Estimating Maintenance Cost of Offshore Substations under Uncertainty: a Markovian Approach

Solène Delannoy-Pavy
Cyrille Vessaire
Manuel Ruiz
RTE (Réseau de Transport d'Electricité), France
{solene.delannoypavy, cyrille.vessaire, manuel.ruiz}@rte-france.com

Vincent Leclère
Axel Parmentier
Ecole des Ponts ParisTech, France
{axel.parmentier, vincent.leclere}@enpc.fr

***Abstract*—Offshore wind is growing in France as part of the green transition. The French Transmission System Operator (RTE) will own and maintain the offshore substations linking wind farms to the onshore grid. Since site access depends on weather and downtime is costly, maintenance planning must balance cost and risk. Penalties scale with curtailed energy, except on a limited number of pre-declared free maintenance days scheduled in advance without weather forecasts. We model the maintenance optimization problem as a Markov Decision Process (MDP). Exact dynamic programming becomes infeasible beyond a few components, so we use weakly coupled MDP methods to derive maintenance strategies. The approach is applied to a case study, using historical weather data to generate scenarios for the future Centre Manche 1 wind farm. As these assets are new and failure data are limited, degradation models remain uncertain; our analysis shows that this uncertainty significantly impacts cost estimates.**

***Index Terms*—Markov Decision Process, Multistage stochastic optimization, Offshore substation maintenance, Weakly coupled dynamic programs**

I. Introduction

The global transition to renewable energy has positioned offshore wind power as a cornerstone of future energy systems, with European targets reaching hundreds of gigawatts by 2050. France's offshore wind development strategy is particularly ambitious, aiming to increase installed capacity from the current 1.5 GW to 45 GW by 2050 [1]. This rapid scaling presents unprecedented challenges for Transmission System Operators (TSOs), who must ensure reliable grid integration while managing substantial operational and financial risks.

The French TSO, RTE, faces unique challenges in connecting these offshore wind farms to the mainland grid. Offshore electrical substations, which collect, transform, and transmit electricity from multiple wind turbines to the onshore grid, represent critical single points of failure. Stations are radially connected to land, without the mesh redundancy present in the rest of the network. Any failure directly impacts wind farm production and triggers substantial financial penalties. Offshore substation maintenance differs from onshore maintenance. Weather-dependent accessibility restricts maintenance operations to periods of favorable conditions, often requiring several consecutive days of calm sea. The novelty of offshore technology means that reliability data are limited to manufacturer-provided Mean Time Between Failures (MTBF) estimates, creating significant uncertainty in degradation modeling. Moreover, the harsh marine environment stresses components in ways that differ from onshore conditions, making historical onshore reliability data inadequate for offshore applications.

Contractual arrangements typically impose penalties proportional to curtailed energy when substations are unavailable, except during pre-declared free maintenance days that are allocated through limited annual quotas. This penalty structure creates a complex optimization problem: maintenance must be scheduled to minimize both failure-induced outages and planned downtime, while accounting for weather uncertainty and the strategic allocation of maintenance quotas. The economic implications are substantial—substation failures can result in penalties worth millions of euros per day, making optimal maintenance planning critical for the TSO.

TSOs must make irreversible strategic investment decisions now—such as substation design selection, redundancy levels, and component specifications—that will determine maintenance costs for the 30-year asset lifetime. For instance, designs incorporating duplicate transformers or enhanced cooling systems can significantly reduce downtime probability and maintenance duration, but at increased capital cost. Quantifying the expected maintenance penalties associated with different design choices is therefore essential for informed investment decisions. This cost reflects the value of well-designed maintenance strategies. However, identifying good maintenance strategies is particularly challenging for offshore substations due to limited operational experience and high uncertainty, as previously discussed. To address this, we propose formulating and solving an optimization problem. While the

Submitted to the 24th Power Systems Computation Conference (PSCC 2026).

model does not need to capture all operational details—given that maintenance strategies can be progressively refined over time—it must adequately represent the key factors that drive major maintenance cost impacts.

