Variance Reduction in Wind Farm Layout Optimization

Bertelsen Gagakuma, Andrew P. J. Stanley, and Andrew Ning

Flight, Optimization, and Wind Laboratory, Department of Mechanical Engineering, Brigham Young University, Provo, UT 84602

Abstract

As demand for wind power continues to grow, it is becoming increasingly important to minimize the risk, characterized by the variance, that is associated with long-term power forecasts. This paper investigates variance reduction in power forecasts from wind farm layout optimization. The problem was formulated as a multi-objective optimization. The $\varepsilon-$constraint method is used to solve the bi-objective problem in a two-step optimization framework where two sequential optimizations are performed. The first is maximizing mean wind farm power alone and second, minimizing the power variance with a constraint on the mean power set to the value from the first optimization. The results show that the variance in power estimates can be greatly reduced, by as much as 30%, without sacrificing mean plant power for the different farm sizes and wind conditions studied. This reduction is attributed to the multi-modality of the design space which allows for unique solutions of high mean plant power at different variances. Thus, wind farms can be designed to maximize power capture with greater confidence.

Keywords: Wind farm layout optimization, Variance reduction, Mean wind power, Wind power variance

1. Introduction

Wind power is one of the fastest-growing sources of electricity. This is evidenced by the steady increase in wind-powered electricity generation since 2001. With commercial wind turbines producing almost 4% of world electricity in 2018[1], wind is becoming a viable and sustainable energy source in the global energy landscape. Although there have been many technological advances in recent times that have made wind energy more favorable, it remains an intermittent energy source, marked by high variability in power generation. This variability is mainly attributed to meteorological fluctuations, and can be classified as either short-term, or long-term.

Short-term variability includes minute-to-minute, hour-to-hour, and day-to-day fluctuations in power production which are consequences of turbulence and transient events. Quantifying short-term variability is crucial for grid management and electricity pricing [1, 2]. It can be managed with appropriate load balancing schemes in grid systems that are fed by different electricity sources to meet demand.

Long-term variability on the other hand involves fluctuations over longer periods such as monthly or seasonal time frames, and years. This slower variation in wind power output is caused by climate effects. Such variability is not important for the day-to-day operation of wind farms, but as the name suggests is crucial for long-term project development and maintenance planning [3, 4].

As wind power continues to grow, it will be important to minimize the risk associated with long-term power forecasts. Because wind is a variable resource, wind energy predictions will unavoidably be subject to epistemic uncertainties. It is, therefore, instructive to study the evolution of key statistical quantities such as variance and different percentiles in farm design to help reduce fluctuations in long-term power forecasts. This will allow wind farm operators and economic planners to make more informed decisions [5].

The energy capture capacity of a site can be improved through wind farm layout optimization. This is a critical step in farm design, as wind turbine assemblies (tower, rotor, generator, controls, etc.) are massive

[1] https://www.iea.org/topics/renewables/wind/

systems that typically remain fixed in place for an operational lifetime. Many groups around the world therefore study this problem of optimal turbine positioning by using past wind data for predicting future power production as smart turbine placement favors superior energy capture. The focus in doing this is usually to maximize power [6], annual energy production (AEP) [7, 8], net present value (NPV) [9], profit [10, 11], or to minimize the cost of energy (COE) [12, 13]. These objectives are fundamentally functions of the mean power output, and therefore do not consider variance. The mean power output is not the only important objective; for power companies it is important to minimize variability as well. A low variance in predicted average power translates to less deviation from mean power output forecasts over time, thereby increasing the precision of estimates.

One approach to reducing the variance in wind farm power predictions is to perform a multi-objective optimization using Pareto optimization. The Pareto approach allows for easy comparison of different objectives and has been used in wind farm design. Mytilinou et al. used Pareto analysis to study the trade-off between location, turbine type, and turbine count to suggest optimum regions for siting wind farms around the UK [14]. Their objectives were life cycle cost (LCC), number of turbines, extracted power, and total installed site capacity. Borrisova and Mustakerov also used such methods to test their multi-objective model and algorithm by studying the effects of different combinations of farm location, turbine counts, and wind conditions on cost of energy and power output in wind farm design [15]. This method can also be used to study the trade-offs between mean power and variance as it offers the best possible compromise between conflicting objectives.

The wind farm layout optimization problem is highly dimensional and in a highly multimodal space [16, 17]. Because of the multimodality, we've seen multiple layout designs that produce nearly identical mean wind power ($\approx$3% of each other) but with variances spread over much wider ranges (up to 30% between extremes). We can take advantage of this multimodality by finding solutions with similar mean performance, but with reduced variance. In other words, we generally don't need to trade any mean power performance in order to reduce the corresponding power variance. We've exploited a similar idea in a past work where for some cases noise in a wind farm could be reduced with no penalty in mean power [18].

As it is desirable to obtain solutions with lower risk, the goal of this paper is to show that we can design wind farm layouts to have lower risk associated with mean power output. We present a two stage process for long-term wind plant power variance reduction in which we integrate statistical constraints into the layout optimization framework and test this on a variety of wind farms with different farm boundaries, wind roses, numbers of turbines, and turbine densities, to search for designs with high power production as well as robustness against changes in wind direction. Our approach requires zero sacrifice in mean plant power production for improving robustness.

The rest of this paper is organized as follows: Section 2 contains a summary of the models and optimization framework for this problem, Section 3 presents the results analysis and discussion, and Section 4 gives the concluding remarks.

2. Methodology

The following is an overview of the different tools and techniques employed in this work, with sections discussing wind farm and flow property modeling, as well as relevant statistical quantities and the optimization framework for this problem.

2.1. Wind Farm Modeling

The 3.35 MW onshore wind turbine developed by the International Energy Agency (IEA) [19] is used in this work. It is comprised of a three-blade rotor, which has a swept diameter of 130 m and a 110 m hub height. The power generator is rated at 3.35 MW.

There were several cases considered in this variance reduction study. A case is made up of a boundary type (circle or square), number of turbines (16, 36, or 64), and wind conditions from locations in Victorville, Redding, or the Princess Amalia wind farm.

The circle and square boundary farms have equal areas in each case. Although operational wind farms rarely have circle, square, or other standard plane geometry boundaries as the ones presented here, these

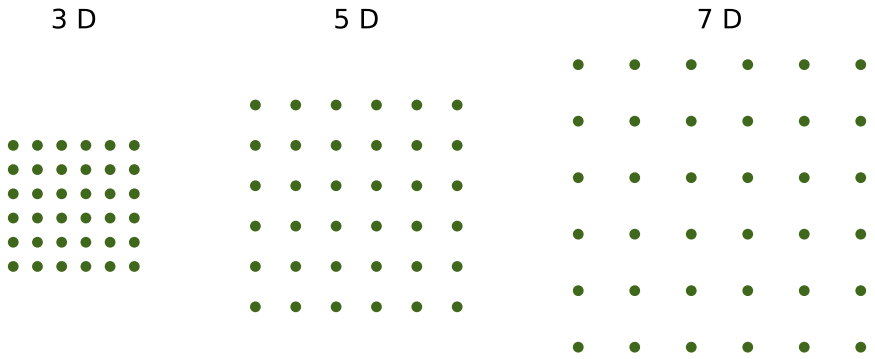

Figure 1: 36-turbine square wind farm at three different turbine densities.

two boundaries are used for demonstrating the intended concepts. With the exception of two cases, the wind farm area is determined by the area required to have the turbines arranged in a perfect square grid with a spacing of five rotor diameters between turbines. The two special cases were performed on 36-turbine circle and square boundary farms with different turbine spacing; cases where the number of turbines is kept constant and the land area varied, based on different initial rotor diameter spacing. Fig. 1 shows this sort of case for a 36-turbine square farm at three different turbine densities which are based on turbine spacings of 3, 5, and 7 rotor diameters.

