

STOCHOS: Stochastic Opportunistic Maintenance Scheduling For Offshore Wind Farms

Petros Papadopoulos, David W. Coit , Ahmed Aziz Ezzat
Department of Industrial & Systems Engineering, Rutgers University

Abstract

Despite the promising outlook, the numerous economic and environmental benefits of offshore wind energy are still compromised by its high operations and maintenance (O&M) expenditures. On one hand, offshore-specific challenges such as site remoteness, harsh weather, transportation requirements, and production losses, significantly inflate the O&M costs relative to land-based wind farms. On the other hand, the uncertainties in weather conditions, asset degradation, and electricity prices largely constrain the farm operator’s ability to identify the time windows at which maintenance is possible, let alone optimal. In response, we propose STOCHOS, short for the stochastic holistic opportunistic scheduler—a maintenance scheduling approach tailored to address the unique challenges and uncertainties in offshore wind farms. Given probabilistic forecasts of key environmental and operational parameters, STOCHOS optimally schedules the offshore maintenance tasks by harnessing the opportunities that arise due to favorable weather conditions, on-site maintenance resources, and maximal operating revenues. STOCHOS is formulated as a two-stage stochastic mixed integer linear program, which we solve using a scenario-based rolling horizon algorithm that aligns with the industrial practice. Tested on real-world data from the U.S. North Atlantic where several offshore wind farms are in-development, STOCHOS demonstrates considerable improvements relative to prevalent maintenance benchmarks, across various O&M metrics, including total cost, downtime, resource utilization, and maintenance interruptions.

1 Introduction

Offshore wind (OSW) is set to play a pivotal role in the global climate change mitigation efforts. As a clean source of energy, it has a significantly smaller environmental footprint than fossil fuels, while it enables access to stronger and steadier winds than those in-land, resulting in a higher and more reliable generation of electricity. Several countries have set multi-year targets for the adoption or the expansion of their OSW portfolios, with recent reports projecting a growth in its global capacity of at least 15 times by 2040 (IEA, 2019).

Despite the promising outlook, OSW still faces substantial technical challenges pertaining to the high expenditures associated with operating and maintaining a fleet of OSW

turbines. Aside from the high upfront capital costs—which are largely due to the foundation construction and underwater cabling—a substantial portion of the total life cycle cost of OSW farms is attributed to operations and maintenance (O&M) activities. Recent estimates suggest that O&M activities performed throughout an OSW farm’s lifetime contribute by more than 30% to the total life cycle costs (Stehly and Beiter, 2020).

1.1 The need for opportunistic maintenance:

The high O&M costs in OSW farms are driven by several unique challenges, including:

(C1) OSW farm accessibility: The ability of the maintenance crew to safely access the OSW farm is frequently prohibited by harsh metocean conditions (e.g., high wind speeds and wave heights) (Gilbert et al., 2021). Our analysis of the metocean conditions in the designated OSW areas in the NY/NJ Bight (where several OSW projects are currently in-development) suggests that an OSW turbine sited at that location can be inaccessible for $\sim 56\%$ of its operational time due to unsafe wind and/or wave conditions (Papadopoulos et al., 2022).

(C2) Transportation and maintenance requirements: OSW turbines are installed at remote locations that are accessed via repair vessels or helicopters, typically chartered from private O&M contractors. Crew and equipment transport alone contributes $\sim 28\text{-}73\%$ of the OSW O&M costs (Carroll et al., 2017; Dalgic et al., 2015). Moreover, the harsh weather conditions at OSW farms further accelerate asset degradation and failure rates (Carroll et al., 2016), which directly translate into frequent maintenance visits. Those high maintenance requirements largely inflate the maintenance setup costs relative to onshore wind farms.

(C3) Revenue losses: Modern OSW turbines have $\sim 2\text{-}3$ times higher capacity than their land-based counterparts (Golparvar et al., 2021). A shutdown of such an ultra-scale turbine results in large revenue losses, especially at times of strong winds or favorable market prices.

Those offshore-specific challenges (C1-C3) make the conventional maintenance scheduling approaches—which may be well-suited for land-based operations—sub-optimal, motivating the need for an “*opportunistic*” approach for OSW maintenance scheduling. The essence of opportunistic maintenance is to incentivize the grouping of maintenance tasks (that would have been otherwise scheduled independently) at times of “maintenance opportunity,” in order to reduce the number of maintenance visits and aggregate them at times when it is most economical for the OSW farm operator (Ren et al., 2021). In light

of $C1-C3$, those opportunities can be one of the following: (i) access-based opportunities: grouping maintenance tasks at times of OSW farm access to leverage the (often scattered) periods of provisional accessibility, hence minimizing maintenance delays, interruptions, and downtime; (ii) resource-based opportunities: grouping maintenance tasks to leverage maintenance resources that are already on-site (e.g., vessels, equipment, crew) in order to share the setup cost among multiple turbines and maximize resource utilization; and (iii) revenue-based opportunities: grouping maintenance tasks at times when production losses would be minimal (e.g., at times of anticipated low winds and/or low market prices).

However, our survey of the literature (see Table 1) reveals that the overwhelming majority of research in wind farm opportunistic maintenance solely focuses on resource-based opportunities, which are often referred to in the literature as “economic dependencies” (Shafiee et al., 2015; Wang et al., 2019; Ding and Tian, 2012; Sarker and Faiz, 2016; Ma et al., 2022; Besnard and Bertling, 2010; Ko and Byon, 2017). A small fraction of those efforts partially account for revenue losses by incentivizing the grouping of maintenance tasks at periods of low power production (Besnard et al., 2009; Yildirim et al., 2017; Song et al., 2018). A separate line of research focuses on forecasting access-based opportunities (Yang et al., 2020; Lubing et al., 2019; Taylor and Jeon, 2018; Zhang et al., 2021), but largely overlooks those based on economic dependencies and/or revenue losses.

Barring few recent efforts (Mazidi et al., 2017; Besnard et al., 2011; Fallahi et al., 2022; Papadopoulos et al., 2022), there is a lack of approaches that holistically account for multiple components of maintenance opportunity. Papadopoulos et al. (2022) show that such a holistic approach can lead to major O&M cost savings due to the inter-dependencies between the three opportunistic components, making their combined impact on O&M costs multiplicative rather than additive. For instance, both access- and revenue-based opportunities are weather-related, while resource-based opportunities are only possible when the OSW farm is accessible. However, a main limitation in the work of Papadopoulos et al. (2022) is its inherent assumption that perfect knowledge about key environmental and operational parameters is available to the maintenance planner.

Table 1: The contribution of STOCHOS to the OSW opportunistic maintenance literature

Research Effort	Opportunities			Uncertainties			
	Access	Revenue	Economic	Access	Production	Elec. prices	Degradation
Besnard et al. (2009)	-	✓	✓	-	-	-	-
Shafiee et al. (2015)	-	-	✓	-	-	-	✓
Wang et al. (2019)	-	-	✓	-	-	-	✓
Ding and Tian (2012)	-	-	✓	-	-	-	✓
Sarker and Faiz (2016)	-	-	✓	-	-	-	✓
Ma et al. (2022)	-	-	✓	-	-	-	✓
Ko and Byon (2017)	-	-	✓	-	-	-	✓
Yildirim et al. (2017)	-	✓	✓	-	✓	-	✓
Mazidi et al. (2017)	✓	✓	✓	-	✓	✓	-
Song et al. (2018)	-	✓	✓	-	-	-	✓
Fallahi et al. (2022)	-	✓	✓	-	✓	✓	✓
Besnard et al. (2011)	✓	✓	✓	✓	✓	-	-
Besnard and Bertling (2010)	-	-	✓	-	-	-	-
Yang et al. (2020)	✓	✓	-	-	-	-	✓
Lubing et al. (2019)	✓	-	-	✓	-	-	✓
Taylor and Jeon (2018)	✓	✓	-	✓	-	-	-
Papadopoulos et al. (2022)	✓	✓	✓	-	-	-	-
STOCHOS	✓	✓	✓	✓	✓	✓	✓

1.2 The value of recognizing uncertainty in OSW maintenance

Several prior studies have advocated the importance of uncertainty modeling in wind farm maintenance planning (Byon et al., 2010; Perez et al., 2015). We argue that, for an opportunistic maintenance approach, the need to recognize uncertainty is even more pressing because the imperfect knowledge about key environmental and operational parameters can largely compromise (or may even reverse) the economic gains of an opportunistic approach.

The main reason is that those parameters, which are uncertain and difficult to predict, are the primary determinants of the “opportunity windows”, i.e. the times at which maintenance grouping is economically desirable, including: *(i)* uncertainty of metocean conditions (wind speed and wave height), which are the main determinants of OSW farm access (access- and resource-based opportunities), as well as the power production (revenue-based opportunities), *(ii)* uncertainty of electricity prices, which govern the generated revenues (revenue-based opportunities), and *(iii)* uncertainty in asset degradation, which affects the turbine’s downtime and production losses (revenue- and resource-based opportunities).

