6th European Conference on Computational Mechanics (ECCM 6)
7th European Conference on Computational Fluid Dynamics (ECFD 7)
1115 June 2018, Glasgow, UK

MULTILEVEL UNCERTAINTY QUANTIFICATION OF A WIND TURBINE LARGE EDDY SIMULATION MODEL

David C. Maniaci[1], Ari L. Frankel[2], Gianluca Geraci[3], Myra L. Blaylock[4], and Michael S. Eldred[3]

[1] Sandia National Laboratories
PO BOX 5800, MS 1124, Albuquerque, NM 87185
dcmania@sandia.gov

[2] Sandia National Laboratories
7011 East Ave, MS 9159, Livermore, CA 94550
alfrank@sandia.gov

[3] Sandia National Laboratories
PO BOX 5800, MS 1318, Albuquerque, NM 87185
ggeraci@sandia.gov
mseldre@sandia.gov

[4] Sandia National Laboratories
7011 East Ave, MS 9957, Livermore CA
mlblayl@sandia.gov

Key words: Uncertainty Quantification, Multilevel Monte Carlo, Wind Energy, Large Eddy Simulation

Abstract. Wind energy is stochastic in nature; the prediction of aerodynamic quantities and loads relevant to wind energy applications involves modeling the interaction of a range of physics over many scales for many different cases. These predictions require a range of model fidelity, as predictive models that include the interaction of atmospheric and wind turbine wake physics can take weeks to solve on institutional high performance computing systems. In order to quantify the uncertainty in predictions of wind energy quantities with multiple models, researchers at Sandia National Laboratories have applied Multilevel-Multifidelity methods. A demonstration study was completed using simulations of a NREL 5MW rotor in an atmospheric boundary layer with wake interaction. The flow was simulated with two models of disparate fidelity; an actuator line wind plant large-eddy scale model, Nalu, using several mesh resolutions in combination with a lower fidelity model, OpenFAST. Uncertainties in the flow conditions and actuator forces were propagated through the model using Monte Carlo sampling to estimate the velocity defect in the wake and forces on the rotor. Coarse-mesh simulations were leveraged along with the lower-fidelity flow model to reduce the variance of the estimator, and the resulting Multilevel-Multifidelity strategy demonstrated a substantial improvement in estimator efficiency compared to the standard Monte Carlo method.

1 Introduction

Uncertainty Quantification (UQ) is critical in order to enable predictive numerical simulations for scientific discoveries and advanced engineering design, and is starting to be part of wind industry engineering practices. However, in the presence of complex high-fidelity simulations and a large number of uncertainty parameters, the computational cost becomes prohibitive for conventional UQ methods, in both research and design applications. In recent years, multi fidelity UQ has been introduced in order to alleviate this issue and it is based on the aggregation of several lower accuracy models with a handful of higher fidelity computations. One method belonging to this class is the multilevel Monte Carlo method (MLMC) [8, 7]. In the MLMC method the goal is to obtain a statistical estimator based on the aggregation of evaluations of the quantity of interest (QoI) over several mesh resolutions. MLMC takes advantage of the convergence of the deterministic scheme in order to build MC estimators that target the coarsest resolution levels and all the subsequent discrepancies between adjacent resolutions. In addition, generalized control variate approaches can also be used in the presence of different fidelities. The combination of control variate and MLMC lead to Multilevel-Multifidelity estimators (MLMF). To demonstrate the use of MLMF methods in quantifying uncertainty in wind energy applications, we consider simulations of a NREL 5MW rotor in an atmospheric boundary layer.

The flow is modeled with Nalu, a large eddy simulation code [3], which uses actuator line models to represent the blades in the present study. The domain is taken to be periodic to emulate the effects of wake interaction with the rotor. Uncertainties in the flow conditions and actuator forces are propagated through the model using Monte Carlo sampling to estimate the velocity defect in the wake and forces on the rotor. A lower fidelity model was implemented, referred to as OpenFAST, which was used in combination with coarse grid Nalu simulations to reduce the variance of the estimator.

2 Problem Description

In a typical wind farm, multiple wind turbines may be placed in multiple rows in order to maintain a compact arrangement. The result of this arrangement is that the front rows of turbines will create a wake from the oncoming wind that will propagate to the wind turbines in back rows. The velocity deficit in these wakes tends to reduce the output power from those turbines, and hence the net wind farm efficiency may be lower than predicted from the nominal efficiency of a single turbine.

It is possible to use computational models to optimize the arrangement of wind turbines to minimize the loss of efficiency from the wake effects. However, models that do not consider variability in the operating conditions of the wind farm may lead to poor wind farm configurations; an optimally designed wind farm under a single nominal set of operating conditions offers no performance guarantees if the conditions deviate on any given day. Even when the turbine is operating within ideal conditions, the generated power is affected by turbulent fluctuations in the wind, changes in the wind shear and veer, atmospheric conditions such as density and stability, and the wakes from other

turbines. Within the turbine itself, there may be misalignment with the primary wind direction, suboptimal blade pitch angles, and mis-tuned tolerances for the rotor and blade controllers. Finally, computational models for wind farms often rely on low order models that cannot capture all of the physics present in the system, and thus add model form uncertainty to the design of such systems.

Uncertainty quantification (UQ) is necessary for treating variability in applications where significant uncertainties are present. In wind energy applications, UQ provides the ability to quantify the expected power output and associated variance from wind farms in the face of the aforementioned uncertainties in operating conditions. UQ may also be used to inform and optimize wind farm configurations that are robust to these variations and may yield superior performance to configurations designed without taking variability into account.

2.1 Computational Approach

As an example problem, we consider the case of a wind turbine in an atmospheric boundary layer and attempt to evaluate the rotor power, thrust, and root-bending moments in the x- and y-directions. To simulate the wind turbine dynamics, we use the OpenFAST software suite developed at the National Renewable Energy Laboratory (NREL) [5]. OpenFAST enables the analysis of complex physical and environment coupling, including turbine controllers, elastic dynamics, and flow-structure interactions with actuator line theory. We use the NREL 5-MW baseline wind turbine [2] as a prototpyical example in our analysis; this is a single three-bladed turbine with rotor diameter 126m and hub height 90m from the ground. The inflow wind for the aerodynamic analysis of the wind turbine may be simulated in two different ways.

The first method is to leverage the multiphysics, massively parallel simulation code Nalu [3]. Nalu is used to perform large eddy simulations (LES) of the atmospheric boundary layer. Nalu solves the Navier-Stokes equations in the low-Mach number approximation with the one-equation, constant coefficient, turbulent kinetic energy model for the subgrid scale stresses [4]. Following [1], the flow-structure interaction from the turbine onto the wind is simulated by adding a body force of the form

$$f_i = \int_0^L F_i(l) g(\vec{r}(l)) dl \tag{1}$$

where f_i is the body force in the momentum equation, l is the distance along an actuator line, F_i are forces computed from common modules of the OpenFAST code-base, and $g(\vec{x}_l)$ is a smoothing kernel of the form

$$g(\vec{r}) = \frac{1}{\pi^{3/2}\epsilon^3} \exp\left(-\frac{|\vec{r}|^2}{\epsilon^2}\right) \tag{2}$$

where ϵ is a characteristic length scale that may be tuned to spread the body forces out over a larger volume. In the simulations performed in this study, ϵ was fixed at 10m. The

Table 1: Cost estimates for Nalu and OpenFAST simulations.

Case	Mesh size	Simulation time (seconds)	CPUs	Cost (CPU-hours)	Cost (relative)
OpenFAST		500	1	0.42	1
Coarse	100x50x50	2000	80	240	576
Medium	200x100x100	2000	160	960	2304
Fine	400x200x200	2000	400	6860	16500
Reference	800x200x200	2000	400	38400	91400

simulation domain was taken to be $[2km, 1km, 1km]$ in the x, y, and z directions, and the wind turbine was placed in the center of the domain. The boundary layer was initialized with a small perturbation in the velocity profile near the ground in order to accelerate the development of turbulence. The domain was taken to be periodic, allowing the wake from the turbine to interact with itself as an approximation to the effects of having many wind turbines upstream.

The second method is to use the OpenFAST module Turbsim to generate realizations of turbulence-like inflow wind combined with the AeroDYN module for solving the wind turbine rotor flow field. This method is the lower fidelity modeling approach and is termed OpenFAST in this paper, despite the use of some OpenFAST modules in the Nalu code. The Kaimal spectrum was used in conjunction with a power-law boundary layer shape to approximate the turbulent fluctuations near the turbine. This approach is not able to capture the wake interaction effects, and would not be an accurate representation of the turbulence or boundary layer structure.

Although the LES-based actuator line approach of Nalu is expected to be more accurate than the OpenFAST approach, the computational cost associated with solving the filtered Navier-Stokes equations on a large domain is substantially higher. Thus we compare the costs of performing an LES on three separate meshes of increasing size against the cost of using OpenFAST. The OpenFAST simulations were run on a single CPU for a total simulation time of 500 seconds. Due to the wake interaction effects and domain size, the Nalu simulations were given a simulation time of 2000 seconds to allow the turbulence to become stationary.

Table 1 reports the relative simulation costs for running OpenFAST and for running Nalu on four meshes of increasing resolution, referred to here as coarse, medium, fine, and reference. The reference case is considered to be high resolution as it has a mesh size comparable to the blade width of the turbine. The OpenFAST cost includes the amount of time required to generate the inflow turbulence realization, but even with this cost, it is orders of magnitude cheaper to run OpenFAST than even the coarse grid LES in Nalu.

3 Multilevel-Multifidelity Strategy

In this section we briefly describe the multilevel, control variate, and multilevel-multifidelity (MLMF) sampling approaches adopted in this work. The main goal of all these strategies is to leverage the computational efficiency of lower accuracy models in order to decrease

the variance of sampling estimators based on a limited number of high-fidelity evaluations.

The sampling strategies adopted in this work are all based on the Monte Carlo (MC) method. In the MC method the expected value $\mathbb{E}[Q]$ for a generic Quantity of Interest (QoI) $Q : \Xi \to \mathbb{R}$ is approximated as

$$\mathbb{E}[Q] = \int_{\Xi} Q(\boldsymbol{\xi}) p(\boldsymbol{\xi}) \, \mathrm{d}\boldsymbol{\xi} \approx \hat{Q}_N^{\mathrm{MC}} = \frac{1}{N} \sum_{i=1}^{N} Q(\boldsymbol{\xi}^{(i)}) = \frac{1}{N} \sum_{i=1}^{N} Q^{(i)}, \tag{3}$$

where N realizations of the vector of random input $\boldsymbol{\xi} \in \mathbb{R}^d$ are drawn according to the joint probability distribution $p(\boldsymbol{\xi})$. For each realization of the vector of random input variable $\boldsymbol{\xi}$, the value of the QoI $Q^{(i)} = Q(\boldsymbol{\xi}^i)$ is computed by solving the system of PDEs describing the problem of interest, *i.e.* the Navier-Stokes equations for the high-fidelity wind turbine model. As well known, the MC estimator is robust and reliable, being unbiased and not affected by the dimensionality d. However, it exhibits a slow rate of convergence. The root-mean-square-error (RMSE) scales linearly with the standard deviation of Q and inversely proportional to $N^{1/2}$. Therefore, in order to obtain accurate numerical estimation it is usually required to evaluate a large number of realization of $Q(\boldsymbol{\xi})$.

A possible mitigation strategy to reduce the overall computational cost of the MC estimator without sacrificing its accuracy is the so-called multilevel MC (MLMC) approach. For a deep review of the method the interested reader can refer to [8]. Here we only present the basic concept to discuss the numerical investigation conducted in this work. The main idea of MLMC is to replace the QoI Q with a sequence of corrections with respect to a coarser solution (for instance, coarser spatial mesh resolutions) applied to the top of the coarsest solution. If we consider the highest resolution level Q_L we can write $Q_L = Q_0 + (Q_1 - Q_0) + \cdots + (Q_L - Q_{L-1})$ and if we define

$$Y_\ell = \begin{cases} Q_\ell - Q_{\ell-1} & \text{for} \quad \ell > 0 \\ Q_0 & \text{for} \quad \ell = 0, \end{cases} \tag{4}$$

we can write more compactly $Q_L = \sum_{\ell=0}^{L} Y_\ell$. The MLMC estimator is formulated by re-writing $\mathbb{E}[Q_L]$, by exploiting the linearity of the expected value, as the sum of the expected values Y_ℓ and adopting an independent MC estimator for them on each level ℓ

$$\hat{Q}_L^{\mathrm{MLMC}} = \sum_{\ell=0}^{L} \frac{1}{N_\ell} \sum_{i=1}^{N_\ell} Y_\ell^{(i)}. \tag{5}$$

The variance of this estimator is known analytically once the variances $\mathbb{V}ar[Y_\ell]$ are estimated, and thus the sample allocation can be optimized in order to obtain a target variance ε^2 with a minimum total computational cost $\mathcal{C} = \sum_{\ell=0}^{L} \mathcal{C}_\ell N_\ell$, where the cost of each realization of Y_ℓ is noted as $\mathcal{C}_\ell$. The resulting optimal allocation of samples per level is given by:

$$N_\ell = \frac{1}{\varepsilon^2} \sum_{k=0}^{L} \sqrt{\mathbb{V}ar\left(Y_k\right) \mathcal{C}_k} \sqrt{\frac{\mathbb{V}ar\left(Y_\ell\right)}{\mathcal{C}_\ell}}. \tag{6}$$

As it can be seen from the previous equation the number of simulations for each level is proportional to $\sqrt{\mathbb{V}ar\left(Y_\ell\right)/\mathcal{C}_\ell}$. Therefore, for a sequence of levels for which $Y_\ell \to 0$ for $\ell \to \infty$, N_ℓ will decrease as well with ℓ, *i.e.* the computational burden is redistributed toward the less expensive coarser levels.

In this work we constrained our sample allocation with the maximum number of available simulations at the highest resolution, *i.e.* $N_L = N_{target}$. Once N_{target} is known the optimal allocation is built backward by using the sequence of ratios (these are independent from the variance ε^2)

$$\tau_{\ell-1} = \frac{N_\ell}{N_{\ell-1}} = \sqrt{\frac{\mathbb{V}ar\left(Y_\ell\right)}{\mathcal{C}_\ell}} \sqrt{\frac{\mathcal{C}_{\ell-1}}{\mathbb{V}ar\left(Y_{\ell-1}\right)}}, \quad \text{for} \quad \ell = L, \ldots, 1. \tag{7}$$

In the presence of numerical models based on different physics, as in the present case where LES are compared to potential flow computations, there is no guarantee that $Y_\ell \to 0$ for $\ell \to 0$, so it is reasonable to rely on a slightly different approach, namely the control variate (CV). We briefly present a formulation in which the statistical properties of the high- and low-fidelity model (HF and LF respectively) are estimated at the same time as introduced in [10, 9]. Also, for an extension to multiple low-fidelity models, the reader can refer to [12]. In the CV approach we can approximate the expected value of a QoI for the high-fidelity model evaluated at the resolution level L (Q_L^{HF}) by adding an unbiased term based on low-fidelity properties to the standard MC estimator

$$\mathbb{E}\left[Q_L^{\mathrm{HF}}\right] \approx \hat{Q}_{L,N_{\mathrm{HF}}}^{\mathrm{CV,HF}} = \hat{Q}_{L,N_{\mathrm{HF}}}^{\mathrm{MC,HF}} + \alpha \left(\hat{Q}_{L,N_{\mathrm{HF}}}^{\mathrm{MC,LF}} - \hat{Q}_{L,N_{\mathrm{LF}}}^{\mathrm{MC,LF}}\right), \tag{8}$$

where $N_{\mathrm{LF}} = N_{\mathrm{HF}} + \Delta_{\mathrm{LF}} = N_{\mathrm{HF}}(1+r)$ and the additional LF evaluation $\Delta_{\mathrm{LF}} = rN_{\mathrm{HF}}$ are drawn independently from the first set N_{HF}. The values for α and the additional parameter $r > 0$ are obtained by minimizing the overall computational cost under the constraint of the variance of the estimator $\mathbb{V}ar\left(\hat{Q}_{L,N_{\mathrm{HF}}}^{\mathrm{MC,HF}}\right)$ being equal to ε^2. We report here the final result obtained through the optimization analytically from the optimality conditions (see [6, 7] for further details):

$$N_{\mathrm{HF}} = \frac{\mathbb{V}ar\left(Q_L^{\mathrm{HF}}\right)}{\varepsilon^2}\left(1 - \frac{r}{r+1}\rho^2\right) \quad \text{and} \quad r = -1 + \sqrt{\frac{\mathcal{C}_{\mathrm{HF}}}{\mathcal{C}_{\mathrm{LF}}}\frac{\rho^2}{1-\rho^2}}, \tag{9}$$

where ρ is the Pearson's correlation coefficient between HF and LF and $\mathcal{C}_{\mathrm{HF}}$ and $\mathcal{C}_{\mathrm{LF}}$ the computational cost of each HF and LF, respectively.

In this work we do not directly employ the CV approach to build a statistical estimator; instead, we combine it with the MLMC approach in order to obtain the so called MLMF

Table 2: Ranges for uniform distributions of the three uncertain variables considered in this study.

Uncertainty	Minimum	Maxmimum
Speed	6.5 m/s	7.5 m/s
Density	0.97 kg/m^3	1.19 kg/m^3
Yaw	-25^o	25^o

estimator. The MLMF estiamtor is obtained as follows. Initially the MLMC estimaor is obtained as in Eq.(5) and afterward each MC term $\frac{1}{N_\ell}\sum_{i=1}^{N_\ell} Y_\ell^{(i)}$ for the corrections is approximated by means of the CV estimator Eq.(8). The final result is similar to the MLMC allocation

$$N_\ell = \frac{1}{\varepsilon^2}\sum_{k=0}^{L}\sqrt{\frac{\mathbb{V}ar\left(Y_k\right)\mathcal{C}_k}{1-\rho_k^2}}\Lambda_k\sqrt{\frac{\mathbb{V}ar\left(Y_\ell\right)\left(1-\rho^2\right)}{\mathcal{C}_\ell}}, \tag{10}$$

where $\Lambda_\ell = 1-\rho_\ell^2 r_\ell/(1+r_\ell)$ denotes the variance reduction obtained at each level ℓ by the CV estimator with respect to the MC. For this estimator we can also compute the ratios τ_ℓ similarly to Eq.(7) in order to build the optimal allocation constrained to the maximum number of realizations at the highest level of resolution of the HF model.

4 Numerical Results

To demonstrate the effectiveness of the MLMF strategy for variance reduction, we consider the case where the wind speed, air density, and yaw angle of the wind turbine are uncertain. The assumed distributions are uniform with bounds given in Table 2.

100 samples each were taken of the coarse-grid LES and OpenFAST simulations (low and high fidelity paired together for the control variate on Y_0), 20 samples each were taken of the medium- and coarse-grid LES (paired together for Y_1), and 7 samples each were taken of the fine- and medium-grids (paired together for Y_2). Due to the cost of the reference mesh, only one sample was taken to benchmark the performance of the other samples, evaluated at the means of the three input random variables.

An snapshot from a fine mesh sample is shown in Figure 1. The wake resulting from the turbine may be seen from the velocity deficit trailing the center of the domain, as well as the extra turbulence production resulting from the flow-structure interaction. The periodic boundary conditions also cause the wake to re-enter the domain upstream, lowering the wind speed at the turbine location when the wake reaches the turbine. Depending on the initialization, turbine parameters, and turbulent fluctuations, the wake may instantaneously veer off in different directions as well. Changes in the turbine yaw angle also cause the wake to change direction.

The following sections will describe the results of the sampling study for key quantities of interest from the turbine operation. In particular, in Sec. 4.1 we analyze and discuss the output of the simulations performed for all the models and the spatial discretizations.

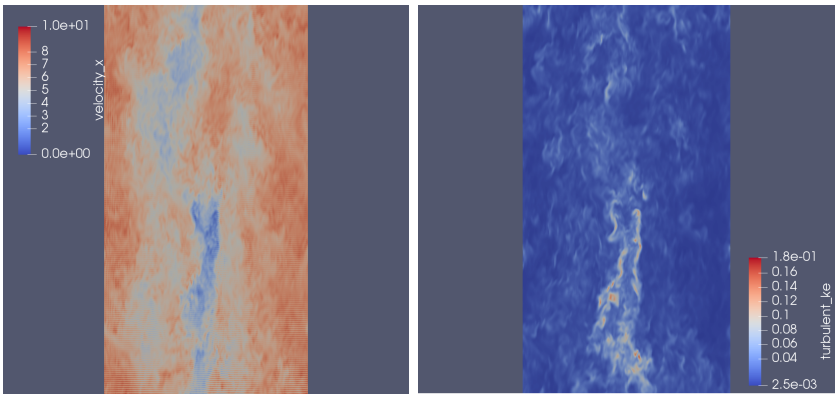

Figure 1: Instantaneous x-direction velocity (left) and turbulent kinetic energy (right) at the z=90m plane for a fine grid sample.

In Sec. 4.2 we present the extrapolation of the performance of all the sampling estimators for the QoIs. For the sake of brevity we present only the results for the power and thrust QoIs. Results for the root bending moments (M_x and M_y) are comparable to Power and Thrust, respectively.

4.1 Pilot Sample Results

In this section we consider the results of the deterministic realizations for both Nalu and OpenFAST. In Figs. 2 and 3 we report the results for power and thrust both in their absolute value for each resolution and their multilevel corrections $Y_\ell = Q_\ell - Q_{\ell-1}$. It is evident from both Figs. 2a and 3a that all the resolutions for Nalu predict values around the nominal value obtained through a reference solution (higher spatial resolution) at the nominal input parameter values. For all the QoI considered, the value computed by OpenFAST is always higher (for power OpenFAST predicts a value almost 4 time larger than Nalu). This is due to the fact that OpenFAST does not resolve the wake induction effect and instead assumes a nominal inflow condition of higher velocity.

When the multilevel corrections are considered (see Figs. 2b and 3b) we observe the expected behavior, *i.e.* the (mean) value of Y_ℓ is decreasing while increasing the level ℓ thanks to the convergence of the Nalu realizations. Also, it is worth noting that the spread of the samples is decreasing while increasing ℓ; therefore, it is reasonable to assume that the variance of Y_ℓ is decreasing as well. This is confirmed from the evaluation of the variance of Y_ℓ for all the QoIs, refer to Table 3.

The evaluation of the performance of the MLMF estimator also requires the estimation of the correlation between Nalu and OpenFAST. In Table 4 we report the correlation between Nalu and OpenFAST for level 0 and 1. Given a higher correlation between the medium mesh results and OpenFast than between the coarse mesh resultsand OpenFast, in the next section we also consider a two-level MLMC estimator and its MLMF counterpart.

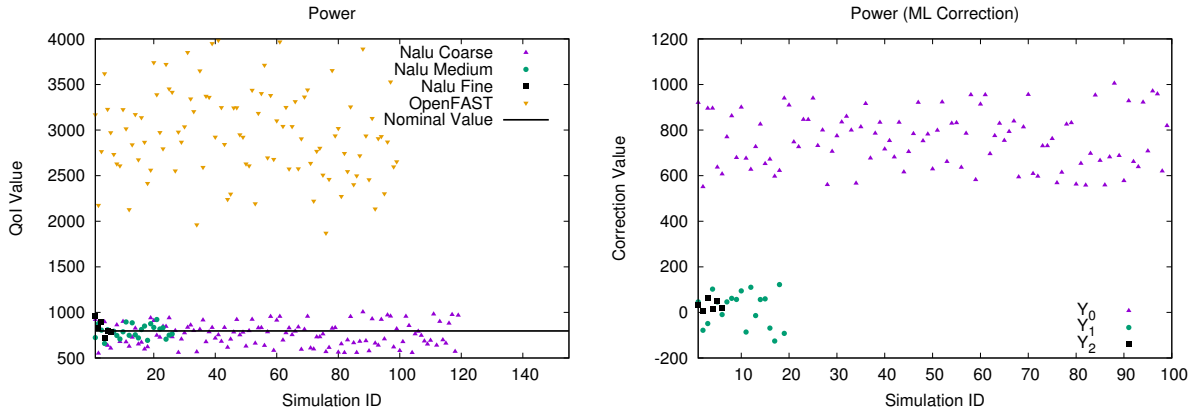

(a) (b)

Figure 2: (a) Nalu and OpenFAST computed values for the power (in kW). The nominal value is computed with Nalu using a finer resolution mesh at nominal conditions for the input parameters. (b) Multilevel corrections ($Y_\ell = Q_\ell - Q_{\ell-1}$) for power.

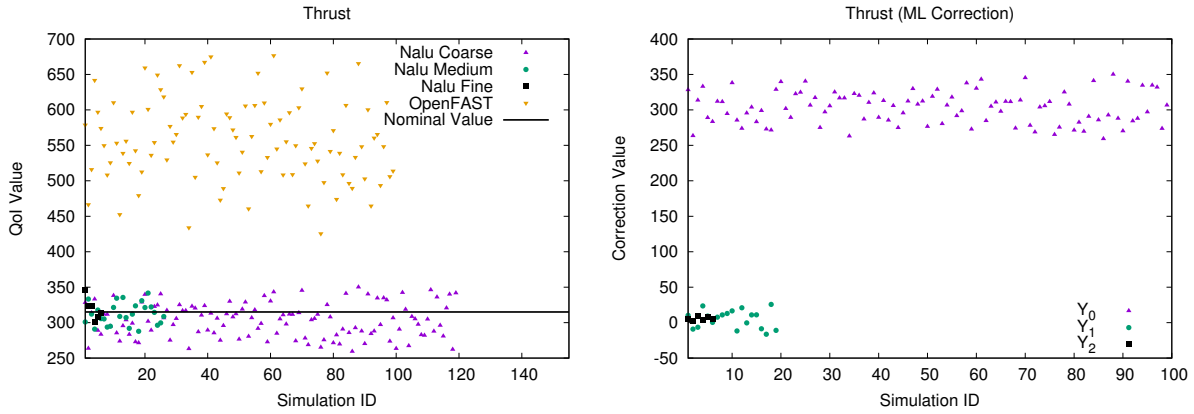

(a) (b)

Figure 3: (a) Nalu and OpenFAST computed values for the thrust (in kN). The nominal value is computed with Nalu using a fine resolution mesh at nominal conditions for the input parameters. (b) Multilevel corrections ($Y_\ell = Q_\ell - Q_{\ell-1}$) for thrust.

Table 3: Variance for the QoI and their ML corrections ($Y_\ell = Q_\ell - Q_{\ell-1}$) per each level ℓ.

	Power		Thrust	
Level	Q_ℓ	Y_ℓ	Q_ℓ	Y_ℓ
0	1.5819e+04		5.5268e+02	
1	5.1535e+03	3.7996e+03	2.6418e+02	9.2135e+01
2	5.7622e+03	506.038	2.2039e+02	13.281

Table 4: Correlation squared between Nalu and OpenFAST.

QoI	ρ_0^2	ρ_1^2
Power	0.27776	0.37157
Thrust	0.43889	0.62452

4.2 Estimator Performance Extrapolation

In this section, given the statistical properties estimated for power and thrust by propagating uncertainties in Nalu and OpenFAST, we are able to extrapolate the behavior of several estimators:

- Standard MC estimator;
- Multilevel $Q_0 + (Q_1 - Q_0) + (Q_2 - Q_1)$, namely MLMC-3l;
- Multilevel $Q_1 + (Q_2 - Q_1)$, namely MLMC-2l;
- MLMF based on MLMC-3l with CV for Q_0, namely MLMF-3l;
- MLMF based on MLMC-2l with CV for Q_1, namely MLMF-2l.

We extrapolate the optimal allocation profile for a higher number of Nalu fine-grid simulations equal to $5, 10, 50, 100, 1000$. For each estimator, once the optimal sample profile has been computed, we also estimated its standard deviation as a measure of its reliability. In Fig. 4 the standard deviation for each estimator is reported as function of the equivalent number of HF realizations needed.

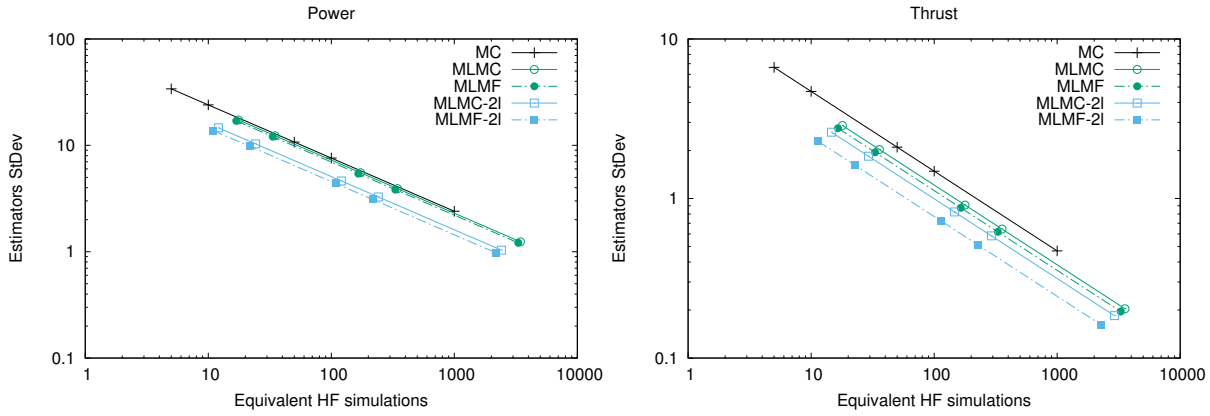

(a) (b)

Figure 4: Extrapolated performance for the MLMC/MLMF estimators for power (a) and thrust (b)

For this problem, the most efficient estimator for both power and thrust is MLMF-2l in which the Nalu coarse-grid simulations are neglected and the CV is introduced at the

Table 5: Samples allocation for MLMC-3l and MLMF-3l.

	Power			Thrust		
	MLMC	MLMF		MLMC	MLMF	
Level	Nalu	Nalu	OpenFAST	Nalu	Nalu	OpenFAST
0	161	137	2040	181	136	2887
1	36	36		34	34	
2	5	5		5	5	

level $\ell = 1$, *i.e.* the medium-grid. Overall, the efficiency of all the MLMC and MLMF estimators is higher for the prediction of the thrust. As an example we report in Table 5 the sample allocation profile for both the QoIs. The sample allocation for MLMF-3l is not altered at levels $\ell = 1, 2$ but only at level $\ell = 0$ where the CV is adopted. The sample allocation is heavily shifted towards OpenFAST and more dramatically for thrust than power given the higher correlation between the numerical codes for this latter QoI.

5 Conclusions

This study served as a demonstration of several estimators for a wind turbine application. The performance of the estimators varied, with the MLMF approach requiring the fewest equivalent samples for a given level of accuracy. These are preliminary results and many improvements are in progress in the utilization of the physics models, OpenFAST and Nalu.

It is anticipated that the uncertainty quantification approach used in this study will be applied to the design optimization of wind plants that have robustness and reliability requirements. Building toward this application, future studies will include problems with a higher number of uncertain parameters, as there are many additional conditions that are of interest in predicting the influence of atmospheric parameters on the performance impacts of wind turbine wakes. To enable such studies, improvements are needed in the coarse mesh resolution accuracy with respect to its relative cost as well as in the lower fidelity model. The primary limitation of the lower fidelity model, OpenFAST, was that it does not model the reduced velocity caused by the wind turbine wake. Reduced fidelity models that include corrections or models for wake effects will be investigated in future work, and are expected to significantly improve the correlation with the higher fidelity model Nalu.

Acknowledgements

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. The views expressed in the article do not nnecessarily represent the views of the U.S. Department of Energy or the United States Government. This work was accomplished through funding from the U.S.

Department of Energy Wind Energy Technology Office.

REFERENCES

[1] Sorensen, J. N. and Shen, W. Z. Numerical modeling of wind turbine wakes. *Journal of Fluids Engineering* (2002) **124(2)**:393-399.

[2] Jonkman, J. and Butterfield, S. and Musial, M. and Scott, G. Definition of a 5-MW reference wind turbine for offshore system development. *Golden: National Renewable Energy Laboratory.* **NREL/TP-500-38060**.

[3] Domino, S. Sierra Low Mach Module: Nalu Theory Manual 1.0. *SAND2015-3107W*, Sandia National Laboratories Unclassified Unlimited Release (UUR). (2015) URL: https://github.com/NaluCFD/NaluDoc

[4] Yoshizawa, A. and Horiuti, K. A statistically-derived subgrid-scale kinetic energy model for the large-eddy simulation of turbulent flows. *J. Phys. Soc. of Japan.* (1985) **54**:2834-2839.

[5] NWTC Information Portal (OpenFAST). URL: https://nwtc.nrel.gov/OpenFAST.

[6] Geraci, G. and Eldred, M. and Iaccarino, G. A multifidelity control variate approach for the multilevel Monte Carlo technique. *CTR Annu. Res. Briefs.* (2015) 169–181

[7] Geraci, G. and Eldred, M. and Iaccarino, G. A multifidelity multilevel Monte Carlo method for uncertainty propagation in aerospace applications. *19th AIAA Non-Deterministic Approaches Conference.* (2017) AIAA 2017-1951

[8] Giles, M. B. Multi-level Monte Carlo path simulation. *Oper. Res.* (2008) **56**, pp. 607–617.

[9] Ng, L.W.T. & Willcox, K.E. Multifidelity approaches for optimization under uncertainty. *Int. J. Numer. Meth. Engng.* (2014) **100**(10), 746–772.

[10] Pasupathy, R., Taaffe, M., Schmeiser, B. W. & Wang, W. Control-variate estimation using estimated control means. *IIE Transactions* (2014) **44**(5), 381–385.

[11] Adams, B.M., Bauman, L.E., Bohnhoff, W.J., Dalbey, K.R., Ebeida, M.S., Eddy, J.P., Eldred, M.S., Geraci, G., Hooper, R.W., Hough, P.D., Hu, K.T., Jakeman, J.D., Maupin, K.A., Monschke, J.A., Rushdi, A., Swiler, L.P., Vigil, D.M., and Wildey, T.M., Dakota, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 6.0 Users Manual *Sandia National Laboratories* **SAND2014-4633.** Updated November 2016 (Version 6.5).

[12] Peherstorfer, B. and Willcox, K. and Gunzburger, M. Survey of multifidelity methods in uncertainty propagation, inference, and optimization. *Massachusetts Institute of Technology* (2016).