Other assumptions made in this analysis are flat terrains in all cases and all turbines are of the same size and characteristics. Also, all wind directions in this work are parallel to the horizontal rotational axis of the wind turbines.

2.2. Wake Model

The wake model used in this work is the Flow Redirection and Induction in Steady-state (FLORIS) engineering wake model [20]. This is a derivative of the Jensen wake model developed by N.O. Jensen [21] which assumes a linearly expanding wake with a velocity deficit that is only dependent on the distance behind a rotor. The FLORIS model defines three separate wake zones with differing expansion and decay rates to more accurately describe the velocity deficit across wind turbine rotors in wake regions. A major advantage of this model over Jensen's is its adaptability for handling of partial wakes and wake mixing which it achieves because of the three-zone divisions. In this work, hub velocity losses due to waking were determined from the square root of the sum of squares of contributing losses. These contributing losses were determined using area weighted averages of the affected wake zones. The FLORIS wake model uses different parameters for characterizing wake deflection, expansion, and velocity which are found in Table 1 of [20].

2.3. Flow Properties

Existing wind direction, speed, and frequency data was used to create a probability density function for wind direction, and a distribution for wind speed as a function of wind direction for each of the three chosen locations (Fig. 2). The locations are Victorville and Redding in California, and the setting of the Princess Amalia wind farm off the coast of the Netherlands. Since the power produced in a wind farm is largely dependent on wind conditions, the three wind roses chosen were dissimilar in order to measure the effectiveness of the variance reduction approach presented here under diverse conditions. The Amalia wind rose has multiple high probability directions with a wind speed profile closely matching the wind rose. Even though the Victorville and Redding wind roses each have two notably dominant wind directions, they

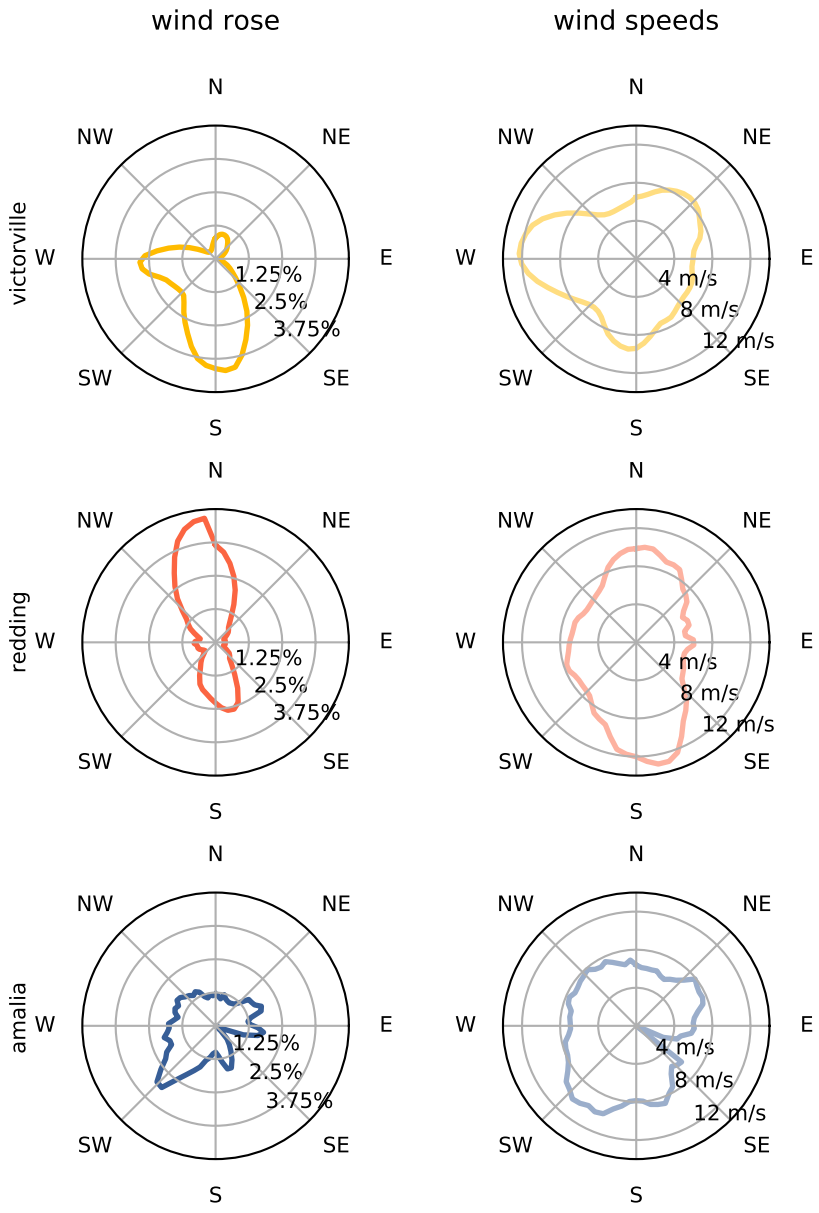

Figure 2: Victorville, Redding, and Amalia, wind roses and directionally averaged wind speeds. The wind roses show the general source wind directions in relation to the frequency with which wind blows from a given direction. The directionally averaged wind speeds represent the magnitude of wind the speed coming from a given direction.

are characterized by different orientations and wind speed profiles. These direction characteristics affect multi-modality in the design space in different ways.

The three different wind roses (Fig. 2) were chosen to study optimization outcomes when this variance reduction concept is applied for different wind conditions. In order to draw equal sample sizes of wind directions and wind speeds over smooth distributions for numerical integration, data from each of these locations was splined using Akima splines [22] to obtain functional forms of wind frequency and speed, both of which depend on wind direction such that $f_{frequency} = f(direction)$, and $f_{speed} = f(direction)$. The resulting wind frequency function was then normalized to give the probability density function of wind frequency. In this work, 70 wind directions were selected between 0 and 360 degrees based on their probability and the corresponding directionally averaged wind speeds were determined from the wind speed function. These speeds represent the magnitude of the wind speed coming from a given direction as seen in Fig. 2. Thus we only consider wind direction uncertainty in this analysis.