This suggests that an opportunistic maintenance approach, despite being desirable in essence, could have conflicting effects: On one hand, if (near) perfect knowledge of the uncertain parameters listed above is available, then the maintenance planner can precisely anticipate the windows of opportunity, and hence, fully harness the economic gains of an opportunistic approach. On the other hand, incorrect information about such parameters can have an adverse impact: Multiple maintenance tasks would be grouped at a time when maintenance may not even be feasible, let alone optimal, leading to significant maintenance delays and excessive downtime. In such situations, even an “unopportunistic” approach that individually schedules maintenance actions would have possibly been more cost-efficient.




Figure 1: (A) Wind speed and wave height data, forecasts, along with 50 scenarios. (B) Accessibility states. (C) Maintenance schedules of a deterministic and a stochastic strategy.

Figure 1 illustrates the value of considering uncertainty in OSW opportunistic maintenance with a realistic example, where the wind speed and wave height data, forecasts, and scenarios are displayed in Figure 1-A. Figure 1-B conveys information about the “access state”; the OSW farm is inaccessible when the wave height and/or the wind speed exceed

their respective safety thresholds. Specifically, the top row of Figure 1-B shows the scenario-based probability of access, whereas the middle and bottom rows show the accessibility as a binary variable informed by the forecasts and real data, respectively. Figure 1-C is a Gantt chart, showing the maintenance schedules obtained by two optimization models: a deterministic (dotted bars) and a stochastic (hashed bars) model. The x -axis denotes the planning horizon (in hours) and the y -axis denotes the wind turbine (WT) index (here, we have 5 WTs). Note that both optimization models are opportunistic, i.e., they both leverage access-, resource-, and revenue-based opportunities. Yet, the main difference is that the former only takes as input the point forecasts of the environmental and operational parameters listed above, while the latter makes use of the full probabilistic distribution of the forecasts to generate uncertainty scenarios as inputs to the stochastic optimization.

In the case of the deterministic model which assigns complete confidence to the point forecasts, *small forecasting errors can incur significant maintenance costs*. Consider the fifth day (96-120-h range) where the point forecasts under-predict the wave height, projecting that the OSW farm would be accessible (while in reality, it is not; see inset in Figure 1-A). Thus, the deterministic approach (mistakenly) perceives it as an optimal time to group three maintenance tasks for WT3, WT4, and WT5 to leverage this access period and concurrent low power production. With access inaccurately estimated, all three maintenance tasks end up being interrupted (See Figure 1-C), leaving the WTs in a non-operational status until the following day, where maintenance is ultimately resumed once access is restored.

The stochastic approach, on the other end, recognizes the forecast uncertainty (by means of the probabilistic scenarios), and hence, the risk of maintenance interruption is foreseen. Here, the model delays the maintenance until day 7, where both wind and wave conditions are expected to be tolerable, resulting in uninterrupted maintenance for all three WTs, and a total cost reduction of 31% relative to the deterministic approach’s schedule.

It is natural to expect that in dynamic operational environments, such as that of OSW farms, similar cases caused by the synergistic effects of parameter uncertainty occur quite often. It is therefore necessary for a maintenance scheduling approach to be able to account, not only for maintenance opportunities, but also for the uncertainties in the environmental and operational parameters that define what constitutes a maintenance opportunity in the first place. Looking at Table 1, one can notice a lack of efforts that attempt to holistically

consider maintenance opportunities *and* the full spectrum of parameter uncertainty in a unified framework. This is the research gap addressed in this work.

1.3 Contributions

To fill this research gap, we propose the STOchastic Holistic Oppportunistic Scheduler (STOCHOS)—In Greek, STOCHOS (written as $\Sigma\tau\acute{o}\chi\omicron\varsigma$) translates to “target,” or “objective.” The contributions of this work are listed as follows:

1. STOCHOS is an offshore-tailored maintenance optimization framework which uniquely
 - (i) leverages the combined maintenance opportunities arising due to accessibility, revenue, and transportation costs, and
 - (ii) accounts for the multi-source uncertainties of metocean conditions, asset degradation, and electricity prices. Compared to the related literature, STOCHOS presents a novel modeling framework for opportunistic maintenance optimization under uncertainty, while its rolling horizon solution algorithm is in line with the industrial practice of OSW maintenance planning.
2. Instead of using temporally invariant distributional assumptions about input parameters, as prevalent in the stochastic optimization literature, we generate scenarios (or trajectories) from a probabilistic forecasting framework. The trajectories recognize the temporal nature of OSW data, and constitute optimal inputs to the stochastic optimization model. We show the economic value of this approach relative to using marginal predictive distributions or point predictions.
3. In terms of practical relevance, we extensively evaluate STOCHOS using real-world data from a unique geographical region in the U.S. North Atlantic where several large-scale projects are currently in-development (BOEM, 2021). We compare STOCHOS against a set of prevalent maintenance benchmarks, and show its superior performance in terms of several O&M metrics. Our data, models, and findings can therefore provide timely and key insights to the operators of those soon-to-be-operational OSW farms.

The remainder of the paper is organized as follows. In Section 2, we present the formulation of STOCHOS. Section 3 presents the probabilistic forecasting and scenario generation approaches, followed by the solution procedure in Section 4. Section 5 presents our results and findings. We summarize our conclusions and future research directions in Section 6.

2 Model Formulation

This section starts with an introduction of notation and modeling assumptions, followed by a detailed description of the mathematical model of STOCHOS.

2.1 Problem description and assumptions

- (A1) We define $\mathcal{I} := \{i \mid i \in \mathbb{Z}^+, 1 \leq i \leq N_{\mathcal{I}}\}$ as the set of offshore WTs that must undergo a maintenance task of τ_i hours (called the repair time). We focus on minor to medium tasks which typically require less than a day to complete (e.g. electrical components, grease/oil cooling liquids, sensors and controls), and comprise $\sim 75\%$ of all maintenance tasks in OSW farms (Carroll et al., 2016; Dinwoodie et al., 2015).
- (A2) The planning horizon is divided into two sub-horizons (Koltsidopoulos Papatzimos, 2020): a day-ahead short-term horizon (STH) and a long-term horizon (LTH). The STH, $\mathcal{T} := \{t \mid t \in \mathbb{Z}^+, 1 \leq t \leq 24 \text{ hours}\}$, is of hourly resolution. The LTH, $\mathcal{D} := \{d \mid d \in \mathbb{Z}^+, 1 \leq d \leq N_{\mathcal{D}} \text{ days}\}$, has a daily resolution, starting the day after the STH up to $N_{\mathcal{D}}$ days ahead. We denote LTH variables and parameters with a superscript L to distinguish them from their STH counterparts.
- (A3) STOCHOS is solved using sample average approximation by optimizing a stochastic MILP (to be presented next) over a random scenario subset, $\mathcal{S} := \{s \mid s \in \mathbb{Z}^+, 1 \leq s \leq N_{\mathcal{S}} \text{ scenarios}\}$. Hereinafter, notations for random variables, real data, point forecasts, and scenarios are denoted as follows. If $z(t)$ denotes an arbitrary random variable in the STH, then z_t denotes its true (actual) value at the t -th hour, $\hat{z}_t$ denotes the corresponding point forecast, while $z_{t,s}$ is the realization of a random scenario s at time t . Replacing t by d extends this notation to an LTH variable $z^L(d)$.
- (A4) We assume that the wind speed, wave height, and electricity prices are uncertain during both the STH and LTH, which are denoted by $\nu(t), \eta(t), \kappa(t)$ for STH, and $\nu^L(d), \eta^L(d), \kappa^L(d)$ for LTH, respectively. For asset degradation, we assume that the OSW turbine is an integrated system with a system-level residual life (RL), denoted by $\lambda^L(i) \forall i \in \mathcal{I}$, which is a turbine-specific random variable defined as the time (in days) at which the i th WT fails. In practice, predictions about $\lambda^L(i)$ can be obtained using

a condition monitoring system that considers the impact of dynamic environmental and operational degradation parameters (Bian et al., 2015). We assume that the RL is only uncertain in the LTH, since the STH, which is the day-ahead horizon, is a small interval of time relative to the typical degradation process of WTs. A WT for which the realized RL is smaller than a day is bound to fail by the end of the STH (unless it is maintained). Maintenance tasks performed before the realized (true) RL are preventive (PM), otherwise they are corrective (CM).

- (A5) The power produced from a WT is primarily a function of the hub-height wind speed, and is estimated using historical data (Ezzat et al., 2018; Ding, 2019). Details of the wind-to-power conversion are discussed in the supplementary materials, SM-1.
- (A6) WTs are accessed by crew transport vessels (CTVs) which are subject to accessibility constraints defined by wind speed and wave height safety thresholds, denoted as ν_{max} and η_{max} , respectively. Maintenance operations are subject to interruptions due to access constraints, and can be resumed once access is restored. To formally account for maintenance interruptions, we associate each maintenance task with a “mission time,” $A(t, i)$ which is a random variable denoting the time needed to complete a task of τ_i hours starting at time t . $A(t, i)$ is a function of τ_i and the access state denoted by $X(t) \in \{0, 1\}$, which, in turn, is a function of the stochastic metocean conditions $\nu(t)$ and $\eta(t)$, and the safety thresholds ν_{max} and η_{max} . Details about the estimation of $X(t)$ and $A(t, i)$ given metocean conditions and maintenance times are provided in the supplementary materials SM-2. Note that while maintenance interruptions are allowed, longer mission times directly translate to prolonged work hours (i.e., higher crew and vessel contracting costs), and longer downtimes (i.e., higher revenue losses).

