

Incorporating High-Fidelity Aerostructural Analyses in Wind Turbine Rotor Optimization

Denis-Gabriel Caprace*, Adam Cardoza†, Andrew Ning‡
Brigham Young University, Provo, Utah 84602

Marco Mangano§, Sicheng He¶, Joaquim R. R. A. Martins||
University of Michigan, Ann Arbor, Michigan 48109

As demand grows for wind turbines with larger blades, the design of future wind turbines must account for multi-physical interactions and an ever-increasing number of design load conditions. One aspect, aerostructural coupling, calls for design tools that are both accurate and computationally efficient. In this paper, we present a combined-fidelity approach that couples high-fidelity computational fluid dynamics and computational solid mechanics simulations, with a conventional aeroelastic turbine modeling tool based on blade element momentum and beam theories. The approach is integrated into a multidisciplinary optimization framework. It takes advantage of the high-fidelity tightly-coupled aerostructural simulations to evaluate the rotor power production and uses conventional unsteady simulations to add structural sizing constraints. We show that the overall method is effective to obtain improved optimal designs that are resistant to extreme and fatigue loads. Finally, we discuss the computational cost and benefits of the proposed approach for the design of wind turbine rotors.

Nomenclature

D	=	structural damage due to fatigue loads
DEL	=	damage-equivalent load
ϵ	=	strain
η	=	safety factor
E	=	Young modulus (material property)
$\mathbf{g}$	=	gravity acceleration vector
F_n, F_t	=	blade normal and tangential forces (in rotor frame)
L^R, L^M	=	load range and average in a cycle
L^{ult}	=	ultimate load (material property)
m	=	Wöhler exponent (material property)
N	=	number of cycles to failure
r, R	=	radial location and rotor radius
V	=	inflow velocity
Y	=	structural yield due to extreme loads
Ω	=	rotor rotation vector

Subscripts/superscripts

EXTR	=	extreme
life	=	lifetime extrapolated or aggregated

*Post-doctoral Research Fellow, Department of Mechanical Engineering; dcaprace@byu.edu. Member AIAA.

†Ph.D. Candidate, Department of Mechanical Engineering. Student Member AIAA.

‡Associate Professor, Department of Mechanical Engineering; Associate Fellow AIAA.

§Ph.D. Candidate, Department of Aerospace Engineering. Student Member AIAA.

¶Post-doctoral Research Fellow, Department of Aerospace Engineering. Member AIAA.

||Professor, Department of Aerospace Engineering. Fellow AIAA.

I. Introduction

SINCE the first industrial use of wind turbines, design trends for new turbines have pushed for ever larger rotors to increase the power harvested from the wind. Nowadays, the blades of utility-scale turbines often reach dimensions superior to 100 m. At those scales, the blade flexibility and the resulting aerostructural coupling considerably influence the turbine operation and lifespan. Accounting for this multidisciplinary interaction —intensified by environmental conditions— is critical in the design of large turbines and requires detailed and comprehensive simulation capabilities. Additionally, the number of operating conditions that must be considered in the design process has also increased over time. Any turbine design must comply with standards, such as IEC 61400-3 [1], that formulate requirements to guarantee the safety and the integrity of the system throughout its life. The standards define a number of design load cases (DLCs) that represent usual and extreme conditions that the turbine may encounter. The design tools for the next generation of turbines must ensure that the turbine is able to withstand all DLCs while enabling the exploration of design spaces of increasing complexity and size. In this paper, we present a new procedure to combine a high-fidelity aerostructural solver with conventional unsteady turbine simulation tools, within a multidisciplinary design optimization (MDO) framework, for the design of a wind turbine rotor.

Recently, MDO has become popular for wind turbines as it offers a systematic approach to integrate various physics and different requirements in the design process. Figure 1 illustrates the features accounted for in the present rotor MDO, focusing on blade aerostructural interactions. On that topic, multiple works have recently demonstrated that passive load alleviation can be achieved through bend-twist coupling, leading to lighter blade designs [2–4].




Figure 1 Wind turbine MDO of a rotor integrating environmental parameters (in green), design variables (in black), output functions (in blue) and other requirements (in pink).

To accommodate MDO, various simulation tools have been considered. On one end of the spectrum, there are conventional tools that run relatively fast but involve rather strong assumptions. They typically rely on the coupling of blade element momentum (BEM) codes and beam solvers which originally enabled the aerostructural optimization of blades [3]. Additional modeling layers including other turbine parts and unsteady simulation capabilities have further enabled comprehensive wind turbine optimizations driven by cost metrics [5–7]. Extension of such conventional tools to the optimization of floating turbines has been proposed by adding the physics of the platform and mooring systems [8].

Increasing the fidelity of the computational models, Bottasso et al. [9, 10] and Bortolotti et al. [11] have gradually introduced 2D and 3D structural finite element (FE) simulations and 2D computational fluid dynamics (CFD) solvers. Their results further highlighted the strong effects of aerostructural coupling that wind turbines are subject to. These effects are best accounted for in a coupled or monolithic approach.

At the far end of the spectrum, we consider “high” fidelity simulation tools based on 3D CFD and computational structural mechanics (CSM) codes. These options are even more computationally expensive but present several benefits for the overall MDO. First, CFD and CSM models typically allow for a fine and accurate description of the rotor geometry, which renders blade shape optimization easier [12]. Typically, CFD offers good solution accuracy over the entire blade radius, including the root and tip regions where the BEM theory typically resorts to empirical corrections [13]. In the latter reference, the high-fidelity optimization gave higher improvements of the objective as compared to optimal solutions obtained with conventional tools. Flow conditions around non-conventional geometries such as curved tips [14] or associated with off-design conditions such as a yawed inflow or the turbulent wake state [15] are also better captured with CFD. Considering aerostructural problems, Mangano et al. [16] provided a preliminary demonstration that the use of tightly coupled high-fidelity simulation tools in an optimization leads to more optimal designs, compared to the historical discipline-per-discipline or loosely coupled approaches. These promising trends motivate us to further develop and expand high-fidelity optimization frameworks.

In the end, wind turbine optimization is subject to a computational trade-off. High-fidelity methods increase accuracy together with computational expenses, whereas conventional tools are more affordable but may lead to conservative or sub-optimal designs because of the underlying assumptions of the models. A possible trade-off resides in the combination of different kinds of tools within the same optimization framework. Such “multi-fidelity” (also termed “mixed-fidelity”) approaches come in different forms. For example, the high-fidelity solver of a given discipline can be accelerated using results from a low-fidelity solver of the same discipline [17, 18]. Different levels of fidelities can also be used for various disciplines in the turbine model, e.g. a simplified aerodynamic solver with 3D structural FE [10] or a beam model with CFD [15]. Following a similar idea, the present work does not aim to replace well-established conventional optimization methods. Instead, we explore how high-fidelity tools can be incorporated with those methods in order to take advantage of both.

Finally, an inherent challenge of MDO resides in the optimization itself. For complex systems such as wind turbines, the number of design variables (DVs) is large. Therefore, gradient-based optimization is more computationally efficient than gradient-free approaches [19]. For this reason, many of the above-mentioned studies adopted the former. However, the robustness and effectiveness of such optimizations greatly rely on the accuracy of the gradients, which often makes simple techniques such as finite-differencing unsuitable. The full power of gradient-based optimization becomes available when exact gradients are available from the simulation tools [20], a capability that we exploit here.

In this work, we propose a new approach for the coupling of a high-fidelity CFD-CSM-based rotor simulation tool with a conventional, comprehensive turbine simulation tool. We develop a multi-fidelity paradigm that we refer to as *combined fidelity*. The approach builds upon the high-fidelity numerical framework developed by Mangano et al. [16]. In their work, they illustrate how the tight CFD-CSM coupling can benefit the blade optimization. However, their application is limited to a single load case. Generally, they suggest that only few design points can be included in a high-fidelity optimization formulation due to high computational expenses. We present an extended numerical framework which raises that limitation, following a treatment of DLC simulations inspired by Ingersoll and Ning [21]. The resulting framework is used for gradient-based optimization of a wind turbine blade with analytic gradients. The preliminary results that we show constitute a first step towards more complex problems.