2.4. Statistics of Interest

Mean plant power is a major statistic of interest as it is used in determining quantities like annual energy production (AEP), net present value (NPV), cost of energy (COE), etc. It is a quantity that depends on the total farm power. The mean, although very useful, is only a measure of some central value and therefore does not paint a complete picture of the power producing capacity of a site. Together with the variance, the mean power can be interpreted to give more meaningful prediction of future power production since variance shows the spread from mean quantities. The variance is thus a measure of risk in stochastic systems. The mean plant power is determined from the power contributions, P_i of all turbines in a farm.

$$P = \sum_{i=1}^{n_{turbs}} P_i \tag{1}$$

$$P_i = \frac{1}{2}\rho A C_p U_i^3 \tag{2}$$

The expression for the power from each turbine, P_i is given by Eq. 2, where ρ is the air density and A is the rotor swept area. The power coefficient, C_P, is the power conversion efficiency of a wind turbine for which a value of 0.458 was used. The effective hub velocity for each turbine is U_i. It is a function of position (x and y) as well as flow parameters and variables (wind direction and speed β and U_∞ respectively). The values of U are determined with the FLORIS wake model mentioned above. Because we only consider wind direction uncertainty in this work, the mean plant power can be expressed as an integral for a random variable (Eq. 3).

$$\mu_P = E[P] = \int_0^{2\pi} P(x, y, \beta, U_\infty)\ \phi(\beta)\ d\beta \approx \sum_{j=0}^{m} P(x, y, \beta_j, U_{\infty,j})\ \phi(\beta_j) \tag{3}$$

In this equation, $\phi(\beta)$ is the probability density function of the wind direction and m is the number of sampled wind directions. The mean power is simply expressed as an integral for determining the expected value of a random variable since the density function of wind direction is known. This integral can be evaluated as that for a discrete random variable using the rectangle rule (far right hand side of Eq. 3) since samples are drawn from the wind distributions in this analysis. Each wind direction β_j has a corresponding directionally averaged free stream wind speed $U_{\infty,j}$ which is determined form the wind speed distribution. The uncertain variable of interest here is the wind direction β, making this a 1-dimensional uncertainty quantification problem.

The variance is evaluated similarly from the integral given in Eq. 4.

$$\sigma_P^2 = \int_0^{2\pi} P^2(x, y, \beta, U_\infty)\ \phi(\beta)\ d\beta - \mu_P^2 \approx \left(\sum_{j=0}^{m} (P^2(x, y, \beta_j, U_{\infty,j})\phi(\beta_j) \right) - \mu_P^2 \tag{4}$$

2.5. Optimization Framework

Multi-objective optimization (MOO) problems are those that involve more than one objective function to be minimized. MOO problems are typically expressed mathematically as

$$\begin{aligned} &\min\ (f_1(x), f_2(x), ..., f_k(x)), \qquad k \geq 2 \\ &\text{s.t.}\ \ x \in \mathbf{X} \end{aligned} \tag{5}$$

where k is the number of objective functions and $\mathbf{X}$ is the set of feasible decision vectors and/or constraint set. Such problems have multiple solutions that quantify the best tradeoff between competing objectives as opposed to global optima found in single-objective optimization problems. Because of the existence of a set of solutions, the concept of dominance is introduced to determine if one solution is better than another. For two solutions x_1 and x_2, x_1 is said to dominate x_2 if x_1 is better than x_2 in all objectives.

Some standard approaches to solving MOO problems include: normalized objective functions, the weighted sum method, $\varepsilon-$constraint method, lexicographic method, multi-objective evolutionary algorithms, and the normal boundary intersection method. Further details about the strengths and weaknesses of these methods can be found in [23, 24]. In this work, however, the $\varepsilon-$constraint method is used due to its simplicity, and effectiveness in both convex and non-convex design spaces.

2.6. Variance Reduction Concept

The approach presented below is a simple yet very rewarding variance reduction concept for power forecasting through wind farm layout optimization. It is a form of the ε-constraint method of multi-objective optimization which involves minimizing just one objective with inequality constraints which are formulated from the other objective functions.

In this work, the main objective to be minimized is the variance. The optimization is carried out in two steps. Step 1 is the common practice of maximizing mean plant power without any constraints set on the variance. This first optimization gives a suitably high mean power value with a corresponding variance that is usually high. Step 2 then aims at reducing the variance associated with the mean power from step 1. This is done by performing a second optimization where the variance is minimized with an inequality constraint set on the mean power to be greater than or equal to the value from step 1. Also, the starting layout in this second optimization is the solution from the step 1 problem. The combined optimization problem is formulated below as:

step 1: Maximizing mean power

$$\begin{aligned} &\text{maximize} && \mu_P^* \\ &\text{w.r.t} && x_i, y_j \quad i,j = 0,1,2,...,n_{turbs} \\ &\text{subject to} && S_{i,j} \geq 2D_{turbine}, \quad i,j = 0,1,2,...,n_{turbs} \\ & && \textit{boundary constraints} \end{aligned} \tag{6}$$

step 2: Variance reduction

$$\begin{aligned} &\text{minimize} && \sigma_P^2 \\ &\text{w.r.t} && x_i, y_j \quad i,j = 0,1,2,...,n_{turbs} \\ &\text{subject to} && \mu_P \geq \mu_P^* \\ & && S_{i,j} \geq 2D_{turbine}, \quad i,j = 0,1,2,...,n_{turbs} \\ & && \textit{boundary constraints} \end{aligned} \tag{7}$$

By doing this, we are able to obtain better solutions, i.e., layouts that produce effectively the same high wind power but with lower variance than in the case of performing a single mean power optimization. The degree of improvement seen after the variance reduction step varies with each case. Naturally, some combinations of wind conditions, farm boundary, and size show greater improvement over others.

The Python Optmization Sparse framework (pyOptSparse) [25] is used with Sparse Nonlinear Optimizer (SNOPT), a gradient-based Fortran software package for solving large-scale nonlinear optimization problems [26], to perform all the optimizations in this work.

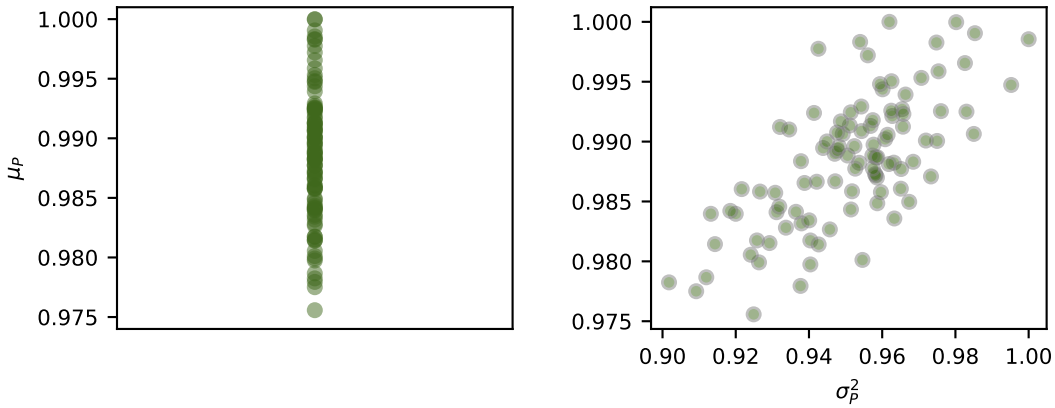

Figure 3: On the left, normalized mean power results from 100 layout optimizations for a 36−turbine, circular boundary farm using the Amalia wind conditions. On the right, mean power and corresponding variance from the same 100 optimizations (each normalized by the highest mean and variance respectively from the step 1 optimizations, where we maximized the mean power only with no constraints on the variance). The difference between maximum and minimum variance values is noticeably larger than that between mean power values.