2.2 The optimization model

There exist different approaches to uncertainty modeling in the maintenance optimization literature. One common approach is to directly encode a set of assumed probability distributions about the uncertain parameters into the optimization model, which often results in nonlinear models that are typically solved using heuristic approaches (Shafiee et al., 2015; Wang et al., 2019; Ding and Tian, 2012; Lu et al., 2017; Sarker and Faiz, 2016). Alter-

natively, stochastic programming uses sampling methods to integrate realizations of those probability distributions to minimize maintenance costs (or similar objectives) (Besnard et al., 2011; Yildirim et al., 2017; Fallahi et al., 2022). STOCHOS belongs to the second cluster, and is formulated as a two-stage stochastic mixed linear integer program (MILP). The following binary variables constitute the decision variables in STOCHOS:

- $m_{t,i}$: A maintenance task for WT i is scheduled to start at hour t in the STH.
- r : A vessel is rented for the STH.
- $m_{d,i,s}^L$: A maintenance task is scheduled for WT i at day d in the LTH, in scenario s .
- $r_{d,s}^L$: A vessel is rented at day d in the LTH, in scenario s .

The objective, shown in (1), is to maximize the profit in the planning horizon, comprising the day-ahead profit l^{STH} , the long-term profit $\{l_d^{LTH}\}_{d \in \mathcal{D}}$, and a stochastic penalty term.

$$\max_{m_{t,i}, m_{d,i,s}^L, r, r_{d,s}^L} \left\{ \overbrace{l^{STH}}^{\text{short-term profit}} + \overbrace{\sum_{d \in \mathcal{D}} l_d^{LTH}}^{\text{long-term profit}} - \underbrace{\frac{1}{N_S} \sum_{s \in \mathcal{S}} \left[\overbrace{\sum_{i \in \mathcal{I}} (U_s \cdot w_{i,s} + Y_s \cdot b_{i,s})}^{\text{prolonged interruptions}} + \overbrace{C_1 \cdot a_s^x + C_2 \cdot a_s^q}^{\text{spot resource contracting}} \right]}_{\text{stochastic penalty term}} \right\}. \quad (1)$$

The short-term profit, l^{STH} , defined as in (2), is the difference between the day-ahead operating revenues and the maintenance costs. The revenue is generated from selling the power output of the i th WT at time t and scenario s , denoted by $p_{t,i,s}$, at an hourly market price $\kappa_{t,s}$. The maintenance costs comprise four components: (1) the repair costs which are dictated by the PM and CM cost coefficients, K (\$/task) and Φ (\$/task), respectively, (2) the daily vessel rental cost with a daily rate of Ω (\$/day), (3) crew-related costs charged at the crew hourly cost rate Ψ (\$/h) where $x_{t,i,s} \in \{0,1\}$ denotes whether a crew is assigned to the i th WT at time t and scenario s , and (4) overtime costs charged at the crew overtime cost rate Q (\$/h) where $q_s \in \mathbb{Z}^+$ denotes the overtime hours worked in the STH. Binary parameter ζ_i indicates the day-ahead operational status of the i -th WT, wherein $\zeta_i = 0$ indicates a failed WT needing a CM. The parameter $\rho_i \in \{0,1\}$ denotes whether a maintenance task has been initiated in a previous day but is yet to be completed (when

$\rho_i = 1$, repair cost goes to zero as it has already been accounted for in a previous day).

$$l^{STH} = - \sum_{i \in \mathcal{I}} \sum_{t \in \mathcal{T}} \left[\overbrace{(1 - \rho_i) \cdot [\zeta_i \cdot K + (1 - \zeta_i) \cdot \Phi] \cdot m_{t,i}}^{\text{repair cost}} - \overbrace{\Omega \cdot r}^{\text{vessel cost}} \right. \\ \left. + \frac{1}{N_S} \sum_{s \in \mathcal{S}} \left[\sum_{i \in \mathcal{I}} \sum_{t \in \mathcal{T}} \left(\underbrace{\kappa_{t,s} \cdot p_{t,i,s}}_{\text{operating revenue}} - \underbrace{\Psi \cdot x_{t,i,s}}_{\text{crew cost}} \right) - \underbrace{Q \cdot q_s}_{\text{overtime cost}} \right] \right]. \quad (2)$$

The long-term profit, in (3), is similarly calculated for daily intervals. The crew work hours are calculated as the mission time $A_{d,i,s}^L$ of the tasks scheduled at that day, $m_{d,i,s}^L$.

$$l_d^{LTH} = \frac{1}{N_S} \sum_{s \in \mathcal{S}} \left\{ \sum_{i \in \mathcal{I}} \left[\underbrace{\kappa_{d,s}^L \cdot p_{d,i,s}^L}_{\text{operating revenue}} - \overbrace{(1 - \rho_i) \cdot [\zeta_{d,i,s}^L \cdot K + (1 - \zeta_{d,i,s}^L) \cdot \Phi] \cdot m_{d,i,s}^L}^{\text{repair cost}} \right. \right. \\ \left. \left. - \underbrace{\Psi \cdot A_{d,i,s}^L \cdot m_{d,i,s}^L}_{\text{crew cost}} \right] - \underbrace{\Omega \cdot r_{d,s}^L}_{\text{vessel cost}} - \underbrace{Q \cdot q_{d,s}^L}_{\text{overtime cost}} \right\} \quad \forall d \in \mathcal{D}. \quad (3)$$

The last term in (1) is a stochastic penalty which entails two sub-terms. The first term, which we call “prolonged interruptions,” represents the cost of maintenance actions that started in the STH, but were not completed, and hence, must be resumed in the LTH (The WT remains non-operational when such event takes place). In that term, $U_s \cdot w_{i,s}$ represent the cost incurred *until* the maintenance is restarted (pre-maintenance), while $Y_s \cdot b_{i,s}$ is the cost incurred *while* the maintenance is ongoing (during maintenance). In specific, $w_{i,s} \in \{0, 1\}$ denotes the occurrence of the interruption event, U_s is the associated pre-maintenance cost, $b_{i,s} \in \mathbb{Z}^+$ denotes the remaining maintenance time that has to be completed in the LTH, and Y_s is the associated cost incurred during the maintenance. The second sub-term penalizes the (unlikely) case when the maximum budget for crew (17) or overtime resources (11) are exceeded. We assign fairly large cost coefficients, C_1 and C_2 , for such exceedances to reflect the unavoidable practice of on-the-spot contracting of additional crew, equipment, or other immediate maintenance resources, when necessary. In our experiments, those exceedances occur less than 0.09% of the time.

Maintenance Constraints: The equality in (4) forces a maintenance task to be scheduled either in the STH or in the LTH.

$$\sum_{t \in \mathcal{T}} m_{t,i} + \sum_{d \in \mathcal{D}} m_{d,i,s}^L = 1 \quad \forall i \in \mathcal{I}, s \in \mathcal{S}. \quad (4)$$

A maintenance task in the STH can only be initiated after the time of first light, t_R , and before the time of last sunlight, t_D , as expressed in (5) and (6), respectively.

$$m_{t,i} \leq \frac{t}{t_R} \quad \forall t \in \mathcal{T}, i \in \mathcal{I}. \quad (5)$$

$$m_{t,i} \leq \frac{t_D}{t} \quad \forall t \in \mathcal{T}, i \in \mathcal{I}. \quad (6)$$

Once a maintenance task is initiated at time t for the i th WT, then it would be under maintenance for a period of time computed as the minimum between the remaining time in the STH, which is $24 - t$, and the mission time $A_{t,i,s}$. This is expressed in (7), where the binary variable $u_{\tilde{t},i,s}$ denotes a turbine under maintenance at time $\tilde{t}$.

$$\sum_{\tilde{t}=t}^{\min\{24,t+A_{t,i,s}\}} u_{\tilde{t},i,s} \geq \min\{24 - t, A_{t,i,s}\} \cdot m_{t,i} \quad \forall t \in \mathcal{T}, i \in \mathcal{I}, s \in \mathcal{S}. \quad (7)$$

If that maintenance task is not completed within the STH, then $b_{i,s} \in \mathbb{Z}^+$, as shown in (8), denotes the remaining time required to complete it in the LTH, while $w_{i,s} \in \{0, 1\}$, as shown in (9), denotes the occurrence of such event and is only set to 1 once $b_{i,s} > 0$, as it is multiplied by M , an arbitrary large number. This is the case where the “prolonged interruption” penalty term in (1) is activated.

$$b_{i,s} \geq \sum_{t \in \mathcal{T}} m_{t,i} \cdot [A_{t,i,s} - 24 + t]^+ \quad \forall i \in \mathcal{I}, s \in \mathcal{S}. \quad (8)$$

$$b_{i,s} \leq M \cdot w_{i,s} \quad \forall i \in \mathcal{I}, s \in \mathcal{S}. \quad (9)$$

The maintenance crew, once dispatched, is occupied until the maintenance is completed or their shift ends by the time of last sunlight, t_D , as shown in (10). An upper bound on the number of available maintenance crews is set to B (crews), as shown in (11). The auxiliary variable $a_s^x \in \mathbb{Z}^+$ is added to the right-hand-side (RHS) of (11) to account for the (unlikely) case in which the crew budget limit is exceeded, but almost always remains zero since it is heavily penalized by C_1 in (1). In practice, this translates to situations where additional crew or equipment is contracted on-the-spot.

