



Optimization of wind plant layouts using an adjoint approach

Ryan N. King^{1,2}, Katherine Dykes², Peter Graf², and Peter E. Hamlington¹

¹University of Colorado, Boulder, Colorado, USA

²National Renewable Energy Laboratory, Golden, Colorado, USA

Correspondence to: Ryan N. King (ryan.king@nrel.gov)

Received: 17 July 2016 – Discussion started: 31 August 2016

Revised: 6 December 2016 – Accepted: 6 January 2017 – Published: 10 March 2017

Abstract. Using adjoint optimization and three-dimensional steady-state Reynolds-averaged Navier–Stokes (RANS) simulations, we present a new gradient-based approach for optimally siting wind turbines within utility-scale wind plants. By solving the adjoint equations of the flow model, the gradients needed for optimization are found at a cost that is independent of the number of control variables, thereby permitting optimization of large wind plants with many turbine locations. Moreover, compared to the common approach of superimposing prescribed wake deficits onto linearized flow models, the computational efficiency of the adjoint approach allows the use of higher-fidelity RANS flow models which can capture nonlinear turbulent flow physics within a wind plant. The steady-state RANS flow model is implemented in the Python finite-element package `FEniCS` and the derivation and solution of the discrete adjoint equations are automated within the `dolphin-adjoint` framework. Gradient-based optimization of wind turbine locations is demonstrated for idealized test cases that reveal new optimization heuristics such as rotational symmetry, local speedups, and nonlinear wake curvature effects. Layout optimization is also demonstrated on more complex wind rose shapes, including a full annual energy production (AEP) layout optimization over 36 inflow directions and 5 wind speed bins.

1 Introduction

Optimizing wind turbine locations within a wind plant is a uniquely challenging problem that combines turbulent flow control with practical engineering challenges concerning the economical development of renewable energy. The problem is further complicated by the strong nonlinear coupling between turbine locations, power production, atmospheric boundary layer turbulence, and mechanical loads on turbine components. Wind plant optimization techniques used in industry often rely on heuristic guidelines and simplified linear flow models that limit computational costs. However, these approaches neglect important turbulent flow physics and can result in the underperformance of wind plants relative to their pre-construction estimates. The reduced power output and increased uncertainty due to optimization with low-fidelity flow models ultimately increases the levelized cost of energy (LCOE) and associated project risk for investors.

Major impediments to improving the LCOE for utility-scale wind plants are the difficulty of performing the turbine layout optimization using fluid dynamical models with sufficient fidelity and performing optimization in high-dimensional design spaces. In order to improve wind plant optimization, we present `WindSE`, a wind plant layout tool that uses a steady-state three-dimensional (3-D) Reynolds-averaged Navier–Stokes (RANS) flow model and gradient-based optimization. This gradient-based approach, made possible by automated adjoint derivations, yields an efficient, high-fidelity, and high-dimensional wind plant optimization framework. The software is part of the National Renewable Energy Laboratory Wind-Plant Integrated System Design and Engineering Model (`WISDEM`TM) (Dykes et al., 2014), which provides a modular framework for comprehensive analysis and optimization of wind plant designs. Through the use of a higher-fidelity 3-D RANS flow model and gradient-based adjoint optimization, `WindSE` is able to

provide optimized layouts with arbitrary topological complexity at relatively low computational cost, while at the same time more accurately capturing turbulent flow effects present in real wind plants.

This paper is organized as follows: background on wind plant flow modeling, layout optimization, and adjoint optimization techniques is provided in Sect. 2; a description of the methods used in *WindSE* for implementing the wind plant optimization problem, flow model, turbine representation, numerics, and optimization algorithm is given in Sect. 3; a demonstration of the flow model capabilities and optimization results for idealized and real world test cases with increasingly complex wind roses, culminating in a full annual energy production (AEP) optimization, is presented in Sect. 4; and the paper is concluded with a discussion of the results and implications for wind plant design in Sect. 5.

2 Background

2.1 Wind plant flow modeling

Utility-scale wind plants in the United States typically involve tens to hundreds of turbines arranged in semi-regular arrays. The layout topology is generally an outcome of optimizing the power production or net capacity factor subject to competing influences from site constraints, the local wind resource, and construction costs. These constraints include the patchwork of viable building areas formed by leases and setbacks due to environmental concerns or physical infrastructure, terrain and soil characteristics like slope or vegetation, turbine manufacturer spacing requirements, and continuity requirements imposed by access roads and electrical connections. This results in a complex design problem, with turbine layouts varying drastically between different geographic locations and exhibiting complex topologies.

The AEP from a wind plant layout has traditionally been, and generally continues to be, assessed using reduced-order linear flow models. Such models estimate the relative wind speed across a site based on the linearized Navier–Stokes equations and treat terrain features as perturbations in boundary conditions. The underlying governing equations, introduced by Jackson and Hunt (1975) and implemented in packages such as *MS3DJH/3R* (Walmsley et al., 1986), are based on analytical perturbation solutions to flow over a low hill. This approach decomposes the terrain into sinusoidal hills and calculates relative speedup effects over each hill. Speedup effects from multiple hills are superimposed to obtain relative velocities over the entire site. The emphasis on relative velocities is motivated by the need to extrapolate from point measurements at meteorological towers to velocity fields covering the entire plant.

A velocity deficit representing the turbine wake is then superimposed on the background wind resource at each turbine location. The *PARK* model developed by Jensen (1983) and the eddy viscosity model developed by Ainslie (1988) are

commonly used, and a comprehensive review of wake modeling can be found in Crespo et al. (1999). Such approaches decouple the wind flow calculation from the wake calculation and use wake models calibrated to a single turbine in isolation. As a result, the wake model approach neglects flow curvature, speed-up effects around turbines, and changes to wakes deep in a plant. This results in known inaccuracies in complex terrain, ad hoc model adjustments for overlapping wakes from multiple turbines, and systemic under-predictions of wake losses and over-predictions of power output (Barthelmie et al., 2009). Despite these limitations, linear flow models and prescribed wake models are commonly used in engineering practice because they are computationally inexpensive and reasonably accurate in simple terrain with few turbines.

Recently, higher-fidelity computational fluid dynamics (CFD) models have been increasingly applied to the study of wind plant performance. Several RANS models have been developed for simulating atmospheric flows in wind plants using actuator disk turbine representations, with a particular emphasis on modified $k - \epsilon$ closures (Cabezón et al., 2011; El Kasmi and Masson, 2008; van der Laan et al., 2015a, b). These approaches greatly improve upon linearized flow models and, given an appropriate model for the turbulent eddy viscosity, better capture turbulent transport effects in the shear layer at the edge of the wake. Large eddy simulation (LES) is also increasingly being applied to wind plant modeling (Calaf et al., 2010; Porté-Agel et al., 2011; Churchfield et al., 2012), as well as to some aspects of wind plant optimization. These simulations are vastly more expensive than RANS approaches, but resolve time-varying turbulent motions up to the filter cutoff scale. This allows for better characterization of wake meandering and unsteady loads on turbines. Finally, wind plant modeling has also been improved through coupling to numerical weather prediction simulations (Fitch et al., 2012; Mirocha et al., 2014). Such simulations are able to incorporate synoptic-scale weather forcing and interactions between the surface, wind plant, and full atmospheric boundary layer.

2.2 Approaches to wind plant optimization

To date, wind plant layout optimization has been performed primarily using linear flow models with analytical wake deficits. The relatively cheap computational cost of such models allows gradient-free methods to be used in the optimization process. A number of layout optimization studies have utilized gradient-free approaches to minimize wake losses (Kuskiak and Song, 2010), maximize net present value (González et al., 2010), and minimize noise propagation (Kwong et al., 2012). Other optimization approaches include particle swarm optimizations that determine turbine layouts and rotor diameters (Chowdhury et al., 2012), extended pattern searches for multimodal layout optimization (Du Pont and Cagan, 2012), and even game-theoretic methods (Mar-

den et al., 2013). An exhaustive review of wind plant optimization efforts has been compiled by Herbert-Acero et al. (2014) that surveys the wide variety of objective functions, flow models, and constraints that have been studied. Results from these prior optimizations are, however, heavily influenced by the use of analytical wake models which do not fully account for nonlinear and turbulent flow physics. Consequently, significant differences in optimal layouts are expected when using higher-fidelity flow models.

Recently, higher-fidelity CFD flow models have been used in a limited range of wind plant optimization applications. The Technical University of Denmark has developed TOPFARM (Larsen et al., 2011), which employs an improved wake model (the dynamic wake meandering model) as well as a parabolic Navier–Stokes solver. TOPFARM uses a hybrid optimization approach that combines sequential linear programming (SLP) with gradient-free genetic algorithms. The genetic algorithm is used to find the neighborhood of the global optimum, and then the gradient-based SLP algorithm completes the optimization. However, the genetic algorithm step penalizes large design spaces, limiting TOPFARM to relatively small wind plants. King et al. (2016) used adjoints of a 2-D RANS flow model to optimize turbine locations and observed that the 2-D nature of the flow solver resulted in substantial flow curvature. Meyers and Meneveau (2012) studied the effects of spacing and alignment in infinite wind plants with turbines arranged in regular gridded turbine arrays. In their study, a pseudospectral LES was used to model atmospheric boundary layer turbulence and complex wake interactions, but the optimized layouts were restricted to grids, as opposed to a more general layout topology optimization where turbines are free to arrange in non-gridded configurations. LES results have also been used to tune linear flow models that were used in optimization of yaw control (Gebraad et al., 2016), turbine layouts (Bokharaie et al., 2016), and coupled layout and yaw optimization (Fleming et al., 2016), or in the tuning of RANS models for wind plant control optimization (Jungo et al., 2016). Wind plant LES has also been used to directly perform adjoint optimization of wind plant controls by adjusting rotor thrust during operation (Goit and Meyers, 2015; Goit et al., 2016). These approaches benefit from high-fidelity CFD and leverage the power of adjoint optimization but have only been applied for fixed layouts. Finally, Funke et al. (2014) performed an adjoint optimization of ocean turbine layouts using an analysis based on the shallow water equations, with ocean turbines represented by increased bottom friction.