3. Results and Discussion

In this section, we show results for variance reduction in wind farm layout optimization using the proposed variance minimization concept. Simulations were carried out for different combinations of farm sizes, boundaries, and wind roses. We also present results for variance reduction using full Pareto optimizations.

3.1. Variance Reduction

As discussed, wind farm layout optimization typically seeks to optimize a metric related to the mean or expected power, like the annual energy production or cost of energy. However, because the design space is multimodal it is good practice to use a multistart procedure, where many random starting points are used and the best result is then taken. In this work we used 100 random starting points in each case to first optimize wind farm layouts for maximum mean power. Those multiple starting locations create a range of values as shown in Fig. 3 for a 36-turbine wind farm with Amalia wind conditions in a circular boundary. The key insight of this work is that those same solutions, which maximize mean power have very different variance. This can be seen in Fig. 3 where the percentage difference between the normalized[2] maximum and minimum mean power is about 2.5%, and that between the variance maximum and minimum is about 10%.

We can exploit this behavior even further by minimizing the variance while constraining the mean power. In Fig. 4 we see that the mean power remains the same, whiles the variance changes after performing a second optimization with constraints set on the mean. The variance values from the μ_P maximization step are noticeably larger than those from the variance reduction step. This is more apparent in the right window where gray arrows show the shift in variance between the first and second optimizations for 10 instances.

This trend of lowered variance after performing two sequential optimizations is observed in all the cases we considered. Table 1 contains the minimum, average, and maximum relative[3] reduction in variance for 18 cases which consist of 16-, 36-, and 64-turbines in both circle and square farm boundaries with the wind conditions from the three locations (Victorville, Redding, and Amalia wind farm). The land area of each wind farm is based on the area of a square grid arrangement of wind turbines with an approximate initial spacing of 5 rotor diameter spans. Thus both circle and square boundaries with the same number of turbines have the same area.

[2]Both mean power and variance are normalized by the highest mean and variance respectively from the step 1 optimizations
[3]The reduction in variance is measured relative to that maximum mean power optimization variance.

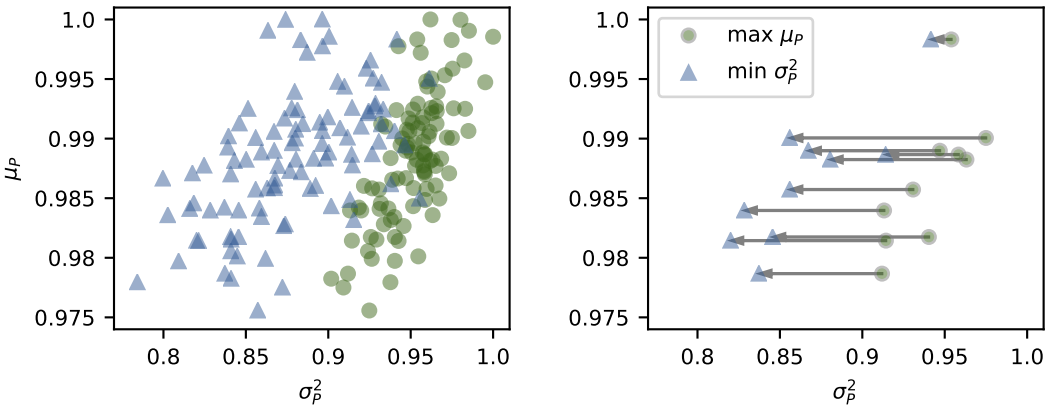

Figure 4: In each subfigure, normalized mean power and variance from the two optimization steps of maximizing mean power, and reducing variance. The second optimization variance results are generally lower than those from the first while the mean power stays the same in both optimizations. On the right, this shift is clearly shown for 10 pairs of optimizations.

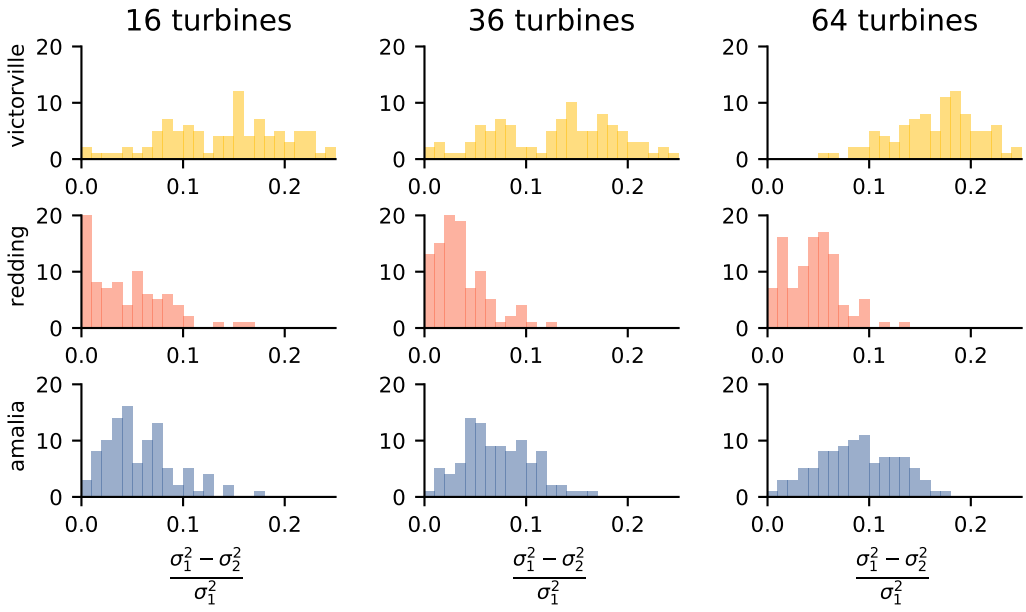

Figure 5: Relative variance decrease for circle farm boundary cases.