$$x_{t,i,s} \geq u_{t,i,s} - \frac{t}{t_D} \quad \forall t \in \mathcal{T}, i \in \mathcal{I}, s \in \mathcal{S}. \quad (10)$$

$$\sum_{i \in \mathcal{I}} x_{t,i,s} \leq B + a_s^x \quad \forall t \in \mathcal{T}, s \in \mathcal{S}. \quad (11)$$

Turbine Availability Constraints: A failed WT (i.e., one which has not been maintained at or before its RL) remains unavailable until a maintenance action is completed. For the STH case, availability can be expressed by the binary variable $y_{t,i,s}$, as in (12).

$$y_{t,i,s} \leq \underbrace{\zeta_i \cdot (1 - \rho_i)}_{\text{WT operational status}} + \underbrace{\frac{24 \cdot \sum_{\tilde{t} \in \mathcal{T}} m_{\tilde{t},i} - \sum_{\tilde{t} \in \mathcal{T}} (\tilde{t} \cdot m_{\tilde{t},i})}{24 - t + g}}_{\text{availability restored after maintenance}} \quad \forall t \in \mathcal{T}, i \in \mathcal{I}, s \in \mathcal{S}. \quad (12)$$

The first term of the RHS in (12) denotes whether the turbine is in a failed state at the beginning of the STH, or if a maintenance task has been initiated in a previous day. In case the RL is reached, then the turbine fails ($\zeta_i = 0$), and can only return to its former operational status once a CM action is performed. This is enforced by the second term in the RHS; if no CM is scheduled, then $m_{\tilde{t},i} = 0$ and the term drops to 0. In contrast, if $m_{\tilde{t},i} = 1$, the term is greater than 1 after the time of maintenance, $\tilde{t}$. An arbitrary small number g avoids division by zero. A similar constraint for the LTH is shown in (13).

$$y_{d,i,s}^L \leq \zeta_{d,i,s}^L \cdot (1 - \rho_i) + \frac{N_{\mathcal{D}} - \sum_{d \in \mathcal{D}} (d \cdot m_{d,i,s}^L)}{N_{\mathcal{D}} - d + g} \quad \forall d \in \mathcal{D}, i \in \mathcal{I}, s \in \mathcal{S}. \quad (13)$$

A WT under maintenance remains unavailable until the task is completed, as in (14).

$$y_{t,i,s} \leq 1 - u_{t,i,s} \quad \forall t \in \mathcal{T}, i \in \mathcal{I}, s \in \mathcal{S}. \quad (14)$$

Vessel Rental Constraints: Vessels are rented daily only if a maintenance task has been scheduled at that day. This is enforced via (15)-(16), for the STH and LTH, respectively.

$$M \cdot r \geq \sum_{t \in \mathcal{T}} \sum_{i \in \mathcal{I}} m_{t,i}. \quad (15)$$

$$M \cdot r_{d,s}^L \geq \sum_{i \in \mathcal{I}} m_{d,i,s}^L \quad \forall d \in \mathcal{D}, s \in \mathcal{S}. \quad (16)$$

Overtime Constraints: The total work hours, as shown in (17), cannot exceed W (hour/crew). Otherwise, overtime, tracked by integer variable q_s , are incurred and compensated at a higher rate determined by Q (\$/hour). An upper bound of H (in hours) limits the number of overtime hours in both the STH and LTH, as in (20) and (21). Similar to a_s^x in (11), a_s^q in (17) accounts for the on-spot contracting of additional crew, at a large cost of C_2 .

A task initiated but not finished in the STH, is prioritized in the first day of the LTH ($d = 1$). This is expressed in (18) where $A_{d=1,i,s}^L \cdot m_{d=1,i,s}^L + b_{i,s}$ denotes the total work hours in the first day of the LTH (the sum of the mission times and the remaining maintenance

time of unfinished tasks). A similar constraint in (19) is imposed for the remaining days of the LTH, i.e. for $d \in \{2, \dots, N_D\}$.

$$\sum_{t \in \mathcal{T}, i \in \mathcal{I}} x_{t,i,s} \leq B \cdot W + q_s + a_s^q \quad \forall s \in \mathcal{S}. \quad (17)$$

$$\sum_{i \in \mathcal{I}} [A_{1,i,s}^L \cdot m_{1,i,s}^L + b_{i,s}] \leq B \cdot W + q_{1,s}^L \quad \forall s \in \mathcal{S}. \quad (18)$$

$$\sum_{i \in \mathcal{I}} [A_{d,i,s}^L \cdot m_{d,i,s}^L] \leq B \cdot W + q_{d,s}^L \quad \forall d \in \{2, \dots, N_D\}, s \in \mathcal{S}. \quad (19)$$

$$q_s \leq H \quad \forall s \in \mathcal{S}. \quad (20)$$

$$q_{d,s}^L \leq H \quad \forall d \in \mathcal{D}, s \in \mathcal{S}. \quad (21)$$

Electricity Generation Constraints: The hourly power output, $p_{t,i,s}$, is computed as a fraction $f_{t,i,s} \in [0, 1]$ of the turbine's rated capacity R (MW), multiplied by the turbine's availability $y_{t,i,s}$, as shown in (22). When the WT is in a failed state, or under maintenance, $p_{t,i,s} = y_{t,i,s} = 0$, and the operator forfeits the associated revenue. The fraction $f_{t,i,s}$ is called *the normalized power level*. Full details of estimating $f_{t,i,s}$ and $f_{d,i,s}^L$ given wind speed conditions are shown in the supplementary materials, SM-1.

$$p_{t,i,s} \leq R \cdot f_{t,i,s} \cdot y_{t,i,s} \quad \forall t \in \mathcal{T}, i \in \mathcal{I}, s \in \mathcal{S}. \quad (22)$$

Likewise, the daily power output, $p_{d,i,s}^L$, is defined in (23), wherein $f_{d,i,s}^L \in [0, 1]$ is a function of the daily average wind speed. The term $m_{d,i}^L \cdot \zeta_{d,i,s}^L \cdot A_{d,i,s}^L / 24$, accounts for the production losses if a preventive maintenance has been scheduled on day d .

$$p_{d,i,s}^L \leq 24 \cdot R \cdot f_{d,i,s}^L \cdot \left(y_{d,i,s}^L - m_{d,i,s}^L \cdot \frac{A_{d,i,s}^L \cdot \zeta_{d,i,s}^L}{24} \right) \quad \forall d \in \mathcal{D}, i \in \mathcal{I}, s \in \mathcal{S}. \quad (23)$$

Curtailling wind power production is common in grid-connected wind farms. This is accounted for by introducing the parameter $C_{t,s} \in [0, 1]$ in (24) which defines the fraction of the farm-level power output that can be injected to the grid at each hourly interval t . $C_{t,s} = 1$ denotes no curtailment—i.e., all power produced is eventually sold.

$$\sum_{i \in \mathcal{I}} p_{t,i,s} \leq \sum_{i \in \mathcal{I}} f_{t,i,s} \cdot R \cdot C_{t,s} \quad \forall t \in \mathcal{T}, s \in \mathcal{S}. \quad (24)$$

3 Probabilistic Forecasting and Scenario Generation

We first introduce the data used in this work and its relevance to the U.S. OSW industry, then present our probabilistic forecasting and scenario generation framework.

3.1 Data description

This work is motivated by the ongoing large-scale OSW developments in the US North Atlantic, and in particular, the NY/NJ Bight, where several Gigawatt-scale OSW projects are in-development (BOEM, 2021). We use real-world OSW data and meteorological forecasts from this region, in proximity to those soon-to-be-operational OSW projects.

Metoccean Data and Forecasts: We use wind speed and wave height measurements collected by the E05 Hudson North buoy between August 2019 and May 2020. This is one of two floating LiDAR buoys recently deployed by the NY State Energy Research & Development Authority (NYSERDA) (NYSERDA, 2021) in the NY/NJ Bight. Co-located with the measurements are numerical weather predictions (NWP) of wind speed and wave height. The wind speed NWP come from RU-WRF, a meso-scale weather model developed by the Rutgers University Center For Ocean Observing Leadership (RUCOOL) to capture the unique physical phenomena in this geographical region (Optis et al., 2020). Wave height NWP come from WAVEWATCH III, a numerical wave model maintained by the National Oceanic and Atmospheric Administration (NOAA). The metoccean data and their NWP are shown in Figures 2-A and 2-B, respectively.

Electricity Price Data and Forecasts: Real day-ahead electricity prices are obtained from PJM’s Data Miner III (PJM, 2021), at node COMED (node id: 33092371). Point forecasts for the same node are obtained by a Lasso Estimated Auto-Regressive (LEAR) model proposed in Lago et al. (2021). The time series of the real electricity prices, the LEAR predictions, and the associated error histograms are shown in Figure 2-C.