The present paper is organized as follows. Section II presents the conventional tool (based on *OpenFAST*) and the high-fidelity tools, all integrated in the gradient-based aerostructural optimization framework *MACH*. In Section III, we detail how the combined fidelity optimization is formulated, with a particular attention to the constraints which involve passing information between different fidelity levels. In Section IV, we present the 10 MW reference turbine model employed as a baseline to our optimization, and we provide some verifications of the consistency of the approach. Finally, we demonstrate the effectiveness of the approach in Section V by analyzing simple optimization results. We show that even though the blade structure is mainly sized by the constraints derived from the various DLCs, MDO provides a convenient way to further explore the aerostructural trends captured with the coupled high-fidelity environment. We close the paper with comments on the numerical behavior and cost of the overall approach.

II. Simulation and Optimization Tools

In this section, we briefly review the numerical simulation tools that we use, classified as “conventional” to designate those commonly used for turbine design, and “high-fidelity” to refer to coupled CFD-CSM-based methods. We add a note on the optimizer.

A. Conventional Tools

Unsteady wind turbine simulations are performed with OpenFAST*, conveniently wrapped within WEIS†. WEIS integrates several capabilities to facilitate the pre- and post-processing of OpenFAST simulations, including the turbine definition, the generation of IEC wind conditions, the tuning of controllers, and the processing of output files. The entire toolset is developed and maintained by the National Renewable Energy Laboratory. We briefly describe it below, and we refer the reader to the software manuals for more details.

OpenFAST is a comprehensive tool for the aero-servo-elastic simulation of wind turbines. We restrict our use to the simulation of rotors. Its interface with WEIS makes it easy to simulate a variety of cases including different velocities and turbulent conditions, as defined in the standard DLCs. WEIS automatically tunes the Reference OpenSource Controller for the input turbine. The rotor aerodynamics is modeled with a dynamic BEM model. The blade structure model uses a beam representation, the properties of which are derived from the description of the inner blade geometry and classical laminate theory.

B. High-fidelity Tools

The simulation tools are comprised of the CFD code ADflow and the CSM code TACS. These codes are tightly coupled through MACH. Again, we provide a concise description here, and we refer the reader to the cited references for implementation details.

ADflow [22, 23] solves the steady-state compressible RANS equations complemented with a Spalart-Allmaras turbulence model, in a finite volume scheme. It utilizes both Newton–Krylov (NK) and approximate NK methods to efficiently converge the flow residuals, with a preconditioner based on the work exposed in [24]. ADflow provides exact sensitivities of aerodynamic functions to design variables, that are computed based on the combination of automatic differentiation and a discrete adjoint method. Advantageously, the computational cost of the adjoint is independent of the number of design variables (but scales with the number of functions of interest, that is, output quantities that are required to form the objective function of the constraints). This enables the efficient treatment of optimization problems with a larger number of design variables than functions of interest, such as the problem at hand.

The *Toolkit for the Analysis of Composite Structures* (TACS) is a finite-element code originally developed to work with shell-structures [25]. TACS computes the global deformation, strain and stresses in a structure under a given loading, and outputs information about material failure and buckling. It also features an adjoint method for the computation of gradients, which makes it well suited for solving structural optimization problems.




Figure 2 XDSM diagram of the MACH framework. [26]

MACH has been extensively used for aero and aerostructural optimization of aerodynamics applications (see the numerous examples listed by Martins [20]), including the aerodynamic shape optimization of wind-turbines [12, 13].

*OpenFAST, v2.6.0, <https://github.com/OpenFAST/openfast> (May 2021).

†Wind Energy with Integrated Servo-control (WEIS), v0.2, <https://github.com/WISDEM/WEIS> (May 2021).

MACH provides a set of methods to facilitate the generation and the deformation of aerodynamic and structural meshes, as part of a comprehensive aerostructural optimization procedure (formalized in the XDSM diagram shown in Fig. 2). In particular, the coupling between ADflow and TACS is handled through a nonlinear Gauss–Seidel method or the coupled NK method. The linear block Gauss–Seidel method or the coupled Krylov method is used to solve the coupled adjoint equation for the derivatives [26], a key feature required for gradient-based optimization.

The current work extends MACH capabilities to support both high and low-fidelity simulations tools, combined in an optimization problem formulated as described in the next section.

C. Optimizer

We employ *SNOPT* which implements a sparse sequential quadratic programming algorithm for the optimization of nonlinear problems [27]. Since this algorithm takes advantage of sparsity, we use the interface to SNOPT provided in *pyOptSparse*[‡] to ease the handling of sparse Jacobian matrices.

III. Combined Fidelity Optimization Framework

The motivation of the present approach is to leverage the relatively fast computational time of unsteady load cases provided by conventional simulation tools, while also relying on the high-fidelity model which provides a more detailed description of the blade and potentially more accurate solutions. Because of the associated cost of unsteady simulations (and current limitations of the tools), we limit the high-fidelity evaluations to steady conditions representative of the rotor power production. The response of the turbine to other potentially unsteady conditions is evaluated through conventional simulations. The purpose of the conventional simulations is to assess if the turbine design can withstand the projected extreme and fatigue loads. Putting together steady high-fidelity results and unsteady loads from conventional simulations, the combined fidelity optimization produces a solution which is compliant with the selected DLCs and takes advantage of realistic trends in the high-fidelity design space.




Figure 3 Block diagram of the combined fidelity approach. The distinct fidelities are shown in blue. Explicit components are shown in green, implicit components in red, and inputs in gray. The sequence WEIS → MACH is placed in an outer iteration loop with feedback shown in purple.

The algorithm is illustrated in Fig. 3. The starting point is an initial blade design, the corresponding aerodynamic and structural meshes, and a set of DLCs that the rotor will be evaluated on. First, the blade information is translated to a beam model, then WEIS performs unsteady simulation of the corresponding DLCs. The internal processing involves generating turbulent inflows, running OpenFAST, and determining extreme and fatigue loads from simulation

[‡]PyOptSparse, v2.6.2, <https://github.com/mdolab/pyoptsparse> (May 2021).

outputs (as detailed in Section III.A). Extreme and fatigue loads are the main information we collect from the unsteady simulations, as they will be used to constrain the high-fidelity optimization. They are then passed to the high-fidelity optimization module, following a procedure that we discuss in Section III.B.

The high-fidelity optimization module hosts the actual MDO, based on a setup which extends the work presented by Mangano et al. [16]. As in the original version, the module features an internal aerostructural kernel that evaluates an objective function and constraints. For instance, this allows us to maximize the torque extraction under tip displacement constraint. Our extension consists in adding additional components to the optimization, to evaluate how the current blade performs under extreme and fatigue loads. Extreme loads may cause the blade to yield, and we constrain the optimization to prevent it (as explained in Section III.C). Similarly, fatigue loads lead to an accumulated damage which we constrain as well (as explained in Section III.D). The aerostructural and the additional structural analyses are arranged in a multi-point optimization, in this case meaning that they are run simultaneously and contribute to constraint evaluations at every iteration of the optimization.

It must be noted that, even though the aerostructural kernel recomputes the loads related to the steady power production at every iteration, the extreme and fatigue loads are frozen during the optimization. However, these loads are also theoretically affected by modifications in the blade design. To resolve this, WEIS and MACH are put in a loop (hereafter referred to as *outer loop*). The sequence "unsteady simulation" and "aerostructural optimization" is repeated until convergence, assessed based on the difference in the previous and updated extreme and fatigue loads.

The computation of the unsteady loads outside of the optimization loop was proposed by Ingersoll and Ning [21]. The authors showed that this approach leads to the same solution as a fully coupled optimization. It has the tremendous advantage that the aerostructural optimization does not require the derivatives of the extreme and fatigue loads with respect to the DVs since the loads are assumed constant. Obtaining such derivatives is generally a computationally expensive operation. The downside is the requirement of an outer loop, where an aerostructural optimization problem must be solved at each iteration. Nevertheless, since the loads are only marginally affected by the blade structural design, we expect the outer loop to converge after a few iterations. A more detailed view of the approach is provided in the Appendix through a complete XDSM diagram (see Fig. 18).

A. Computation of Extreme and Fatigue Loads

The simulation of each DLC produces time series of loads L . Each simulation typically captures 600 seconds of physical time. We first process the time series to determine the short-term damage and extreme loads experienced by the turbine during the simulations. Then, we aggregate these values among all the evaluated load cases and extrapolate them to obtain the expected lifetime damage and the extreme load.

We define the extreme load as the result of an extreme event that will not occur more often than every 50 years. There are various techniques to estimate them based on a limited set of simulated time series [28] ; we proceed as follows. We monitor the aerodynamic loads in the time series and we build their binned probability distribution. Each bin counts the number of times that a load with a certain intensity occurs in the time series, divided by the total number of samples in the simulation. We aggregate the distributions from all the simulations by performing a weighted sum of the binned distributions. The weight is chosen based on the expected percentage of time the turbine will operate in the corresponding conditions, assuming a turbine availability of 1. Here, the probability p_j that the turbine operates in certain condition is taken from the wind probability distribution. The result is a discrete probability density function (*pdf*) under the form of a histogram, as shown in Fig. 4.

To predict an event with a return period longer than the duration of the simulation, we fit an analytic probability distribution to the histogram. For aerodynamic loads, we determined that a Gaussian distribution provided the best match. The mean and variance of the fitted distribution are tuned to minimize the least square error with the data. The fit is performed on the logarithm of the probability of exceedance ($1 - pdf$), since we are mainly interested in matching the tail of the distribution. Finally, based on the fitted distribution, we determine the extreme load L^{EXTR} as the load which gives the probability

$$P[L > L^{\text{EXTR}}] = 1 - \frac{dt}{50 \text{ year}}, \quad (1)$$

where dt is the sampling period, that is, the time step of the conventional simulation. An example of the typical quality of the fit is presented in Fig. 5, together with the result of the extrapolation. This process is repeated for every spanwise station along the blade, so as to obtain the spanwise distribution of extreme loads.




Figure 4 Probability density function of the aerodynamic loads measured at $r/R = 0.5$. The tail of the distribution is fitted with a normal distribution (dashed line).




Figure 5 Probability of exceedance of the aerodynamic loads measured at $r/R = 0.5$, and extrapolation of the extreme load. The “x” symbol shows the load with a return period equal to 50 years.

Blade fatigue originates from the gravitational and the turbulent conditions experienced by the blade. Veers [29] and Ravikumar et al. [30] recently reviewed the state of the art of fatigue computation for wind turbines. In this work, we follow the guidelines set forth in the MLife documentation [31]. Our approach relies on the computation of Damage-Equivalent Loads (DEL), that is, a set of steady loads which lead to the same damage as the entire load spectrum the turbine will experience over its lifetime. We here design the blades for a 20 years lifetime duration.

Our simplified implementation works as follows. Rainflow counting is applied to every time series, which decomposes the load signal into a succession of load cycles. The i^{th} cycle in the j^{th} time series has a load range $L_{i,j}^R$ and average load $L_{i,j}^M$. We assume a power law to model the S-N curve, that is, the relation between the loading conditions and the number of cycles to failure N :

$$N_{ij} = \left(\frac{L^{\text{ult}} - |L_{i,j}^M|}{\frac{1}{2}L_{i,j}^R} \right)^m, \quad (2)$$

where L^{ult} is the ultimate load and m is the Wöhler exponent, both being material properties. We work with $m = 10$, a typical value for composite materials. Following Palmgren-Miner’s rule, we assume the structural damage caused by each cycle adds up linearly,

$$D_j = \left(\sum_i \frac{1}{N_{i,j}} \right). \quad (3)$$

We then aggregate the lifetime damage D by summing the damages accumulated over the short-term time series rescaled by a factor, f_j , defined as

$$f_j = \frac{p_j T^{\text{life}}}{T_j}, \quad (4)$$

where T_j is the duration of the time series and T^{life} is the turbine lifetime chosen for design. Hence, the lifetime damage reads:

$$D = \sum_j f_j D_j = \sum_j f_j \left(\sum_i \frac{1}{N_{i,j}} \right) = \left(\frac{L^{\text{ult}}}{\text{DEL}^{\text{life}}} \right)^{-m}, \quad (5)$$

where we introduced the lifetime damage-equivalent load. Assuming that one can use a Goodman correction for each load cycle to define an equivalent load range $L_{i,j}^{R0}$ such that $N_{i,j} = \left(\frac{L^{\text{ult}}}{\frac{1}{2} L_{i,j}^{R0}} \right)^m$, we finally obtain the expression:

$$\text{DEL}^{\text{life}} = \frac{1}{2} \left(\sum_j f_j \left(\sum_i (L_{i,j}^{R0})^m \right) \right)^{1/m}. \quad (6)$$

We follow this process and compute damage-equivalent normal and tangential aerodynamic loads for every spanwise station of the blade.

B. Transfer of Extreme and Fatigue Loads to the High-fidelity Module

One unique aspect of this work is that we aim to transfer extreme and fatigue loads from the conventional model to the high-fidelity model. We identify two ways to accomplish this. The first one consists in working with the beam internal forces and moments, which is the approach commonly used in the literature. The ensuing damage and yield constraints in the optimization then still rely on a beam representation. We will refer to it as the “conventional implementation”. The second way is based on external loads instead: aerodynamics loads, gravitational loads, and inertial loads due to rotation. By applying the extreme and fatigue components of each of these loads to the FE model, we can derive the corresponding constraints. We will refer to it as the “combined fidelity implementation”.

In this work, we chose to explore the second option as the use of the high-fidelity structural model may potentially yield higher accuracy predictions. However, since the results from conventional simulations are only a function of the radius, we need to extrapolate the extreme and fatigue load distributions so that they can be used with the 3D models. Hereafter, we explain how we reconstruct the 3D aerodynamic loads.

The procedure that we developed uses a precomputed reference aerodynamic solution to extrapolate the load in the chordwise direction and reconstructs the load over the entire airfoil. A similar technique was proposed by Bottasso et al. [32]. The reference aerodynamic solution is computed using CFD for the baseline blade at rated power. It provides a stress distribution σ^{ref} in $[N/m^2]$ over the surface of the blade, expressed as a function of the radial location r and the curvilinear coordinate along the airfoil s (Fig. 6b). The corresponding reference spanwise distributions of normal and tangential loads (Fig. 6a), in $[N/m]$, are

$$F_n^{\text{ref}}(r) = \int_{\partial s} \left(\sigma^{\text{ref}}(r, s) \cdot \hat{\mathbf{e}}_n \right) ds, \quad (7)$$

$$F_t^{\text{ref}}(r) = \int_{\partial s} \left(\sigma^{\text{ref}}(r, s) \cdot \hat{\mathbf{e}}_t \right) ds, \quad (8)$$

where ∂s denotes the closed contour of the airfoil, and $\hat{\mathbf{e}}_n, \hat{\mathbf{e}}_t$ are unit normals in the normal and tangential direction respectively.

On the other hand, from the conventional simulation, we obtain target spanwise load distributions, $F_n^{\text{target}}, F_t^{\text{target}}$, also in $[N/m]$, which are only function of the radius (Fig. 7a). We define a scaling factor as $\lambda(r) = \frac{F_n(r)}{F_n^{\text{ref}}(r)}$. The 3D extrapolated stress distribution (Fig. 7b) is then obtained as

$$\sigma = \lambda(r) \sigma^{\text{ref}}(r, s). \quad (9)$$

By definition, the resulting spanwise normal force distribution $F_n = F_n^{\text{target}}$ matches the target. However, since every stress vector is multiplied by a scalar, the procedure likely introduces a small discrepancy between the target tangential load distribution F_t^{target} and the extrapolated one,

$$F_t(r) = \lambda(r) \int_{\partial s} \left(\sigma^{\text{ref}}(r, s) \cdot \hat{\mathbf{e}}_t \right) ds. \quad (10)$$




Figure 6 Reference aerodynamic solution




Figure 7 Extrapolated aerodynamic loads. The solid line shows the 3D-extrapolated loading, and the dashed line the target loading.

We favored maintaining a consistent solution over the airfoil, in place of multiplying by different factors the normal and tangential component of the stresses.

Eventually, the loads extrapolated on the aerodynamic surface mesh are transferred to the structural mesh using the standard routine as described by Kenway et al. [26].

We note that, as long as the aerodynamic shape does not change during the optimization, the reference solution does not need to be recomputed between each outer iteration of the combined fidelity approach. Otherwise, the reference could be updated based on the most recent aerodynamic solution computed within the aerostructural kernel. This would ensure a better consistency of the solution of an aerostructural optimization with shape variables.

C. Definition of the Yield Constraint

We define both the conventional and combined fidelity implementations of the yield constraint. Note that MACH currently only supports the latter. Future versions will include the conventional implementation as well.

Conventional implementation In this work, we consider only structural yield caused by extreme loads computed

from the conventional simulations. Extreme values of the beam internal forces and moments are obtained using the method outlined in the previous section. Their distribution over the blade radius is used to determine the corresponding strains ϵ^{EXTR} . We assume that the local strain ϵ at a given location (x, y) in the beam cross-section can be obtained from the edgewise bending moment, flapwise bending moment and radial load, resp. M_x , M_y and F_z , as

$$\epsilon = - \left(\frac{M_x}{EI_{xx}} y - \frac{M_y}{EI_{yy}} x + \frac{F_z}{EA} \right), \quad (11)$$

where E is the Young's modulus of the cross section, I_{ii} is the mass moment of inertia about axis i , and A is the cross sectional area (see also [33, eq.11.10]). Finally, we constrain the strain calculated at the location of the spar caps (pressure side and suction side), and we obtain a yield constraint:

$$Y = \frac{\eta_Y \epsilon^{\text{EXTR}}}{\epsilon_{\max}} \leq 1, \quad (12)$$

where $\epsilon_{\max}$ is the yield strain of the material or laminate, and $\eta_Y = 1.35$ is a safety factor to account for uncertainty in the load estimation process and in the material properties.

Combined fidelity implementation The yield constraints in the MACH optimization are obtained through an evaluation of the FE model. We apply the 3D reconstructed aerodynamic loading to the outer skin of the blade model. Since gravitational and rotational loads are absent from the aerodynamic component, we add them as separate external loads acting on the whole structure. The FE simulation returns the value of a chosen failure criterion at every node of the structural mesh. Typically, TACS implements the von-Mises criterion for alloys and the Tsai–Wu criterion for composites. As we will see, the present work exclusively uses the former. Locally, we thus obtain the yield constraint as

$$Y = \eta_Y \mathcal{F}(\mathbf{x}, F_n^{\text{EXTR}}, F_t^{\text{EXTR}}, \mathbf{g}, \mathbf{\Omega}) \leq 1, \quad (13)$$

where $\mathcal{F}$ stands for the evaluation of the failure criterion through the current FE model. It depends on the location of the node $\mathbf{x}$, the extreme aerodynamic loads, the gravity vector and the rotation rate vector. We use the same safety factor $\eta_Y = 1.35$.

From the perspective of the optimization problem, the criterion is aggregated to reduce the number of function evaluations that need to be differentiated using the adjoint (and thus to reduce computational time). In our optimization, we perform a Kreisselmeier–Steinhauser aggregation and obtain three separate constraints: one for the whole suction side of the blade, one for the whole pressure side, and one for all shear webs. We refer to Kennedy and Martins [25] for details on the theoretical foundation and implementation of failure criteria in TACS.

The advantage of applying the extreme loads to the high-fidelity structural model is that it benefits from the theoretically higher accuracy of the FE representation. Compared to conventional implementations based on strain, the present method relaxes the assumptions of linear relationship between stress and strain. Furthermore, it can accommodate multi-axial stresses. This may be prominent for composite blade models.

D. Definition of the Damage Constraint

Akin to the yield constraint, MACH currently only supports the combined fidelity implementation, even though we here define the conventional implementation as well.

Conventional implementation From the damage-equivalent beam internal forces and moments, we determine a damage-equivalent strain ϵ^{life} at the location of the spar caps from Eqs. (11). The latter is then linked to the number of cycles to failure using Eqs. (2) and (5):

$$N^{\text{life}} = \left(\frac{\epsilon_{\max}}{\epsilon^{\text{life}}} \right)^m, \quad (14)$$

where we use the fact that the DEL are established for a zero-average load. At a given location of the blade structure, the damage constraint is:

$$D = \left(\frac{\eta_D \epsilon^{\text{life}}}{\epsilon_{\max}} \right)^m \leq 1, \quad (15)$$

where we cover for the various uncertainties with a single safety factor $\eta_D = 1.35$. The latter is applied to the loads and not to the damage, as proposed by Veers [29]. Indeed, considering the nonlinearity introduced with the exponent, a safety factor on the loads has a much larger influence, and thus leads to a safer design.

Combined fidelity implementation We note that the conventional expression of the damage, Eq. (15), reduces to a failure criterion to the power m . Therefore, we express the high-fidelity damage constraint as

$$D = \left(\eta_D \mathcal{F}(\mathbf{x}, F_n^{\text{life}}, F_t^{\text{life}}, \mathbf{g}, \mathbf{\Omega}) \right)^m \leq 1. \quad (16)$$

Just as for the conventional implementation, we evaluate the response of the model to DELs. As a reminder, the DELs ($F_n^{\text{life}}, F_t^{\text{life}}$ in the above equation) are static loads that lead to the same damage as the entire spectrum of unsteady loads. We therefore evaluate the FE solution with the damage-equivalent aerodynamic loads as static external loads. We opt for applying gravitational and rotational loads as well, to represent the average loading of the blade in operation.

A similar procedure was proposed by Musial et al. [34] to determine loads for load-based fatigue tests of a blade. In this type of blade experiment, a set of external damage-equivalent loads are applied as static loads to a full-scale blade. Our approach is comparable, except that we perform the “test” numerically using a FE structural simulation.

Eventually, it must be noted that the validity of the damage constraint greatly depends on the assumptions made in the damage law (Miner’s rule), and on the material model (S-N curve). This is a known limitation of the present handling of fatigue [35]. Other high-fidelity methods exist to determine the blade fatigue. However, they usually involve unsteady FE simulations [36] which is currently impractical for MDO.

E. Discussion

The underlying idea of the combined fidelity approach is to use the most appropriate fidelity to each component in the process. Even though we refer to the CFD-CSM-based model as “high-fidelity”, we do not consider that the conventional approach has a low or lower fidelity. In fact, in some respect, it can be considered as a higher fidelity model than the coupled CFD-CSM model. The former enables comprehensive dynamic simulations of the whole turbine whereas the latter is limited to steady simulations of the rotor. On the other hand the conventional method introduces simplifications originating in the foundation of the BEM and beam theories. These assumptions are valid for standard blade designs, but our purpose with the MDO of turbine blades is to identify new or better solutions that may arise when those assumptions are relaxed.

For operating conditions where multiple models are valid (i.e. a standard blade design in steady flow conditions), there should be sufficient agreement between the models to guarantee the consistency of the approach. In the next section, we verify this assertion through a series of comparisons, for the aerodynamic and structural models separately. Note, that unlike other mixed-fidelity approaches [17, 18], we do not try to correct the results from one fidelity level with results from the other. This can only be done when both fidelities are valid over the entire range of operating conditions, which is not the case in the present approach.

IV. Turbine Definition and Modeling

Our baseline blade configuration is inspired from the 10 MW reference turbine from the Technical University of Denmark (DTU) [37]. We work with the blade outer shape modified by Madsen et al. [13] to leave more room for optimization. Additionally, the internal geometry in our model is greatly simplified: the blade skin and the three shear webs are made of a single layer of plain material. For the illustration of the combined fidelity approach in this study, we make up a fictitious isotropic material with properties inspired from an aluminium alloy, except for the Wöhler exponent which we chose as $m = 10$. Future extensions of the present work will consider actual composite materials, and laminates with realistic layups.

Figure 8 shows the blade outer mold line, and the corresponding surface aerodynamics and structural meshes used in the rest of this work. In the reference aerodynamic mesh (level L0), each blade has approximately 128 cells over the span and 256 elements over the chord. The aerodynamic volume mesh is extruded from the surface mesh using a hyperbolic procedure with 128 levels, leading to a total number of 14.2M cells. The outer boundary is located approximately at a distance of $20R$ from the blade. Coarser mesh levels are obtained by merging 8 adjacent cells together. Applied recursively, we obtain mesh levels L1, L2 and L3.

The structural mesh has a unique resolution with a total number of 208k cells. The three blades are separate and treated as clamped at the root. Cells are grouped in an arbitrary number of panels (in the present case, 117 per blade, as shown in the figure). The properties of each panel can be used as a design variable in the optimization. The baseline configuration considers a uniform thickness distribution of 15 mm for all panels.

Mangano et al. [16] give a more detailed description of the rotor geometry and modeling that we employ. Hereafter, we verify that the predictions of the high-fidelity aerodynamic and structural models are consistent with those of




Figure 8 Aerodynamic surface level 0 mesh (top) and structural mesh (bottom). Colors on the structural mesh show the division in panels, each possibly with a different properties.

conventional simulation tools.

A. Aerodynamics

We compare aerodynamic quantities predicted by the various models at hand. A sufficient agreement is required to maintain the consistency of the combined fidelity approach. To isolate the aerodynamic contributions only, simulations are run for a rigid blade with the baseline airfoil, chord and twist distributions. We use steady uniform inflow conditions.

Figure 9 shows the power coefficient as a function of the inflow velocity, as computed by the BEM code (AeroDyn) implemented in OpenFAST, by ADflow, and by EllipSys3D. The latter is a finite volume incompressible CFD solver. EllipSys3D results were provided to us by Madsen et al. [13] as a courtesy and are here considered as an external reference. Note that the results from both CFD codes are obtained on the same L0 mesh.

We observe a satisfactory agreement between the results from the three methods. The match between the conventional and the high-fidelity results is highly dependent on the airfoil polars used in the BEM. The polar data was obtained from 2D RANS simulations with EllipSys and was provided in the definition of the reference turbine [37]. This may explain the slightly better agreement of the BEM results with EllipSys3D. We do not expect the agreement to deteriorate much when changing the blade design as long as the airfoil shape remains unchanged.

As noted by Madsen et al. [13], the discrepancies between ADflow and EllipSys3D can be attributed to different turbulence models, compressibility effects, and order of the numerical schemes. A mesh convergence study is also provided in the same reference.

B. Structures

We work with two distinct models of the blade structure: the conventional beam model, and the 3D shell FE element model. Each model encodes the blade material properties, panel thicknesses and geometry.

To verify that the models are consistent with each other, we compare their predictions when the baseline constant-thickness structure is subject to a given loading. Since the evaluation of yield and damage eventually matters most in the optimization, we compare the output for those quantities directly.

For the evaluation of failure due to yield, we compute the structural response of a blade operating in under-rated conditions (region II). First, we simulate the rotor in a steady inflow at a velocity of 9 m/s and blade pitch angle of 0°




Figure 9 Aerodynamic torque and thrust of the baseline turbine, obtained with AeroDyn (+), ADflow (-■-) and EllipSys3D (-▲-). Both CFD curves result from simulations with the L0 mesh.

with OpenFAST. We isolate the instantaneous blade loads measured at an azimuthal angle of 270° (i.e., when the blade is horizontal). The corresponding beam internal forces and moments are used in the evaluation of the conventional yield constraint (Eq. (12)). For the high-fidelity model, the technique presented in Section III.B is employed to transfer the aerodynamic loads to the 3D mesh. We then apply them as a static external loading to the FE model (Eq. (13)).

The failure due to yield Y at the location of the spar caps along the radius are compared in Fig. 10. For this relatively simple load condition, the distributions agree qualitatively well. Both models exhibit similar overall trends, with minimum and maximum values of the same order of magnitude. Future work will focus on seeking validation of these curves with comparison against strain measured in actual blades. For composite blades, we expect that the assumption of purely axial stresses in the conventional model will lead to inaccuracies that the combined fidelity implementation may help overcome.




Figure 10 Evaluation of the failure due to yield in the spar caps. Conventional formulation evaluated on the pressure side (—●—) and suction side (—○—); combined fidelity evaluation over the pressure side (- - -) and suction side (- - -).

We perform a similar comparison for the damage. DELs are gathered from a single 10 min OpenFAST simulation performed with a 9 m/s inflow velocity and "normal" turbulent conditions (as defined in the IEC standards [1]). The loads

are then scaled to represent the damage equivalent to one year of operation in those conditions. The damage-equivalent beam loads lead straight to the conventional damage constraint (Eq. (15)). The aerodynamic DELs are transferred to the 3D mesh and then used in the evaluation of the high-fidelity damage constraint (Eq. (16)).

The results are compared in Fig. 11. We observe a similar agreement as for the failure due to yield. Discrepancies may originate in the different treatments of the gravitational and rotation loads in each model, the aerodynamic load transfer from 1D to 3D, or the different implementations of the failure itself. Again, in future work, we will proceed with a more careful validation of these quantities.




Figure 11 Evaluation of the failure due to fatigue damage in the spar caps. Conventional formulation evaluated on the pressure side (—●—) and suction side (—●—); combined fidelity evaluation over the pressure side (---) and suction side (---).

V. Optimization Results

In this section, we present the results of a structural optimization and an aerostructural optimization using the combined fidelity approach. The first case demonstrates the viability of the approach showing the convergence of the outer iterations, whereas the second case explores the numerical behavior of the approach when it uses multiple structural and aerostructural evaluations in the loop. Mangano et al. [16] present more examples of high-fidelity aerostructural optimization of wind turbine blade based on MACH.

The following optimizations consider the DLCs listed in Table 1, referring to the nomenclature of the standards [1]. The number of velocities and random seeds is also shown, which amounts to a total of 132 simulations that are performed at every outer iteration of the approach.

DLC	Inflow velocities [m/s]	N seeds	Mode
1.2	5, 7, 9, 11, 13, 15, 17, 19, 21, 23, 25	6	Fatigue
1.3	5, 7, 9, 11, 13, 15, 17, 19, 21, 23, 25	6	Extreme

Table 1 List of DLCs and related conditions considered in the unsteady simulations. Cases marked with ‘Extreme’ contribute to the yield constraint, and cases with ‘Fatigue’ to the damage constraint.

Even though the following results only account for two DLCs, the approach lends itself to considering more cases. For instance, fatigue loads can be aggregated across various DLCs. For extreme loads from different DLCs, one can elect to choose the worst case and use it as an input to a single failure evaluation in the optimization. The alternative is to add separate high-fidelity failure evaluations for each DLC, taking advantage of the multi-point formulation of the optimization.

A. Structural Optimization

The high-fidelity optimization is set up to minimize the blade structural mass, starting from the initial constant-thickness blade design. For the purpose of making a simple demonstration, the design variable vector only includes the thickness of the 117 panels of the model. In future work, we will use a more realistic composite model where other parameters will be optimized, such as the fiber orientation in the composite laminates, the fraction of each ply in the layup, etc. In the present case, all parameters other than panel thicknesses are fixed, including the rotor radius, the blade shape, the pitch angle, and rotor rotation rate. We formulate both the constraint on yield and damage, as explained in Section III. No additional constraints are considered. Note that there is no high-fidelity aerostructural simulation involved in this first case.




Figure 12 Convergence history of the first optimization in the combined fidelity structural optimization. The target optimality is reached after 265 major iterations.

The convergence history of the first outer iteration in the optimization is presented in Fig. 12. While the initial design is not feasible, the optimizer reaches a feasible solution after around 20 iterations, at the cost of an increase in mass. In later iterations, the mass decreases as the optimization progresses. Convergence is achieved when SNOPT's optimality metric reaches 10^{-4} . Figure 12 also shows the value of the damage and yield constraints aggregated over the entire blade pressure side. In this case, only the yield constraint is active in the spar caps at the optimal solution. Generally, the number of included DLCs and the way the fatigue and extreme loads are computed affect which constraint is active at the optimum. Also, further analysis of our results shows that both aggregated constraints on yield and damage are active in the webs, which confirms the importance of considering them simultaneously in the optimization.

Convergence of the outer loop is obtained after 4 outer iterations, as can be seen in Fig. 13. The first iteration requires the largest number of major iterations. For subsequent ones, the optimization restarts from the last design point of the previous iteration and uses the updated extreme and fatigue loads. Only the first two iterations provide a significant decrease in the objective function. This whole behavior is due to the damage-equivalent loads being only marginally affected by the blade structural design. Figure 14 confirms that trend, showing that the DELs do not change by more than 4% relative to the first simulation, and quickly converge.




Figure 13 Convergence of the outer loop in a combined fidelity structural optimization, with relative variation of the blade mass with respect to the baseline (left), and number of major optimization iterations per outer iteration in the combined fidelity approach (right).




Figure 14 Relative change in damage-equivalent normal and tangential aerodynamic loads over the course of the outer iterations (here labeled i), with respect to the output of the first evaluation.

The final blade design is presented in Fig. 15. Compared to the baseline model where all panels have a thickness of 15 mm, an outstanding feature of the optimal design is the increased thickness of the spar caps on both sides of the blade. This is consistent with engineering common sense as the spar caps are subject to the largest stresses. A larger thickness is used for the spar cap on the pressure side to compensate for relatively higher damage and yield values. This feature was already apparent in Figs. 10 and 11. The mass was globally reduced with respect to the baseline configuration as many of the other panels from the skin and the webs (not shown) are thinner. Note that the baseline panel distribution was purely arbitrary. A finer panel distribution could lead to further mass reduction, without adding cost to the optimization thanks to the adjoint formulation.

B. Aerostructural Optimization

Starting from the optimal solution found for the blade structure in the previous section, we complete an additional optimization of a slightly modified problem. In addition to the yield and damage constraints, we add a torque constraint derived from the coupled aerostructural solver. The resulting problem thus consists of two structural simulations and one aerostructural simulation which are performed every time the optimizer needs to evaluate the objective and the constraints. The CFD simulations inside the coupled aerostructural kernel are here performed with a L3 mesh to reduce the computational burden for the present demonstration.




Figure 15 Optimal panel thickness distribution over the blade suction side (top) and pressure side (bottom) obtained from the structural combined fidelity optimization.

Comparing the loads in normal operation and those related to extreme events and fatigue, it is clear that the blade structure is mainly sized by extreme and fatigue considerations. However, the design space is potentially open for improvements on the aerostructural design while maintaining compliance with the yield and damage constraints. The purpose of this case is to show that the optimizer is capable of adjusting aerodynamic design variables while keeping the structural constraints active.

As a surrogate to more complex aerostructural problems, we here consider the problem of finding the blade pitch angle that will provide a given rotor torque. The mass is kept as the objective. The pitch angle is added to the set of design variables, and we set a torque constraint such that rotor power must be arbitrarily increased by 5% with respect to the initial design. The convergence history of this simplified optimization problem is presented in Fig. 16.




Figure 16 Convergence history of the aerostructural optimization, starting from the optimal solution of the structural problem and adding a torque constraint.

The figure shows that the optimizer eventually finds a feasible solution that still satisfies structural constraints. Note that, because we consider the coupled aerostructural problem, the torque has a non-zero sensitivity to the panel thicknesses. We can see that, towards the beginning of the optimization, the optimizer moderately increases the pitch angle but also adapts the thicknesses. As a result, by iteration 25, the solution almost reached the required torque, but the damage constraint is violated. This is because the optimizer exploited the torsional flexibility of the blade. As the optimization progresses, a better solution is found that uses a higher pitch and meets all constraints.

Extending the current formulation with more complex aerostructural problems such as those considered by Mangano et al. [16] and beyond will allow us to explore more thoroughly the design space and the underlying aerostructural coupling. This constitutes the next major step in our work.

C. Note on the Computational Cost

Computational cost is one of the main barriers to a more extensive use of high-fidelity simulation for the design and optimization of wind turbines. We provide some insights into the cost of the various steps involved in the present approach.

We break down the computational time of the structural optimization presented above in Fig. 17. The timings shown in this figure result from computations run on the *marylou9* supercomputer from the BYU Fulton Supercomputing Lab. We used a single compute node equipped with Intel Broadwell 2.4 GHz CPUs. The execution times are only to be used for a comparison of their relative orders of magnitude as they are not representative of the average behavior of the codes on multiple nodes. The total time for WEIS includes all the OpenFAST simulations, the related pre-processing (computation of turbulent inflows) and post-processing (rainflow counting, load extrapolation and aggregation). Each of these operations are executed in serial for a given inflow condition, but the various load cases are simulated and processed in parallel. The total computational cost of the conventional simulations mostly depends on the number of simulated DLCs and the related conditions and seeds, on the chosen length of the unsteady simulations and on the complexity of the model.

On the other hand, the primary drivers of the cost of the high-fidelity optimization are the chosen mesh resolution, the number of differentiated functions evaluated at each iteration, and the convergence of the optimization itself. Both aerodynamic and structural high-fidelity codes provide a MPI-parallel implementation.




Figure 17 Execution time of the combined fidelity optimization. The total times for WEIS and MACH (blue colors) are broken down into the most time-consuming operations. For WEIS, the total times required for the preprocessing, the simulation and the postprocessing are reported. For MACH, we measured the time spent in the evaluation of the structural solutions used for each constraint separately. The number of conventional simulations or function calls is indicated on top of the bars.

In the first outer iteration, the cost of the high-fidelity optimization is dominant. However, the ratio becomes more favorable for later iterations. The number of function calls from the optimizer drastically reduces in the second and subsequent iterations since they restart from the output of the previous iteration. On the other hand, the number of unsteady simulations remains the same over the outer iterations. Note that some computational expenses are saved in WEIS by computing the turbulent inflow conditions only once in the first iteration (prepro.) and reusing them after. This also highlights the advantage of the combined fidelity approach. We can see that performing the OpenFAST simulations and the related processing within the optimization loop would lead to a much higher computational cost, since it would require the evaluation of all DLCs with the conventional model at every function call from the optimizer.

Finally, from the above example of aerostructural optimization, we evaluate the average execution time of the

structural simulation and the aerostructural simulation. On average, one call to the structural simulation takes 0.057 CPU.h, whereas it takes 0.475 CPU.h for the aerostructural simulation to execute. As anticipated, the aerostructural simulation is far more expensive and most of the computational time is dedicated to converging the CFD solution. Additionally, the necessity of resolving the tight aerostructural coupling with an iterative method increases the overall computational cost. From this observation, it is apparent that a proper computational load balancing is required at the level of the function call from the optimizer. Since the structural evaluation of the damage and yield constraints are cheaper than the aerostructural simulation, we allocate more computational resources to the aerostructural simulation (roughly proportionally to the average timings given above). Doing so avoids wasting computational resource in the situation where some processors need to wait for the completion of the task performed by another set of processors, before the computation can continue.

VI. Conclusions

In this paper, we have presented a combined fidelity approach to the design of wind turbines and showcased its use for the optimization of a rotor. The approach relies on a coupled CFD-CSM solver to perform simulation that captures the tight aerostructural coupling inherent to modern flexible blades. While the optimization of the power production incorporates high-fidelity steady-state simulations, structural sizing constraints are added to make sure the optimal design complies with DLCs defined in the standards. These extra constraints are derived from loads obtained from conventional unsteady turbine simulations, which are performed by OpenFAST in the present work.

We presented in detail the combined fidelity algorithm, the derivation of the extreme and fatigue loads from the unsteady simulations, the technique to transfer them to the high-fidelity model, and the subsequent formulation of the yield and damage constraints. For a simple blade model inspired from the DTU 10 MW reference design, we evaluated the consistency of the approach by comparing predicted aerodynamic quantities. We also verified that the newly formulated constraints agreed with their conventional implementation. Finally, we considered a structural optimization problem. We demonstrated that the outer loop required in the combined fidelity approach converges in only a few iterations, since the precomputed loads are only marginally affected by the blade design itself. We performed a subsequent optimization combining both structural and aerostructural solvers, and we measured and compared the related computational costs.

In a concurrent study [16], we concluded that our MDO framework provides an easy and efficient way to handle problems with hundreds of design variables. Additionally, the tight aerostructural coupling captured with the high-fidelity simulations can lead to improved designs as compared to sequential aerodynamic and structural optimizations. The work that we presented here enables the extension of such an MDO with many load cases by incorporating both conventional and high-fidelity tools in the design process. Even though the cost of high-fidelity MDO is probably too high for preliminary design applications, our framework can provide valuable design insights into the aerostructural behavior of large flexible blades.

Future work will consider more complex aerostructural optimization problems such as improving the model and increasing the number of design variables (e.g., considering airfoil shape). This should open more possibilities for the optimizer to arrive at novel designs. Integrating more turbine components in the optimization such as the tower and the nacelle should also enable the identification of more global optima, which could be done through the combined fidelity approach. Regarding the combined fidelity approach itself, we will work on improving the constraints involving multiple fidelities through more DLCs and specific treatment of the corresponding loads. For instance, the formulation of the fatigue constraint could be revised to incorporate a better knowledge of the fatigue of composite materials, including a better treatment of average loads and their influence on the accumulated damage. A proper validation of the high-fidelity damage and yield constraints is also required to more clearly identify their benefits over the conventional formulation. Even though the present application was limited to rotor blades, the initial results showed promising trends that we will continue to explore in future work.

Appendix

Figure 18 presents a XDSM diagram of the combined fidelity approach with details on the various components and information they exchange. Starting in the top-left corner, steps 1 to 4 include all the treatments necessary to generate the conventional turbine model using the set of high-fidelity design variables, to perform the conventional simulations with OpenFAST and to post-process their output. The load transfer procedure outlined in Section III.B takes place in step 4.

The remainder of the operations are gathered under the optimization loop managed in step 5. The aerostructural

kernel covers steps 7 to 10, where the tight coupling of the CFD solver ADflow and the structural solver TACS is visible. The iterative method required to converge the coupling is handled by the multidisciplinary analysis (MDA) module, step 7. The additional structural analyses and the subsequent evaluations of the combined fidelity yield and damage constraints take place in steps 11 and 12, respectively. Finally, the feedback of the outer iteration loop is illustrated with the link between steps 13 and 0.

Acknowledgments

This research is supported by the Department of Energy (DOE) Advanced Research Projects Agency-Energy (ARPA-E) Program award DE-AR0001186 entitled “Computationally Efficient Control Co-Design Optimization Framework with Mixed-Fidelity Fluid and Structure Analysis.” The authors thank DOE ARPA-E Aerodynamic Turbines Lighter and Afloat with Nautical Technologies and Integrated Servo-control (ATLANTIS) Program led by Dr. Mario Garcia-Sanz. Special thanks to the entire ATLANTIS Team for their support.

The authors are grateful for the computing resources provided by the Office of Research Computing at Brigham Young University and by the Texas Advanced Computing Center (TACC) Stampede2 High Performance Computing system via the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by the National Science Foundation grant number ACI-1548562.

References

- [1] “Wind Turbines – Part 3: Design requirements for offshore wind turbines,” Tech. Rep. IEC 61400-3, International Electrotechnical Commission, Geneva, Switzerland, 2009.
- [2] Bottasso, C., Campagnolo, F., Croce, A., and Tibaldi, C., “Optimization-based study of bend–twist coupled rotor blades for passive and integrated passive-active load alleviation,” *Wind Energy*, Vol. 16, No. 8, 2013, pp. 1149–1166. <https://doi.org/10.1002/we.1543>, URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/we.1543>.
- [3] Zahle, F., Tibaldi, C., Pavese, C., McWilliam, M. K., Blasques, J. P. A. A., and Hansen, M. H., “Design of an Aeroelastically Tailored 10 MW Wind Turbine Rotor,” *Journal of Physics: Conference Series*, Vol. 753, 2016, p. 062008. <https://doi.org/10.1088/1742-6596/753/6/062008>.
- [4] McWilliam, M. K., Zahle, F., Dicholkar, A., Verelst, D., and Kim, T., “Optimal aero-elastic design of a rotor with bend-twist coupling,” *Journal of Physics: Conference Series*, 2019. URL https://repository.lboro.ac.uk/articles/journal_contribution/Optimal_aero-elastic_design_of_a_rotor_with_bend-twist_coupling/9572537.
- [5] Gray, J. S., Hearn, T. A., Moore, K. T., Hwang, J., Martins, J. R. R. A., and Ning, A., “Automatic Evaluation of Multidisciplinary Derivatives Using a Graph-Based Problem Formulation in OpenMDAO,” *15th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, 2014. <https://doi.org/10.2514/6.2014-2042>, URL <https://arc.aiaa.org/doi/abs/10.2514/6.2014-2042>.
- [6] Dykes, K., Ning, A., King, R., Graf, P., Scott, G., and Veers, P. S., “Sensitivity Analysis of Wind Plant Performance to Key Turbine Design Parameters: A Systems Engineering Approach,” *32nd ASME Wind Energy Symposium*, 2014. <https://doi.org/10.2514/6.2014-1087>, URL <https://arc.aiaa.org/doi/abs/10.2514/6.2014-1087>.
- [7] Ashuri, T., Zaaijer, M., Martins, J., van Bussel, G., and van Kuik, G., “Multidisciplinary design optimization of offshore wind turbines for minimum levelized cost of energy,” *Renewable Energy*, Vol. 68, 2014, pp. 893–905. <https://doi.org/10.1016/j.renene.2014.02.045>, URL <https://www.sciencedirect.com/science/article/pii/S0960148114001360>.
- [8] Gaertner, E. M., and Lackner, M. A., “Aero-elastic design optimization of floating offshore wind turbine blades,” *2018 Wind Energy Symposium*, 2018. <https://doi.org/10.2514/6.2018-2015>, URL <https://arc.aiaa.org/doi/abs/10.2514/6.2018-2015>.
- [9] Bottasso, C. L., Campagnolo, F., and Croce, A., “Multi-disciplinary constrained optimization of wind turbines,” *Multibody System Dynamics*, Vol. 27, No. 1, 2012, pp. 21–53. <https://doi.org/10.1007/s11044-011-9271-x>.
- [10] Bottasso, C. L., Bortolotti, P., Croce, A., and Gualdoni, F., “Integrated aero-structural optimization of wind turbines,” *Multibody System Dynamics*, Vol. 38, No. 4, 2016, pp. 317–344. <https://doi.org/10.1007/s11044-015-9488-1>.
- [11] Bortolotti, P., Bottasso, C. L., and Croce, A., “Combined preliminary–detailed design of wind turbines,” *Wind Energy Science*, Vol. 1, No. 1, 2016, pp. 71–88. <https://doi.org/10.5194/wes-1-71-2016>, URL <https://wes.copernicus.org/articles/1/71/2016/>.
- [12] Dhert, T., Ashuri, T., and Martins, J. R. R. A., “Aerodynamic shape optimization of wind turbine blades using a Reynolds-averaged Navier–Stokes model and an adjoint method,” *Wind Energy*, Vol. 20, No. 5, 2017, pp. 909–926. <https://doi.org/10.1002/we.2070>, URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/we.2070>.

- [13] Madsen, M. H. A., Zahle, F., Sørensen, N. N., and Martins, J. R. R. A., “Multipoint high-fidelity CFD-based aerodynamic shape optimization of a 10 MW wind turbine,” *Wind Energy Science*, Vol. 4, No. 2, 2019, pp. 163–192. <https://doi.org/10.5194/wes-4-163-2019>, URL <https://wes.copernicus.org/articles/4/163/2019/>.
- [14] Madsen, M. H. A., Zahle, F., Horcas, S. G., Barlas, T., and Sørensen, N. N., “CFD-based curved tip shape design for wind turbine blades,” *Wind Energy Science Discussions*, Vol. 2021, 2021, pp. 1–48. <https://doi.org/10.5194/wes-2021-115>, URL <https://wes.copernicus.org/preprints/wes-2021-115/>.
- [15] Heinz, J., Sørensen, N., and Zahle, F., “Fluid-structure interaction computations for geometrically resolved rotor simulations using CFD,” *Wind Energy*, Vol. 19, No. 12, 2016, pp. 2205–2221. <https://doi.org/10.1002/we.1976>.
- [16] Mangano, M., He, S., Liao, Y., Caprace, D.-G., and Martins, J. R. R. A., “Towards Passive Aeroelastic Tailoring of Large Wind Turbines Using High-Fidelity Multidisciplinary Design Optimization,” *AIAA Scitech Forum*, 2022.
- [17] McWilliam, M., Lawton, S., and Crawford, C., “Towards a Framework for Aero-elastic Multidisciplinary Design Optimization of Horizontal Axis Wind Turbines,” *51st AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, 2017. <https://doi.org/10.2514/6.2013-200>, URL <https://arc.aiaa.org/doi/abs/10.2514/6.2013-200>.
- [18] Jasa, J., Bortolotti, P., Zalkind, D., and Barter, G., “Effectively using multifidelity optimization for wind turbine design,” *Wind Energy Science Discussions*, Vol. 2021, 2021, pp. 1–22. <https://doi.org/10.5194/wes-2021-56>, URL <https://wes.copernicus.org/preprints/wes-2021-56/>.
- [19] Ning, A., and Petch, D., “Integrated design of downwind land-based wind turbines using analytic gradients,” *Wind Energy*, Vol. 19, No. 12, 2016, pp. 2137–2152. <https://doi.org/10.1002/we.1972>, URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/we.1972>.
- [20] Martins, J. R. R. A., “Perspectives on aerodynamic design optimization,” *AIAA Scitech 2020 Forum*, 2020. <https://doi.org/10.2514/6.2020-0043>, URL <https://arc.aiaa.org/doi/abs/10.2514/6.2020-0043>.
- [21] Ingersoll, B., and Ning, A., “Efficient incorporation of fatigue damage constraints in wind turbine blade optimization,” *Wind Energy*, Vol. 23, No. 4, 2020, pp. 1063–1076. <https://doi.org/10.1002/we.2473>, URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/we.2473>.
- [22] Kenway, G. K. W., Mader, C. A., He, P., and Martins, J. R. R. A., “Effective Adjoint Approaches for Computational Fluid Dynamics,” *Progress in Aerospace Sciences*, Vol. 110, 2019, p. 100542. <https://doi.org/10.1016/j.paerosci.2019.05.002>.
- [23] Mader, C. A., Kenway, G. K. W., Yildirim, A., and Martins, J. R. R. A., “ADflow: An Open-Source Computational Fluid Dynamics Solver for Aerodynamic and Multidisciplinary Optimization,” *Journal of Aerospace Information Systems*, Vol. 17, No. 9, 2020, pp. 508–527. <https://doi.org/10.2514/1.1010796>.
- [24] Yildirim, A., Kenway, G. K. W., Mader, C. A., and Martins, J. R. R. A., “A Jacobian-free approximate Newton–Krylov startup strategy for RANS simulations,” *Journal of Computational Physics*, Vol. 397, 2019, p. 108741. <https://doi.org/10.1016/j.jcp.2019.06.018>.
- [25] Kennedy, G. J., and Martins, J. R., “A parallel finite-element framework for large-scale gradient-based design optimization of high-performance structures,” *Finite Elements in Analysis and Design*, Vol. 87, 2014, pp. 56–73. <https://doi.org/10.1016/j.finl.2014.04.011>, URL <https://www.sciencedirect.com/science/article/pii/S0168874X14000730>.
- [26] Kenway, G. K. W., Kennedy, G. J., and Martins, J. R. R. A., “Scalable Parallel Approach for High-Fidelity Steady-State Aeroelastic Analysis and Adjoint Derivative Computations,” *AIAA Journal*, Vol. 52, No. 5, 2014, pp. 935–951. <https://doi.org/10.2514/1.J052255>.
- [27] Gill, P. E., Murray, W., and Saunders, M. A., “SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization,” *SIAM Rev.*, Vol. 47, No. 1, 2005, pp. 99–131. <https://doi.org/10.1137/S0036144504446096>.
- [28] Moriarty, P. J., Holley, W., and Butterfield, S., “Extrapolation of extreme and fatigue loads using probabilistic methods,” Tech. Rep. NREL/TP-500-34421, National Renewable Energy Lab., Golden, CO (US), 2004.
- [29] Veers, P., “Fatigue loading of wind turbines,” *Wind Energy Systems*, edited by J. D. Sørensen and J. N. Sørensen, Woodhead Publishing Series in Energy, Woodhead Publishing, 2011, Chap. 5, pp. 130–158. <https://doi.org/10.1533/9780857090638.1.130>, URL <https://www.sciencedirect.com/science/article/pii/B9781845695804500053>.
- [30] Ravikumar, K., Subbiah, R., Ranganathan, N., Bensingh, J., Kader, A., and Nayak, S. K., “A review on fatigue damages in the wind turbines: Challenges in determining and reducing fatigue failures in wind turbine blades,” *Wind Engineering*, Vol. 44, No. 4, 2020, pp. 434–451. <https://doi.org/10.1177/0309524X19849851>.

- [31] Hayman, G., “MLife theory manual for version 1.00,” *National Renewable Energy Laboratory (NREL)*, 2012.
- [32] Bottasso, C. L., Campagnolo, F., Croce, A., Dilli, S., Gualdoni, F., and Nielsen, M. B., “Structural optimization of wind turbine rotor blades by multilevel sectional/multibody/3D-FEM analysis,” *Multibody System Dynamics*, Vol. 32, No. 1, 2014, pp. 87–116. <https://doi.org/10.1007/s11044-013-9394-3>.
- [33] Hansen, M., *Aerodynamics of wind turbines*, 2nd ed., Earthscan, London, UK, 2008.
- [34] Musial, W., Clark, M., Egging, N., Zuteck, M., Musial, W., Clark, M., Egging, N., and Zuteck, M., “Comparison of strength and load-based methods for testing wind turbine blades,” *35th Aerospace Sciences Meeting and Exhibit*, 1997. <https://doi.org/10.2514/6.1997-957>, URL <https://arc.aiaa.org/doi/abs/10.2514/6.1997-957>.
- [35] Sutherland, H. J., “On the Fatigue Analysis of Wind Turbines,” Tech. rep., Sandia National Laboratory, United States, 1999. <https://doi.org/10.2172/9460>, URL <https://www.osti.gov/biblio/9460>.
- [36] Hu, W., Choi, K. K., Zhupanska, O., and Buchholz, J. H. J., “Integrating variable wind load, aerodynamic, and structural analyses towards accurate fatigue life prediction in composite wind turbine blades,” *Structural and Multidisciplinary Optimization*, Vol. 53, No. 3, 2016, pp. 375–394. <https://doi.org/10.1007/s00158-015-1338-5>.
- [37] Bak, C., Zahle, F., Bitsche, R., Kim, T., Yde, A., Henriksen, L. C., Natarajan, A., and Hansen, M., “Description of the DTU 10 MW reference wind turbine,” Tech. Rep. DTU Wind Energy Report-I-0092, Danish Technical University, DTU Wind Energy Roskilde, DK, 2013.