy and LM Wind Power, and use it to demonstrate that passive load alleviation blade technologies could result in up to 3% increase in annual energy production without incurring increases in platform loads. The project succeeded in implementing a complete tool chain for designing blades in an integrated manner, and this design framework is now actively being used and developed at DTU Wind Energy and at LM Wind Power. The project also produced two blade designs that showed that it was possible to increase the annual energy production more than 3% with respect to the reference blade. However, the most promising blade, SRB2, was designed with loosened structural geometric constraints, which means that more work is needed to develop a structural design with on one hand the desirable load alleviating characteristics, while at the same time being manufacturable. The SRB6 design included a number of additional manufacturing constraints and model constraints on material fatigue damage, as well as platform fatigue loads based on a frequency domain approach. The results showed that the new model additions still suffer from robustness challenges, which adversely affected the stability of the numerical optimizations.

Below is a summary of the main results achieved in the project:

- The multidisciplinary blade design framework HAWTOpt2 was developed to meet the requirements of LM. The framework was installed and is operational at LM. HAWTOpt2 enables us to explore advanced blade designs that incorporate passive load alleviation, and includes modelling of key design driving constraints. Compared to other available design tools, HAWTOpt2 is amongst the most sophisticated, owing to a large extent to the advanced aeroelastic and structural codes HAWC2, HAWCStab2, and BECAS used as the backbone of the framework.
- Implementation of new geometric parameterisation suitable for commercial blades: The parameterisation allows to quickly develop a blade structural geometry according to simple geometric rules, and was used for all designs developed in the project.
- Development of a new material fatigue model: Modern flexible wind turbine blades are often limited by fatigue in the composite materials used in the blades. A numerically efficient model for evaluating fatigue load capacity in the integrated design framework was therefore implemented.
- Implementation of a frequency domain based fatigue loads model: This model allows us to evaluate fatigue loads on the wind turbine in a numerically efficient manner. The model was previously implemented, but has in this project been refined and extended to also estimate extreme loads based on normal operation with extreme turbulence.
- Development of inverse design method for blades: To establish the reference blade used in the project, a method was developed to create a detailed internal structural design based on an aeroelastic definition of the blade. This method solves a feasibility problem to match pre-defined structural characteristics by sizing material thicknesses and internal structural geometry.
- Development of a cross-sectional mesher: The project faced challenges to accurately model the geometric complexity of an industrial grade blade. It was therefore decided to pursue the development of a new cross-sectional meshing tool that could accurately model geometric details not possible with the meshing used. The tool was demonstrated to achieve this goal, but has not yet been incorporated in the design framework.
- Development of a multi-fidelity methodology to increase accuracy of loads constraints: A key challenge in integrated design is to capture the design driving loads correctly, while at the same time keep down computational time. This development aimed at addressing this through a multi-fidelity approach, where detailed load calculations were used in an outer loop to correct a lower fidelity load estimator inside the optimization loop. The methodology was demonstrated on a sample structural design problem and showed promising results, but faced challenges in the application to loads correction, due to one one hand implementational difficulties in HAWTOpt2 as well as difficulties in computing gradients of the quantities of interest with sufficient accuracy.

- Development of a method to include multiple airfoil families and add-on devices in the integrated design tool: A multidimensional interpolation scheme was developed for including multiple airfoil families, airfoils with add-ons or varying degree of soiling in the optimization. It was demonstrated on a simple example, but not used for the final SRB designs due to time constraints.
- Development of a series of aeroelastically tailored blade designs: HAWTOpt2 was applied to design a number of blades aiming to maximise annual energy production subject to constraints to not exceed the loads of the baseline design.
  - The SRB2 design achieved a 3.2% increase in AEP while also reducing tower bottom extreme loads by up to 10%. However, blade root edgewise and torsional loads increased, and the materials were allowed to be unrealistically thin. Also, the design did not include constraints on material fatigue.
  - The SRB6 design was designed with realistic material thickness constraints, and achieved a 2.4% increase in AEP, including constraints on material fatigue and blade root edge and torsional loads that were now not exceeded.
- LM Blades Interface: LM developed an interface between HAWTOpt2 and their internal design tool LMBlades which was used to develop blade designs based on SRB2 and SRB6 according to LM design standards, materials and airfoil geometries, which was not shared with DTU due to confidentiality.
- Detailed Design: LM carried out a series of detailed designs of blades based on SRB2 taking into account LMs detailed manufacturing constraints and design rules. It was found that the SRB2_lmstruct_rev1 blade achieved 2.8% increase in AEP, but due to slight increases in stiffnesses and mass, minor increases in platform loads were observed. By loosening up some manufacturing constraints to lower the torsional stiffness an AEP increase of 3.4% was achieved. Work needs to be done in order to find a way to realize this.
- Commercial milestone on a demonstration project: LM decided not to pursue a demonstration project directly based on the SRB designs. Regarding the software developed in the project, the HAWTOpt2 design tool will remain available to LM after the finalization of the project. LM and DTU plan to collaborate on further development of the design framework, but the exact terms of this collaboration has not been settled.

Bibliography

[1] http://becas.dtu.dk/.

[2] http://www.hawc2.dk/.

[3] http://pyopt.org/.

[4] http://openmdao.org, 2018.

[5] N. M. Alexandrov, R. M. Lewis, C. R. Gumbert, L. L. Green, and P. A. Newman. Approximation and model management in aerodynamic optimization with variable-fidelity models. *Journal of Aircraft*, 38(6):1093–1101, Nov. 2001. ISSN 0021-8669. doi: 10.2514/2.2877. URL http://dx.doi.org/10.2514/2.2877.

[6] J. P. Blasques and M. Stolpe. Multi-material topology optimization of laminated composite beam cross sections. *Composite Structures*, 94(11):3278 – 3289, 2012. ISSN 0263-8223. doi: http://dx.doi.org/10.1016/j.compstruct.2012.05.002.

[7] P. Blasques. User ' s Manual for BECAS A cross section analysis tool for anisotropic and inhomogeneous beam sections of arbitrary geometry e Pedro Blasques. 2012.

[8] DTU. Beam cross section analysis software. http://www.becas.dtu.dk/, 2018. URL http://www.becas.dtu.dk/.

[9] M. Hansen and L. Henriksen. *Basic DTU Wind Energy controller*. DTU Wind Energy, Denmark, 2013. Contract Number: EUDP project Light Rotor; Project Number 46-43028-Xwp3 The report number is DTU-Wind-Energy-Report-E-0028 and not DTU-Wind-Energy-Report-E-0018 as stated inside the report.

[10] M. Hansen, K. Thomsen, A. Natarajan, and A. Barlas. *Design Load Basis for onshore turbines - Revision 00*. DTU Wind Energy, Denmark, 2015.

[11] J. Jonkman, S. Butterfield, W. Musial, and G. Scott. Definition of a 5-mw reference wind turbine for offshore system development. Technical report, National Renewable Energy Laboratory, 2009.

[12] G. R. Pirrung. *Computation of fatigue and extreme loads using a linear model*. DTU Wind Energy Report-I-0761, Roskilde, Denmark, 2018.

[13] C. Tibaldi, L. C. Henriksen, M. H. Hansen, and C. Bak. Wind turbine fatigue damage evaluation based on a linear model and a spectral method. *Wind Energy*, 19(7):1289–1306. doi: 10.1002/we.1898.