Asset Degradation Data and Forecasts: Unlike the above-listed parameters (metoccean conditions and electricity prices), actual turbine degradation data requires access to O&M records. Currently, there are no operational OSW farms in this region. So, we assume a set of observed RLs, $\{\lambda_i\}_{i \in \mathcal{I}}$, which are unknown to the planner, who only has access to a set of (imperfect) RL predictions $\{\hat{\lambda}_i\}_{i \in \mathcal{I}}$ provided by a condition monitoring system. In

this study, true and predicted RLs are carefully selected in light of the summary statistics derived in previous studies on offshore WT reliability analyses (Carroll et al., 2016), thus ensuring a faithful replication of real-world WT failures (More details in Section 5.1).




Figure 2: Time series of the data and forecasts of wind speed (A), wave height (B) and electricity price (C). Histograms of the forecast residuals are shown to the right.

3.2 Uncertainty modeling and scenario generation

Modeling uncertainty in scenario-based stochastic optimization requires probabilistic characterizations of the input parameters. A prevalent approach in the literature is to impose a distributional assumption on the marginal distribution of the forecast residuals at each lead time. Albeit convenient, this approach is sub-optimal when temporal dependencies exist, as in the case of OSW data. Figure 3 depicts the density of the forecast residuals at four time lags for wind speed (top), wave height (middle), and electricity price (bottom).

Instead, a more powerful approach is to fully describe the density of the temporal process governing each input parameter. Scenarios generated from such approach would preserve the temporal dependence and constitute the optimal input to a stochastic program (Pinson, 2013). Specifically, we define $z(t)$ as the random variable denoting either the wind speed $\nu(t)$, wave height $\eta(t)$, or electricity price $\kappa(t)$ (we treat the RL predictions differently as discussed at the end of this section). We also denote by $\mathbf{z} = [z_1, \dots, z_{tc}]^T$ the set of historical




Figure 3: Density plots of the point forecast residuals versus their lagged versions for wind speed (A), wave height (B) and electricity price (C) time series forecasts.

measurements of $z(t)$ up to the current time t_c . Similarly, $\hat{\mathbf{z}} = [\hat{z}_1, \dots, \hat{z}_{t_c}]^T$ are the raw point predictions available to the planner (RU-WRF for $\nu(t)$, WAVEWATCH III for $\eta(t)$, and LEAR for $\kappa(t)$). Then, we propose the following probabilistic forecasting framework, comprising three independent terms:

$$z(t) = \mu(t) + \omega(t) + \epsilon(t), \quad (25)$$

where $\mu(t)$ is a deterministic mean function defined as the raw point predictions available to the maintenance planner (that is, $\mu(t) = \hat{z}_t \forall t$). The term $\omega(t)$ is a zero-mean temporal Gaussian Process (GP) with its pairwise covariance denoted by $\sigma_{t,t'} = \text{Cov}\{\omega(t), \omega(t')\}, \forall t, t'$, while $\epsilon(t)$ is the zero-mean Gaussian noise, such that $\boldsymbol{\epsilon} = [\epsilon_1, \dots, \epsilon_{t_c}]^T \sim \mathcal{N}(0, \delta \mathbf{I})$, where $\mathbf{I}$ is the $t_c \times 1$ identity matrix, and δ is the noise variance.

Let us use $\mathbf{C}$ to denote the $t_c \times t_c$ covariance matrix whose (t, t') -th entry is defined as $\sigma_{t,t'} + \delta \mathbb{I}(t = t')$, where $\mathbb{I}(\cdot)$ is the indicator function. Those entries are estimated using a stationary kernel $C(\cdot, \cdot)$ which encodes the temporal dependence. Several mathematically permissible choices for $C(\cdot, \cdot)$ are possible and we use the squared exponential kernel, for which the hyperparameters are estimated by maximizing the log-likelihood of the GP. Using the estimated parameters, we can obtain the estimated covariance matrix, $\hat{\mathbf{C}}$.

Putting together all the pieces, one can fully characterize the predictive distribution of the random variable $z^*(t_c + h)$ which denotes the forecast of $z(t_c + h)$. By virtue of the GP, this distribution is Gaussian with its mean and variance as in (26) and (27), respectively.

$$\bar{z}_{t_c+h} = \mathbb{E}[z^*(t_c + h)|\mathbf{z}, \hat{\mathbf{z}}] = \hat{z}_{t_c+h} + \hat{\mathbf{k}}^T \hat{\mathbf{C}}^{-1}(\mathbf{z} - \hat{\mathbf{z}}), \quad (26)$$

$$\bar{\sigma}_{t_c+h}^2 = \mathbb{V}[z^*(t_c + h)|\mathbf{z}, \hat{\mathbf{z}}] = \hat{\alpha} - \hat{\mathbf{k}}^T \hat{\mathbf{C}}^{-1} \hat{\mathbf{k}}, \quad (27)$$

where $h \in \{1, \dots, H\}$ is the forecast lead time, $\hat{z}_{t_c+h}$ is the raw forecast (e.g. from RU-WRF for wind speed) at $t_c + h$. The $t_c \times 1$ vector $\hat{\mathbf{k}}$ contains the pairwise covariances between $[1, \dots, t_c]^T$ and $t_c + h$, computed using $C(\cdot, \cdot)$. A similar model to (25) is developed for the LTH (in daily resolution), and similar expressions to those in (26) and (27) are derived.

The advantages of this probabilistic framework are two-fold. First is its ability to effectively account for the temporal dependencies in OSW data. This is obvious in how the predictive mean in (26) depends on the correlation between the target forecast and the historical observations. This is also apparent in how the predictive variance in (27) is reduced from the marginal variance when information about strongly correlated values in the process history is available. This is in contrast to distributional assumptions which treat each lead time individually, thus overlooking the value added by the temporal dependence.

The second advantage is our ability to draw random temporal trajectories (or scenarios) by sampling from the joint multivariate Gaussian distribution of the forecasts, as shown in (28). The resulting trajectories, which naturally encode the temporal dependence in them, are optimal inputs to the stochastic program, as opposed to marginal predictive densities.

$$\mathbf{z}^* = [z^*(t_c + 1), \dots, z^*(t_c + H)]^T \sim \mathcal{N}(\bar{\mathbf{z}}, \bar{\mathbf{C}}), \quad (28)$$

where $\bar{\mathbf{z}} = [\bar{z}_{t_c+1}, \dots, \bar{z}_{t_c+H}]^T$ is the $H \times 1$ vector of the GP predictive means, and $\bar{\mathbf{C}}$ is the correspondent $H \times H$ predictive covariance matrix.

Figure 4 shows a fan chart of 1,000 trajectories of wind speeds drawn randomly on a select day using (28). Similar trajectories can be drawn for the wave height and electricity prices. While the above two advantages are largely acknowledged in the wind forecasting literature (Pinson, 2013; Ezzat et al., 2019), limited research has bridged such probabilistic forecasting models with the stochastic maintenance optimization realm. One of the key outcomes of our analyses in this paper is to demonstrate the “economic value” realized by an effective uncertainty modeling of temporal dependencies in input parameters.




Figure 4: Fan charts of 1,000 random scenarios (trajectories) on 09/21/2019.

Table 2: Average solution time in seconds, for different number of scenarios (N_S) and number of wind turbines (N_I), using an optimality gap of 0.1%.

		N_S			
		25	50	100	200
N_I	5	7.1	14.3	23.3	96.9
	15	24.6	92.5	157.2	592.6
	50	199.1	631.1	1,861.6	2,531.4

For turbine RL estimates, we assume a simpler probabilistic model where the prediction error from the i th WT is assumed to follow a Weibull distribution with $\hat{\lambda}_i$ and ξ_i^λ as scale and shape parameters, respectively, such that the forecast variable $\lambda_i^* \sim \text{Weibull}(\hat{\lambda}_i, \xi_i^\lambda)$. The predictive mean of λ_i^* is denoted by $\bar{\lambda}_i = \hat{\lambda}_i \Gamma(1 + \frac{1}{\xi_i^\lambda})$.

4 Solution Procedure and Computational Efficiency

We solve STOCHOS using the Gurobi 9.0 solver with a Python interface in a server with two 14-core Intel Xeon CPUs with a base frequency of 2.60 GHz and 128 GB of memory, for a relative optimality gap of 0.1%. Table 2 shows the solution time (in seconds) as a function of the number of scenarios ($N_s = 25, 50, 100$, and 200) and the number of WTs ($N_I = 5, 15$ and 50). Even with 200 scenarios and 50 WTs, STOCHOS is solved in under 45 minutes. In practice, wind farm operators only need to run STOCHOS once per day, so they can use 200 scenarios (or even more), although our experiments suggest that the solution quality does not significantly change beyond 50 scenarios.

In actual operations, STOCHOS would be solved at the end of each working day to inform the day-ahead maintenance plan (Koltsidopoulos Papatzimos, 2020). The rolling

horizon procedure presented in Algorithm 1, explains the details of such implementation, starting from parameter and data input, to training the probabilistic forecasting models of Section 3.2, followed by the scenario generation step, in which N_s trajectories for the wind speed, wave height, and electricity price are generated, together with N_s turbine-specific RL prediction samples. Those scenarios are then used to determine scenario-specific power output predictions, accessibility estimations, and projected mission times. STOCHOS is then solved by minimizing the objective function in (1), returning an hourly day-ahead maintenance schedule ($m_{t,i}$ and r), and a daily schedule for the long-term horizon ($m_{d,i,s}^L$ and $r_{d,s}^L$). Once a solution has been obtained, the planning horizon is shifted by one day into the future. We call this as a single optimization “roll,” and the whole process is repeated until all maintenance tasks are completed. Finally, all STH schedules from all rolls are patched representing the executed hourly maintenance schedule for the whole planning horizon. This is the hourly schedule that is implemented by following STOCHOS’ solutions, and hence, is used to assess its performance relative to other benchmarks.