In the present study, we use adjoint techniques to enable gradient-based optimization of wind turbine locations within a plant, subject to realistic turbulent flow fields. A steady 3-D RANS flow solver is employed as a first-principles model that can predict new turbulent flow physics, rather than prescribing fixed wake behaviors as in linear flow models. The RANS model provides an accurate model of a neutral atmospheric boundary layer at moderate computational cost and

without requiring calibration using LES. The RANS model is also amenable to automatic differentiation and gradient-based optimization, as explained in the next section. This gradient-based approach permits the use of high-dimensional control spaces, thereby providing optimized layouts of arbitrary complexity (i.e., optimized layouts are not restricted to grids or any other regular arrangement). We further demonstrate layout optimization for the full plant AEP based on a real site wind rose, going beyond the uniform speed optimization considered previously (King et al., 2016). The resulting optimization framework thus represents a novel application of adjoint techniques to the optimization of utility-scale wind plants using a higher-fidelity CFD flow model.

2.3 Adjoint techniques for efficiently calculating gradients

The greater expense of CFD flow models such as RANS and LES compared to linearized flow models requires the use of an efficient optimization technique that minimizes the number of flow model evaluations. Gradient-based optimization methods are promising candidates for CFD-driven wind plant optimization since they require orders of magnitude fewer function evaluations than gradient-free techniques. However, finding these gradients can be challenging when using complex flow models or when there are many control variables. Calculating such gradients with a finite-difference approach requires a function evaluation for each control variable, making this approach prohibitively expensive for optimizing utility-scale plants with many turbines.

In the present optimization framework, the necessary gradients are obtained relatively inexpensively by using adjoint optimization techniques. A comprehensive review of discrete techniques for calculating gradients of engineering design problems, including the adjoint approach, can be found in Martins and Hwang (2013). The adjoint approach allows one to calculate gradients at a cost that scales with the dimension of the objective function rather than with the number of design variables. For wind plant optimizations with scalar objective functions, this means that gradients can be found at a fixed cost regardless of the number of control variables.

Adjoint systems arise naturally from consideration of dynamical systems whose behavior can be described by differential equations, and they have a rich history in fluid dynamics and sensitivity analysis. Reviews of adjoint techniques in various flow control and optimization applications have been published by Jameson (2003), Giles and Pierce (2000), Giannakoglou and Papadimitriou (2008), and Luchini and Bottaro (2014). At its core, the adjoint approach reverses the propagation of information in a dynamical system. Given an ordinary or partial differential equation governing the evolution of a dynamical system and an objective function measuring a quantity of interest, the adjoint approach produces a differential equation whose states evolve backward in time and indicates where a perturbation would maximally influence the

objective function. By reversing the flow of information, the adjoint reveals what is effectively the optimal open-loop control input.

The adjoint operator is defined by the bilinear identity $\langle \mathbf{A}u, v \rangle = \langle u, \mathbf{B}v \rangle$, where operator \mathbf{B} is adjoint to operator \mathbf{A} . This holds if \mathbf{A} is a continuous differential operator, in which case \mathbf{B} is found through integration by parts, or if \mathbf{A} is a matrix, in which case $\mathbf{B} = \mathbf{A}^*$, where \mathbf{A}^* is the complex conjugate transpose of \mathbf{A} . This bilinear identity reveals that the adjoint operator is implicitly defined by the forward model, but in order for the adjoint problem to be well posed, additional constraints are required. Because the adjoint travels backwards in time, these constraints are terminal conditions rather than initial conditions, and they come from the specific objective function under consideration as well as from the states produced by the forward model. The RANS model considered in this study is steady state in time and consequently the adjoint system is also steady state, equivalent to solving a single time step in the dynamical systems interpretation.

The following general framework illustrates the computational advantages of the adjoint approach. Consider a dynamical system with governing equations that can be expressed in residual form as $\mathcal{F}(\mathbf{u}, \mathbf{m}) \equiv 0$, where \mathcal{F} is a vector-valued differential equation (e.g., the RANS equations), $\mathbf{u} \in \mathbb{R}^n$ is the system state vector (e.g., the flow field velocities), and $\mathbf{m} \in \mathbb{R}^m$ is a vector of control variables (e.g., the wind turbine coordinates). Additionally, consider a scalar objective functional $J(\mathbf{u}, \mathbf{m}) \in \mathbb{R}$ that measures a quantity of interest (e.g., the LCOE). Many engineering problems can be formulated as constrained optimization problems that seek optimal states and control parameters to minimize J , namely

$$\begin{aligned} \min_{\mathbf{u}, \mathbf{m}} \quad & J(\mathbf{u}, \mathbf{m}) \\ \text{subject to} \quad & \mathcal{F}(\mathbf{u}, \mathbf{m}) = 0 \\ & h(\mathbf{m}) = 0 \\ & g(\mathbf{m}) \leq 0, \end{aligned} \quad (1)$$

where h and g are additional equality and inequality constraints on the control parameter \mathbf{m} , such as upper and lower bounds on a control input (e.g., wind plant site boundaries). In common engineering problems, \mathcal{F} is a partial differential equation (PDE) that is expensive to evaluate, and the dimensions of both the control space and state space are high. Efficiently solving this PDE-constrained optimization problem requires algorithms that scale well to high dimensions and that minimize the number of evaluations of \mathcal{F} .

A common approach to solving a PDE-constrained problem is to minimize a reduced functional $\hat{J}(\mathbf{m}) \equiv J[\mathbf{u}(\mathbf{m}), \mathbf{m}]$. This formulation takes advantage of the fact that the PDE constraint $\mathcal{F}(\mathbf{u}, \mathbf{m}) = 0$ implicitly defines the state \mathbf{u} in terms of the control \mathbf{m} . Minimizing the reduced functional is equivalent to minimizing the original functional if the governing equations \mathcal{F} produce a unique state \mathbf{u} for a

given set of controls \mathbf{m} , and if $\mathcal{F}[\mathbf{u}(\mathbf{m}), \mathbf{m}]$ is assumed to be continuously invertible so that $\mathbf{u}(\mathbf{m})$ is continuously differentiable (Hinze et al., 2009).

Gradient-based optimization algorithms can approach second-order convergence to local minima and can minimize the evaluations of \mathcal{F} , but require the gradient of the objective functional with respect to all of the control parameters. This gradient, $d\hat{J}/d\mathbf{m}$, is given by the chain rule as

$$\frac{d\hat{J}}{d\mathbf{m}} = \frac{dJ[\mathbf{u}(\mathbf{m}), \mathbf{m}]}{d\mathbf{m}} = \frac{\partial J}{\partial \mathbf{m}} + \frac{\partial J}{\partial \mathbf{u}} \frac{d\mathbf{u}}{d\mathbf{m}}. \quad (2)$$

Since $J(\mathbf{u}, \mathbf{m})$ is generally a user-defined function, $\partial J/\partial \mathbf{m} \in \mathbb{R}^{1 \times m}$ and $\partial J/\partial \mathbf{u} \in \mathbb{R}^{1 \times n}$ are straightforward to determine analytically. However, $d\mathbf{u}/d\mathbf{m} \in \mathbb{R}^{n \times m}$ is expensive to compute for high-dimensional control and state spaces. A finite-difference approach to calculating this gradient would require m evaluations of \mathcal{F} , which is intractable for many engineering problems.

In the adjoint approach, $d\mathbf{u}/d\mathbf{m}$ is eliminated from Eq. (2) by taking the derivative of the PDE constraint $\mathcal{F}(\mathbf{u}, \mathbf{m}) = 0$ with respect to \mathbf{m} , resulting in the tangent linear system

$$\frac{\partial \mathcal{F}}{\partial \mathbf{m}} + \frac{\partial \mathcal{F}}{\partial \mathbf{u}} \frac{d\mathbf{u}}{d\mathbf{m}} = 0. \quad (3)$$

Solving for $d\mathbf{u}/d\mathbf{m}$ in Eq. (3) and substituting into Eq. (2) then yields

$$\frac{d\hat{J}}{d\mathbf{m}} = \frac{\partial J}{\partial \mathbf{m}} - \underbrace{\frac{\partial J}{\partial \mathbf{u}} \left[\frac{\partial \mathcal{F}}{\partial \mathbf{u}} \right]^{-1}}_{\Psi^T} \frac{\partial \mathcal{F}}{\partial \mathbf{m}}, \quad (4)$$

where we have introduced the new *adjoint* variable Ψ . This variable maps source perturbations in $\mathcal{F}(\mathbf{u}, \mathbf{m}) = 0$ to sensitivities of J . By definition, it is governed by

$$\left[\frac{\partial \mathcal{F}}{\partial \mathbf{u}} \right]^T \Psi = \left[\frac{\partial J}{\partial \mathbf{u}} \right]^T. \quad (5)$$

With a solution of the forward model \mathcal{F} and the adjoint Ψ , the derivative of the objective function expressed in Eq. (2) can then be calculated as

$$\frac{d\hat{J}}{d\mathbf{m}} = \frac{\partial J}{\partial \mathbf{m}} - \Psi^T \frac{\partial \mathcal{F}}{\partial \mathbf{m}}. \quad (6)$$

Because Eq. (5) is independent of \mathbf{m} , the gradient can be calculated at a fixed cost regardless of the dimension of \mathbf{m} . This enables efficient gradient-based optimization for systems governed by computationally demanding PDEs. Equation (5) provides an intuitive interpretation of the adjoint as a linear mapping between source perturbation in the governing equations and changes in the objective function. The adjoint thus provides the optimal control input and directly encodes the sensitivities of \hat{J} . Alternatively, the adjoint can also be interpreted as a Lagrange multiplier enforcing the PDE constraint \mathcal{F} .

3 Methodology

In the present study, wind plant layout optimization is approached as a PDE-constrained optimization problem using the adjoint theory developed in the previous section. The wind plant power output is maximized with gradient-based optimization techniques, subject to a PDE constraint corresponding to the RANS equations, which are used to predict turbulent flow within the plant.