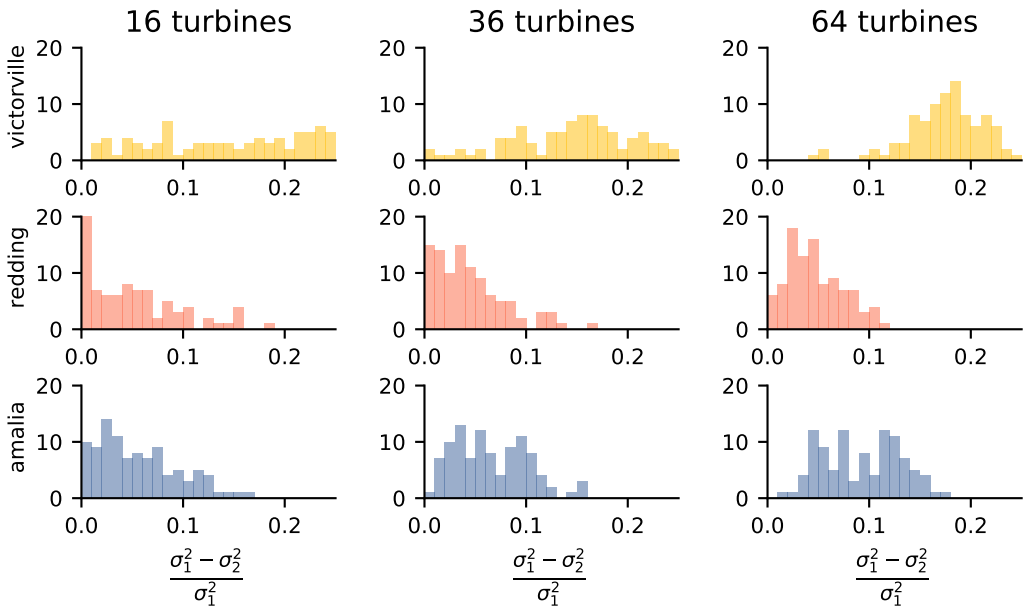

Figure 6: Relative variance decrease for square farm boundary cases.

There is nonzero variance decrease in all our analyzed cases although some reductions are smaller than others. Also, a wider range of variance reduction is observed in some cases over others. In Fig. 5, we observe that the Victorville and Amalia wind roses yield wider ranges of relative variance reduction compared to the bidirectional Redding cases. The Victorville farms generally produce the highest variance decreases of up to 20%. The Redding cases produce lower variance reductions in comparison to the other two wind roses. This is because multidirecional wind roses further increase the occurrence of local minima in the design space. Thus, allowing for more distinct solutions that yield high mean power at different variances.

The square cases (Fig. 6) show similar trends as the circle ones with the Victorville and Amalia farms again favoring variance reduction over the Redding farms.

The wind farm layout optimization process creates a pattern for arranging turbines in within a boundary that promotes high power output. Fig. 7 shows optimized layouts for 36-turbine circle and square farm boundaries respectively. The mean power stays the same for the two optimization steps as previously mentioned, whereas the variances[4] after the second optimization are noticeably less than those in first. Each layout solution is indicative of the shape of the wind rose used.

[4]Presented here as the standard deviation for consistency of units.

boundary	wind rose		16	36	64
		min	0.520	0.230	5.53
	Victorville	avg	14.6	12.7	17.2
		max	29.9	24.8	27.9
			0.00340	0.0244	0.377
circle	Redding		3.82	3.49	4.56
			16.7	12.2	13.6
			0.517	0.949	0.434
	Amalia		5.74	7.14	8.81
			17.4	16.1	17.6
		min	1.23	0.508	4.54
	Victorville	avg	16.8	15.0	17.8
		max	33.3	30.3	27.6
			0.00660	0.0270	0.141
square	Redding		4.31	4.46	4.69
			18.2	16.1	11.1
			0.00830	0.533	1.46
	Amalia		5.61	6.63	9.30
			16.2	15.8	17.6

Table 1: Minimum, average, and maximum % change in plant power variance for 16−, 36−, and 64-turbine circle and square boundary wind farms and all three wind roses from 100 individual layout optimizations using the proposed two step variance reduction approach.

In the Victorville and Amalia wind roses, we observe that turbines are more evenly spread within the farm boundaries. This is a consequence of their wind direction and wind speed profiles. The diagonally symmetric bidirectional Redding wind rose has turbines at the boundary edges, and in the middle of the farm as expected. Turbines are arranged within farm boundaries such that rotor-on-rotor wake interactions are minimized in order to promote energy capture.

Lastly, the square farms generally yield higher mean power than their circle counterparts for each of the wind conditions even though the land areas are equal. This is because the angled corners in the square boundaries offer better turbine separation, which in turn reduces the occurrence of wakes.

This is the main advantage of applying this concept to the wind farm layout optimization problem–there is no reduction in mean power for a corresponding decrease in variance. The approach yields positive results because the wind farm layout optimization design space is highly multi-modal. Hence, although multi-modality makes optimization problems more difficult, it can be exploited to minimize variance in long-term wind power forecasting.

3.2. Variance Reduction with Different Turbine Densities

We also studied a few cases using the proposed concept with different turbine densities, meaning the same number of turbines are confined to different land areas. The presented optimization results are for the Amalia wind rose, and 36-turbine farms. Wind farms with six different turbine densities were studied, with average turbine spacings of 3–8 rotor diameters[5].

[5]Fig. 1 shows this sort of arrangement for three square boundary farms which are spaced at 3, 5, and 7-rotor diameters.