5 Experimental Results

This section starts with the experimental setup, followed by our results and analyses.

5.1 Experimental Setup

We run our optimization model for 100 experiments representing different weather conditions and electricity prices. For each experiment, Algorithm 1 is implemented, and a maintenance schedule is obtained. To be able to conduct extensive numerical evaluations, we set $N_S = 50$ scenarios and consider $N_T = 5$ WTs with a planning horizon of 20 days.

Data and forecasts used in this case study are described in detail in Section 3.1. True RLs (unknown to the planner ahead of time) are set at $\lambda_1 = 2.0$, $\lambda_2 = 6.8$, $\lambda_3 = 11.5$, $\lambda_4 = 16.2$, $\lambda_5 = 21.0$ days. Point predictions of those RLs (available to the planner) are set at $\hat{\lambda}_1 = 4.0$, $\hat{\lambda}_2 = 6.1$, $\hat{\lambda}_3 = 13.2$, $\hat{\lambda}_4 = 6.8$, $\hat{\lambda}_5 = 23.8$ days, reflecting different error magnitudes (low for WTs 2, 3, and 5, relatively high for WTs 1, and 4). Maintenance times are set at $\tau_1 = 11$, $\tau_2 = 5$, $\tau_3 = 6$, $\tau_4 = \tau_5 = 4$ hours. Maintenance times, as well as true and predicted RLs are set in light of the summary statistics for minor repairs

Algorithm 1 A rolling-based solution procedure for STOCHOS

```

1: Input set dimensions  $N_{\mathcal{I}}, N_{\mathcal{D}}, N_{\mathcal{S}} \rightarrow \#$  of WTs,  $\#$  of days, and  $\#$  of scenarios
2: Input parameters  $K, \Phi, \Psi, \Omega, Q, R, B, W, \tau_i, \nu_{max}, \eta_{max} \rightarrow$  operational parameters
3: Initialize  $\theta_i = 1 \quad \forall i \in \mathcal{I} \rightarrow$  parameter denoting maintenance requirement
4: Initialize  $\rho_i = 0, \zeta_i = 1 \quad \forall i \in \mathcal{I} \rightarrow$  parameters denoting WT operational status
5: Set the roll counter  $j = 0$ 
6: while  $\sum_{i \in \mathcal{I}} \theta_i > 0$  do
7:   Set  $\mathcal{T} = \{24 \cdot j, \dots, 24 \cdot (j + 1)\}; \mathcal{D} = \{j + 1, \dots, N_{\mathcal{D}} + j\}$ 
8:   Input observations and predictions for  $\nu(t), \eta(t), \kappa(t)$ , and  $\lambda(i) \quad \forall i$ .
9:   Train the STH probabilistic models of Section 3.2 for  $\nu(t), \eta(t)$ , and  $\kappa(t)$ .
10:  Train the LTH probabilistic models of Section 3.2 for  $\nu^L(d), \eta^L(d), \kappa^L(d), \lambda^L(i) \quad \forall i$ .
11:  Sample  $N_{\mathcal{S}}$  scenarios  $\{\nu_{t,s}\}_{t \in \mathcal{T}}^{s \in \mathcal{S}}, \{\eta_{t,s}\}_{t \in \mathcal{T}}^{s \in \mathcal{S}}$ , and  $\{\kappa_{t,s}\}_{t \in \mathcal{T}}^{s \in \mathcal{S}} \rightarrow$  STH scenarios for wind
    speed, wave height & price
12:  Sample  $N_{\mathcal{S}}$  scenarios  $\{\nu_{d,s}^L\}_{d \in \mathcal{D}}^{s \in \mathcal{S}}, \{\eta_{d,s}^L\}_{d \in \mathcal{D}}^{s \in \mathcal{S}}$ , and  $\{\kappa_{d,s}^L\}_{d \in \mathcal{D}}^{s \in \mathcal{S}} \rightarrow$  LTH scenarios for wind
    speed, wave height & price
13:  Sample  $N_{\mathcal{S}}$  scenarios  $\{\lambda_{i,s}^L\}_{i \in \mathcal{I}, s \in \mathcal{S}} \rightarrow$  scenarios for turbine RL
14:  Evaluate  $\{f_{t,s}\}_{t \in \mathcal{T}}^{s \in \mathcal{S}}$  and  $\{f_{d,s}^L\}_{d \in \mathcal{D}}^{s \in \mathcal{S}} \rightarrow$  Speed-to-power conversion; SM-1
15:  Evaluate  $\{A_{t,i,s}\}_{t \in \mathcal{T}}^{s \in \mathcal{S}}$  and  $\{A_{d,i,s}^L\}_{d \in \mathcal{D}}^{s \in \mathcal{S}} \rightarrow$  Computing STH and LTH mission times; SM-2
16:  Solve STOCHOS for  $N_{\mathcal{S}}$  scenarios
17:  Return the decisions for the  $j$ th roll:  $\mathcal{S}_j^{STH} = \{m_{t,i}, r\}$  and  $\mathcal{S}_j^{LTH} = \{m_{d,i,s}^L, r_{d,s}^L\}$ .
18:  Evaluate  $\mathcal{S}_j^{STH}$  under real operational parameters
19:  for  $i \in \mathcal{I}$  do
20:    if  $\sum_{t \in \mathcal{T}} m_{t,i} > 0$  then
21:      if  $b_i > 0$  then
22:        Set  $\rho \leftarrow 1; \tau_i \leftarrow b_i$  (maintenance resumed if unfinished in the STH)
23:      else
24:        Set  $\theta_i \leftarrow 0$  (maintenance no longer required if completed in the STH)
25:      end if
26:    end if
27:  end for
28:  Update roll counter  $j = j + 1$ 
29:  Update RLs & WT status  $\lambda_i = (\lambda_i - 1, 0)^+; \hat{\lambda}_i = (\hat{\lambda}_i - 1, 0)^+; \zeta_i = 1 - \mathbb{I}(\lambda_i = 0) \quad \forall i$ 
30: end while
31: return the final executed schedule by STOCHOS,  $\mathcal{S}^* = \{\mathcal{S}_1^{STH}, \dots, \mathcal{S}_{j-1}^{STH}\}$ 

```

Table 3: Operational parameters for our experimental setup

Notation	Parameter	Value
K	Cost of preventive maintenance (PM)	\$4,000
Φ	Cost of corrective maintenance (CM)	\$10,000
C_1, C_2	Spot contracting costs	\$1,000
B	Number of available crews	2 crews
Ψ	Crew hourly rate	\$250/hour
Q	Overtime hourly rate	\$125/hour
W	Maximum number of non-overtime work hours	8 hours
H	Maximum number of overtime work hours	8 hours
Ω	Daily vessel rental cost	\$2,500/day
ν_{max}	Wind speed safety threshold	15 m/s
η_{max}	Wave height safety threshold	1.8 m
t_R	Time of first light	6:00 am
t_D	Time of last sunlight	9:00 pm

reported by Dinwoodie et al. (2015); Carroll et al. (2016). Those, together with the real-world data about the wind speed, wave height, and electricity price, are used as input to the probabilistic forecasting and scenario generation framework described in Section 3.2.

Table 3 shows our selection for the remaining parameters: K and Φ are chosen following Yildirim et al. (2017), B , Ψ , and Q are selected based on the recommendations of Maples et al. (2013), while C_1 , C_2 are assigned large values relative to “regular” maintenance resource parameters, B and Ψ . Medium-sized crew transport vessels (CTVs), which are common for minor to medium repairs in OSW turbines, are used, wherein Ω , ν_{max} , and η_{max} are chosen according to the analyses in Dalgic et al. (2013); Anderberg (2015). Curtailment is assumed to be 0%, and hence, we set $C_{t,s} = 1 \forall t, s$.

5.2 Numerical Results

The performance of STOCHOS is compared against the following set of benchmarks:

1. Holistic Opportunistic Strategy with Perfect Knowledge (**PK-HOST**): A deterministic variant of STOCHOS which assumes perfect knowledge of all parameters: metocean data, electricity prices, and RL predictions. This benchmark assesses how far STOCHOS is from the best possible performance attained had the planner known, with full certainty, all environmental and operational conditions.
2. Holistic Opportunistic Strategy with Point Forecasts (**PF-HOST**): A deterministic version of STOCHOS which uses the raw point forecasts of all the uncertain parameters: RU-WRF for wind speed, WAVEWATCH III for wave height, LEAR for

electricity prices, and point RL predictions. This benchmark shows the value of uncertainty modeling in identifying the “right” windows of opportunity, and its relevance to fully harnessing the benefits of opportunistic maintenance grouping.