3.1 Optimization problem definition

We seek to maximize steady-state power output from N different turbines experiencing K different inflow wind states (each state is defined by a wind speed and direction) by controlling the 2-D Cartesian positions of the turbines, $\mathbf{x} = (x_1 \dots x_N)$ and $\mathbf{y} = (y_1 \dots y_N)$. The Reynolds-averaged velocity field $\bar{\mathbf{u}}_i \in \mathbb{R}^3$ and pressure field $\bar{p} \in \mathbb{R}$ are taken to be the states (collectively denoted \mathbf{u}) and the turbine coordinates are the control variable $\mathbf{m} = [\mathbf{x}, \mathbf{y}]^T \in \mathbb{R}^{2N}$. This leads to the following optimization problem:

$$\begin{aligned}
 \min_{\mathbf{u}, \mathbf{m}} \quad & J[\mathbf{u}(\mathbf{m}), \mathbf{m}] = -\sum_{k=1}^K \sum_{n=1}^N \alpha_k \frac{1}{2} \rho A_n \\
 & c'_{p,n} \beta_n^{-1} \varphi_n \|\bar{\mathbf{u}}(\mathbf{m}) \cdot \hat{\mathbf{n}}_k\|^3 \\
 \text{subject to} \quad & \bar{u}_j \frac{\partial \bar{u}_i}{\partial x_j} = -\frac{1}{\rho} \frac{\partial \bar{p}}{\partial x_i} + \nu \frac{\partial^2 \bar{u}_i}{\partial x_j^2} - \frac{\partial \tau_{ij}}{\partial x_j} \\
 & + \frac{1}{\rho} \sum_{n=1}^N f_{AD,n} \hat{\mathbf{n}}_k \\
 & \frac{\partial \bar{u}_i}{\partial x_i} = 0 \\
 & f_{AD,n} = \frac{1}{2} \rho A_n c'_{t,n} \beta_n^{-1} \varphi_n \|\bar{\mathbf{u}}(\mathbf{m}) \cdot \hat{\mathbf{n}}_k\|^2 \\
 & \tau_{ij} = -2\nu_T \bar{S}_{ij} \\
 & \nu_T = \ell_{\text{mix}}^2 (2\bar{S}_{ij} \bar{S}_{ij})^{1/2} \\
 & \bar{\mathbf{u}}_1 = u_{in}(z) \text{ on } \Gamma_{\text{in,top,bottom}} \\
 & L_{x,\text{lower}} < x_n < L_{x,\text{upper}} \\
 & L_{y,\text{lower}} < y_n < L_{y,\text{upper}} \\
 & (x_i - x_j)^2 + (y_i - y_j)^2 > D_{\text{min}}^2 \\
 & \forall i, j = 1 \dots N, i \neq j,
 \end{aligned} \tag{7}$$

where $J[\mathbf{u}(\mathbf{m}), \mathbf{m}]$ is the negative of the total power production, α_k are weights for each inflow state from the wind speed distribution and wind rose describing the site climatology, A_n is the turbine rotor area, and ρ is the air density. The PDE constraints are the steady 3-D RANS and continuity equations which govern the evolution of $\bar{\mathbf{u}}_i$ and \bar{p} . Boundary conditions corresponding to the log law are applied on inflow, top, and bottom boundaries, and outflow boundary conditions are applied on the sides and exit of the domain as further described in Sect. 3.3. The RANS equations have an

additional body force term $f_{AD,n}$ that represents the force imparted on the flow by each wind turbine. The Boussinesq hypothesis is used to close the deviatoric Reynolds stress term τ_{ij} by making use of an eddy viscosity ν_T . The eddy viscosity is calculated with a mixing length model (Wilcox, 2006), where the mixing length ℓ_{mix} is taken to be one-eighth the vertical distance from the bottom surface. A mixing length model was chosen because of its simplicity, differentiability, and familiarity to the wind energy community from use in traditional wake models (Ainslie, 1988) and in recent RANS models developed for real-time wind plant controls (Boersma et al., 2016). The power and thrust coefficients for turbine n , denoted $c'_{p,n}$ and $c'_{t,n}$, respectively, are derived from actuator disk theory and are described in the next section. The variables φ_n and β_n are geometric smoothing kernels and normalization constants, respectively, and are also described further in the next section. The site constraints are imposed as lower, $L_{x,\text{lower}}$ and $L_{y,\text{lower}}$, and upper, $L_{x,\text{upper}}$ and $L_{y,\text{upper}}$, bounds on the turbine \mathbf{x} and \mathbf{y} locations. An inter-turbine minimum spacing constraint is further enforced with a minimum spacing D_{min} of 3 times the rotor diameter (RD) used in the AEP simulations. The turbines are assumed to always yaw into the wind and the turbine body force is thus directed into the incoming wind by the unit vector $\hat{\mathbf{n}}_k$. Note that the subscript n in the above system of equations represents one of the N different turbines and the subscript k denotes one of the K different wind states included in the analysis. As described in the previous section, a reduced functional $\hat{J}(\mathbf{m}) = J[\mathbf{u}(\mathbf{m}), \mathbf{m}]$ is used in the adjoint optimization approach.

We stress that the strength of this approach is not the particular form of the RANS closure model, but rather its embedding in an adjoint optimization framework. Different objective functions, turbulence closures, or turbine representations can be easily implemented in this framework and still benefit from the adjoint approach.

3.2 Wind turbine representation

Turbines in the simulations are represented as non-rotating actuator disks using actuator disk theory, as described in standard wind energy texts (Burton et al., 2011). The power, P , and thrust force, T , generated by an actuator disk are given in terms of a power coefficient, c_p , a thrust coefficient, c_t , and an upstream reference velocity, u_{ref} , as

$$P = \frac{1}{2} \rho A c_p u_{\text{ref}}^3, \quad T = \frac{1}{2} \rho A c_t u_{\text{ref}}^2, \tag{8}$$

where power and thrust coefficients are given in terms of an axial induction factor a as

$$c_p = 4a(1-a)^2, \quad c_t = 4a(1-a). \tag{9}$$

The axial induction factor describes the velocity at the rotor disk, $u_{\text{rotor}} = u_{\text{ref}}(1-a)$, as a fraction of the upstream reference velocity. The reference velocity is typically taken to be

the far-field upstream velocity for a turbine in isolation, but as discussed by Sanderse et al. (2011), the determination of u_{ref} for waked turbines or in complex terrain is more difficult. We adopt the same approach used by Calaf et al. (2010) and Meyers and Meneveau (2012) and define modified power and thrust coefficients, c'_p and c'_t , respectively, that are based on the local rotor disk velocity u_{rotor} rather than u_{ref} . These modified coefficients are

$$c'_p = \frac{c_p}{(1-a)^3}, \quad c'_t = \frac{c_t}{(1-a)^2}, \quad (10)$$

resulting in power and thrust given by

$$P = \frac{1}{2} \rho A c'_p u_{\text{rotor}}^3, \quad T = \frac{1}{2} \rho A c'_t u_{\text{rotor}}^2. \quad (11)$$

For this study, turbine operating parameters of $c_t = 3/4$, $c_p = 0.34$, and $a = 1/4$ are used for all turbines, consistent with values used in previous CFD studies (Calaf et al., 2010; Meyers and Meneveau, 2012; Jimenez et al., 2007). This results in modified power and thrust coefficients of $c'_p = 0.81$ and $c'_t = 4/3$.

Standard actuator disk theory assumes that the rotor disk is uniformly loaded, but this introduces singularities at the rotor edge. To ensure that the thrust force and power production are continuously differentiable, as well as to avoid numerical instabilities, the turbine force and power production are smoothly distributed across the rotor swept area. This differentiability with respect to position is crucial for gradient-based layout optimization, and smoothly distributing rotor forces is common practice in actuator disk (Wu and Porté-Agel, 2010) and actuator line (Churchfield et al., 2012) implementations. This is accomplished here using a geometric smoothing kernel φ . The following multivariate exponential distribution is used as the smoothing kernel in order to create a rotor with almost compact support that is still continuously differentiable:

$$\varphi_n = \varphi(x, y, z; x_n, y_n, z_n) = \exp \left\{ - \left[\frac{x - x_n}{w} \right]^\gamma - \left[\frac{(y - y_n)^2 + (z - z_n)^2}{r^2} \right]^\gamma \right\}, \quad (12)$$

where $\varphi_n(x, y, z; x_n, y_n, z_n)$ is the distribution for the n th turbine centered on (x_n, y_n, z_n) , r is the rotor radius, w is the rotor half-width, and γ controls the sharpness of the rotor edge. Finally, we also introduce a constant β_n given by

$$\beta_n = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \varphi_n(x, y, z; x_n, y_n, z_n) dx dy dz, \quad (13)$$

which is used to normalize the values of the smoothed actuator disk. Both 1-D and 2-D prototypes of this smoothing function are shown in Fig. 1a and b for $\gamma = 6$ and $w = r/4$ with $r = 40$ m, and a hub-height slice of the actuator disk representation of a small wind plant is shown in Fig. 1c.

3.3 Simulation setup

The 3-D computational domain used in the simulations has horizontal dimensions of $2.4 \text{ km} \times 2.4 \text{ km}$ and a vertical dimension of 640 m (equivalent to $30 \times 30 \times 8 \text{ RD}$, with $\text{RD} = 80 \text{ m}$), as shown in Fig. 2. Dirichlet boundary conditions are used to prescribe the wind speed and direction on inflow, top, and bottom boundaries, corresponding to the planes at $x = -1.2 \text{ km}$, $z = 640 \text{ m}$, and $z = z_0$, respectively. The inflow velocity profile is specified according to a neutral logarithmic velocity profile

$$u_{\text{in}}(z) = \frac{u_*}{\kappa} \ln \left(\frac{z}{z_0} \right), \quad (14)$$

where z is the vertical coordinate, z_0 is the roughness height, u_* is the friction velocity, and $\kappa = 0.4$ is the von Kármán constant. The roughness height is taken to be $z_0 = 0.04 \text{ m}$, corresponding to open and relatively smooth terrain, and u_* is found by solving for the desired hub-height velocity. A no-slip boundary condition is used on the bottom boundary and the velocity along the top boundary is u_{in} (640 m) from Eq. (14). On outflow boundaries (i.e., the planes at $x = 1.2$, $y = -1.2$, and $y = 1.2 \text{ km}$), a standard “do-nothing” finite-element outflow boundary condition (Heywood et al., 1996) is applied. The initial velocities at all locations in the domain are given by the logarithmic profile in Eq. (14), and the initial relative pressure is assumed to be 0 bar everywhere.