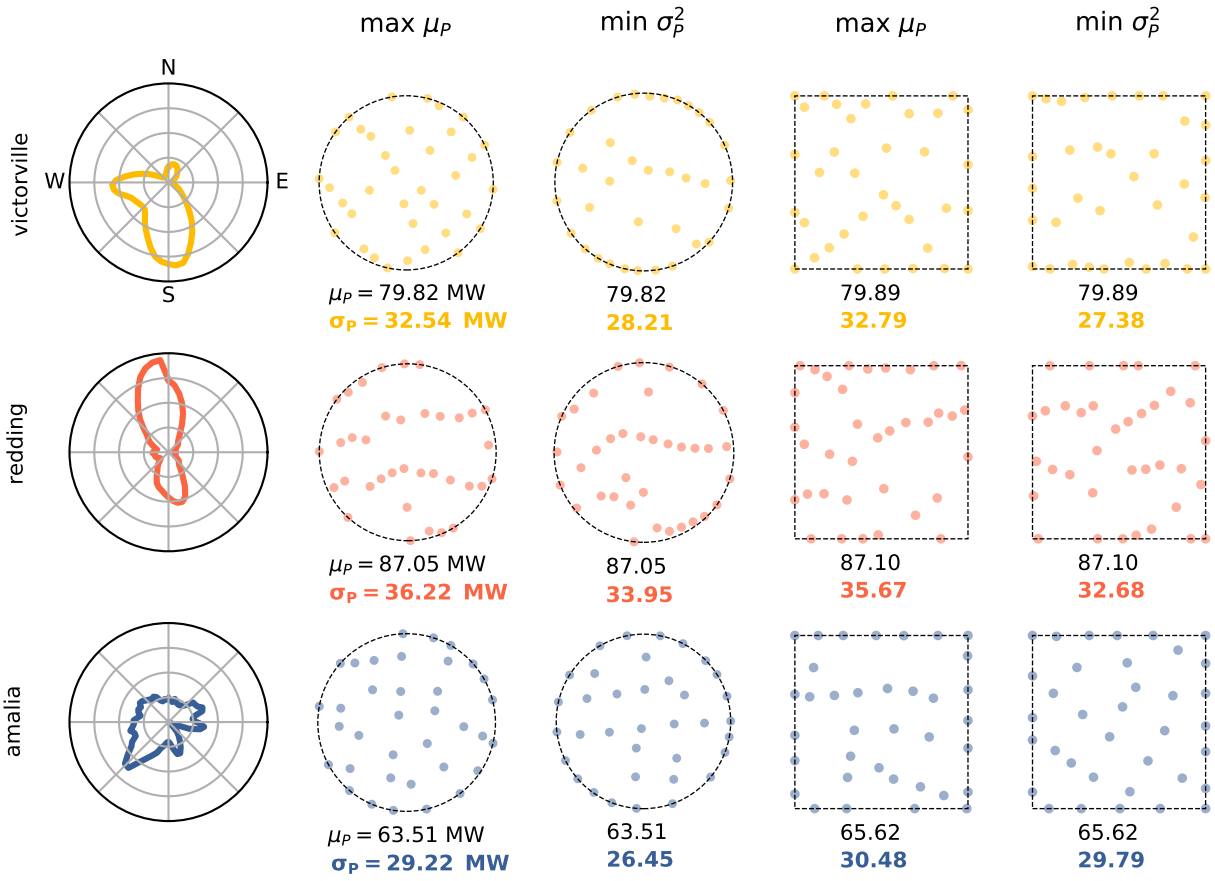

Figure 7: Layout optimization solutions for 36-turbine circle and square farms using the Victorville, Redding, and Amalia wind conditions. Step 1 and 2 solutions are shown with their corresponding mean power and standard deviation.

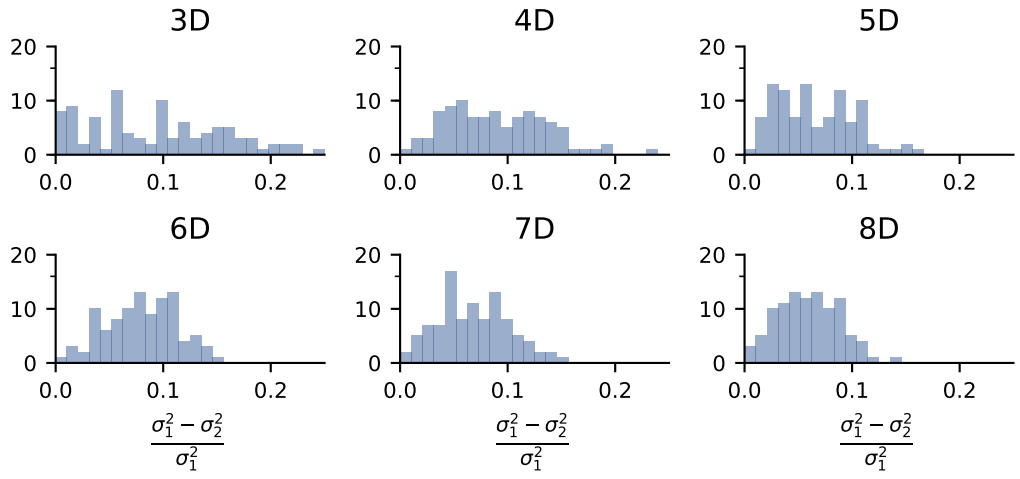

Figure 8: Relative variance decrease for different turbine densities based on 3–8 rotor diameter spacings for circle Amalia farms.

The relative differences in variance between the two optimization steps for the different turbine densities for a circle boundary farm are presented in Fig. 8. The denser (5D spacing and below) farms yield a wider range of relative variance decrease over the sparser ones (6D spacing and above). These denser farms also show higher average variance decreases, with the largest being approximately 9% for the farms with 3-rotor diameter average spacing. The lowest average variance decrease is just over 5% as seen in the farms with 8-rotor diameter average spacing. Fig. 9 shows the variance reduction results for varied turbine spacing in square wind farms. As with the circular farms, the average variance reduction is inversely proportional to turbine density. However, even for the farms with larger average turbine spacing, a large variance reduction was achieved.

The denser wind farms achieve larger variance reduction than the sparser farms. This is because of the strong wake effects that exist in denser wind farms. When the turbines are close together, the turbine wakes have not had a chance to recover before reaching downstream turbines. Thus, as the wind changes direction the power production changes more drastically, as turbines move in and out of the strong wakes of upstream turbines. In wind farms where the turbines are spaced farther apart, wakes have more time to recover before interacting with downstream turbines. This means that the power production is not affected as much by turbine wakes.

3.3. Comparison to Pareto Optimization

The previous sections demonstrated that after optimizing a wind farm layout for maximum power production, the variance can be significantly reduced by re-optimizing the layout for minimum variance. By constraining the mean power, this variance reduction can be achieved without any sacrifice to mean performance. This process is very simple, and can be done using existing wind farm models and optimization frameworks for very little additional effort. However, if one has the computational resources and is willing, further insight can be gained by performing a full multi-objective optimization and finding the Pareto front of the design space.

Pareto optimal results offer the best possible compromise between two objectives. A Pareto front is computationally expensive to generate, requiring thousands of individual optimizations. The black points in Fig. 10 represent the approximate Pareto front for a 36-turbine wind farm with square boundaries, five rotor diameter average turbine spacing, and using the Amalia wind data. These data are close to the lowest variance solution that can be found with certain constraints on the mean power. They were generated by constraining the mean wind farm power to different values, and minimizing the associated power variance. As can be seen, a point on the Pareto front cannot achieve lower variance for the same mean power. The

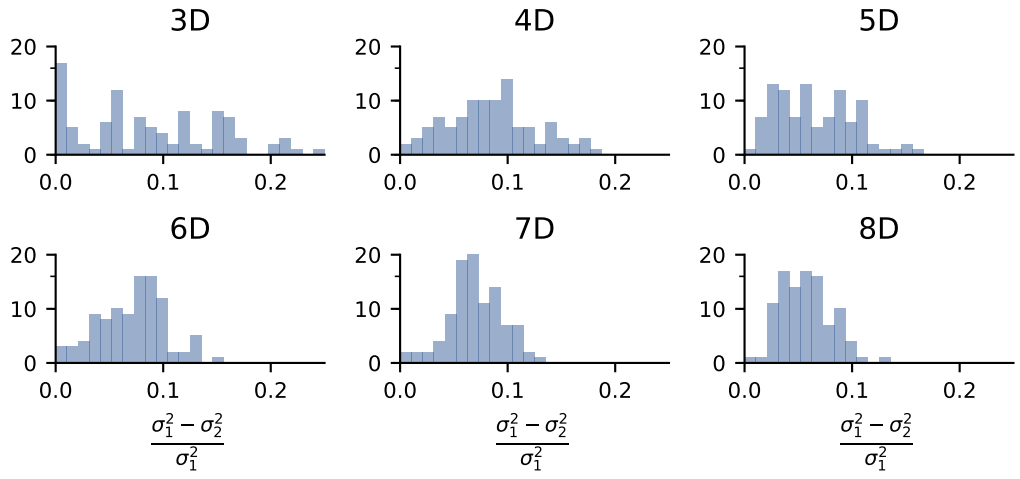

Figure 9: Relative variance decrease for different turbine densities based on 3–8 rotor diameter spacings for square Amalia farms.

objectives of maximum mean and minimum variance conflict, meaning that to improve one objective you must sacrifice some of the other. For this wind farm, and the others that we considered during this study, the Pareto front is relatively flat near the upper right hand corner. This is significant, as it means that if one is willing to sacrifice a small amount of mean power production they are able to greatly reduce the variance. For the case shown in Fig. 10, a 2% sacrifice in the mean power can result in more than 25% decrease in variance. Much larger reductions in variance can be achieved, but they come at a higher cost in mean power.