3. Condition-Based Strategy (**CBS**): A deterministic strategy that relies on point RL predictions, $\{\hat{\lambda}_i\}_{i \in \mathcal{I}}$, wherein maintenance is scheduled prior to the predicted RLs (with a three-day buffer time), or upon the occurrence of an unexpected failure.
4. Corrective Maintenance Strategy (**CMS**): Maintenance tasks are only performed reactively, or post-failure. Both CBS and CMS are prevalent benchmarks in the academic literature and practice, and hence are used herein to showcase the combined benefit of accounting for opportunities and parameter uncertainties.

Figure 5 illustrates the boxplots of the final costs (across 100 experiments) achieved by STOCHOS and the four benchmarks, namely: PK-HOST, PF-HOST, CBS, and CMS. Looking at Figure 5, we make few immediate observations: First, STOCHOS is the closest to PK-HOST with a relative cost difference (in terms of median costs) of 10.69%, while PF-HOST comes as a far second with a relative cost difference of 18.14%. This demonstrates how accounting for uncertainty makes the maintenance planner as close as possible to the “utopian” case of having perfect knowledge of information.

Second, STOCHOS improves upon PF-HOST by 6.30%. Another interesting observation is that the interquantile range (the difference between the third and first quantiles) of STOCHOS is 24.16% smaller than that of PF-HOST, suggesting that STOCHOS is not only better “on average,” but is also *consistently* better, and hence, a more reliable strategy. PF-HOST experiences multiple extreme cases where the performance is significantly worse than other strategies (e.g., the motivating example in Figure 1). This again confirms the value of accounting for uncertainty in terms of both average performance and reliability, relative to assuming full confidence in the available forecasts. Third, both STOCHOS and PF-HOST significantly outperform “non-opportunistic” approaches like CBS and CMS, further confirming the importance of opportunistic maintenance planning in OSW farms.

We summarize key O&M metrics in Table 4, including: number of vessel rentals, downtime, accessibility downtime (i.e. how much downtime is attributed to incorrect assessment of accessibility), production losses (in MWh) and revenue losses (in \$K). We also report




Figure 5: Boxplots of the total costs across 100 optimization experiments corresponding to different metocean conditions and electricity price profiles. Solid circles (and the numbers on top of them) denote median costs (in \$K). Numbers below the boxplots are the percentage improvements (%) of STOCHOS relative to each benchmark.

the number of PM tasks (higher the better) and CM tasks (lower the better), as well as the number of maintenance interruptions. Finally, we report the total maintenance cost, as well as the increase in cost relative to the optimal solution. Across all metrics, STOCHOS significantly outperforms all benchmarks (excluding the non-realistic case of PK-HOST).

5.3 The value of probabilistic forecasting

We would like to demonstrate the value of our probabilistic forecast and scenario generation approach presented in Section 3.2. To do so, we add two additional benchmarks:

1. Holistic Opportunistic Strategy with Calibrated Point Forecasts (**CPF-HOST**), which is similar to PF-HOST, except that it uses the predictive means of the probabilistic models as point predictions, instead of the raw point forecasts as in PF-HOST.
2. Stochastic Holistic Opportunistic Scheduler with Marginal Densities (**MD-STOCHOS**), which is similar to STOCHOS except that it uses scenarios generated from Gaussian distributions estimated via the residuals of the raw point forecasts of Section 3.1. This approach is prevalent in the classical stochastic optimization literature.

Table 4: O&M metrics. Bold-faced values denote best performance (PK-HOST excluded)

	PK-HOST	STOCHOS	PF-HOST	CBS	CMS
Number of vessel rentals	2.36	2.18	2.26	4.92	6.72
Total downtime [h]	39.52	44.20	47.30	97.20	142.03
Accessibility downtime [h]	9.52	14.20	17.30	67.20	112.03
Production loss [MWh]	101.21	141.05	154.59	597.68	1037.24
Revenue loss [\$K]	5.10	7.10	7.83	29.85	51.64
Total PM tasks	5.00	4.53	4.48	2.35	0.00
Total CM tasks	0.00	0.47	0.52	2.65	5.00
Maintenance interruptions	0.66	0.43	0.64	1.23	1.83
Avg. total cost [\$K]	38.92	43.43	45.00	86.90	127.91
Cost inc. from optimal [\$K]	0.10	4.60	6.14	48.07	89.08
Median total cost [\$K]	35.86	39.69	42.36	84.11	127.16

Table 5: The economic value of our probabilistic forecast framework.

	STOCHOS	PF-HOST	CPF-HOST	MD-STOCHOS
Median cost [\$K]	39.69	42.36	42.18	44.57
% Improvement [%]	-	6.30	5.91	10.95

Table 5 shows the median cost of those two benchmarks, in addition to that of STOCHOS and PF-HOST. The key interesting finding from this analysis is to show how an inadequate approach of handling uncertainty via marginal densities can lead to decisions that are even worse than deterministic strategies. Another interesting finding is that if the maintenance planner decides not to use the rich information in our probabilistic models (perhaps because of reluctance to change in business practice), then using the single-valued predictive means from our probabilistic models as input to the deterministic approach provides a modest but noticeable improvement in cost reduction. This is evident in how CPF-HOST improves upon PF-HOST. Finally, STOCHOS, with its adequate treatment of uncertainty and calibrated forecasts is able to achieve the maximal economic gain relative to both its deterministic and stochastic variants: PF-HOST, CPF-HOST, and MD-STOCHOS.

5.4 Sensitivity analysis and practical insights

To examine the effect of key operational parameters on the optimal solution, we conduct an experiment where we vary the number of WTs ($N_{\mathcal{I}} = 5, 15$ and 50), maintenance crews ($B = 1, 2$ and 3) and vessel costs ($\Omega = \$1,500, \$2,500$ and $\$3,500$). Figure 6 shows the results, where we fix all experimental parameters and data as those of Section 5.1 (except the parameter being varied, be it $N_{\mathcal{I}}$, B , or Ω). The only exception is that we use of $B = 4$




Figure 6: Costs resulting from the sensitivity analysis of the number of WTs (A), the number of available maintenance crews (B) and the transportation cost (C).

crews (instead of $B = 2$) only when we vary the number of WTs (Figure 6A), since larger number of WTs (i.e. maintenance tasks) naturally requires higher maintenance resources.

Figure 6A shows a nonlinear relationship between the number of WTs and the maintenance cost per turbine. Analyzing the resulting schedules in more depth provides a possible explanation: with 5 WTs and $B = 4$, the maintenance schedule is mainly dictated by the WT which is bound to fail first. In this case, WT1 with the shortest RL, is prioritized to avoid its failure. STOCHOS then finds it more economical to group the remaining tasks at the same day to share the setup cost already incurred by the maintenance of WT1, despite the production losses being relatively high. Going from 5 WTs to 15 WTs (with $B = 4$) creates more room for cost savings, as not all maintenance tasks can be scheduled in the same day, which incentivizes the search for a better opportunity. With $N_T = 50$ WTs, the schedule becomes too congested and some compromises must be made in order to complete all tasks on time and avoid the high costs of turbine failures.

Figure 6B shows that a too small maintenance budget (e.g., only having access to $B = 1$ crew), results in higher maintenance costs, mainly due to prolonged maintenance schedules. Gradually increasing B parallelizes the maintenance workload and allows the sharing of other on-site maintenance resources (e.g., vessel rentals). This results in more efficient (i.e. shorter) maintenance schedules, although at the expense of higher crew costs. Finally, Figure 6C shows a fairly linear relationship between the total cost and the vessel rental costs. While the number of vessel rentals may be reduced in response to the increase in vessel costs, it causes an increase in the revenue losses and repair costs, since transportation-based opportunities are now weighted more heavily than their revenue-based counterparts.

Before we conclude, we would like to furnish few insights for OSW practitioners. First,

one of our key findings in this work is the significance of accessibility information in OSW maintenance planning. Our research therefore calls upon both the research community and practitioners to develop more advanced, OSW-tailored access forecasting solutions (Gilbert et al., 2021), similar to those historically developed for classical maintenance parameters (e.g., asset degradation). Second, this work amplifies the recent calls in the academic community to shift away from deterministic to probabilistic forecasting and decision-theoretic frameworks (Pinson, 2013). This is evident by the substantial improvements attained by STOCHOS relative to deterministic strategies like PF-HOST or CBS which solely rely on single-valued predictions in making critical OSW maintenance decisions.

6 Conclusion

The unique challenges and uncertainties in OSW farms motivate the need for an offshore-tailored approach for maintenance optimization. To that end, we proposed STOCHOS, short for the stochastic holistic opportunistic scheduler—a maintenance scheduling approach that harnesses the maintenance opportunities arising due to favorable weather conditions, on-site maintenance resources, and maximal operating revenues, while adequately accounting for key operational and environmental uncertainties in the planning horizon. Our results showed that STOCHOS outperforms several prevalent benchmark strategies, across several key O&M metrics. Future work will look into two broad research questions: (1) Given the same optimization model, can we seek more powerful probabilistic representations that can achieve maximal economic gains?; and (2) Given the same forecasting model, how can we improve our optimization models to accommodate more intricate representation of the complex failure modes in wind turbines, as well as the impact of vessel routing and crew logistics in large-scale wind farms.

Supplementary Materials

SM-1 describes how to statistically obtain wind power estimates given hub-height wind speed forecasts. SM-2 details the estimation of mission times given accessibility information.