A coarse computational mesh of finite elements is generated for the entire domain and then further refined within a circle that circumscribes the site boundaries, as shown in Fig. 2. For all tests presented here, the site boundary is assumed to be a square of side 18 RD centered in the computational domain. This area corresponds to a 6 RD spacing in streamwise and spanwise directions if 16 turbines are arranged in a regular grid. The resolution of the finite elements used in the simulation can be quantified by the radius of a circle that circumscribes a single tetrahedral element, termed the circumradius. The mesh is refined such that the turbine rotor diameter is at least 4 times larger than the circumradius of the finite elements within the wind plant site boundaries. The mesh is also stretched by a factor of 1.2 in the vertical direction in order to increase resolution near the bottom boundary. This results in a mesh with approximately 200 000 degrees of freedom.

To confirm that the domain is sufficiently large for the number of turbines, we examined the blockage ratio. In wind tunnel studies, the blockage ratio is typically defined as the ratio of the total tunnel cross-sectional area to the total rotor disk area. A worst-case blockage scenario would involve all 16 turbines impeding the flow with no overlap between rotors. This worst-case results in a blockage ratio of 5.2 %, which is well below the threshold of 10 % blockage that requires a correction in wind tunnel studies of horizontal-axis wind turbines (Chen and Liou, 2011).




Figure 1. A continuously differentiable modified exponential distribution is used to smoothly distribute the turbine forces and power production over the rotor swept area. The smoothing function is demonstrated for an 80 m rotor diameter turbine with $\gamma = 6$ for 1-D and 2-D prototypes in panels (a) and (b) by taking conditional distributions of $\varphi(x, y, z)$ in Eq. (12). The edge of the rotor disk is represented by vertical dashed red lines in panel (a). Panel (c) is a hub-height slice showing how a wind plant is represented as a summation of smoothed actuator disks.

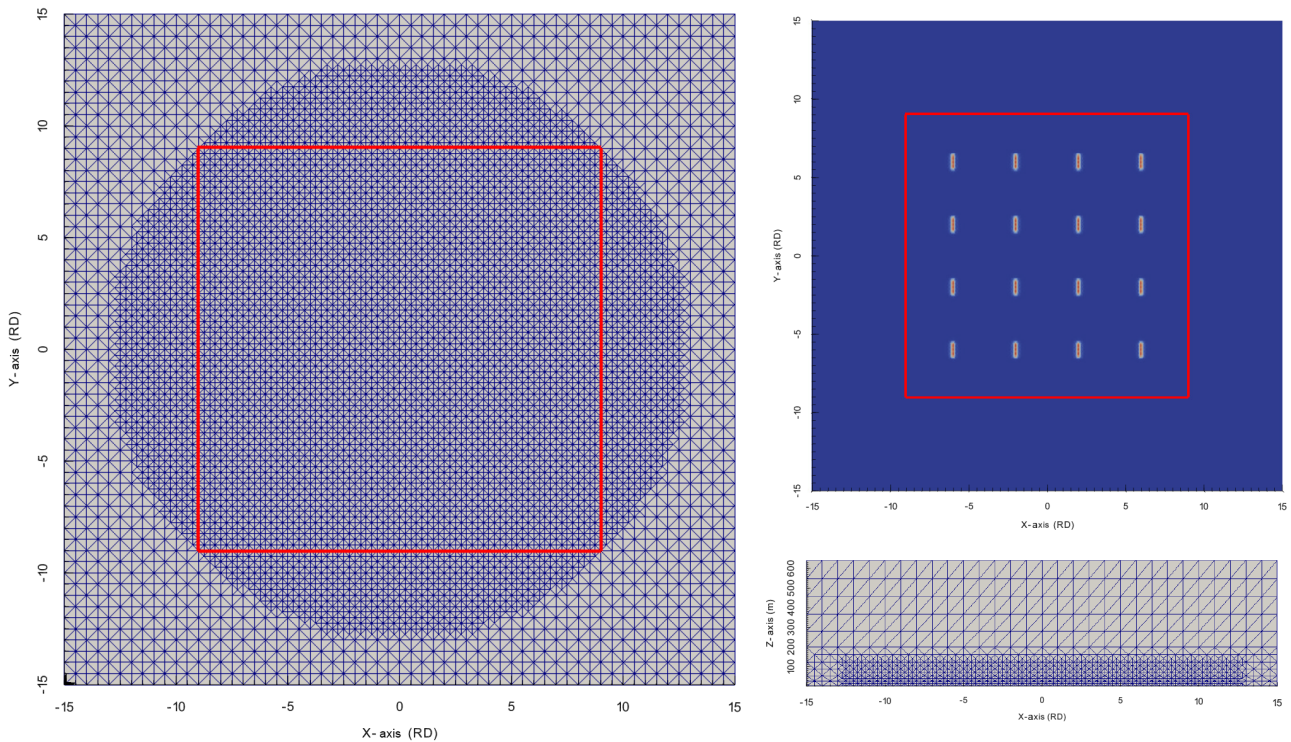


Figure 2. Plan view of the computational mesh at hub height showing additional refinement around the wind turbine area (left), schematic of the computational domain showing the turbine locations used to initialize each optimization (top right), and side view of the mesh showing refinement below 2 rotor diameters (bottom right). The wind plant site constraint is shown as a red square in the plan view, and horizontal dimensions are normalized by the rotor diameter $RD = 80$ m.

Simulations are performed for a range of different inflow directions and inflow speeds, corresponding to both idealized and real-world wind roses and wind speed distributions. Steady-state solutions are found for each of the K wind states. We assume that the turbines are always yawed into the

wind such that the rotor disk is perpendicular to the inflow direction. The computational domain and boundary conditions are assumed constant and the control variables containing the turbine locations are defined in a fixed reference frame. Different wind directions are modeled by applying a rota-

tion to the turbine coordinates corresponding to the inflow angle when calculating actuator disk forces. This rotation is included in the adjoint calculation and the resulting gradients are with respect to changes in the reference frame turbine coordinates. A weighted sum of the power production for each wind state (i.e., speed and direction) is performed based on the site wind speed distribution and wind rose, and the adjoint gradient is calculated over the total power output, taking into account the layout rotation. Because the boundary conditions are part of a well-posed PDE used as an optimization constraint, this approach of rotating the layout is preferred to explicitly changing the inflow direction because the PDE constraint in the optimization is kept constant.

3.4 Gradient-based layout optimization process

Figure 3 shows a schematic of the multi-state optimization workflow. The layout is optimized over all inflow states simultaneously in a multilevel optimization process with respect to the total power output rather than in a sequential optimization process over each inflow separately. The basic optimization process is the following:

1. Begin with an initial layout \mathbf{m}_i , which can be either a gridded or random layout of the N turbines.
2. Perform flow-field simulations for each of the K desired wind states (corresponding to different wind speeds and directions obtained from either real or idealized wind roses and speed distributions).
3. Calculate the negative of the wind plant power, J_k , for each of the K wind states.
4. Calculate the objective function by taking a weighted sum $J = \sum \alpha_k J_k$ over all K wind states, where α_k is the relative probability of each wind state obtained from either a real or idealized wind rose and speed distribution.
5. Compute adjoint simulations for the forward simulations, and calculate the gradient of the reduced functional $d\hat{J}/d\mathbf{m}$.
6. Use the gradient to perform a gradient-based optimization of the layout to obtain \mathbf{m}_{i+1} and go to step 2 above.

This process is repeated until the change in the objective function or its gradient falls below a user-defined threshold, which typically occurs in 30–50 iterations.

It should be noted that, because the optimization algorithm is gradient-based, it finds local rather than global minima in the total objective function J . The layout optimization problem can have many local minima, particularly with just a few inflow wind states. Starting with a regular gridded layout slightly smaller than the site constraint was generally observed to produce reasonable results since none of the variables are initially constrained and a gridded layout is

often a “good enough” guess to be in the radius of convergence to the global optimum. This initial layout, used for all simulations described in the following, is shown in Fig. 2. However, gradient-based methods are still fundamentally local searches and cannot provide strong assurances of finding a global minima. We note that with many inflow states (i.e., for large K), the optimization problem actually becomes more convex and convergence is achieved in fewer iterations. Additionally, running the optimization from many different starting configurations generated by random sampling or Latin hypercube samples can be used to further characterize the robustness of the minima.

3.5 Numerical implementation

The WindSE flow solver is implemented in a software package called FEniCS (Logg et al., 2012), which automates the solution of PDEs using the finite-element method. FEniCS is written in Python and can be easily integrated with other Python-based systems engineering tools like WISDEM. The FEniCS project is based on the DOLFIN problem-solving environment and connects a number of useful components for the automatic discretization and solution of finite-element problems. These components include a form language that allows users to specify equations in variational form using a syntax that closely resembles their mathematical description, automated compilers that generate finite-element forms for a chosen basis, and just-in-time compilation to C++ to enhance computational speed. FEniCS can interface to common HPC libraries such as PETSc and Trilinos for numerical linear algebra, ParMETIS and SCOTCH for domain decomposition, and MPI and OpenMP for parallelization. FEniCS has been extensively tested and validated on a number of computational problems in solid and fluid mechanics, eigenvalue problems, and coupled PDEs (Logg et al., 2012).