For wind farms, the region of the Pareto front of most interest is the upper right corner, the solutions with high mean power. The right sub-figure in Fig. 10 shows a zoomed in portion of this upper right corner. When performing an optimization to maximize the mean power of the wind farm (step 1 of the method presented in this paper), the solution is somewhere near this corner of the Pareto front. A few of these solutions are shown in Fig. 10 with the green dots. If one could easily find a point on the Pareto front then the two-step optimization process presented in this paper would be unnecessary, and reductions in variance would require sacrifices in mean power. However, because of the complexity of the design space, generally hundreds of optimizations are required to find just one point on the Pareto front. When step 2 of our method is applied, new solutions are found with lower variance. The solutions from step 1 are pushed closer to the Pareto optimal, for very little added computation. Rather than requiring tens or hundreds of optimizations to approximate a point on the Pareto front, we can get close with just two. Because this upper right corner of the Pareto front is relatively flat, our method can produce large variance reduction, even for the higher mean power solutions.

In the context of full multi-objective optimization, the method we present in this paper is extremely useful. Rather than seek to fully explore the design space, and trade-offs that exist between the mean and variance, one can simply apply our method. This will produce a wind farm layout with high mean power, as well as variance that approaches the Pareto optimal solution, with greatly reduced computational expense.

4. Conclusion

This work explored a simple, yet efficient approach to reducing variance in wind farm layout optimization. Various cases were evaluated that demonstrate how effective the approach is for different farm boundaries, turbine counts, and wind conditions. The approach led to solutions with lower variance and at no expense to the mean plant power.

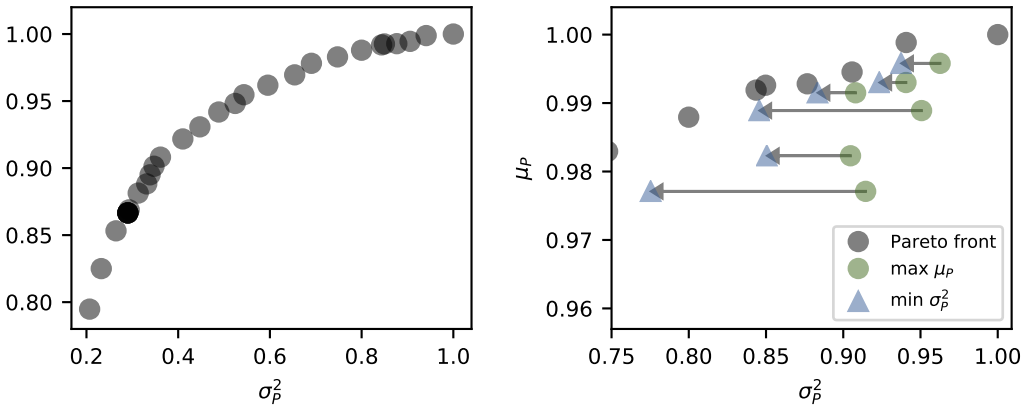

Figure 10: On the left, the Pareto front for the multi-objective wind farm layout optimization problem of high mean power versus low variance on 36-turbine square boundary farm using the Amalia wind conditions. On the right, a zoomed in view on the upper right corner of the full Pareto front. Also shown are six pairs of points from mean power and variance optimizations.

The mean plant power is a major statistic of interest during the design and development stages of wind farms as it is used in determining quantities of interest such as AEP, COE, NPV, etc. In general, high values of mean power are sought in layout optimization for power forecasting purposes. Turbine placement is a significant step in wind farm design since intelligently situating turbines can improve the energy capture capacity of a site. Because mean power does not provide information about the potential power fluctuations in a wind farm, it is instructive to study the associated variance (standard deviation) to get a sense of the risk associated with this quantity in power forecasting. A low variance will translate to better stability of forecasts and thus variance should be lowered where possible.

The two step approach for reducing variance presented in this paper uses constraints on the mean power to minimize the variance during wind farm layout optimization. It involves two optimizations where the first is a layout optimization which only considers maximizing mean plant power. The second step is to re-optimize the wind farm layout to minimize variance. In this step, the mean power is constrained to be greater than or equal to the mean power achieved in the first step. This means that a wind farm layout with both high mean power and reduced power variance can be found. The success of our presented two-step process is enabled by the multi-modality of the wind farm layout design space. Because the wind farm layout optimization design space is highly multi-modal, many solutions can be found that yield comparable power outputs but very different power variances. Our method finds the solutions that correspond to the lower variances.

The results for all the different cases studied show that this technique is effective at reducing the wind farm power variance while maintaining the initial high mean power. In every case that we considered, there was a noticeable average decrease between step one and step two of our method. We also observe from this analysis that variance reduction in wind farm layout optimization improves for regions where the wind conditions are characterized by high variability. This is observed in the results from the Victorville and Amalia cases that showed higher average variance reduction. For the 64-turbine wind farms, the Vitorville wind farms achieved over 17% average variance reduction, and the Amalia wind farms over 8% average variance reduction. The Redding wind farms experienced a lower average variance reduction, of about 4.5% for both circle and square farm boundaries.

For the cases with different turbine densities, there was an overall decrease in variance after performing the two sequential optimizations. The denser wind farms experienced greater power variance reductions between steps one and two of our method. For both wind farm boundaries, the average variance decrease in the 3D farm was highest, about 9%. The wind farms with turbines spaced farther apart still noticeably reduced the variance between the two steps, although by a smaller percentage. The largest turbine spacing

of 8 rotor diameters achieved an average variance reduction of about 5%. This indicates that our method will therefore be more beneficial to situations where land is scarce and turbines would need to be built close together. However, even wind farms with turbines spaced far apart can benefit greatly from our presented methodology.

This concept is meant to be used to improve the accuracy of long-term power forecasting through wind farm layout optimization. Our method presents an extremely simple and easy to implement method that significantly reduces wind farm power variance. In this paper, we only consider uncertainty in wind direction. This was done for simplicity in presenting our new concept. Future studies could consider uncertainty in wind speed, and other sources of uncertainty such as wake model parameters. They could also apply our method to currently operating wind farms, and comparing the variance reduction that can be achieved while maintaining the mean power production of the existing layout.

Competing Interests

The authors declare no competing interests.

Author Contributions

BG led the research, including writing model code, running optimizations, and writing the paper. APJS wrote model code, created figures, helped develop methodology, and provided editing for the paper. AN developed ideas and methodology, provided feedback throughout the entire process, and provided editing for the paper.

Funding

This work was partially supported by the National Science Foundation under Grant No. 1539384. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

[1] R. Barthelmie, F. Murray, S. Pryor, The economic benefit of short-term forecasting for wind energy in the uk electricity market, Energy Policy 36 (2008) 1687 – 1696. URL: http://www.sciencedirect.com/science/article/pii/S0301421508000323. doi:https://doi.org/10.1016/j.enpol.2008.01.027.

[2] K. D. Orwig, M. L. Ahlstrom, V. Banunarayanan, J. Sharp, J. M. Wilczak, J. Freedman, S. E. Haupt, J. Cline, O. Bartholomy, H. F. Hamann, B. Hodge, C. Finley, D. Nakafuji, J. L. Peterson, D. Maggio, M. Marquis, Recent trends in variable generation forecasting and its value to the power system, IEEE Transactions on Sustainable Energy 6 (2015) 924–933. doi:10.1109/TSTE.2014.2366118.

[3] L. Yu, S. Zhong, X. Bian, W. E. Heilman, Temporal and spatial variability of wind resources in the united states as derived from the climate forecast system reanalysis, Journal of Climate 28 (2015) 1166–1183. URL: https://doi.org/10.1175/JCLI-D-14-00322.1. doi:10.1175/JCLI-D-14-00322.1. arXiv:https://doi.org/10.1175/JCLI-D-14-00322.1.

[4] V. Torralba, F. J. Doblas-Reyes, D. MacLeod, I. Christel, M. Davis, Seasonal climate prediction: A new source of information for the management of wind energy resources, Journal of Applied Meteorology and Climatology 56 (2017) 1231–1247. URL: https://doi.org/10.1175/JAMC-D-16-0204.1. doi:10.1175/JAMC-D-16-0204.1. arXiv:https://doi.org/10.1175/JAMC-D-16-0204.1.

[5] K. Methaprayoon, W. Lee, C. Yingvivatanapong, J. Liao, An integration of ANN wind power estimation into UC considering the forecasting uncertainty, 2005, pp. 116 – 124. doi:10.1109/ICPS.2005.1436364.

[6] F. Kaminsky, R. Kirchhoff, L. Sheu, Optimal spacing of wind turbines in a wind energy power plant, Solar Energy - SOLAR ENERG 39 (1987) 467–471. doi:10.1016/0038-092X(87)90053-3.

[7] B. Hrafnkelsson, G. Oddsson, R. Unnthorsson, A method for estimating annual energy production using monte carlo wind speed simulation, Energies 9 (2016) 286. doi:10.3390/en9040286.

[8] R. Aghabi Rivas, J. Clausen, K. S. Hansen, L. Jensen, Solving the turbine positioning problem for large offshore wind farms by simulated annealing, Wind Engineering 33 (2009) 287–297. doi:10.1260/0309-524X.33.3.287.

[9] K. Attias, S. Ladany, Optimal econoimc layout of turbines on windfarms, Wind Engineering 30 (2006) 141–151. doi:10.1260/030952406778055045.

[10] U. A. Ozturk, B. Norman, Heuristic methods for wind energy conversion system positioning, Electric Power Systems Research 70 (2004) 179–185. doi:10.1016/j.epsr.2003.12.006.
[11] P. Crosby, Application of a monte carlo optimization technique to a cluster of wind turbines, Journal of Solar Energy Engineering, Transactions of the ASME 109 (1987) 330–336. doi:10.1115/1.3268225.
[12] G. Marmidis, S. Lazarou, E. Pyrgioti, Optimal placement of wind turbines in a wind park using monte carlo simulation, Renewable Energy 33 (2008) 1455–1460. doi:10.1016/j.renene.2007.09.004.
[13] S. Grady, M. Hussaini, M. Abdullah, Placement of wind turbines using genetic algorithms, Renewable Energy 30 (2005) 259–270. doi:10.1016/j.renene.2004.05.007.
[14] D. Borissova, I. Mustakerov, Designing of wind farm layout by using of multi-objective optimization, International Journal of Mathematical Models and Methods in Applied Sciences 11 (2017) 290–295.
[15] V. Mytilinou, A. Kolios, A multi-objective optimisation approach applied to offshore wind farm location selection, Journal of Ocean Engineering and Marine Energy 3 (2017). doi:10.1007/s40722-017-0092-8.
[16] J. Thomas, A. Ning, A method for reducing multi-modality in the wind farm layout optimization problem, Journal of Physics: Conference Series 1037 (2018) 042012. doi:10.1088/1742-6596/1037/4/042012.
[17] B. L. Du Pont, J. Cagan, An extended pattern search approach to wind farm layout optimization, volume 134, 2010. doi:10.1115/DETC2010-28748.
[18] E. Tingey, J. Thomas, A. Ning, Wind farm layout optimization using sound pressure level constraints, 2015, pp. 154–159. doi:10.1109/SusTech.2015.7314339.
[19] P. Bortolotti, H. C. Tarres, K. L. Dykes, K. Merz, L. Sethuraman, D. Verelst, F. Zahle, Iea wind tcp task 37: Systems engineering in wind energy - wp2.1 reference wind turbines (2019). doi:10.2172/1529216.
[20] P. Gebraad, F. W. Teeuwisse, J. W. Wingerden, P. Fleming, S. D. Ruben, J. Marden, L. Pao, Wind plant power optimization through yaw control using a parametric model for wake effects—a cfd simulation study, Wind Energy (2014). doi:10.1002/we.1822.
[21] N. Jensen, A note on wind generator interaction, 1983.
[22] H. Akima, A new method of interpolation and smooth curve fitting based on local procedures, J. ACM 17 (1970) 589–602. doi:10.1145/321607.321609.
[23] K.-H. Chang, Chapter 19 - multiobjective optimization and advanced topics, in: K.-H. Chang (Ed.), e-Design, Academic Press, Boston, 2015, pp. 1105 – 1173. URL: http://www.sciencedirect.com/science/article/pii/B9780123820389000193. doi:https://doi.org/10.1016/B978-0-12-382038-9.00019-3.
[24] R. Marler, J. Arora, Survey of multi-objective optimization methods for engineering, Structural and Multidisciplinary Optimization 26 (2004) 369–395. doi:10.1007/s00158-003-0368-6.
[25] R. E. Perez, P. W. Jansen, J. R. R. A. Martins, pyOpt: A Python-based object-oriented framework for nonlinear constrained optimization, Structural and Multidisciplinary Optimization 45 (2012) 101–118. doi:10.1007/s00158-011-0666-3.
[26] P. Gill, W. Murray, M. Saunders, Snopt: An sqp algorithm for large-scale constrained optimization, SIAM Journal on Optimization 12 (2002) 979–1006. doi:10.2307/20453604.