A. Literature Review

Maintenance optimization for wind energy systems has received considerable attention, with established literature on onshore [2], [3] and offshore [4] wind turbines. However, these studies focus on individual turbine maintenance rather than the centralized substation infrastructure that connects entire wind farms. The specific challenges of offshore substation maintenance, including weather-dependent accessibility, combined with penalty-driven cost structures and quota-based maintenance planning—remain largely unexplored in the academic literature.

This paper addresses this gap by proposing a comprehensive decision support methodology for offshore substation maintenance planning under uncertainty. Our approach is based on a Markov Decision Process (MDP) formulation [5], a well-established framework for sequential decision-making under uncertainty that has proven effective in maintenance optimization [6]. The MDP framework naturally accommodates the stochastic nature of component degradation, weather patterns, and the multistage decision structure inherent in maintenance planning. The state of the substation is defined by the joint states of its individual components. A major challenge arises from the exponential growth of the state space with the number of components, making it computationally intractable to solve the MDP exactly using dynamic programming. This can be addressed using approximate dynamic programming [7]. Specific algorithms have been proposed for weakly coupled [8], [9] and decomposable MDPs [10], which model multi-component systems requiring coordinated or global actions. We explore the fluid approximation introduced in [10], a solution method designed for decomposable MDPs that yields high-quality solutions.

B. Contributions

The key contributions of this work are threefold: i) A novel multi-horizon MDP formulation that captures the temporal structure of offshore maintenance planning, including operational maintenance and strategic quota allocation; ii) A fluid-approximation-based policy and numerical results showing its performance; iii) A sensitivity analysis with respect to the choice of the degradation model.

C. Paper Structure

The remainder of this paper is organized as follows. Section II introduces the optimization model. Section III presents a fluid-approximation-based policy for addressing it. Section IV reports computational results obtained from a medium-scale substation model and analyzes the sensitivity of cost estimates to uncertainties in the degradation model. Finally, Section V concludes the paper and outlines potential directions for future research.

II. Optimization model

A. Problem statement

We formulate the offshore substation maintenance optimization problem as a finite-horizon MDP with time-dependent components. The MDP incorporates the following key modeling assumptions:

Assumption 1: Time is discretized into two-month periods. To benefit from free maintenance days in a two-month period, the TSO must declare them in advance, at the start of that given period. Therefore, we assume that maintenance is scheduled at the beginning of each period.

Assumption 2: Both the wind park's production and substation accessibility are weather-dependent.

Assumption 3: Maintenance is planned without consideration of weather forecasts, reflecting the fact that forecasts beyond two weeks are not available.

Assumption 4: Component replacement may span multiple days, during which the component remains powered down from the planned maintenance start date until its completion.

Assumption 5: Maintenance operations may fall behind schedule when accessibility is poor, which can cause additional costs. However, all maintenance operations scheduled at the start of a two-month period must be completed within that same period.

To model maintenance scheduling in line with Assumption 1, we introduce *scheduling stages* of two months. We assume that the substation's degradation state is fully observable. In our formulation, the system state at the start of each scheduling stage consists of the degradation state of each component of the substation and the number of remaining free maintenance days. An action for a scheduling stage specifies both the maintenance tasks scheduled during the stage and the set of pre-declared free maintenance days for that stage. The objective is to minimize the long-run expected cost produced by the maintenance scheduling strategy. The system dynamics and stage costs induced by an action are summarized in Fig. 1; further details follow in the next paragraphs.

B. Time, state and action space

A timestep T in the MDP corresponds to a scheduling stage. We denote by $\overline{\mathrm{T}}$ the total number of scheduling stages considered. Since producer contracts typically last 30 years and each year contains 6 scheduling stages (two months each), we set $\overline{\mathrm{T}} = 180$ scheduling stages. The degradation of the substation and the weather are modeled with a *daily time step* t. For notational clarity, we assume that each scheduling stage has the same number of time steps which we denote by $\bar{\mathrm{t}}$. We write $\mathcal{T}_T = [\bar{\mathrm{t}}T : \bar{\mathrm{t}}(T+1) - 1]$ the set of daily time steps associated with scheduling stage T.