The description of finite-element problems as variational forms in FEniCS lends itself to highly abstracted algorithmic differentiation. The software package `dolfin-adjoint` (Farrell et al., 2013) performs a high-level algorithmic differentiation of a forward problem implemented in FEniCS and can derive both discrete adjoint and tangent linear models. The discrete adjoint and tangent linear models are important in the gradient-based optimization algorithms that are used in data assimilation, optimal control, and error estimation. The algorithmic differentiation routines in `dolfin-adjoint` act on the forward problem discrete equations described in variational form and stored in memory. The corresponding discrete adjoint equations are derived in `dolfin-adjoint` and passed back to FEniCS as an additional PDE problem. This approach operates at a higher level of abstraction than traditional algorithmic differentiation, which typically treats forward models as a series of elementary instructions involving native operations in low-level code like addition and multiplication. This higher level of abstraction gives `dolfin-adjoint` greater flexibility

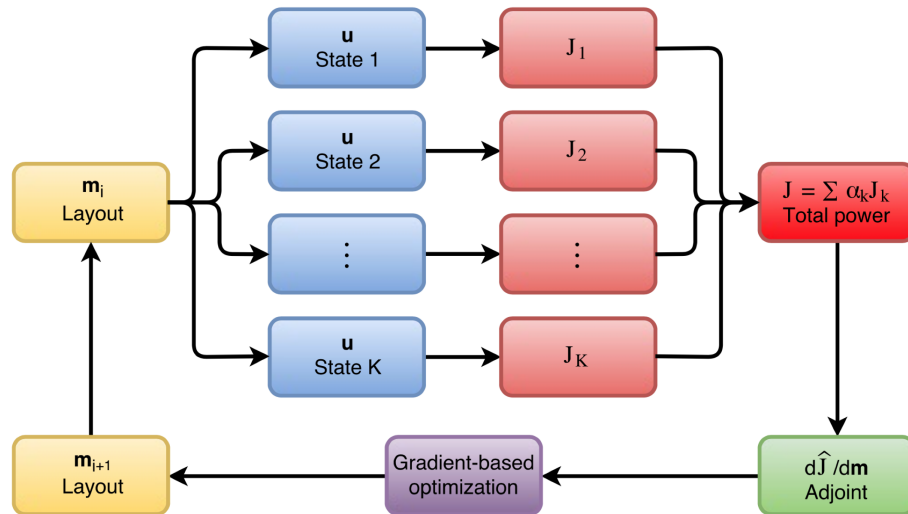


Figure 3. Schematic of the multilevel optimization process for a wind plant with K wind states (i.e., wind speeds and directions).

and automation across a wide range of PDE applications because it avoids differentiating across low-level code where the mathematical and implementation details have been inter-mixed. Moreover, `dolfin-adjoint` can be implemented on unsteady and nonlinear PDEs, and can also be run in parallel. It can directly interface to the optimization algorithms in SciPy and also contains routines for checking the correctness of adjoint gradients and checkpointing.

The 3-D RANS and continuity equations that form the PDE constraint in Eq. (7) are solved with a nonlinear Newton solver in a coupled fashion using a mixed finite-element space with piecewise linear elements for both the velocity and pressure fields. To satisfy the Ladyzhenskaya–Babuška–Brezzi (LBB) (or inf-sup) compatibility condition (Brezzi and Fortin, 1991) with equal-order basis functions, we augment the momentum equation with an additional pressure-stabilized Petrov–Galerkin term that weights the residual of the momentum equation by the gradient of the pressure test function. This pressure-based stabilization alleviates the saddle-point nature of the equal-order finite-element problem (Donéa and Huerta, 2003) but still vanishes for the exact solution to the momentum equation. Each nonlinear solve is initialized with the base logarithmic velocity profile and the relative residual is converged to below 10^{-7} with Newton’s method. The Newton solver uses Jacobians derived automatically within `FEniCS` and linear systems are solved directly with the sparse, parallel solver MUMPS (Amestoy et al., 2000). The choice of equal-order piecewise linear mixed finite-element spaces differs from previous studies on wind and ocean turbine layout optimization in `FEniCS` (King et al., 2016; Funke et al., 2014) that used Taylor–Hood mixed finite-element spaces that are piecewise quadratic for the velocity field. The lower-order representation in this study was necessary when implementing a 3-D solver to keep the to-

tal degrees of freedom sufficiently low so that a direct linear algebra solver could be used.

The gradients obtained from `dolfin-adjoint` are used to optimize turbine locations with Python’s SciPy implementations of the sequential least-squares programming (SLSQP) or limited-memory Broyden–Fletcher–Goldfarb–Shanno (BFGS) algorithm with bounds (L-BFGS-B) algorithms. SLSQP is a gradient-based optimization algorithm that can also handle constraints (Nocedal and Wright, 1999). SLSQP minimizes a quadratic approximation to the objective function at each optimization iteration, with a linear approximation of the constraints. The L-BFGS-B algorithm (Byrd et al., 1995) is a limited-memory version of the popular BFGS algorithm that approximates the inverse Hessian matrix used in quasi-Newton methods. We use the L-BFGS-B algorithm for simpler test cases without inter-turbine spacing constraints as in Sect. 4.2 and 4.3. In Sect. 4.4 we consider a real-world AEP optimization over a full wind rose and do enforce inter-turbine spacing constraints, which requires the use of the SLSQP algorithm. Gradients of the objective function are provided by `dolfin-adjoint` and gradients of the minimum turbine spacing constraint are calculated analytically.

The forward and adjoint problems are parallelized with MPI and can be run on a desktop or in a high performance computing environment. The discrete adjoint calculation is automatically parallelized by `dolfin-adjoint` if the forward model is run in parallel, which drastically simplifies code development.

4 Results

In the following, we present results for several different layout optimization cases. First, we provide results for standard test cases of flow past a single turbine and flow through a




Figure 4. Flow past a single turbine obtained using the 3-D RANS flow solver. The top panel shows the velocity at hub height and the bottom two panels show profiles of the velocity deficit in horizontal and vertical planes passing through the center of the wind turbine rotor. Velocity deficits are relative to the respective profiles 3 RD upstream of the turbine and the velocities are normalized by the incoming hub-height velocity. Axes are in units of rotor diameter $RD = 80$ m.




Figure 5. Hub-height velocity fields from the 3-D RANS solver used in *WindSE* show good qualitative agreement with time-averaged LES velocity fields of similar wind plants reported in the literature (Wu and Porté-Agel, 2013), even in the case of very deep wind plants. Velocities are normalized by incoming hub-height speed of 8 m s^{-1} and distances are normalized by the 80 m rotor diameter.

very deep wind plant in order to demonstrate that the RANS flow solver accurately captures wind turbine wakes, thereby providing confidence that subsequent layout optimizations are performed according to the correct flow physics. Second, we optimize a 16-turbine wind plant using wind roses with evenly weighted wind directions and a constant wind speed of 8 m s^{-1} in order to demonstrate new layout insights and optimization heuristics when accounting for nonlinear flow effects with a high-fidelity model. Third, we optimize a 16-turbine wind plant using unevenly weighted wind roses that exhibit complex directional preferences, again with a constant wind speed of 8 m s^{-1} . Finally, we perform a full AEP optimization using data from the M2 meteorological tower

at the National Wind Technology Center (NWTC) to demonstrate the capabilities of *WindSE* when optimizing layouts for conditions that are representative of real wind plants. In the full AEP optimization, both the wind direction and wind speed are varied, resulting in approximately 100 total wind states over which the optimization is performed. In all cases, we assume a turbine with a rotor diameter of $RD = 80$ m operating with fixed power and thrust coefficients (as described in Sect. 3.2) and a site constraint that is a square of side 18 RD centered in the computational domain.

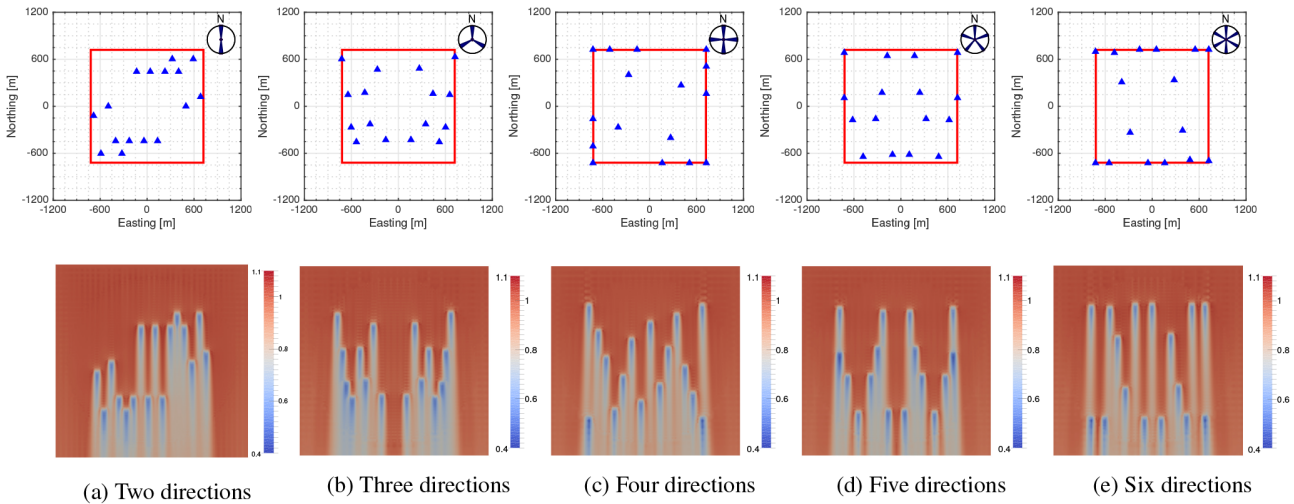


Figure 6. Optimal layouts (top row) and flow fields (bottom row) for five test cases with an increasing number of evenly weighted inflow directions. The wind roses in the upper right corner of each layout plot show the inflow directions used in the optimization. The blue triangles show the optimized turbine locations and the red square indicates the site boundaries. The velocities are normalized by the hub-height inflow velocity, which is 8 m s^{-1} in these simulations. Each of the flow fields are shown for a wind that blows from the north.

Table 1. Relative efficiency of power production from the test cases compared to 16 turbines with no wake effects. Results from a “naïve” case with two parallel turbine rows and no layout optimization are also included to demonstrate the improvements achieved by the layout optimization.

All unawaked	Two-dir. naïve	Two-dir. opt.	Three-dir. opt.	Four-dir. opt.	Five-dir. opt.	Six-dir. opt.
1.0	0.833	0.9854	0.953	0.960	0.948	0.974

4.1 RANS model testing

As a test of the qualitative performance of the RANS flow solver, WindSE was used to simulate flow fields for both a single turbine and a deep wind plant. Figure 4 shows vertical and horizontal velocity profiles in the wake of a single wind turbine. Consistent with theoretical expectations and previous results, a logarithmic velocity profile is observed in the undisturbed flow upstream of the turbine, a Gaussian velocity deficit is observed in the turbine wake, and a gradual wake recovery is observed with increasing distance downstream from the turbine. Moreover, a slight speedup around the edges of the wake is observed near the turbine rotor – such speedups are not captured by traditional linear wake flow models and appear here due to the use of the higher-fidelity RANS flow solver, which more accurately captures nonlinear flow physics. Figure 5 shows the velocity field from a simulation of a wind plant with 10 turbine rows perpendicular to the incoming wind direction. This flow field qualitatively agrees with time-averaged LES results reported by Wu and Porté-Agel (2013).

The results presented in this section are only intended to demonstrate the qualitative agreement of the RANS solver with prior high-fidelity studies and analytical wake theory. This demonstrates that the present flow solver captures a rea-

sonable level of fidelity to introduce the adjoint optimization framework. A detailed turbulence model verification and validation is beyond the scope of this study as we are instead focused on the integration of the model into a flexible and automated adjoint optimization framework. A comprehensive study on the implementation of more sophisticated turbulence models within the WindSE framework is left for future research.

4.2 Layout optimization for evenly weighted wind roses

In order to demonstrate many of the characteristics of optimal layouts and flow fields found using WindSE, Fig. 6 shows results obtained for wind roses with an increasing number of evenly weighted inflow directions. In all cases, the wind speed was assumed constant at 8 m s^{-1} . During the optimization, the turbines are constrained to stay within an $18 \text{ RD} \times 18 \text{ RD}$ region in the center of the computational domain (see also Fig. 2). This region serves as the site boundary in the present tests. The turbines are initialized on a regular grid that spans two-thirds of the site constraints (as shown in Fig. 2) and no inter-turbine spacing constraint was applied. Convergence was achieved in approximately 30–50 iterations of the optimization loop shown in Fig. 3.




Figure 7. Spanwise velocities produced by the optimized layout found in the five-inflow-direction case shown in Fig. 6d. The spanwise velocity is normalized by the hub-height inflow velocity, indicating strong curvature at the rotor and a slight overall deflection of the flow away from the center of the wind plant.

Figure 6 shows that optimal layouts are symmetric about a rotation or reflection when the wind directions are evenly weighted. For the two-direction and six-direction cases shown in Fig. 6a and e, the layout is symmetric about a 180° rotation and for the four-direction case shown in Fig. 6c, the layout is symmetric about a 90° rotation. For odd numbers of inflow directions, Fig. 6b and d show that optimal layouts are symmetric about a horizontal reflection across the north–south axis. The rotational and reflectional symmetries of the layouts for evenly weighted, uniform speed wind roses are useful heuristics for checking the flow solver and optimization results.

The flow fields shown in Fig. 6 indicate that the RANS solver captures flow curvature due to nonlinear turbulent transport effects and pressure increases upwind of the rotor disk. These effects are neglected when superimposing a wake deficit on a flow field obtained from a linear model, as is typically done in many industry-standard optimization frameworks. The flow curvature results in local speedup effects between two turbines, or just outside a wake, that downwind turbines take advantage of in strongly directional wind roses. This effect is further demonstrated in Fig. 7, which shows nonzero spanwise velocities generated by the flow curving around the turbines. This curvature is responsible for the “staggered” appearance of many of the layouts in Fig. 6 where the optimizer takes advantage of these local speedups. Because the power production scales with the cube of the wind speed, these small speedups can have a strong nonlinear effect on the power output. This speedup effect due to flow curvature is particularly enhanced in 2-D, but is still present in the 3-D simulations.

Flow curvature also affects the propagation and interaction of the turbine wakes. Wakes near the edge of the plant

are slightly deflected away from the plant center and reflect the overall spreading of the flow streamlines. This curvature can be observed near the edges of the plant in the cases shown in Fig. 6. Such curvature is again not captured by engineering wake models which prescribe wakes that always travel perpendicular to the rotor. Additionally, the wakes are pinched, curved, or merged when encountering speedups around downwind turbines or other wakes. The RANS flow solver accounts for the effects of other turbines and their wakes on the expansion and dissipation of wakes beyond what is accounted for in prescribed wake models. The deflection and curvature of wakes likely has important ramifications for yaw control strategies that attempt to steer wakes away from downwind turbines.

It is emphasized that the results in Fig. 6 show that the optimizer does not place turbines in straight rows perpendicular to a single predominant wind direction or in a regularly spaced grid for these evenly weighted wind rose cases. Instead, the turbines are placed closer together and staggered to take advantage of local speedups between laterally placed turbines. This is a different strategy than the maximized spacing found when optimizing with linear flow models and prescribed wakes, and is likely sensitive to the number and size of wind direction bins.

We compared the optimized layout for two inflow directions shown in Fig. 6a to a case with two parallel rows of turbines aligned perpendicular to the inflow directions and with maximal spacing between the rows as allowed by the site constraints. This “naïve” strategy is often found when using linear flow models. The optimized layout increases power production by 18.4% over the “naïve” layout due to the strongly directional speedups in this case. We also examined the power production of the test cases when normalized by the total power available if all turbines were unwaked, shown in Table 1. The optimized layouts substantially reduce, but do not entirely eliminate, the wake losses.

4.3 Layout optimization with unevenly weighted wind roses

In the previous section, each inflow direction was given an equal weight (i.e., α_k) in creating the total objective function J . However, real-world wind roses are seldom so simple and typically have several preferred wind directions, with many other less dominant directions. Here we demonstrate optimization of a 16-turbine layout using the same computational domain and setup as in the previous section, but with unevenly weighted wind roses that have two and three dominant directions. Once again, in both cases we assume a constant 8 m s^{-1} wind speed and do not enforce inter-turbine spacing constraints.

The first wind rose considered has two prominent wind directions spaced 107° apart, with a random normal distribution of directions around the primary axes. The wind rose is binned into 36 directions and is roughly shaped like a




Figure 8. Optimal layout (left) and flow field (right) for the boomerang test case with the wind rose binned into 36 inflow directions and a uniform hub-height velocity of 8 m s^{-1} . The wind rose is shown in the upper right corner of the left panel, and the site boundary is shown by a solid red line. Turbine locations are denoted by blue triangles. The flow field in the right panel is normalized by the incoming hub-height velocity of 8 m s^{-1} and shows results for an inflow wind from the south.




Figure 9. Optimal layout (left) and flow field (right) for the NWTC M2 8 m s^{-1} test case with the wind rose binned into 36 inflow directions. The wind rose is shown in the upper right corner of the left panel, and the site boundary is shown by a solid red line. Turbine locations are denoted by blue triangles. The flow field in the right panel is normalized by the incoming hub-height velocity of 8 m s^{-1} and shows results for an inflow wind from the west.

boomerang, as shown in Fig. 8. Compared to the initial uniform gridded layout (see Fig. 2), the optimized layout shown in Fig. 8 improves power production by 9.4%. Despite the uneven weighting of the wind rose, the resulting flow field in the right panel of Fig. 8 once again conforms to many of the heuristics outlined in the previous section, including turbines that take advantage of local speedups between upstream turbines and slight flow curvature at the edge of the plant.

The second wind rose considered is given by the directional distribution of 8 m s^{-1} wind speeds from the M2 mast at NWTC (Jager and Andreas, 1996), as shown in Fig. 9. The wind rose is constructed from publicly available data recorded over the 2015 calendar year. For this wind speed, the wind rose has three prominent directions roughly aligned with inflow from the north, south, and west, along with

many other much less prominent inflow directions. The wind rose is again binned into 36 directions (giving $K = 36$) and the optimized layout and resulting flow field are shown in Fig. 9. Compared to the initial regular gridded layout shown in Fig. 2, the power production of the optimized layout is improved by 7.0%. Once again, despite the much more topologically complex nature of the optimized layout shown in Fig. 9, heuristics such as local speed-ups, staggered spacing, and flow curvature are shown to be important in determining the final layout. It is also emphasized that for such a complex wind rose, the resulting layout bears little resemblance to a regular grid where the turbines would be spaced as far from each other as possible.

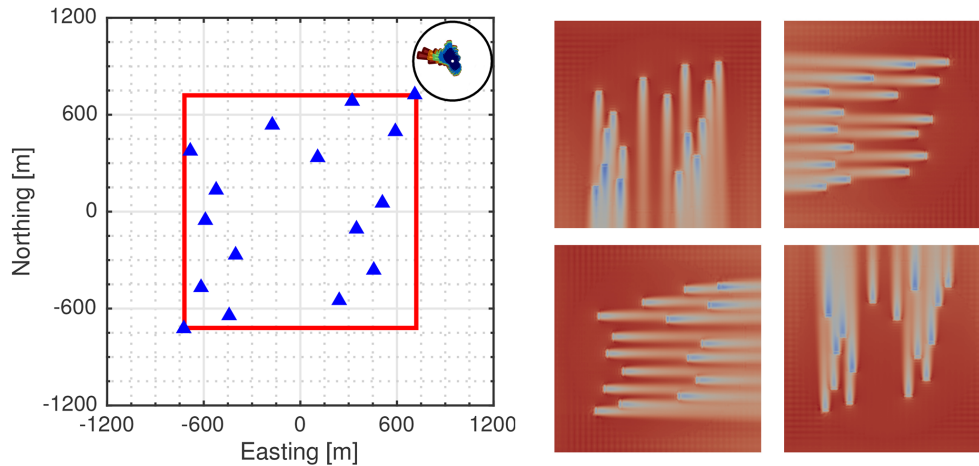


Figure 10. Optimal layout (top left) and flow fields resulting from a full AEP optimization using data from the M2 mast at NWTC. Differences in both wind speed and wind direction are accounted for in the optimization, and the wind rose used is shown in the top right corner of the layout plot. In the top left panel, the site boundary is shown by a solid red line and turbine locations are denoted by blue triangles. Flow fields are shown for inflow winds from the north (top left), east (top right), south (bottom right), and west (bottom left).