Data Availability Statement

Details to reproduce the results of this work are included in the reproducibility report. The codes and data are available on GitHub: <https://github.com/petros-pap/STOCHOS>.

Acknowledgements

This work is partially supported by the National Science Foundation (ECCS-2114422).

References

- Anderberg, C. (2015). Challenges of achieving a high accessibility in remote offshore wind farms. Master’s thesis. Chalmers University of Technology.
- Besnard, F. and L. Bertling (2010). An approach for condition-based maintenance optimization applied to wind turbine blades. *IEEE Transactions on Sustainable Energy* 1(2), 77–83.
- Besnard, F., M. Patriksson, A. Stromberg, A. Wojciechowski, and L. Bertling (2009). An optimization framework for opportunistic maintenance of offshore wind power system. In *2009 IEEE Bucharest PowerTech*, pp. 1–7. IEEE.
- Besnard, F., M. Patriksson, A. Strömberg, A. Wojciechowski, K. Fischer, and L. Bertling (2011). A stochastic model for opportunistic maintenance planning of offshore wind farms. In *2011 IEEE Trondheim PowerTech*, pp. 1–8.
- Bian, L., N. Gebraeel, and J. P. Kharoufeh (2015). Degradation modeling for real-time estimation of residual lifetimes in dynamic environments. *IIE Transactions* 47(5), 471–486.
- BOEM (2021). Lease and grant information. Bureau of Ocean Energy Management. Available at: <https://www.boem.gov/renewable-energy/lease-and-grant-information>.
- Byon, E., L. Ntaimo, and Y. Ding (2010). Optimal maintenance strategies for wind turbine systems under stochastic weather conditions. *IEEE Transactions on Reliability* 59(2), 393–404.
- Carroll, J., A. McDonald, I. Dinwoodie, D. McMillan, M. Revie, and I. Lazakis (2017). Availability, operation and maintenance costs of offshore wind turbines with different drive train configurations. *Wind Energy* 20(2), 361–378.
- Carroll, J., A. McDonald, and D. McMillan (2016). Failure rate, repair time and unscheduled o&m cost analysis of offshore wind turbines. *Wind Energy* 19(6), 1107–1119.
- Dalgic, Y., I. Lazakis, I. Dinwoodie, D. McMillan, and M. Revie (2015). Advanced logistics planning for offshore wind farm operation and maintenance activities. *Ocean Engineering* 101, 211–226.
- Dalgic, Y., I. Lazakis, and O. Turan (2013). Vessel charter rate estimation for offshore wind O&M activities. In *International Maritime Association of Mediterranean IMAM 2013*.

- Ding, F. and Z. Tian (2012). Opportunistic maintenance for wind farms considering multi-level imperfect maintenance thresholds. *Renewable Energy* 45, 175–182.
- Ding, Y. (2019). *Data Science for Wind Energy*. New York: Chapman and Hall/CRC.
- Dinwoodie, I., O. E. Endrerud, M. Hofmann, R. Martin, and I. B. Sperstad (2015). Reference cases for verification of operation and maintenance simulation models for offshore wind farms. *Wind Engineering* 39(1), 1–14.
- Ezzat, A., M. Jun, and Y. Ding (2018). Spatio-temporal asymmetry of local wind fields and its impact on short-term wind forecasting. *IEEE Transactions on Sustainable Energy* 9(3), 1437–1447.
- Ezzat, A., M. Jun, and Y. Ding (2019). Spatio-temporal short-term wind forecast: A calibrated regime-switching method. *The Annals of Applied Statistics* 13(3), 1484–1510.
- Fallahi, F., I. Bakir, M. Yildirim, and Z. Ye (2022). A chance-constrained optimization framework for wind farms to manage fleet-level availability in condition based maintenance and operations. *Renewable and Sustainable Energy Reviews* 168, 112789.
- Gilbert, C., J. Browell, and D. McMillan (2021). Probabilistic access forecasting for improved offshore operations. *International Journal of Forecasting* 37(1), 134–150.
- Golparvar, B., P. Papadopoulos, A. Ezzat, and R. Wang (2021). A surrogate-model-based approach for estimating the first and second-order moments of offshore wind power. *Applied Energy* 299, 117286.
- IEA (2019). Offshore wind outlook 2019. Technical report, International Energy Agency, Paris, France. Available at: <https://www.iea.org/reports/offshore-wind-outlook-2019>.
- Ko, Y. M. and E. Byon (2017). Condition-based joint maintenance optimization for a large-scale system with homogeneous units. *IIE Transactions* 49(5), 493–504.
- Koltsidopoulos Papatzimos, A. (2020). *Data-driven Operations & Maintenance for Offshore Wind Farms: Tools and Methodologies*. Ph. D. thesis, University of Exeter.
- Lago, J., G. Marcjasz, B. De Schutter, and R. Weron (2021). Forecasting day-ahead electricity prices: A review of state-of-the-art algorithms, best practices and an open-access benchmark. *Applied Energy* 293, 116983.
- Lu, Y., L. Sun, J. Kang, H. Sun, and X. Zhang (2017). Opportunistic maintenance optimization for offshore wind turbine electrical and electronic system based on rolling horizon approach. *Journal of Renewable and Sustainable Energy* 9(3), 033307.
- Lubing, X., R. Xiaoming, L. Shuai, and H. Xin (2019). An opportunistic maintenance strategy for offshore wind turbine based on accessibility evaluation. *Wind Engineering* 44(5), 455–468.
- Ma, Y., W. Zhang, and J. Branke (2022). Multi-objective optimisation of multifaceted maintenance strategies for wind farms. *Journal of the Operational Research Society*, 1–16.
- Maples, B., G. Saur, M. Hand, R. Van De Pietermen, and T. Obdam (2013). Installation, operation, and maintenance strategies to reduce the cost of offshore wind energy. Technical report, NREL, Golden, CO, USA.

- Mazidi, P., Y. Tohidi, and M. A. Sanz-Bobi (2017). Strategic maintenance scheduling of an offshore wind farm in a deregulated power system. *Energies* 10(3), 313.
- NYSERDA (2021). NYSERDA floating LiDAR buoy data. New York State Energy Research & Development Agency. Available at: <https://oswbuoysny.resourcepanorama.dnvg1.com>.
- Optis, M., A. Kumler, G. N. Scott, M. C. Debnath, and P. J. Moriarty (2020). Validation of RU-WRF, the custom atmospheric mesoscale model of the Rutgers Center for Ocean Observing Leadership. Technical report, NREL, Golden, CO, USA.
- Papadopoulos, P., D. W. Coit, and A. A. Ezzat (2022). Seizing opportunity: Maintenance optimization in offshore wind farms considering accessibility, production, and crew dispatch. *IEEE Transactions on Sustainable Energy* 13(1), 111–121.
- Perez, E., L. Ntamo, and Y. Ding (2015). Multi-component wind turbine modeling and simulation for wind farm operations and maintenance. *Simulation* 91(4), 360–382.
- Pinson, P. (2013). Wind energy: Forecasting challenges for its operational management. *Statistical Science* 28(4), 564–585.
- PJM (2021). PJM data directory. PJM Interconnection LLC. Available at: https://dataminer2.pjm.com/feed/da_hrl_lmps/definition.
- Ren, Z., A. S. Verma, Y. Li, J. J. Teuwen, and Z. Jiang (2021). Offshore wind turbine operations and maintenance: A state-of-the-art review. *Renewable and Sustainable Energy Reviews* 144, 110886.
- Sarker, B. and T. Faiz (2016). Minimizing maintenance cost for offshore wind turbines following multi-level opportunistic preventive strategy. *Renewable Energy* 85, 104–113.
- Shafiee, M., M. Finkelstein, and C. Bérenguer (2015). An opportunistic condition-based maintenance policy for offshore wind turbine blades subjected to degradation and environmental shocks. *Reliability Engineering & System Safety* 142, 463–471.
- Song, S., Q. Li, F. A. Felder, H. Wang, and D. W. Coit (2018). Integrated optimization of offshore wind farm layout design and turbine opportunistic condition-based maintenance. *Computers & Industrial Engineering* 120, 288–297.
- Stehly, T. J. and P. C. Beiter (2020). 2018 cost of wind energy review. Technical report, National Renewable Energy Lab (NREL), Golden, CO, USA.
- Taylor, J. and J. Jeon (2018). Probabilistic forecasting of wave height for offshore wind turbine maintenance. *European Journal of Operational Research* 267(3), 877–890.
- Wang, J., X. Zhao, and X. Guo (2019). Optimizing wind turbine’s maintenance policies under performance-based contract. *Renewable Energy* 135, 626–634.
- Yang, L., G. Li, Z. Zhang, and X. Ma (2020). Operations & maintenance optimization of wind turbines integrating wind and aging information. *IEEE Transactions on Sustainable Energy* 12(1), 211 – 221.
- Yildirim, M., N. Gebraeel, and A. Sun (2017). Integrated predictive analytics and optimization for opportunistic maintenance and operations in wind farms. *IEEE Transactions on Power Systems* 32(6), 4319–4328.
- Zhang, H., J. Yan, S. Han, L. Li, Y. Liu, and D. Infield (2021). Uncertain accessibility estimation method for offshore wind farm based on multi-step probabilistic wave forecasting. *IET Renewable Power Generation* 15(13), 2944–2955.