State. The substation is divided into a set $\mathcal{C}$ of components such that the degradations of two components are independent. A component may be composed of several submodules; for example, the "converter" component is composed of several thousand "thyristor" submodules. $\mathcal{C}$ depends on the design of the substation. Only components whose degradation can

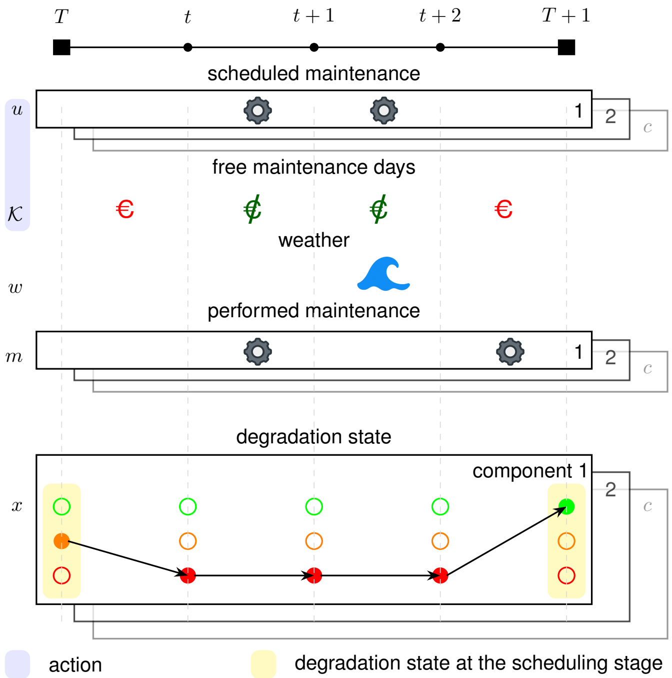

Fig. 1: Dynamics between two scheduling stages.

The first two rows represent the maintenance schedule by component and the free maintenance days decided at the beginning of the scheduling stage. The middle row shows the performed maintenance that is affected by the weather. The bottom row shows the dynamics of the degradation state.

impact transmission capacity are included—typical examples for an HVDC (High Voltage Direct Current) substation are transformers, converters, converter cooling systems, air conditioning units for the conversion hall, and busbars. We represent the overall condition of the substation by a *joint degradation state* x, which captures the degradation levels of all relevant components. Each component $c \in \mathcal{C}$ is modeled independently using a discrete degradation scale with n_c possible states, ranging from state 1 (new) to state n_c (failure). For example, in Fig. 1, component 1 has $n_1 = 3$ states: state 1 (green), state 2 (orange), and state 3 (red). The joint degradation state $x = (x_c)_{c \in \mathcal{C}} \in \mathcal{X}$, where $\mathcal{X} = \bigotimes_{c \in \mathcal{C}} [1 : n_c]$, is then defined as the cartesian product of the individual component states. This makes a total of $n = \prod_{c \in \mathcal{C}} n_c$ possible joint states. We denote by k the *number of free maintenance days remaining* at the beginning of the scheduling stage. The *state* of the MDP is the pair (x, k).

Action. At the start of each scheduling stage, the TSO selects which maintenance tasks to schedule during that period and the planned start dates for those tasks. For notational simplicity, we focus on full repairs only, although the model can be extended to allow partial repairs. We enforce no workforce capacity constraint; multiple components may be maintained simultaneously. The model could also be adapted to account for these constraints. For each daily time step $t \in \mathcal{T}_T$ and component $c \in \mathcal{C}$, we define $u_{t,c} = 1$ if a maintenance on component c is scheduled to begin on day t, and 0 otherwise. The *scheduled maintenance* for stage T is the set of binary variables

$$\mathbf{u_T} = (u_{t,c})_{t \in \mathcal{T}_T, c \in \mathcal{C}} \in \{0,1\}^{\mathcal{T}_T \times \mathcal{C}}. \tag{1}$$

Additionally, the TSO must schedule a set of *free maintenance days* $\mathcal{K}_T \in \mathcal{P}(\mathcal{T}_T)$ on stage T, where $\mathcal{P}(\mathcal{T}_T)$ is the power set of $\mathcal{T}_T$. The *action* of the MDP at scheduling stage T is the pair $(\mathbf{u_T}, \mathcal{K}_T) \in \{0,1\}^{\mathcal{T}_T \times \mathcal{C}} \times \mathcal{P}(\mathcal{T}_T)$.

C. Dynamics between two scheduling stages

Let T be a scheduling stage and k_T be the number of free maintenance days remaining at scheduling stage T. Let K denote the *annual quota of free maintenance days* that the TSO is contractually allowed to use. The *Quota dynamics* is

$$k_{T+1} = \begin{cases} K & \text{if } T+1 \equiv 0 \pmod 6, \\ k_T - |\mathcal{K}_T| & \text{otherwise.} \end{cases} \tag{2}$$

This constraint ensures that k_T is reset to K every year, i.e., every 6 scheduling stages. To ensure that the contractual quota is not exceeded, we add the domain constraint $k_T \geq 0$. We introduce a finite *set of weather scenarios* $\mathcal{W}_T$. A weather scenario $w \in \mathcal{W}_T$ is a time series of length $\bar{t}$ describing two correlated daily variables: *accessibility* and *production*. At daily time step $t \in \mathcal{T}_T$ under scenario w, the substation's accessibility is represented by $h_t^w \in \{0,1\}$ where $h_t^w = 1$ indicates that the substation is accessible at day t (and 0 otherwise). This parameter is obtained by thresholding the wave height: if the daily wave height exceeds the threshold, the substation is considered inaccessible. The production at daily time step t under scenario w is denoted by $p_t^w \in \mathbb{R}^+$. Following Assumption 3, maintenance is planned without any knowledge of accessibility. Thus, the scheduled maintenance introduced in (1) does not depend on the weather. But the weather influences the actual end date of maintenance. We introduce the *performed maintenance* at time t in weather scenario w on component c, it is a binary variable $m_{t,c}^w \in \{0,1\}$ that equals 1 if c is maintained at t and 0 otherwise. If weather conditions are unfavorable, an ongoing maintenance operation may be paused until the site becomes accessible again. Under Assumption 5, however, any maintenance that has been scheduled at the start of the period cannot be canceled. To minimize costs, work is resumed and continued as soon as access permits. On Fig. 1, a maintenance on component 1 has been scheduled at daily time step t. The scheduled maintenance requires two workdays on the substation, and two free maintenance days have been scheduled. Maintenance starts on day t as scheduled. But the substation is not accessible on day $t+1$, so the maintenance is postponed to day $t+2$. The TSO pays a penalty to maintain the substation on day $t+2$. To represent this in the model, we introduce a set of *counters* $q_t^w = (q_{t,c}^w)_{c \in \mathcal{C}}$. The counter $q_{t,c}^w \in \mathbb{N}$ tracks the number of maintenance days remaining for component c at daily time step t under scenario w (with $q_{t,c}^w = 0$ when no maintenance is in progress). Let $d_c \in \mathbb{N}$ be the *replacement time* of component c (in workdays). We now state the *Accessibility constraints*. The dynamics of $q^w_{t,c}$ for all $c \in \mathcal{C}$ and $w \in \mathcal{W}$ is

$$q^w_{t+1,c} = q^w_{t,c} - m^w_{t,c} + d_c u^w_{t+1,c} \quad \forall t \in \mathcal{T}_T, \tag{3a}$$
$$q^w_{\bar{\mathrm{t}}T,c} = d_c u^w_{\bar{\mathrm{t}}T,c}. \tag{3b}$$

At the start of a scheduled maintenance, the counter is increased by the required number of workdays and reduced by one after each maintenance day performed. The maintenance is performed as soon as the substation is accessible, thus $m^w_{t,c}$ satisfies the relation

$$m^w_{t,c} = \begin{cases} h^w_t & \text{if } q^w_{t,c} \geq 1, \\ 0 & \text{otherwise.} \end{cases} \quad \forall t \in \mathcal{T}_T, \quad \forall c \in \mathcal{C} \tag{4}$$

A *margin* M is imposed to prevent scheduling maintenance during the final days of the stage, ensuring completion beforehand

$$u^w_{t,c} = 0 \quad \forall t \in [\bar{\mathrm{t}}(T+1) - M : \bar{\mathrm{t}}(T+1) - 1]. \tag{5}$$

Degradation state dynamics are modeled at the component level, as the degradations of different components are assumed independent. For each component $c \in \mathcal{C}$, let $p_c(x_c, x'_c)$ denote the probability of transitioning from state x_c to state x'_c without maintenance. We assume that a component cannot spontaneously improve its condition, so $p_c(x_c, x'_c) = 0$ if $x'_c < x_c$. The transitions of a component's degradation state depends on the maintenance performed. For component c, the *transition matrix* associated with a given performed maintenance $m \in \{0; 1\}$ is a right stochastic matrix $P^c(m) \in [0, 1]^{n_c \times n_c}$. Applying no maintenance yields random transitions while performing a full repair restores the component to the "new" state, thus we have

$$P^c(0)_{x_c, x'_c} = p_c(x_c, x'_c), \tag{6a}$$
$$P^c(1)_{x_c, x'_c} = \mathbb{1}_{\{x'_c = 1\}}. \tag{6b}$$

Note that the performed maintenance is weather-dependent, thus the degradation state evolution over the scheduling stage is weather-dependent. Consequently, we denote x^w_t the degradation state at daily time step t in weather scenario w.

D. Objective and Bellman equation

The objective of the problem is to minimize expected penalties that depend on the production of the wind park, substation capacity, and free maintenance days. Following Assumption 4, the *capacity of the substation* is a function $C : \mathcal{X} \times \mathbb{N} \longrightarrow \mathbb{R}^+$ of the degradation state of the substation and the counter tracking the number of maintenance days still to be performed. This function depends on the substation design. For example, if the transformers are redundant, a maintenance or a failure in one of the transformers does not cause a capacity drop. A fuller example is given in Annex A. For the sake of clarity, this paper assumes that the model objective is equal to the penalties due to capacity loss. The model could easily be extended to include other costs (*e.g.,* components, maintenance, or transport costs). Further, we assume that the penalty cost per unit of curtailed production is 1. More precisely, the cost over one daily time step t of scheduling stage T in weather scenario w is

$$\mathbb{1}_{\{t \notin \mathcal{K}_T\}} \left(p^w_t - C(x^w_t, q^w_t)\right)^+. \tag{7}$$

The indicator ensures that the cost is 0 on free maintenance days. The second factor ensures that the cost is proportional to the curtailed power otherwise. Let us denote V_T the *value function* at scheduling stage T. $V_T(x_{\bar{\mathrm{t}}T}, k_T)$ represents the minimum expected cost from scheduling stage T to the end of the horizon, starting from state $(x_{\bar{\mathrm{t}}T}, k_T)$. $V_{\overline{\mathrm{T}}}$ equals 0 for each state, and the value function satisfies the Bellman equation

$$V_T(x_{\bar{\mathrm{t}}T}, k_T) = \min_{\mathbf{u_T}, \mathcal{K}_T} \quad \frac{1}{|\mathcal{W}_T|} \sum_{w \in \mathcal{W}_T} \mathbb{E}\Bigg[\sum_{t \in \mathcal{T}_T} \mathbb{1}_{\{t \notin \mathcal{K}_T\}} \times \left(p^w_t - C(x^w_t, q^w_t)\right)^+ + V_{T+1}\left(x^w_{\bar{\mathrm{t}}(T+1)}, k_{T+1}\right) \Bigg] \tag{8}$$

$$\text{subject to} \quad \text{Quota dynamics (2)}$$
$$\text{Accessibility contraints (3), (4), (5).}$$

The expectation is over the stochastic degradation state dynamics, which start deterministically at $x_{\bar{\mathrm{t}}T}$ in all weather scenarios; transition probabilities in weather scenario w then depend on the performed maintenance in w. We assume that the system starts in the deterministic state where $x = \mathbf{1}_{\mathcal{C}}$ and $k = K$, where $\mathbf{1}_{\mathcal{C}}$ is the vector of length $|\mathcal{C}|$ with all components equal to one. This means that all components are new and no free maintenance days have yet been used. $V_0(\mathbf{1}_{\mathcal{C}}, K)$ is an estimation of the maintenance cost. We emphasize that the maintenance of the components is coupled due to the cost function and free maintenance days.

III. Policy based on fluid approximation

We define a *policy* π as a sequence of decision rules $\pi_0, \pi_1, \ldots, \pi_{\overline{\mathrm{T}}-1}$, where each π_T is a mapping from states to actions. Cost estimation requires computing the *optimal policy* π^*. The actions associated with π^* are those that attain the minimum in Bellman equation (8); that could be rewritten as a Mixed-Integer Linear Program (MILP) and solved with a standard MILP solver. However, the exponential size of the state space prevents solving the problem using Dynamic Programming (DP) when modeling realistic substations. Any MDP can be formulated as a Linear Program (LP), whose solution from a given initial state yields an optimal action. Unfortunately, since the variables encode the probability of being in a given state, the size of the LP is exponential in the number of components. Like for DP, this makes it intractable in our case. Fortunately, our MDP is structured as a *weakly coupled MDP*. This can be exploited to build an outer approximation of the polytope of reachable probabilities given feasible policies, which is of tractable size. Based on it, we design a heuristic: for any given state, we build the approximation, solve it, and take the first decision corresponding to the solution. The rest of the section describes the approximation.

Following [10], we introduce a *fluid approximation* of the problem at the current state $(\mathbf{x}, \mathbf{k})$ and scheduling stage T. It approximates the MDP at stages $H \in [T : T + \overline{\mathrm{H}}]$ where $\overline{\mathrm{H}}$ is a given horizon. Typically, we consider 6 stages for one year. We denote the *action space* at the scheduling stage H by $\mathcal{A}(H) = \{0,1\}^{\mathcal{T}_H \times \mathcal{C}} \times \mathcal{P}(\mathcal{T}_H)$. Notably, the number of variables in the fluid approximation grows only linearly with the number of components $|\mathcal{C}|$. While the dynamics induced by actions are decomposable by component, the cost is not. We therefore introduce a decomposable surrogate cost function that approximates the true cost.

Decomposable surrogate cost. We associate a *subsystem* with each component c of the substation. The *subsystem state* is the pair (x_c, k) and we denote by $\mathcal{S}^c = [1 : n_c] \times [0 : K]$ the *subsystem state space*. We define a surrogate cost function that is additive in subsystem states. For each component $c \in \mathcal{C}$ and subsystem state $s = (x_c, k) \in \mathcal{S}^c$, let $\ell^c_{s,a}(H)$ denote the *subsystem surrogate cost*. This is the expected cost over the scheduling stage H when action a is taken, component c is in degradation state x_c, and all other components are new at the beginning of the stage, averaged over all weather scenarios. This cost is precomputed. The approximate cost resulting from taking action a is $\sum_{c \in \mathcal{C}} \ell^c_{s,a}(H)$.

Marginal probability flow. Such as in [10], the decision variables correspond to the marginal probabilities of each subsystem being in each of its possible states while each possible action is taken at each stage. Let $\nu^c_{s,a}(H)$ be the probability that at scheduling stage $H \in [T : T + \overline{\mathrm{H}}]$ component $c \in \mathcal{C}$ is in state $s \in \mathcal{S}^c$ and action $a \in \mathcal{A}(H)$ is chosen. Let $p^c_{s,s',a}(H)$ denote the probability that component c transitions from state s to state s' between the scheduling stages H and $H+1$ when action $a \in \mathcal{A}(H)$ is taken. For all $c \in \mathcal{C}$, $H \in [T : T + \overline{\mathrm{H}} - 1]$, $s' \in \mathcal{S}^c$, we express the temporal evolution of these probabilities through *Flow constraints* inspired by [10], where p is a parameter and ν a variable:

$$\sum_{a \in \mathcal{A}(H+1)} \nu^c_{s',a}(H+1) = \sum_{s \in \mathcal{S}^c} \sum_{a \in \mathcal{A}(H)} p^c_{s,s',a}(H) \nu^c_{s,a}(H). \tag{9}$$

Let $A_a(H)$ be a variable representing the probability that action $a \in \mathcal{A}(H)$ is taken at time H. We have the *Consistency constraints* from [10]

$$\sum_{s \in \mathcal{S}^c} \nu^c_{s,a}(H) = A_a(H) \quad \forall H \in [T : T + \overline{\mathrm{H}}],\ c \in \mathcal{C}, \quad \forall a \in \mathcal{A}(H). \tag{10}$$

The following constraint from [10] ensures that the degradation state at scheduling stage T is $\mathbf{x}$ and the number of free maintenance days remaining is $\mathbf{k}$,

$$\sum_{a \in \mathcal{A}} \nu^c_{s,a}(T) = \mathbb{1}_{s=(\mathbf{x_c},\mathbf{k})} \quad \forall c \in \mathcal{C}, s \in \mathcal{S}^c. \tag{11}$$

We also introduce $\mu_k(H)$, the probability that the number of free maintenance days remaining at scheduling stage H is k, and enforce the *Non-negativity of probabilities constraints* from [10]

$$\begin{aligned}
\nu^c_{s,a}(H) \geq 0 \qquad & \forall H \in [T : T + \overline{\mathrm{H}}],\ c \in \mathcal{C},\ s \in \mathcal{S}^c, \quad \forall a \in \mathcal{A}(H) && \text{(12a)}\\
A_a(H) \geq 0 \qquad & \forall H \in [T : T + \overline{\mathrm{H}}],\ a \in \mathcal{A}(H) && \text{(12b)}\\
\mu_k(H) \geq 0 \qquad & \forall H \in [T : T + \overline{\mathrm{H}}],\ k \in [0 : K]. && \text{(12c)}
\end{aligned}$$

Finally, we add two *Additional constraints* coupling the subsystems, which are specific to our problem. The first one is the consistency constraint

$$\sum_{\substack{x_c \in [1:n_c] \\ a \in \mathcal{A}(H)}} \nu^c_{(x_c,k),a}(H) = \mu_k(H) \quad \forall H \in [T : T + \overline{\mathrm{H}}] \quad \forall c \in \mathcal{C},\ k \in [0 : K]. \tag{13}$$

The second one prevents using more free maintenance days than are available:

$$\nu^c_{(x_c,k),(u,\mathcal{K})}(H) = 0, \quad \forall H \in [T : T + \overline{\mathrm{H}}],\ c \in \mathcal{C}, \quad \forall x_c \in [1 : n_c],\ k \in [0 : K], \quad \forall u \in \{0,1\}^{\mathcal{T}_H \times \mathcal{C}}, \mathcal{K} \in \mathcal{P}(\mathcal{T}_H),\ |\mathcal{K}| > k. \tag{14}$$

Fluid approximation. The fluid approximation is then formulated as a linear program:

$$\underset{\nu, A, \mu}{\text{minimize}} \sum_{H \in [T:T+\overline{\mathrm{H}}]} \sum_{c \in \mathcal{C}} \sum_{s \in \mathcal{S}^c} \sum_{a \in \mathcal{A}(H)} \ell^c_{s,a}(H) \nu^c_{s,a}(H) \tag{15}$$

$$\begin{aligned}
\text{subject to} \quad & \text{Flow constraints (9)}\\
& \text{Consistency constraints (10), (11)}\\
& \text{Non-negativity of probabilities (12)}\\
& \text{Additional constraints (13), (14).}
\end{aligned}$$

Fluid policy. The fluid heuristic policy takes as parameters the surrogate subsystem costs ