Objectives and Constraints for Wind Turbine Optimization

S. Andrew Ning*
Email: andrew.ning@nrel.gov
Rick Damiani
Patrick J. Moriarty

National Wind Technology Center
15013 Denver West Parkway
Golden, Colorado, 80401

Efficient extraction of wind energy is a complex, multidisciplinary process. This paper examines common objectives used in wind turbine optimization problems. The focus is not on the specific optimized designs, but rather on understanding when certain objectives and constraints are necessary, and what their limitations are. Maximizing annual energy production, or even using sequential aero/structural optimization, is shown to be significantly suboptimal compared to using integrated aero/structural metrics. Minimizing the ratio of turbine mass to annual energy production can be effective for fixed rotor diameter designs, as long as the tower mass is estimated carefully. For variable diameter designs, the predicted optimal diameter may be misleading. This is because the mass of the tower is a large fraction of the total turbine mass, but the cost of the tower is a much smaller fraction of overall turbine costs. Minimizing the cost of energy is a much better metric, though high fidelity in the cost modeling is as important as high fidelity in the physics modeling. Furthermore, deterministic cost of energy minimization can be inadequate, given the stochastic nature of the wind and various uncertainties associated with physical processes and model choices. Optimization in the presence of uncertainty is necessary to create robust turbine designs.

Nomenclature

AEP	annual energy production
BOS	balance-of-station
COE	cost of energy
D	rotor diameter
I	area moment of inertia
J	objective
M_b	bending moment
N	number of cycles
O&M	operation and maintenance
S_f	fatigue stress
S_{plan}	planform area
TCC	turbine capital costs
V_{e50}	50-year extreme wind speed
V_{hub}	hub speed
V_{in}	cut-in speed
V_{out}	cut-out speed
V_{tip}	tip speed
Ω	rotor rotation speed
δ	blade-tip deflection
ε_{50}	strain at 50-year extreme wind condition
ε_{cr}	critical buckling strain
ε_{ult}	ultimate strain
γ	safety factor
λ	tip-speed ratio
ω	natural frequency of structure
σ	stress
θ	airfoil twist angle
bm	buckling margin
c	chord
c_{set}	a repeatedly used set of constraints
i_{aero}	an index proportional to blade mass
m	mass
$rating$	machine rating
t	spar cap thickness
t/c	airfoil thickness to chord ratio
x	design variables
0	quantity of the reference model
$\bar{x}$	a quantity normalized relative to the reference design (e.g., $\bar{x} = x/x_0$)

1 Introduction

Increasing global energy requirements and a greater awareness of the benefits of renewable energy sources have driven renewed interest in wind energy. Harvesting wind en-

*Address all correspondence to this author.

ergy efficiently is a complex process that requires a multidisciplinary effort in wind turbine design, site selection, and plant layout. Many trade-offs exist in aerodynamic performance, structural efficiency, land-use footprint, operational versus manufacturing and maintenance costs, and so on. Multidisciplinary optimization and uncertainty analysis are important tools to evaluate design choices and further improve the economics of wind energy. A number of previous studies have examined optimization of wind turbines using a wide variety of approaches [1–8]. In these studies, design variables ranged from parameterizing only the rotor blades to parameterizing complete turbines. Model fidelity included simple analytic models, time-domain unsteady aeroelastic calculations, and three-dimensional computational fluid dynamics with structural finite element analyses. Objectives included maximum annual energy production (AEP), multiobjective maximum power and minimum blade root bending moment, and minimum cost of energy. Optimization approaches included gradient-based methods, direct-search methods, multilevel methods, and optimization under uncertainty.

This paper approaches the turbine optimization problem with a different focus than the previous studies. The purpose of this study is not to demonstrate a specific methodology, or to present optimized designs per se, but rather to understand how different choices in the optimization problem and model choices impact the quality of the solutions. Primarily, the goal is to better understand the appropriateness of various commonly used optimization objectives. To facilitate this understanding, the models used should: capture the fundamental trade-offs in the physics, execute rapidly to allow for a wide range of design studies, and converge robustly and with high accuracy to allow for fair comparisons in the designs. To this end, simple physics-based models were developed that produce smooth output (continuously differentiable) to allow for reliable gradient estimation. The optimization studies of this work concentrate on the design of the rotor blades, but the impact on resizing the rest of the turbine as well as plant-level costs was included. The work conducted was part of a larger effort at the National Renewable Energy Laboratory (NREL) to apply systems engineering techniques to wind energy applications [9].

The following section describes the methodology, which includes the rotor aerodynamic analysis, rotor structural analysis, cost model, reference model, and optimization strategy used. Next, a number of optimization studies are presented. These studies examined the important considerations in maximizing annual energy production, minimizing the ratio of turbine mass to annual energy production, and minimizing cost of energy. Finally, the conclusions of the study are summarized.

2 Methodology

Although the impact on the entire turbine was considered, the focus of this study was on optimizing the rotor blades. Thus, the main thrust of the methodology was the development of appropriate aerodynamic and structural analysis tools for the blades. The effect of hub, nacelle, and foundation sizing was handled through simpler scaling relations. A modified version of the NREL cost and scaling model [10] was used to predict the cost of energy. The physics-based models were implemented in C++ and Fortran, and were linked together in a common framework in Python. This approach retained most of the speed advantage of the compiled languages, but allowed for high flexibility in an object-oriented environment. The computational efficiency and flexibility were important to allow for both a large number and wide variety of studies. Only a high-level overview of the methodology is contained in this section, further details are available in a conference proceedings version of this paper [11].

2.1 Rotor Aerodynamics

2.1.1 Blade Element Momentum Method

The rotor aerodynamic analysis was based on blade element momentum (BEM) theory. Using BEM theory in a gradient-based rotor optimization problem can be challenging because of occasional convergence difficulties of the BEM equations. Ning [12] recently developed a new solution methodology to the BEM equations that offers guaranteed and efficient convergence properties. When used with a continuously differentiable description of the airfoil force coefficients, not only is the solution guaranteed to be found, but the resulting solution is continuously differentiable. This behavior allows gradient-based algorithms to be used to solve the rotor optimization problem much more effectively than with traditional BEM solution approaches. Implementation of this methodology is contained in an open-source code called CCBlade. Details of the methodology are discussed in Ref. 12.

2.1.2 Airfoil Section Analysis

Two-dimensional airfoil data were corrected for rotational effects using the Du-Selig method [13] for lift and the Eggers method [14] for drag. Next, airfoil data were extrapolated to $\pm 180^\circ$, using Viterna's method [15]. Finally, for each section, a two-dimensional cubic B-spline[1] was fit to the lift and drag data separately as a function of the Reynolds number and angle of attack. A small amount of smoothing was used on each spline to reduce high-frequency noise that could cause artificial multiple solutions (0.1 for lift, 0.001 for drag). Many BEM implementations use linear interpolation to estimate lift and drag coefficients; however, such an approach is unsuitable for gradient-based optimization because it introduces discontinuities in the derivatives.

2.1.3 Rotor Aerodynamic Analysis

The aerodynamic power predicted from the BEM method was modified to account for losses in the drivetrain. The drivetrain efficiency (η_{dt}) was assumed to vary with the aerodynamic power normalized by the rated power

[1] DIERCKX package, http://www.netlib.org/dierckx/

($\overline{P} = P_{aero}/P_{rated}$) as

$$\eta_{dt} = 1.0 - (a/\overline{P} + b) \tag{1}$$

where $a = 0.0129$, and $b = 0.0851$. This drivetrain efficiency curve is for a three-stage geared design and comes from an NREL study [16] using WindPACT data [17]. The net power produced is $P = \eta_{dt} P_{aero}$. Note that the maximum drivetrain efficiency is 91.5%. Including this loss had a significant effect on total annual energy production and the estimated rated speed. Although this is less important in comparing relative aerodynamic performance between designs, it does have a significant effect on overall cost of energy.

All studies in this paper focus on variable-speed, variable-pitch machines. The optimal tip-speed ratio for operation in Region 2 (below maximum rotation speed) was determined externally as part of the optimization problem. If the maximum rotation speed was reached before rated power, rotation speed was held at the maximum value in Region 2.5. In Region 3, blade pitch was varied to feather in order to maintain rated power.

AEP was computed using a Rayleigh distribution with a mean speed of 10 m/s, as specified in the International Electrotechnical Commission (IEC) standard for Class I turbines [18]. Losses caused by wake interference from other turbines in the wind farm and losses caused by electrical grid unavailability were estimated simply by using an array loss factor and an availability factor. These studies assumed that the grid was available 95% of the year, and that 10% of the potential wind farm AEP was lost because of wake interference effects.

2.2 Rotor Structure

2.2.1 Beam Finite Element Analysis

A beam finite element code, called pBEAM (polynomial beam element analysis module), was developed for the structural analysis. The methodology uses Euler-Bernoulli beam elements with 12 degrees of freedom (three translational and three rotational at each end of the element); the basic theory is described in any standard finite element textbook (see Yang [19]). A unique feature of the code is that section properties can be described as polynomials of any order between nodes. This means that, rather than using precomputed structural matrices for an assumed distribution and fixed shape-functions, the matrix coefficients are re-computed for each geometry. However, because the distribution is a polynomial, this can be done analytically using polynomial integration.

2.2.2 Sectional Composite Analysis

An existing National Wind Technology Center code, PreComp, was used to estimate equivalent sectional inertial and stiffness properties of composite blades [20]. PreComp uses modified classic laminate theory combined with a shear-flow approach. In addition to inertial and stiffness properties, its computation of the elastic center was used in this analysis. PreComp requires the geometric description of the blade (e.g., chord, twist, section profile shapes, and web locations), along with the internal structural layup (e.g., laminate schedule, orientation of fibers, and laminate material properties). It allows for high flexibility in the specification of the composite layup both spanwise and chordwise.

All loads and inertial properties were transferred to the elastic center and principal axes of the section. For each airfoil section, the maximum strain location was assumed to occur in the outer skin layer at the chordwise location of maximum airfoil thickness. Stress was estimated using a smeared effective modulus of elasticity based on classical laminate theory. The smeared modulus calculation ignored laminate shear and bending moment effects (the latter would be zero for a symmetric laminate), which is a good approximation for slender turbine blades.

2.2.3 Additional Structural Considerations

In addition to the capabilities of the finite element analysis, two additional structural considerations specific to the rotor problem were added. The first was a panel buckling calculation and the second was a fatigue cycle estimation. Both considerations can be particularly important for very large blades [21]. This analysis considered only spar-cap buckling; however, detailed design should also consider trailing-edge panel buckling. The panel buckling calculation used the simple method suggested by Bir [22]. The calculation results from an eigenanalysis of a flat panel that is loaded edgewise. The derivation assumes that the width of the panel is much smaller than the length of the panel, which is true for the rotor blades (except near the tip). It is also assumed conservatively that the longitudinal edges are simply supported, while the chordwise edges are free.

Finally, a fatigue calculation for the blade root was included. A full lifetime fatigue analysis can be quite complex; however, a simplistic assessment can be done using the edgewise gravity loads, because these loads fully reverse every blade rotation. Although edgewise gravity loads are not always the dominant load condition, they can be more significant than flapwise aerodynamic loads in determining the fatigue strength of very large blades [21]. This constraint is an extreme simplification, and has since been replaced by better analysis methods. However, for the purposes of this study it was useful to prevent unrealistically small chord sizes at the blade root—effectively it provided a minimum chord size at the root. The blade root typically requires thick laminates to stiffen the connection to the pitch bearing, and without a proper constraint there is sometimes a large incentive for the optimizer to decrease the chord to impractical values at the blade root.

To aid the fatigue estimation, it is assumed that the S-N curve for the root section can be parameterized as

$$S_f = aN^{-1/b} \tag{2}$$

where b is assumed to be 10 (a typical value for glass-reinforced composite materials [23]).The maximum stress at the root of the blade caused by only the gravity loads was

computed as described above. Because this loading is fully reversed, the stress value at the 3 o'clock azimuth can be used directly as the damage equivalent load for the S-N curve. An average rotation speed was estimated by computing the expected value of the rotor speed using the wind speed distribution

$$\overline{\Omega} = \int_{V_{in}}^{V_{out}} \Omega(V) f(V) dV \tag{3}$$

where $f(V)$ is a probability distribution function of the wind speeds (a Rayleigh distribution with a mean wind speed of 10.0 m/s was used, as discussed previously). A 20-year lifetime of continuous rotation was assumed. The value for a in Eqn. (2) was calibrated using the loading conditions for the baseline rotor so that, at 20 years, the root stress had a 10% margin relative to the fatigue stress (i.e., $\sigma_{\text{root}}/S_f = 0.9$). The margin was added to avoid overly constraining the problem because the S-N curve was only defined relative to the baseline.

2.3 Cost Model

The cost model was based primarily on the NREL cost and scaling model [10], with a few modifications that were found to be important for this study. First, the rotor mass was computed from the structural model instead of the scaling law. Blade cost was a linear function of blade mass, with coefficients derived from an internal NREL study [24].

Next, the tower mass estimation was replaced. The current cost and scaling relationship scales linearly with the square of the rotor diameter (for a fixed hub height). Although this is reasonable under the assumption that the entire turbine scales proportionally, it was not reasonable for many of the design studies in this paper. Because tower mass is such a large fraction of the total turbine mass, it is important to capture changes in tower mass in a more realistic manner. Although capturing these changes does not affect the conclusions of this paper (because we only compared the relative value of different objectives), obtaining a more realistic tower mass estimate is desirable to produce a more accurate estimate of the various metrics.

A full tower model was developed for this study. However, as stated previously, the scope of this paper was intentionally limited to only include rotor design variables. Thus, a simpler scaling relationship for tower sizing was desired. Using physics-based scaling arguments [11], the mass of the tower was estimated as

$$\frac{m_{tower}}{m_{tower0}} = \max\left(\sqrt{0.9}, \frac{[m_{RNA}g + TH/r_{mid}]}{[m_{RNA}g + TH/r_{mid}]_0}\right) \tag{4}$$

where m_{RNA} is the mass of the rotor-nacelle-assembly, g the acceleration of gravity, T the thrust load, H the height of the tower, and r_{mid} the radius of the tower at half-height. In implementation, the function max is not used, instead a small cubic spline is used to maintain a differentiable transition between the two regions.

Although this is not a universal scaling law, it should work well for cylindrical shell designs that are scaled relative to a baseline design at a constant taper ratio and constant diameter-to-thickness ratio. Strictly speaking, r_{mid} is a function of the tower mass, given those scaling assumptions. However, including that variation creates much more complexity in the mathematical expression and has little effect as the change in radius is relatively small. For simplicity, we kept it fixed at the radius of the reference design that is discussed in the following section ($r_{mid} = 2.4675$). The factor of 0.9 comes from the assumption that the moment of inertia at the critical section of the tower cannot shrink more than 10% relative to the baseline design. The tower-top deflection, column buckling stress, and first natural frequency are approximately proportional to I^{-1}, $I^{-1.25}$, $I^{0.5}$ respectively. Thus, constraining the moment of inertia is a useful approximation for these unmodeled tower constraints.

The tower cost estimate was unmodified from the simple linear scaling with mass used in the NREL cost and scaling model. Hub mass and cost and nacelle mass and cost were also unmodified. The nacelle calculation already scales with relevant variables for this study, such as the hub thrust, torque, and maximum rotor rotation speed.

Finally, a modified balance-of-station (BOS) cost model was utilized. NREL recently developed a new BOS model for land-based wind turbines [25]. The previous BOS model scales simply based on machine rating, and does not accurately scale up to the large machine ratings of today. The new version uses a substantially different bottom-up approach by estimating component quantities and sizes. Operation and maintenance costs remain unchanged from the NREL cost and scaling model.

2.4 Reference Geometry

The reference geometry used in this study was based on the NREL 5-MW reference wind turbine [26]. The airfoils and chord schedule used in the development of the blade model were adopted from the Dutch Offshore Wind Energy Convertor project [27, 28]. To arrive at a description of the structural stress/strain field under different loading conditions, a composite-material layup was defined at 38 span sections. A preliminary version of the initial layup was provided by Sandia National Laboratories [29], as a NuMAD [30] property file. The layup was constructed so as to reproduce the prescribed stiffness/inertia properties along the blade span as close as possible, while still satisfying structural constraints. Material properties were largely taken from the U.S. Department of Energy (DOE)/Montana State University (MSU) Composite Material Fatigue Database [23] and a Sandia large rotor study [21], but were modified somewhat to represent generic prestacked laminates. The various materials included: GelCoat, glass fabrics (Unidirectional (E-LT-5500, [0]2), SNL TRIAX ([±45]2[0]2), Saertex Double-Dias (DB, [±45]4), carbon fabrics (generic unidirectional with an effective thickness of 0.47 mm), generic foam, and epoxy resins. The structural layup was not intended to be a fully assessed blade design, but a good starting

point for comparative studies (such as the study described in this paper). Unidirectional carbon was chosen to both reduce the tip deflection and potential tower strike, and to match the stiffness distribution of the original NREL 5-MW model.

Starting from the NREL 5-MW reference design, a parameterization was needed for subsequent optimization. Ideally, the rotor geometry should be parameterized with a small number of terms for optimization efficiency, yet still allow for significant flexibility in describing the geometry. The chord was parameterized with five variables, shown in Fig. 1. The first and last chord positions were fixed at the root and tip, respectively. The third chord variable was fixed at 62.6% of the blade length. The radial location of the second chord variable (typically the location of the maximum chord) was itself a variable (r_2). An Akima spline was fit to the four radial/chord pairs to compute the chord at any other radial station.[2] The airfoil shapes were fixed; therefore, changes in chord were directly proportional to changes in airfoil thickness (fixed t/c for a given section). A fixed t/c does limit the freedom available to the structural design, but reliable aerodynamic airfoil data for very thick airfoils was not available at the time of the study.

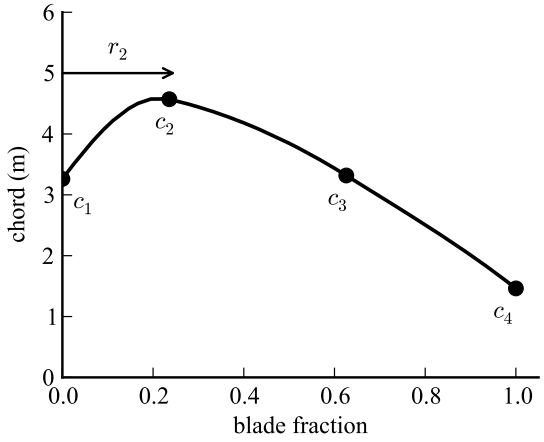

Fig. 1. Rotor blade chord distribution parameterization

The twist distribution was parameterized in a similar manner (Fig. 2). The inboard portion of the blades has cylindrical sections, which are invariant to twist. Thus, a constant value was used—up to 16.7% of the blade length where the first airfoil was defined. Twist was defined at four linearly spaced radial points from this point to the blade tip, and was fit with an Akima spline.

[2]An Akima spline was chosen because of its robustness to outliers. If one of the chord variables differs significantly in magnitude from the others (which can happen during the course of an optimization), then oscillations are produced for many spline types. This may cause some sections to have negative chord values, which is nonphysical and will prevent the analysis from running properly. An Akima spline prevents these types of oscillations. A simple bound constraint on chord is sufficient to prevent intermediate designs with negative chord, as opposed to a nonlinear constraint on chord that would be required if using a cubic spline or Bezier curve.

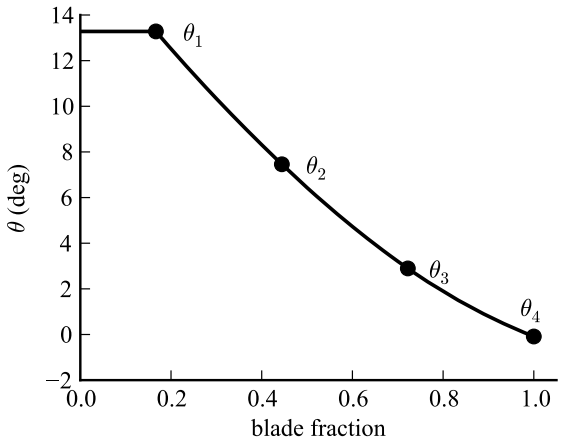

Fig. 2. Rotor blade twist distribution parameterization

For both the chord and twist distribution, parameters were chosen to best fit the reference model's chord and twist. The fit was found by solving the optimization problem

$$\begin{aligned} &\text{minimize} && ||\text{chord}_{\text{NREL}} - \text{chord}_{\text{Akima}}(x))||_2^2 \\ &\text{with respect to} && x = \{r_2, c_1, c_2, c_3, c_4\} \end{aligned} \tag{5}$$

where the chord array was defined at 38 span stations. A corresponding minimization problem was used for twist. These parameterizations led to a baseline model that was very similar to but not exactly the same as the NREL 5-MW reference model.

The reference model's spar cap was significantly thicker over the first half-meter to stiffen the connection to the pitch bearing. Because of the abrupt change in thickness, rather than attempting to fit a spline over the entire span, the spar cap thickness over the first half-meter of the blade length was assumed fixed. For simplicity, the spar cap thickness over the rest of the cylindrical section was assumed constant but not fixed (from $r = 2$ to $r = 11.75$ – note that these locations were normalized by the blade length to accommodate resizing of the rotor diameter). Over the outer portion, spar cap thickness was parameterized in the exact same way as the twist distribution (defined at four linearly spaced stations and fit with an Akima spline), except for the spar cap thickness at the blade tip, which was fixed to prevent unrealistically small thicknesses. The parameterization of the spar cap thickness is shown in Fig. 3.

The baseline spar cap thickness was not resized using a best fit to the reference model. Because a preliminary version of the 5-MW reference structural layup was used, the structure was found to buckle at the extreme load case. Though the buckling model used in this analysis was simplistic, the results were consistent with a linear-buckling analysis conducted in ANSYS, which also predicted buckling in the spar cap and trailing-edge panels. Using an infeasible design as the baseline would lead to rather unfair comparisons. Instead, the spar cap was resized to satisfy the buckling constraint everywhere along the span. The optimization problem

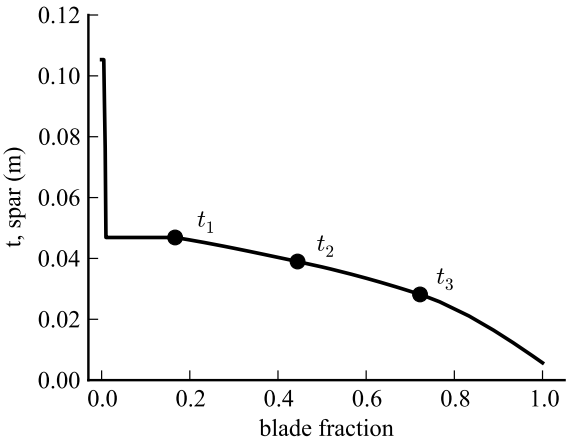

Fig. 3. Parameterization of spar cap thickness distribution

Table 1. Various normalization and reference quantities for the baseline design

AEP	20.6×10^6 kWh
Blade mass	18,246 kg
Turbine mass	657,180 kg
Tip deflection (at rated speed)	2.44 m
Cost of energy	0.0498 $/kWh

was defined as

$$\begin{array}{ll} \text{minimize} & m_{\text{rotor}}(x) \\ \text{with respect to} & x = \{t_1, t_2, t_3\} \\ \text{subject to} & bm(x)_j > 0,\ j = 1, \ldots, n \end{array} \tag{6}$$

where *bm* is the buckling margin computed at every section in the structure at the extreme load condition. (See Section 2.5 for a description of the buckling margin calculation – note that a slightly larger safety factor $\gamma_f = 1.36$ was used to ensure that the baseline design was strictly feasible.)

For the remainder of the paper, all reference quantities refer to the modified baseline model, as opposed to the original NREL 5-MW reference model. Some of the relevant reference quantities for the baseline model are defined in Tab. 1.

2.5 Optimization Strategy and Constraints

Several optimizers, both commercial and open source, were tested on the optimization problems described in this paper. The active-set algorithm of MATLAB's function *fmincon*(), which uses a sequential quadratic programming method, was found to be the most robust for these particular problems and was used for all reported results.

Gradients were estimated using central differencing, and a multistart approach was used to increase the likelihood of finding the global optimum. Early on, forward differencing was used to estimate the gradients, and the multistart approach was found to be necessary, because some starting points terminated prematurely. However, the better gradient estimates provided by central differencing resulted in improved convergence behavior, and even with multiple starting points, the same optimal solution was always found.

All objectives and constraints were normalized to be of order one for improved scaling. For example, an objective minimizing the ratio of turbine mass to AEP is implemented as

$$\min \frac{m/m_0}{AEP/AEP_0} \tag{7}$$

In the following optimization problems, these normalization constants are not explicitly written to reduce clutter; the objective is simply denoted as minimum m/AEP. All solutions were converged to a function tolerance of 1×10^{-6} and a constraint tolerance of 1×10^{-5}. Bound constraints were set large enough to never be active, unless otherwise noted.

A real turbine must be designed to meet a very large number of structural constraints [18], but only a handful of representative cases were used for this purpose. First, an ultimate strength analysis was performed at an extreme load condition. The 50-year extreme wind condition is defined as $V_{e50} = 1.4 V_{ref}$ ($V_{ref} = 50$ m/s for class I turbines) [18]. The distributed weight loads were added to the aerodynamic loads at 0 degrees pitch and the 3 o'clock azimuthal position, which is the worst case for the edgewise loads.

A maximum strain condition was used

$$\varepsilon_{50} \leq \frac{1}{\gamma_f \gamma_m} \varepsilon_{ult} \geq -\varepsilon_{50} \tag{8}$$

where the partial safety factor for loads (γ_f) was set at 1.35, and the partial safety factor for materials (γ_m) was set at 1.1 per the IEC requirements [18]. (For actual designs the material and loading partial safety factors are more involved than using a simple number, and should include additional knockdowns for uncertainty in material properties, expected environmental conditions, and other areas of uncertainty. For the purposes of this paper, this additional complexity is unnecessary.) Because the spar cap is primarily carbon, representative numbers of ultimate strain in tension and compression are about 2% and 1%, respectively. Conservatively, the smaller value of 1% was used as the ultimate strain (ε_{ult}). Only the strain from a few representative sections was used to reduce the number of constraints. The sections were biased inboard because those were the critical locations. The sections were at 0%, 11.1%, 30%, and 63.3% of the blade length on both the upper and lower surface. Although using a large number of constraints is not problematic, it is also not necessary because the chord and thickness must vary smoothly due to the parameterization, which leads to relatively smooth variations in the strain.

Panel buckling was estimated using the simplified method described in Section 2.2.3. The buckling margin was

computed as

$$bm = \frac{\varepsilon_{50}\gamma_f - \varepsilon_{cr}}{\varepsilon_{ult}} > 0 \tag{9}$$

where both values are negative because the structure is in compression for a buckling condition, and the ultimate strain is only used for scaling purposes. Similar to the strain case, only a few representative stations were used. These stations were at 16.7%, 36.7%, 63.3%, and 83.3% of the blade length on the upper surface only.

Although the constraints on extreme loading are useful when sizing the structure, without additional constraints the optimum structures would be far too flexible to be practical. As a result, constraints on the maximum tip deflection and natural frequency of the blade were added to ensure adequate stiffness. The deflection of the structure was computed at rated speed in the 3 o'clock azimuth position (worst case). Because the conceptual rotor design did not often have a defined substructure model, we could not use blade-strike as the constraint criteria. Instead, we assumed that the baseline rotor was designed with adequate stiffness and therefore constrained the deflection to be within 10% of the baseline tip deflection.

$$\delta < 1.1\delta_0 \tag{10}$$

To avoid structural resonance issues, the first natural frequency (and thus all natural frequencies) of the blade should be above the maximum rotor blade passing frequency

$$\omega_1 > \gamma_{freq}(3\Omega_{rated}) \tag{11}$$

where the safety factor γ_{freq} was set to be 1.1. The maximum rotation speed occurs at rated speed, and the factor of 3 came from the number of blades.

The fatigue strength at the root was computed for a 20-year lifetime, as discussed in Section 2.2.3. The fatigue constraint was imposed as

$$-S_f < \sigma_{\text{root-gravity}} < S_f \tag{12}$$

Finally, a constraint on the maximum tip speed was imposed as a surrogate for a noise constraint. This constraint was not imposed by the optimizer, but was implemented directly into the analysis.

The nominal optimization problem is outlined below:

$$\begin{aligned} \underset{x}{\text{minimize}} \quad & J(x) \\ \text{subject to} \quad & \text{ultimate strain (Eqn. (8))} \\ & \text{spar cap buckling (Eqn. (9))} \\ & \text{tip deflection at rated speed (Eqn. (10))} \\ & \text{blade natural frequency (Eqn. (11))} \\ & \text{blade root fatigue due to gravity (Eqn. (12))} \\ & \text{max tip speed (imposed directly in analysis)} \end{aligned}$$

Because this set of constraints is repeated in subsequent optimization problems, it will be referred to as $c_{set}(x)$, where the constraints are reorganized as needed so that feasibility occurs when $c_{set}(x) < 0$. Design variables are summarized in Tab. 2. Not all design variables were used in every problem. As stated earlier, airfoil thickness was changed through the chord distribution, as the nondimensional airfoil shapes were fixed.

Table 2. Design variables used in the optimization problems

Description	Name (# of variables)
Chord (Fig. 1)	$\{c\}$ (5)
Twist (Fig. 2)	$\{\theta\}$ (4)
Spar cap thickness (Fig. 3)	$\{t\}$ (3)
Tip-speed ratio in Region 2	λ (1)
Rotor diameter	D (1)
Machine rating	$rating$ (1)

3 Results

The studies in this paper examine the impact of using the following objectives to optimize a wind turbine rotor: maximizing annual energy production (AEP), minimizing the ratio of turbine mass to AEP, and minimizing cost of energy. Fundamental differences and important considerations of each case are highlighted.

3.1 Maximum Annual Energy Production

Although typically the most appropriate metric for wind turbine optimization is minimizing the cost of energy (COE), many studies focus on optimizing the aerodynamic performance of a wind turbine rotor, either by maximizing power at a fixed speed, or maximizing AEP for a given wind distribution. Pure aerodynamic optimization may be done for a variety of reasons, such as: an appropriate structural model is not available to the designer, an appropriate cost model is not available, the organization separates aerodynamic and structural design, or high-fidelity tools are used and for computational efficiency the aerodynamic and structural optimizations are decoupled. This section examines when pure aerodynamic optimization might be a suitable practice. In other words, under what circumstances does maximizing AEP, or sequentially maximizing AEP and then minimizing mass, yield a good design? A "good design" in this study means that it achieves a similar cost of energy as the minimum COE design.

Unfortunately, maximizing AEP without at least some consideration of the rotor mass is not a particularly well-defined problem. Part of the difficulty with using maximum AEP as an objective is that a multitude of solutions exist that

produce essentially the same AEP (in other words, many local optima exist). To clarify, consider the related optimization problem

$$\begin{array}{ll} \text{maximize} & AEP(x) \\ \text{with respect to} & x = \{\{c\},\{\theta\},\lambda\} \\ \text{subject to} & c_{set} < 0 \\ & \overline{m_{blade}} = m_c \end{array} \tag{13}$$

where m_c is a constant mass value. The result of repeating this optimization problem at several discrete constraints on the blade mass m_c is shown in Fig. 4. Essentially the same AEP can be achieved across a wide range of designs with very different masses. Point M1 is noted on the figure as a design with the same mass as the baseline design. Clearly, point M1 is superior to all points to the right of it, because it achieves essentially the same AEP, but with significantly reduced mass.

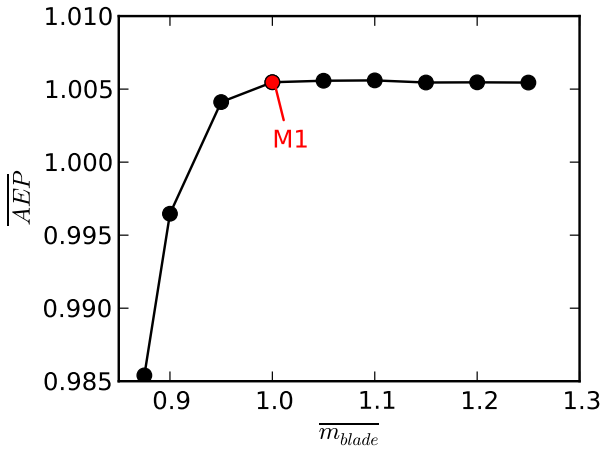

Fig. 4. Maximum annual energy production for different designs. Each design was constrained to have a different blade mass. Point M1 is highlighted and corresponds to the maximum AEP solution, with the mass constrained such that $\overline{m_{blade}} = 1$.

Although design M1 is superior to other designs in Fig. 4 with the same AEP, we presume that structural calculations are not available to a designer attempting to maximize AEP (otherwise a different objective would be used). Clearly, a more reasonable strategy is to maximize the AEP while constraining the mass using a surrogate. Unfortunately, choosing an appropriate surrogate constraint is not straightforward. One possibility is to constrain the root bending moment to be the same as the reference design. However, this is often not a helpful constraint. For example, if we define the problem as

$$\begin{array}{ll} \text{maximize} & AEP(x) \\ \text{with respect to} & x = \{\{c\},\{\theta\},\lambda\} \\ \text{subject to} & V_{tip} < V_{tip_{max}} \\ & M_{b\text{root}} < M_{b\text{root}_0} \end{array} \tag{14}$$

then the solution is exactly the same as if the root bending moment constraint was not included. This is because maximum AEP designs tend to decrease in root bending moment even without the constraint. Every point in Fig. 4 actually has a lower root bending moment than the reference design. For example, Fig. 5 compares the flapwise loads on the baseline design and design M1. The optimized design has larger loading inboard but decreased loading outboard; the net effect is that the root bending moment decreases slightly. Thus, including a root bending moment constraint would have no effect on the optimal solution.

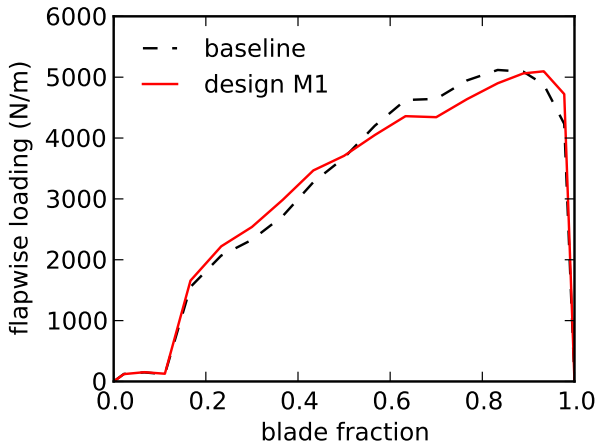

Fig. 5. Flapwise loading for the baseline design and the maximum AEP solution with fixed mass (design M1). The root bending moment for the optimized design decreases, even without a root bending moment constraint.

A number of other potential surrogate constraints were evaluated. One of the more useful ones comes from conceptual aircraft design, where the wing weight is assumed to consist of a portion that scales with the planform area and a portion that scales with the required loading. The portion of the mass that scales with the bending loads can be estimated as [31]

$$m \propto i_{aero} \equiv \int \frac{M_b}{t} dr \tag{15}$$

The assumptions are more of a stretch for wind turbine blades, because the composite structure can vary significantly in effective density, and the blade structure is generally not fully stressed. Even with these limitations, this

metric is still likely to be more useful than root bending moment, because it captures the penalty of a shrinking structural box. Note that, while a uniform t/c is often assumed for aircraft design, an estimate for the actual airfoil t/c distribution is necessary to use this metric for wind turbine applications, because the variation is typically very significant.

Constraining the planform area S_{plan} and the mass index i_{aero} should help constrain the mass of the optimized design. If manufacturing or transportation limitations are known, a constraint on the maximum chord may be used instead of, or in addition to, a constraint on the area. However, even all these constraints are generally insufficient to prevent impractical solutions. Typically, the optimal solution is to decrease the root chord in exchange for a larger chord at the point of the maximum chord. This is slightly better aerodynamically, but is considerably worse structurally. To prevent this penalty, we added a constraint on the stress at the root of the blade. For a thin shell circular section at the blade root, the stress is equal to

$$\sigma_{root} = \frac{M_b r}{\pi r^3 t} \tag{16}$$

where r is the radius of the circular section and t is the shell thickness. In this optimization problem, the internal structure was fixed so

$$\sigma_{root} \propto \frac{M_b}{r^2} \tag{17}$$

The "AEP 1st" optimization problem was done sequentially as follows

$$\begin{aligned}
&\text{maximize} && AEP(x) \\
&\text{with respect to} && x = \{\{c\},\{\theta\},\lambda\} \\
&\text{subject to} && V_{tip} < V_{tip_{max}} \\
&&& i_{aero} < i_{aero_0} \\
&&& S_{plan} < S_{plan_0} \\
&&& (M_b/r^2)_{root} < (M_b/r^2)_{root_0}
\end{aligned} \tag{18}$$

$$\begin{aligned}
&\text{minimize} && m(x) \\
&\text{with respect to} && x = \{t\} \\
&\text{subject to} && c_{set}(x) < 0
\end{aligned}$$

The resulting design achieved a relatively high AEP, though not quite as high as the maximum (0.39% increase compared to 0.56%); however, it achieved this AEP with a slight decrease in blade mass (0.5%). The net result was a 0.3% reduction in cost of energy. Even though this is a much better solution than simply maximizing AEP, it is still far from a minimum COE design. This is because the minimum COE solution does not just avoid a mass increase, but achieves a significant mass decrease. Because the blade shape is already dictated by the aerodynamic optimization, the potential to minimize mass is limited to only changing the internal structure. When the aerodynamic and structural disciplines are coupled in shaping the blade, much greater net benefits are possible.

If the disciplines must be decoupled, an alternative approach is to allow the structural analysis to dictate the blade shape and the aerodynamic analysis to dictate the airfoil shape (only the twist distribution for this problem). Then the structural analysis must be repeated one final time to ensure that the constraints have been satisfied. This approach was used in the "mass 1st" problem, which was defined as

$$\begin{aligned}
&\text{minimize} && m(x) \\
&\text{with respect to} && x = \{\{c\},\{t\}\} \\
&\text{subject to} && c_{set}(x) < 0
\end{aligned}$$

$$\begin{aligned}
&\text{maximize} && AEP(x) \\
&\text{with respect to} && x = \{\{\theta\},\lambda\} \\
&\text{subject to} && V_{tip} < V_{tip_{max}}
\end{aligned} \tag{19}$$

$$\begin{aligned}
&\text{minimize} && m(x) \\
&\text{with respect to} && x = \{\{c\},\{t\}\} \\
&\text{subject to} && c_{set}(x) < 0
\end{aligned}$$

Minimizing mass, even when subject to the constraints, is not a particularly well-defined problem. The solution depends on the choice of variable bounds. In the "mass 1st" problem, the tip chord shrunk to the lower bound (0.5 m) and the position of the maximum chord (r_2 in Fig. 1) moved to its right bound (40% blade fraction). Moving to the extreme bounds allowed the optimizer to minimize mass near the root where the structure was heavy, while still keeping a large enough chord at the root to satisfy the fatigue constraint. For the "mass 1st" problem, the mass of the blades decreased significantly (6.6%), but the AEP also decreased (0.9%). There was a net decrease of 0.35% in cost of energy, but this is not a particularly useful design technique because the results are sensitive to the choice of bounds on the chord distribution.

We compared these designs to the more appropriate objective—the minimum cost-of-energy ("min COE") design

$$\begin{aligned}
&\text{minimize} && COE(x) \\
&\text{with respect to} && x = \{\{c\},\{\theta\},\{t\},\lambda\} \\
&\text{subject to} && c_{set}(x) < 0
\end{aligned} \tag{20}$$

The AEP, turbine mass, and cost of energy of the three optimized designs discussed in this section are compared in Fig. 6. The "AEP 1st" design maximizes AEP first and then minimizes mass (Eqn. (18)). The AEP optimization has constraints on the planform area and bending moment distribution. The "mass 1st" design minimizes mass first and then maximizes AEP (Eqn. (19)). Changes in the chord distribution are dictated by the structural optimization. The "min

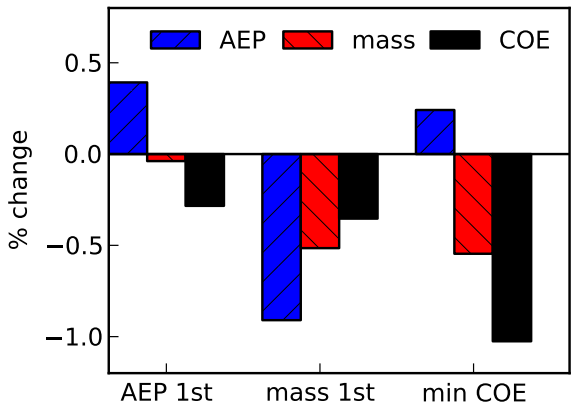

Fig. 6. Comparison between sequential aerodynamic and structural optimizations and an integrated aerodynamic and structural optimization. The percent change in mass is relative to the total turbine mass.

COE" design minimizes the cost of energy, which is the actual target objective (Eqn. (20)). Although only the rotor mass is optimized, the percent change in total turbine mass is shown in this figure (using constant mass values for the rest of the turbine). The mass of the rest of the turbine is added because the rotor mass changes by a relatively large percentage and would overwhelm the scale of the figure. The masses for the rest of the turbine come from those defined for the NREL 5-MW reference model [26] (m_{hub} = 56,780 kg, $m_{nacelle}$ = 240,000 kg and m_{tower} = 347,460 kg). Finally, the chord and twist distributions of the optimized designs are compared in Fig. 7.

With the inclusion of reasonable constraints to limit mass changes, maximizing annual energy production by shaping the blade and subsequently minimizing mass by changing the internal structure can lead to designs that decrease the cost of energy. However, the impact is much smaller than it could be, because minimum cost-of-energy designs tend to decrease mass until structural constraints become active. This is very difficult to approximate without including a structural model and a corresponding integrated metric into the optimization.

On the other hand, with reasonable bounds on the design variables, minimizing the mass by changing the blade's external and internal structure and subsequently maximizing aerodynamic performance by changing the airfoils and tip-speed ratio can also lead to designs that decrease the cost of energy. However, aerodynamic performance suffers and the result is significantly suboptimal compared to designs that integrate the aerodynamic and structural analyses and use a relevant combined metric.

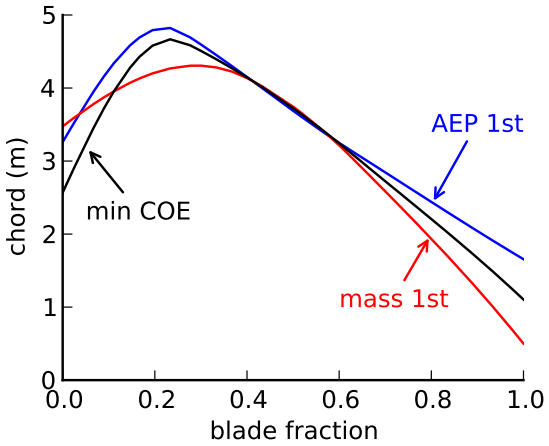

(a) Chord distribution along blade

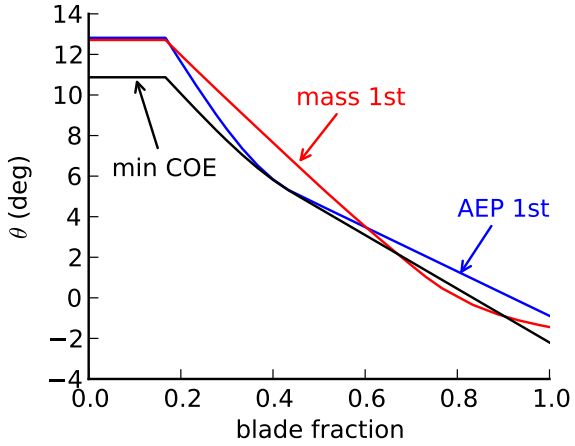

(b) Twist distribution along blade

Fig. 7. Chord and twist distribution for the three designs that were examined

3.2 Minimum Turbine Mass/AEP

Clearly, it is beneficial to directly integrate the aerodynamic and structural analyses in the optimization problem. Although maximizing AEP and minimizing mass sequentially is not particularly effective, an appropriate combined aero/structural metric can be useful. For designs with fixed materials, a reasonable choice is to minimize the ratio of the turbine mass to the annual energy production $m_{turbine}/AEP$. This objective can be motivated by the cost-of-energy equation, which includes capital costs and operating expenses in the numerator and AEP in the denominator

$$\text{COE} = \frac{\text{FCR} \cdot (\text{TCC} + \text{BOS}) + \text{O\&M}}{\text{AEP}} \tag{21}$$

where FCR is a fixed charge rate. In the absence of a detailed cost model, the assumption might be made that the capital costs and operating expenses will scale proportionally with the overall turbine mass.

For the remainder of this section, the mass of the turbine is simply denoted as m, rather than $m_{turbine}$. To compare the effectiveness of minimizing m/AEP with minimizing COE,

a number of studies were examined. The first study compared optimal turbines using both objectives at a fixed rotor diameter and a fixed machine rating (5-MW). The rotor diameter was varied at several different values to observe the trends in performance. The optimization problem was defined as

$$\begin{array}{ll} \text{minimize} & COE(x;D) \text{ or } m(x;D)/AEP(x;D) \\ \text{with respect to} & x = \{\{c\},\{\theta\},\{t\},\lambda\} \\ \text{subject to} & c_{set}(x) < 0 \end{array} \tag{22}$$

Although two separate objectives were used, for comparison purposes, both final designs were evaluated using the cost-of-energy metric. The difference in cost of energy between the two objectives as a function of rotor diameter is shown in Fig. 8a, and the variation in m/AEP is shown in Fig. 8b.

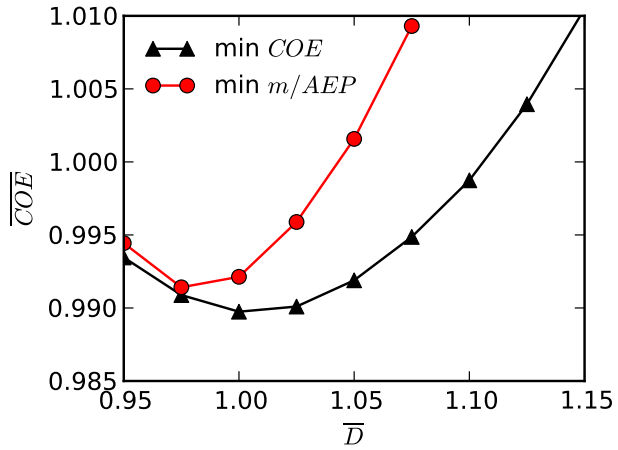

(a) Minimum m/AEP solutions are compared to minimum cost-of-energy solutions (both evaluated using COE)

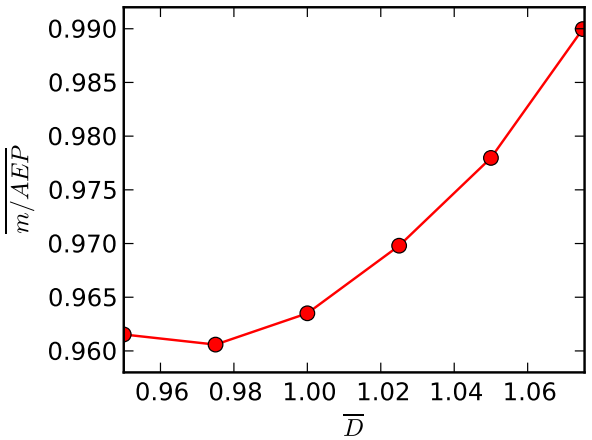

(b) Variation in m/AEP approximately predicts the correct optimal diameter

Fig. 8. Designs with minimum m/AEP as a function of rotor diameter

The metric m/AEP resulted in significantly suboptimal designs (suboptimal in terms of COE—they were the optimal solution for the chosen metric of minimum m/AEP). Both metrics found very similar designs for diameters smaller than the reference design, but for larger diameter rotors, the minimum m/AEP designs had about a 1% higher cost of energy. The reason for this discrepancy is that the metric m/AEP overemphasizes the role of the tower mass. Figure 9a shows the relative contribution to total mass from the different components of the turbine and Fig. 9b shows the relative contribution to total cost for the baseline design. The tower dominates the mass of the turbine at about 53%. However, it contributes a much smaller fraction to the total costs (about 9%). These relative contributions are typical of land-based turbines [32]. Thus, when the objective is to minimize m/AEP, the contribution of the tower plays a disproportionate role. At least, a fairly accurate optimal rotor diameter is predicted using this metric (Fig. 8b), which is not always the case with m/AEP metrics.

The difference between the minimum COE and minimum m/AEP designs can be seen more clearly by examining the chord and twist distributions as the rotor diameter increases. Figure 10 compares the optimal chord and twist distribution for the two metrics at $D/D_0 = 1.05$. The minimum m/AEP design attempts to sacrifice rotor performance to reduce thrust and decrease tower mass. We observe that the minimum m/AEP design uses a thinner structure outboard to reduce thrust (and power) and unloads the tip significantly to further reduce thrust. To compensate for the reduced structure, the design has a much higher spar cap thickness outboard. Compared to the minimum COE design, the minimum m/AEP design produces 7.7% less thrust at rated speed, which allows for a 6.7% lighter tower. At the same time, the annual energy production decreases by 1.9%. Because the contribution of the tower mass is so large, the net effect is still a reduction in m/AEP. In this case, minimizing m/AEP is not a good surrogate for minimizing cost of energy (though it is better than just maximizing AEP).

To improve the predictive ability of m/AEP, other alternatives may be considered. For example, if tower mass is disproportionately emphasized in minimizing m/AEP, perhaps the tower mass could be fixed. However, ignoring the tower completely and minimizing m_{rotor}/AEP is not a useful option. Similar to the previous example, this metric overemphasizes the effect of the rotor mass and thus leads to optimal solutions that reduce rotor mass at great expense to aerodynamic performance. Instead, including a fixed tower and nacelle mass is a much more reasonable approach. We defined

$$m_{fixed} = m_{blades} + m_{other} \tag{23}$$

where the mass of the blades was estimated as before, but the contributions of the other components were fixed at the values reported for the NREL 5-MW reference model [26] (m_{hub} = 56,780 kg, $m_{nacelle}$ = 240,000 kg, and m_{tower} =

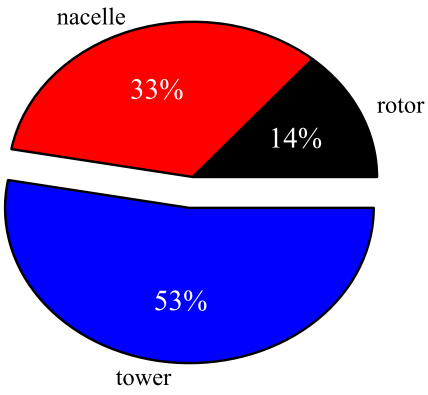

(a) Relative contribution to total mass

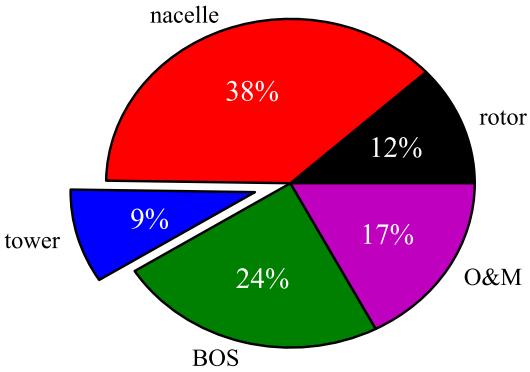

(b) Relative contribution to total cost

Fig. 9. Relative contributions to total mass and total cost of the baseline design. Cost contributions already include the fixed charge rate and tax rate. The tower contribution is of particular note.

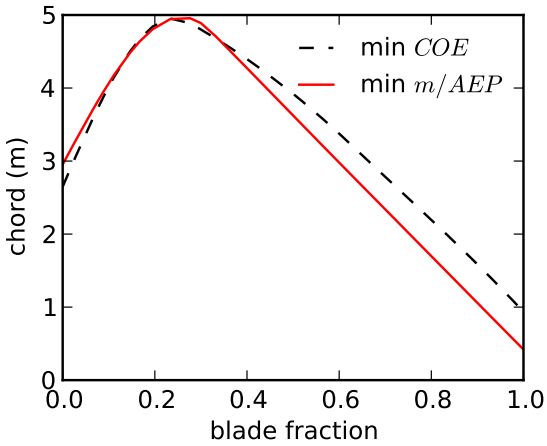

(a) Comparison between chord distributions

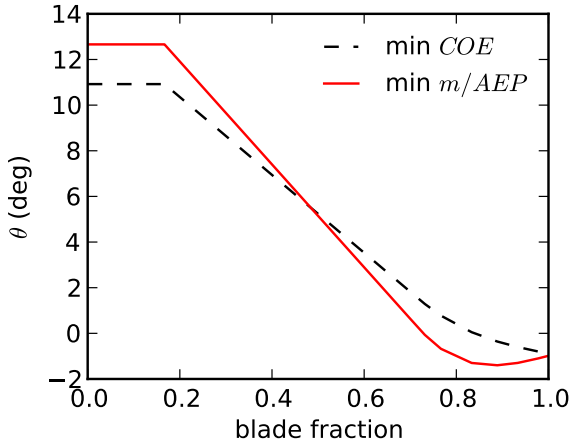

(b) Comparison between twist distributions

Fig. 10. Comparison of chord and twist distributions for minimum COE and minimum m/AEP design at $D/D_0 = 1.05$

347,460 kg). The optimization problem was then defined as

$$\begin{array}{ll} \text{minimize} & m_{fixed}(x;D)/AEP(x;D) \\ \text{with respect to} & x = \{\{c\},\{\theta\},\{t\},\lambda\} \\ \text{subject to} & c_{set}(x) < 0 \end{array} \tag{24}$$

The minimum m_{fixed}/AEP designs are compared to the minimum COE designs as a function of rotor diameter in Fig. 11a. Although the first design was not optimized for minimum COE, both designs were evaluated for cost of energy to facilitate a comparison. We note that this objective produced essentially the same designs as the minimum COE designs, which is somewhat fortuitous. The constant nacelle and tower masses played a fractional role that was similar to the role of the balance-of-station (BOS) and operation and maintenance (O&M) costs in estimating cost of energy (though they were not precisely constant). Figure 9 shows that the rotor mass and rotor costs consume similar fractions of total mass and total cost, respectively.

However, there is still a problem with the metric m_{fixed}/AEP. The use of this metric implies that a cost-of-energy estimate is absent. Figure 11b shows the variation in m_{fixed}/AEP as a function of rotor diameter. We see that the trend is completely opposite to that of the cost of energy. Although the trend may not always reverse, the relative change between designs will certainly differ between the two metrics. As the rotor diameter increases, the annual energy production decreases faster than m_{fixed} increases. The large discrepancy occurs because only the rotor mass is changing in m_{fixed}, and it is a relatively small component of the total mass. On the other hand, as the rotor diameter increases, turbine capital costs generally increase at a much faster rate than m_{fixed} (BOS and O&M costs also increase somewhat). These results suggest that m_{fixed}/AEP may work well for fixed diameter optimizations, but for variable diameter studies, this metric will not predict the correct rotor diameter, and may, in fact, suggest very different optimal diameters.

We also explored other approaches designed to prevent overincentivizing a reduction in tower mass by reformulating the optimization problem [31]. None of the approaches allowed for designs that were similar to minimum cost-of-energy designs, and simultaneously predicted correct opti-

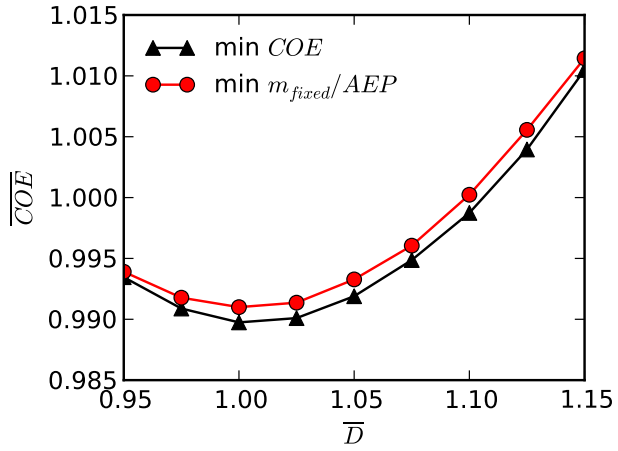

(a) Minimum m_{fixed}/AEP solutions are compared to minimum cost-of-energy solutions (both evaluated using COE)

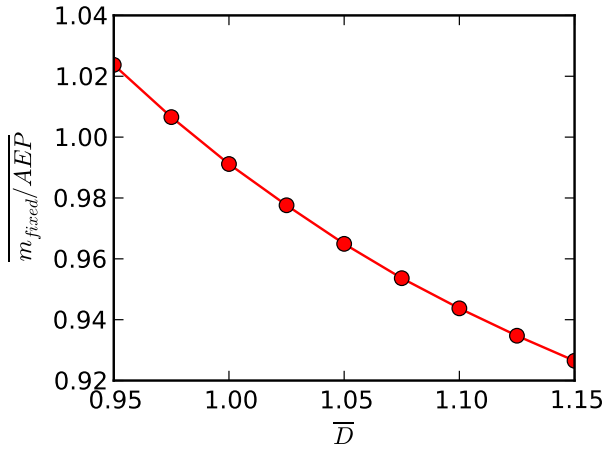

(b) Variation in m_{fixed}/AEP predicts wrong trend

Fig. 11. Designs with minimum m_{fixed}/AEP as a function of rotor diameter

mal diameters. However, the metric m/AEP can still be useful for designs with a fixed rotor diameter and power rating. Because the mass of the tower is such a large portion of turbine mass, and tower cost is not nearly as large of a fraction of total cost, the problem must be carefully designed to prevent the optimizer from taking advantage of this difference.

3.3 Minimum Cost of Energy

Though cost of energy is clearly a superior metric, a deterministic model may still be insufficient to produce robust designs. Even a relatively simple example can illustrate this. Consider that all designs discussed in this paper have been optimized for a fixed wind speed distribution (Rayleigh distribution with a mean wind speed of 10 m/s). However, even though a manufacturer designs a wind turbine for a particular wind power class, the turbine may end up being installed across a wider range of wind inflow conditions. We considered a scenario where the wind turbine was instead designed from the outset for use in a wide range of wind power classes (mean wind speeds of 6.4–11.9 m/s at 50 m above the ground). The turbine was optimized to minimize the expected value of the cost of energy across those wind speeds, and the 50-m wind speed was used as a proxy for the hub height wind speed. We gave each wind speed an equal weight (uniform distribution), though any distribution could be used just as easily. The optimization problem was defined as

$$\begin{aligned} &\text{minimize} && < COE(x;\overline{V}_{hub}) > \\ &\text{where} && \overline{V}_{hub} \sim \mathcal{U}(6.4, 11.9) \\ &\text{with respect to} && x = \{\{c\},\{\theta\},\{t\},\lambda, D, rating\} \\ &\text{subject to} && c_{set}(x) < 0 \end{aligned} \tag{25}$$

The cost of energy as a function of wind speed is shown for the optimal solution as compared to the point design, which minimizes COE at $\overline{V}_{hub} = 10$ m/s in Fig. 12. While the point design has a slightly lower cost of energy at the design point of $V_{hub} = 10$ m/s, the robust design has a much lower cost of energy at the lower wind speeds. The optimal blade shapes are similar; the main difference is that the robust design has a 16% lower machine rating to achieve a higher capacity factor at the lower wind speeds. This also allows the robust design to use a 3.4% larger radius and still satisfy the structural constraints. The net result is that the robust design achieves a 1.2% lower average cost of energy than the point design. Although there are clear benefits to site-specific design, even within a single site the variation in wind conditions may have a considerable impact.

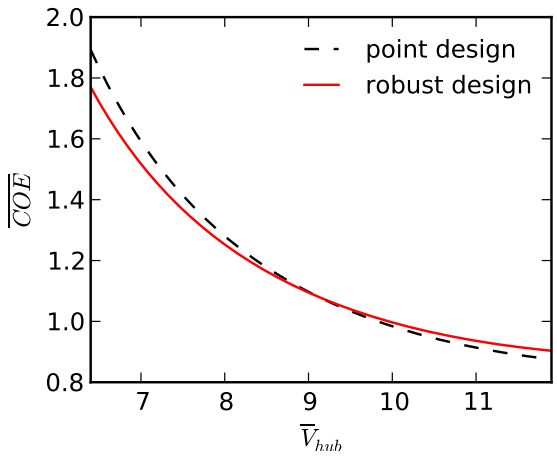

Fig. 12. Cost of energy as a function of hub wind speed for a point design that was optimized at 10 m/s, as compared to a robust design that was optimized to minimize the expected value of the cost of energy across the wind speeds

Structural sizing constraints are affected by uncertainty associated with manufacturing, loading conditions, and safety factors; performance metrics are affected by variability in the wind speed distributions, availability, and other losses. Characterizing the uncertainty of the model inputs and performing optimization under uncertainty is important to achieve robust designs.

4 Conclusions

In this paper, we discussed the limitations of various objectives that are commonly used in wind turbine optimization problems. The specific optimized designs produced in this study are not of primary interest, as they will vary depending on the assumptions and fidelity of the model. However, the relative differences in performance between different objectives highlight the following fundamental conclusions.

First, maximizing annual energy production typically leads to significantly suboptimal designs (in terms of cost of energy), even when the internal structure is subsequently optimized. Part of the difficulty is that similar aerodynamic performance can be achieved with designs that have very different masses. Appropriate aerodynamic surrogates for mass and structural limitations are helpful, but are still considerably inferior to a true structural model. Shaping the blade to minimize mass and subsequently optimizing airfoil sections for maximum aerodynamic performance is not any better. Although both approaches lead to decreases in cost of energy, they are inferior to metrics that combine the aerodynamic and structural performance.

Second, minimizing the ratio of turbine mass to annual energy production can be a useful metric, but only for certain design problems and only if used with care. Even if a nacelle and tower model are not used, a constant estimate of their mass must be included, otherwise potential decreases in rotor mass are overemphasized, which can lead to extremely suboptimal aerodynamic performance. As the rotor diameter (and thrust) change, exactly how the tower mass is estimated can substantially affect the optimal result. If a fixed tower mass is used, then the optimization works well at a fixed rotor diameter, but for a variable-diameter design, it predicts a very inaccurate optimal diameter. This error arises because, without changing the size of the nacelle or drivetrain, the ratio of turbine mass to annual energy production decreases much faster than it otherwise should. On the other hand, when the tower is allowed to resize, the problem must be constructed very carefully. The difficulty is that tower mass consists of a large portion of total mass, but tower cost is a rather small contribution to total cost. Thus, using m/AEP as the objective significantly overemphasizes the role of the tower. Without careful construction of the problem, the objective m/AEP overincentivizes the optimizer to decrease tower mass at the expense of aerodynamic performance. With carefully constrained designs, this metric can work quite well for a fixed rotor diameter, but it may still lead to incorrect optimal diameters for variable-diameter designs. For variable machine rating problems (or designs where the material is varied), this metric is not helpful at all.

Finally, minimum cost of energy is the appropriate metric to balance aerodynamic and structural performance with plant-level and operational costs. However, the fidelity of the cost model can dramatically affect the results. Along with increased fidelity in the physics, increased fidelity in cost modeling is needed. Furthermore, simply minimizing cost of energy may not be the most appropriate metric, because turbines need to be designed to perform well in a variety of conditions. Even within a single site, environmental conditions may vary significantly. Many of these inputs are inherently stochastic, and uncertainty exists in operational and model parameters. These considerations suggest that simply minimizing cost of energy will lead to inferior designs, and that optimization under uncertainty is particularly important given the stochastic nature of the wind.

This paper has highlighted many of the important design considerations in choosing appropriate objectives for wind turbine optimization problems. However, many opportunities exist to improve upon the insights discussed here. Rather than using scaling relationships, full physics-based models for the drivetrain and tower can be used. This will lend additional degrees of freedom to the optimization problem, allow for more rigorous sizing constraints, and lead to a better understanding of the trade-offs in rotor aerodynamic performance and turbine weight. The capital cost models used in this study are very simplistic; more detailed cost models are needed to better capture the effect of materials and manufacturing costs that are the result of changes to the structural ply schedule and blade shape. More thorough studies combining optimization with uncertainty quantification are needed, along with a better understanding of the nature of the uncertainties associated with the environmental conditions, physical processes, and cost metrics. Some of these potential improvements are currently being investigated.

Acknowledgments

The authors gratefully acknowledge George Scott of NREL for providing a Python version of NREL's cost and scaling model, and Katherine Dykes of NREL for providing the blade cost scaling relationship and suggesting the inclusion of drivetrain efficiency losses. This work was supported by the U.S. Department of Energy under Contract No. DE-AC36-08GO28308 with the National Renewable Energy Laboratory. Funding for the work was provided by the DOE Office of Energy Efficiency and Renewable Energy, Wind and Water Power Technologies Office.

References

[1] Fuglsang, P., and Madsen, H., 1999. "Optimization method for wind turbine rotors". *Journal of Wind Engineering and Industrial Aerodynamics,* **80**, pp. 191–206.

[2] Diveux, T., Sebastian, P., Bernard, D., Puiggali, J., and Grandidier, J., 2001. "Horizontal axis wind turbine systems: Optimization using genetic algorithms". *Wind Energy,* **4**, pp. 151–171.

[3] Fuglsang, P., and Thomsen, K., 2001. "Site-specific design optimization of 1.5–2.0 MW wind turbines". *Journal of Solar Energy Engineering,* **123**(4), p. 296.

[4] Fuglsang, P., Bak, C., Schepers, J., Bulder, B., Cockerill, T., Claiden, P., Olesen, A., and van Rossen, R., 2002. "Sitespecific design optimization of wind turbines". *Wind Energy,* **5**(4), pp. 261–279.

[5] Kenway, G., and Martins, J., 2008. "Aerostructural shape optimization of wind turbine blades considering site-specific winds". In Proc. of 12th AIAA/ISSMO

Multidisciplinary Analysis and Optimization Conference.

[6] Petrone, G., de Nicola, C., Quagliarella, D., Witteveen, J., Axerio-Cilies, J., and Iaccarino, G., 2011. "Wind turbine optimization under uncertainty with high performance computing". In AIAA Applied Aerodynamics Conference, AIAA 2011-3806.

[7] Maki, K., Sbragio, R., and Vlahopoulos, N., 2011. "System design of a wind turbine using a multi-level optimization approach". *Renewable Energy,* **43**(2012), pp. 101–110.

[8] Bottasso, C., Campagnolo, F., and Croce, A., 2010. "Computational procedures for the multi-disciplinary constrained optimization of wind turbines". *Scientific Report DIA-SR*, January, pp. 10–02.

[9] Dykes, K., Meadows, R., Felker, F., Graf, P., Hand, M., Lunacek, M., Michalakes, J., Moriarty, P., Musial, W., and Veers, P., 2011. Applications of systems engineering to the research, design, and development of wind energy systems. Tech. Rep. NREL/TP-5000-52616, National Renewable Energy Laboratory, December.

[10] Fingersh, L., Hand, M., and Laxson, A., 2006. Wind turbine design cost and scaling model. Tech. Rep. NREL/TP-500-40566, National Renewable Energy Laboratory, Golden, CO, December.

[11] Ning, A., Damiani, R., and Moriarty, P., 2013. "Objectives and constraints for wind turbine optimization". In 31st ASME Wind Energy Symposium.

[12] Ning, A., 2013. "A simple solution method for the blade element momentum equations with guaranteed convergence". *Wind Energy*.

[13] Du, Z., and Selig, M., 1998. "A 3-D stall-delay model for horizontal axis wind turbine performance prediction". In 1998 ASME Wind Energy Symposium, no. AIAA-1998-21.

[14] Eggers Jr, A. J., Chaney, K., and Digumarthi, R., 2003. "An assessment of approximate modeling of aerodynamic loads on the UAE rotor". In Aerospace Sciences Meeting and Exhibit, no. AIAA-2003-0868.

[15] Viterna, L., and Janetzke, D., 1982. Theoretical and experimental power from large horizontal-axis wind turbines. NASA TM-82944, National Aeronautics and Space Administration, Cleveland, OH. Lewis Research Center, September.

[16] Maples, B., Hand, M., and Musial, W., 2010. Comparative assessment of direct drive high temperature superconducting generators in multi-megawatt class wind turbines. Tech. Rep. NREL/TP-5000-49086, National Renewable Energy Laboratory, October.

[17] Bywaters, G., John, V., Lynch, J., Mattila, P., Norton, G., Stowell, J., Salata, M., Labath, O., Chertok, A., and Hablanian, D., 2004. Northern power systems WindPACT drive train alternative design study report. Tech. Rep. NREL/SR-500-35524, National Renewable Energy Laboratory, October.

[18] International Electrotechnical Commission, 2005. "International Standard IEC 61400-1". *Wind turbine generation systems*.

[19] Yang, T. Y., 1986. *Finite Element Structural Analysis*. Prentice-Hall.

[20] Bir, G., 2005. User's guide to PreComp. Tech. rep., National Renewable Energy Laboratory, September.

[21] Griffith, D. T., and Ashwill, T. D., 2011. The Sandia 100-meter all-glass baseline wind turbine blade: SNL100-00. SAND2011-3779, Sandia National Laboratories, Albuquerque, June.

[22] Bir, G. S., 2001. "Computerized method for preliminary structural design of composite wind turbine blades". *Transactions of the ASME,* **123**(4), pp. 372–381.

[23] Mandell, J., and Samborsky, D., 1997. DOE/MSU composite material fatigue database: Test methods, materials, and analysis. Tech. Rep. Contractor Report SAND97-3002, Sandia National Laboratories, Albuquerque, NM.

[24] Dykes, K., 2014. Development of wind turbine component mass-based cost models. NREL Technical Report (forthcoming).

[25] Maples, B., Hand, M., and Saur, G., 2014. Land-based wind plant balance of station cost and scaling model. NREL Technical Report (forthcoming).

[26] Jonkman, J., Butterfield, S., Musial, W., and G., S., 2009. Definition of a 5-MW reference wind turbine for offshore system development. Tech. Rep. NREL/TP-500-38060, National Renewable Energy Laboratory, Golden, CO, Feb.

[27] Lindenburg, C., 2002. Aeroelastic modeling of the LMH64-5 blade. Tech. Rep. DOWEC-02-KL-083/0, DOWEC 10083_001, Energy Research Center of the Netherlands, December.

[28] Kooijman, H., Lindenburg, C., Winkelaar, D., and van der Hooft, E., 2003. DOWEC 6 MW pre-design: Aero-elastic modeling of the DOWEC 6 MW pre-design in phatas. Tech. Rep. DOWEC 10046_009, Energy Research Center of the Netherlands, September.

[29] Resor, B., 2012. Definition of a 61.5-meter wind turbine blade reference model. Personal communication regarding a draft Sandia National Laboratories report.

[30] Berg, J., and Resor, B., 2012. Numerical manufacturing and design tool (NuMAD V2.0) for wind turbine blades: User's guide. Tech. Rep. SAND2012-728, Sandia National Laboratories, Albuquerque, NM.

[31] Ning, A., and Kroo, I., 2010. "Multidisciplinary considerations in the design of wings and wing tip devices". *Journal of Aircraft,* **47**(2), p. 534.

[32] Tegen, S., Hand, M., Maples, B., Lantz, E., Schwabe, P., and Smith, A., 2012. 2010 cost of wind energy review. Tech. Rep. NREL/TP-5000-52920, National Renewable Energy Laboratory, Golden, CO, April.