4.4 Layout optimization based on annual energy production

As a final demonstration of the power and flexibility of *WindSE*, we optimize a 16-turbine wind plant based on the full AEP from real-world site data. In this case, wind states corresponding to both different wind speeds and wind directions are considered, instead of simply considering a single uniform wind speed as in the tests described in the previous two sections. Data are once again used from the M2 mast at NWTC (Jager and Andreas, 1996), and distributions are formed by binning the data into 36 wind directions and 5 wind speed classes centered on 4, 6, 8, 10, and 12 m s^{-1} , giving a total of $K = 180$ possible wind states in the analysis. Wind states that occurred less than 0.2% of the time were neglected in the optimization, reducing the total number of states considered to 99. We further enforce a minimum inter-turbine spacing D_{\min} of 3 times the rotor diameter.

As shown in Fig. 10, the wind rose for the M2 mast is predominately distributed along the west-northwest direction, with secondary influences from the north and south. The highest wind speeds are also observed for winds from the west-northwest, and so it can be anticipated that a full AEP optimization will result in a layout that is preferentially suited for winds that blow from this direction. This is indeed the case, as shown in the final optimized layout in Fig. 10. The turbines are loosely arranged in two north-south rows that result in relatively large separations between upstream and downstream turbines when the wind is from the west-northwest. As with other tests for evenly and unevenly weighted wind roses with uniform wind speeds, the turbines in the full AEP optimization are staggered with respect to each other in order to take advantage of local speedups between upwind turbines. The resulting optimized layout im-

proves AEP by 8.6% compared to the initial regular gridded layout shown in Fig. 2.

It should be noted that despite the predominant high-speed winds from the west-northwest, the turbines in the north-south rows shown in Fig. 10 do not each fall perfectly on the site boundaries. That is, if only winds from the west-northwest were included in the analysis, one might naïvely place all turbines along the upstream and downstream site boundaries in order to maximize the separation between turbines in the direction of the dominant wind. However, since other wind directions are also included in the analysis, the final optimal layout is more complicated and ensures some degree of turbine staggering in the east-west direction to take advantage of local speedups and minimize wake losses when the wind blows from the north or south. It is also emphasized that simply by accounting for the full AEP in the optimization, the optimal layout in Fig. 10 is substantially different from the optimal layout shown in Fig. 9, where only a single uniform wind speed was considered.

5 Summary and conclusions

WindSE represents a fundamentally new approach to utility-scale wind plant layout optimization that implements adjoint optimization of turbine locations in a flexible CFD framework. The steady-state 3-D RANS flow model advances the fidelity of the fluid dynamics in wind plant optimization simulations, but also requires gradient-based optimization techniques because of the expense of solving the RANS PDE constraint. The adjoint optimization framework in *WindSE* provides these gradients, and further enables efficient, high-dimensional optimization of very large control spaces. This enables layout optimization with a first-

principles flow model without running expensive LES for tuning purposes. The results presented in this paper are achieved at a relatively low computational cost as all optimization results were obtained on a single workstation with a six-core Intel Xeon processor and 32 GB of memory.

The results presented in this paper show that the nonlinear flow effects leading to wake curvature and local speedups are significant when optimizing over a few prominent wind directions. We find consistent rotational symmetry in the optimal layouts with evenly weighted inflow directions, suggesting that evenly weighted wind roses may be useful diagnostic tests for wind plant optimization. As the wind rose is refined into more bin directions, the optimizer is able to take advantage of prominent wind directions and increase energy production. As the number of wind direction and inflow speed combinations increases, WindSE is able to perform a full AEP optimization and achieve sizable gains of almost 9% compared to initial gridded layouts. In the full AEP optimization, the optimizer emphasizes the high speed winds from the west-northwest, since these winds contain the most energy. However, rather than aligning the turbines into rows perpendicular to the incoming wind, the turbines are offset within the general row structure. This is beneficial when the wind is blowing parallel to the row, which is the most common secondary wind direction. It also allows the turbines to take advantage of slight speedup effects around the edges of the wake from upstream turbines. The gradient-based techniques used in this study are inherently constrained to a local search, however, and further research is needed to assess what types of initial layouts are needed to draw conclusions about global optima.

A number of different future studies utilizing WindSE can be imagined. Because the adjoint approach is insensitive to the number of control variables, coupled optimization of turbine locations, hub height, rotor diameter, control settings, etc. can be considered. The flow curvature captured in WindSE is observed to deflect or modify turbine wakes and will likely have important implications in yaw control optimization applications. The high level of abstraction and automation also makes WindSE a useful framework for studying the effects of different turbulence closure models on optimal layouts. We further intend to compare layouts optimized using WindSE and linear flow models by testing both layouts using LES. Additionally, the finite-element method used in WindSE is, in principle, also well suited to examining layout optimization in the presence of terrain-induced complex flows. Finally, the integration of WindSE within NREL systems engineering tools will enable economic analysis and a consideration of LCOE alongside AEP in future optimization studies.

Data availability. Simulations were performed using the WindSE code that will be made available through NREL's WISDEM software tool (Dykes et al., 2014), and can be accessed at <https://nwtc.nrel.gov/WISDEM>.

The NWTC wind data used for layout optimizations in Sect. 4.3 and 4.4 are available at https://www.nrel.gov/midc/nwtc_m2/ (Jager and Andreas, 1996).

Competing interests. The authors declare that they have no conflict of interest.

Acknowledgements. This work was supported by award UGA-0-41026-70 through the Alliance Partner University Program in partnership with the National Renewable Energy Laboratory.

Edited by: J. Peinke

Reviewed by: three anonymous referees

References

- Ainslie, J.: Calculating the flowfield in the wake of wind turbines, *J. Wind Eng. Ind. Aerod.*, 27, 213–224, doi:10.1016/0167-6105(88)90037-2, 1988.
- Amestoy, P. R., Duff, I. S., and L'Excellent, J. Y.: Multifrontal parallel distributed symmetric and unsymmetric solvers, *Comput. Method. Appl. M.*, 184, 501–520, doi:10.1016/S0045-7825(99)00242-X, 2000.
- Barthelmie, R. J., Hansen, K., Frandsen, S. T., Rathmann, O., Schepers, J. G., Schlez, W., Phillips, J., Rados, K., Zervos, A., Politis, E. S., and Chaviaropoulos, P. K.: Modelling and measuring flow and wind turbine wakes in large wind farms offshore, *Wind Energy*, 12, 431–444, doi:10.1002/we.348, 2009.
- Boersma, S., Gebraad, P., Vali, M., Doekemeijer, B., and van Wingerden, J.: A control-oriented dynamic wind farm flow model: “WFSim”, *J. Phys. Conf. Ser.*, 753, 032005, doi:10.1088/1742-6596/753/3/032005, 2016.
- Bokharaie, V. S., Bauweraerts, P., and Meyers, J.: Wind-farm layout optimisation using a hybrid Jensen-LES approach, *Wind Energ. Sci.*, 1, 311–325, doi:10.5194/wes-1-311-2016, 2016.
- Brezzi, F. and Fortin, M. (Eds.): *Mixed and Hybrid Finite Element Methods*, Springer Series in Computational Mathematics, Springer New York, New York, NY, 15, ISBN 978-1-4612-7824-5, 1991.
- Burton, T., Sharpe, D., Jenkins, N., and Bossanyi, E.: *Wind energy handbook*, John Wiley & Sons, Chichester, West Sussex, 2nd Edn., 2011.
- Byrd, R. H., Lu, P., Nocedal, J., and Zhu, C.: A Limited Memory Algorithm for Bound Constrained Optimization, *SIAM J. Sci. Comput.*, 16, 1190–1208, doi:10.1137/0916069, 1995.
- Cabezón, D., Migoya, E., and Crespo, A.: Comparison of turbulence models for the computational fluid dynamics simulation of wind turbine wakes in the atmospheric boundary layer, *Wind Energy*, 14, 909–921, doi:10.1002/we.516, 2011.
- Calaf, M., Meneveau, C., and Meyers, J.: Large eddy simulation study of fully developed wind-turbine array boundary layers, *Phys. Fluids*, 22, 015110-1–015110-16, doi:10.1063/1.3291077, 2010.
- Chen, T. and Liou, L.: Blockage corrections in wind tunnel tests of small horizontal-axis wind turbines, *Exp. Therm. Fluid Sci.*, 35, 565–569, doi:10.1016/j.expthermflusci.2010.12.005, 2011.

- Chowdhury, S., Zhang, J., Messac, A., and Castillo, L.: Unrestricted wind farm layout optimization (UWFLO): Investigating key factors influencing the maximum power generation, *Renew. Energ.*, 38, 16–30, doi:10.1016/j.renene.2011.06.033, 2012.
- Churchfield, M. J., Lee, S., Michalakes, J., and Moriarty, P. J.: A numerical study of the effects of atmospheric and wake turbulence on wind turbine dynamics, *J. Turbul.*, 13, 1–32, doi:10.1080/14685248.2012.668191, 2012.
- Crespo, A., Hernández, J., and Frandsen, S.: Survey of modelling methods for wind turbine wakes and wind farms, *Wind Energy*, 2, 1–24, doi:10.1002/(SICI)1099-1824(199901/03)2:1<1::AID-WE16>3.0.CO;2-7, 1999.
- Donéa, J. and Huerta, A.: *Finite element methods for flow problems*, Wiley, Chichester, Hoboken, NJ, 2003.
- Du Pont, B. L. and Cagan, J.: An Extended Pattern Search Approach to Wind Farm Layout Optimization, *J. Mech. Design*, 134, 081002, doi:10.1115/1.4006997, 2012.
- 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, American Institute of Aeronautics and Astronautics, 1087, 1–25 doi:10.2514/6.2014-1087, 2014.
- El Kasmi, A. and Masson, C.: An extended model for turbulent flow through horizontal-axis wind turbines, *J. Wind Eng. Ind. Aerod.*, 96, 103–122, doi:10.1016/j.jweia.2007.03.007, 2008.
- Farrell, P. E., Ham, D. A., Funke, S. W., and Rognes, M. E.: Automated Derivation of the Adjoint of High-Level Transient Finite Element Programs, *SIAM J. Sci. Comput.*, 35, C369–C393, doi:10.1137/120873558, 2013.
- Fitch, A. C., Olson, J. B., Lundquist, J. K., Dudhia, J., Gupta, A. K., Michalakes, J., and Barstad, I.: Local and Mesoscale Impacts of Wind Farms as Parameterized in a Mesoscale NWP Model, *Mon. Rev.*, 140, 3017–3038, doi:10.1175/MWR-D-11-00352.1, 2012.
- Fleming, P. A., Ning, A., Gebraad, P. M. O., and Dykes, K.: Wind plant system engineering through optimization of layout and yaw control: Wind plant system engineering, *Wind Energy*, 19, 329–344, doi:10.1002/we.1836, 2016.
- Funke, S., Farrell, P., and Piggott, M.: Tidal turbine array optimisation using the adjoint approach, *Renew. Energ.*, 63, 658–673, doi:10.1016/j.renene.2013.09.031, 2014.
- Gebraad, P. M. O., Teeuwisse, F. W., van Wingerden, J. W., Fleming, P. A., Ruben, S. D., Marden, J. R., and Pao, L. Y.: Wind plant power optimization through yaw control using a parametric model for wake effects—a CFD simulation study: Wind plant optimization by yaw control using a parametric wake model, *Wind Energy*, 19, 95–114, doi:10.1002/we.1822, 2016.
- Giannakoglou, K. C. and Papadimitriou, D. I.: Adjoint methods for shape optimization, in: *Optimization and computational fluid dynamics*, Springer, 79–108, http://link.springer.com/chapter/10.1007/978-3-540-72153-6_4, last access: 8 December 2008.
- Giles, M. B. and Pierce, N. A.: An introduction to the adjoint approach to design, *Flow Turbul. Combust.*, 65, 393–415, 2000.
- Goit, J., Munters, W., and Meyers, J.: Optimal Coordinated Control of Power Extraction in LES of a Wind Farm with Entrance Effects, *Energies*, 9, 29, doi:10.3390/en9010029, 2016.
- Goit, J. P. and Meyers, J.: Optimal control of energy extraction in wind-farm boundary layers, *J. Fluid Mech.*, 768, 5–50, doi:10.1017/jfm.2015.70, 2015.
- González, J. S., Gonzalez Rodriguez, A. G., Mora, J. C., Santos, J. R., and Payan, M. B.: Optimization of wind farm turbines layout using an evolutive algorithm, *Renew. Energ.*, 35, 1671–1681, doi:10.1016/j.renene.2010.01.010, 2010.
- Herbert-Acero, J., Probst, O., Réthoré, P.-E., Larsen, G., and Castillo-Villar, K.: A Review of Methodological Approaches for the Design and Optimization of Wind Farms, *Energies*, 7, 6930–7016, doi:10.3390/en7116930, 2014.
- Heywood, J. G., Rannacher, R., and Turek, S.: Artificial boundaries and flux and pressure conditions for the incompressible Navier-Stokes equations, *Int. J. Numer. Meth. Fl.*, 22, 325–352, 1996.
- Hinze, M., Pinnau, R., Ulbrich, M., and Ulbrich, S.: *Optimization with PDE Constraints*, Mathematical Modelling: Theory and Applications, Springer, 2009, 23, 2009.
- Iungo, G. V., Viola, F., Ciri, U., Leonardi, S., and Rotea, M.: Reduced order model for optimization of power production from a wind farm, American Institute of Aeronautics and Astronautics, 2200, 1–9, doi:10.2514/6.2016-2200, 2016.
- Jackson, P. S. and Hunt, J. C. R.: Turbulent wind flow over a low hill, *Q. J. Roy. Meteor. Soc.*, 101, 929–955, doi:10.1002/qj.49710143015, 1975.
- Jager, D. and Andreas, A.: NREL National Wind Technology Center (NWTTC): M2 Tower; Boulder, Colorado (Data), https://www.nrel.gov/midc/nwtc_m2/ (last access: 12 May 2016), 1996.
- Jameson, A.: Aerodynamic shape optimization using the adjoint method, Lectures at the Von Karman Institute, Brussels, <http://aero-comlab.stanford.edu/Papers/jameson.vki03.pdf> (last access: 26 July 2013), 2003.
- Jensen, N. O.: A note on wind generator interaction, Risø-M No. 2411 Risø National Laboratory Roskilde, 1–16, http://orbit.dtu.dk/files/55857682/ris_m_2411.pdf (last access: 28 February 2017), 1983.
- Jimenez, A., Crespo, A., Migoya, E., and Garcia, J.: Advances in large-eddy simulation of a wind turbine wake, *J. Phys. Conf. Ser.*, 75, 012041, doi:10.1088/1742-6596/75/1/012041, 2007.
- King, R., Hamlington, P., Dykes, K., and Graf, P.: Adjoint Optimization of wind Farm Layouts for Systems Engineering Analysis, American Institute of Aeronautics and Astronautics, vol. AIAA 2016-2199, doi:10.2514/6.2016-2199, 2016.
- Kusiak, A. and Song, Z.: Design of wind farm layout for maximum wind energy capture, *Renew. Energ.*, 35, 685–694, doi:10.1016/j.renene.2009.08.019, 2010.
- Kwong, W. Y., Zhang, P. Y., Romero, D., Moran, J., Morgenroth, M., and Amon, C.: Wind farm layout optimization considering energy generation and noise propagation, in: *ASME 2012 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, American Society of Mechanical Engineers, DETC2012-71478, 323–332, doi:10.1115/DETC2012-71478, 2012.
- Larsen, G. C., Aagaard Madsen, H., Troldborg, N., Larsen, T. J., Réthoré, P.-E., Fuglsang, P., Ott, S., Mann, J., Buhl, T., Nielsen, M. et al.: TOPFARM-next generation design tool for optimisation of wind farm topology and operation, Tech. Rep. Risø-R-1805(EN), Risø DTU National Laboratory for Sustainable Energy, <http://forskingsbasen.deff.dk/Share.external?sp=S14b3124a-fef2-4c0b-8f80-6dae7c10ddd2&sp=Sdtu> (last access: 28 February 2017), 2011.
- Logg, A., Mardal, K.-A., and Wells, G. (Eds.): *Automated Solution of Differential Equations by the Finite Element Method*, Lec-

- ture Notes in Computational Science and Engineering, Springer Berlin Heidelberg, Berlin, Heidelberg, 84, <http://link.springer.com/10.1007/978-3-642-23099-8>, 2012.
- Luchini, P. and Bottaro, A.: Adjoint Equations in Stability Analysis, *Annu. Rev. Fluid Mech.*, 46, 493–517, doi:10.1146/annurev-fluid-010313-141253, 2014.
- Marden, J. R., Ruben, S. D., and Pao, L. Y.: A Model-Free Approach to Wind Farm Control Using Game Theoretic Methods, *IEEE T. Contr. Syst. T.*, 21, 1207–1214, doi:10.1109/TCST.2013.2257780, 2013.
- Martins, J. R. R. A. and Hwang, J. T.: Review and Unification of Methods for Computing Derivatives of Multidisciplinary Computational Models, *AIAA J.*, 51, 2582–2599, doi:10.2514/1.J052184, 2013.
- Meyers, J. and Meneveau, C.: Optimal turbine spacing in fully developed wind farm boundary layers, *Wind Energy*, 15, 305–317, doi:10.1002/we.469, 2012.
- Mirocha, J. D., Kosovic, B., Aitken, M. L., and Lundquist, J. K.: Implementation of a generalized actuator disk wind turbine model into the weather research and forecasting model for large-eddy simulation applications, *Journal of Renewable and Sustainable Energy*, 6, 013104, doi:10.1063/1.4861061, 2014.
- Nocedal, J. and Wright, S. J.: Numerical optimization, Springer, New York, ISBN: 978-0-387-30303-1, 1999.
- Porté-Agel, F., Wu, Y.-T., Lu, H., and Conzemius, R. J.: Large-eddy simulation of atmospheric boundary layer flow through wind turbines and wind farms, *J. Wind Eng. Ind. Aerod.*, 99, 154–168, doi:10.1016/j.jweia.2011.01.011, 2011.
- Sanderse, B., Pijl, S., and Koren, B.: Review of computational fluid dynamics for wind turbine wake aerodynamics, *Wind Energy*, 14, 799–819, doi:10.1002/we.458, 2011.
- van der Laan, M. P., Sørensen, N. N., Réthoré, P.-E., Mann, J., Kelly, M. C., Troldborg, N., Hansen, K. S., and Murcia, J. P.: The $k-\epsilon$ model applied to wind farms, *Wind Energy*, 18, 2065–2084, doi:10.100/we.1804, 2015a.
- van der Laan, M. P., Sørensen, N. N., Réthoré, P.-E., Mann, J., Kelly, M. C., Troldborg, N., Schepers, J. G., and Macheaux, E.: An improved $k-\epsilon$ model applied to a wind turbine wake in atmospheric turbulence, *Wind Energy*, 18, 889–907, doi:10.1002/we.1736, 2015b.
- Walmsley, J. L., Taylor, P. A., and Keith, T.: A simple model of neutrally stratified boundary-layer flow over complex terrain with surface roughness modulations (MS3DJH/3R), *Bound.-Lay. Meteorol.*, 36, 157–186, doi:10.1007/BF00117466, 1986.
- Wilcox, D. C.: Turbulence modeling for CFD, DCW Industries, La Cãnada, Calif, 3rd Edn., 2006.
- Wu, Y.-T. and Porté-Agel, F.: Large-Eddy Simulation of Wind-Turbine Wakes: Evaluation of Turbine Parametrisations, *Bound.-Lay. Meteorol.*, 138, 345–366, doi:10.1007/s10546-010-9569-x, 2010.
- Wu, Y.-T. and Porté-Agel, F.: Simulation of Turbulent Flow Inside and Above Wind Farms: Model Validation and Layout Effects, *Bound.-Lay. Meteorol.*, 146, 181–205, doi:10.1007/s10546-012-9757-y, 2013.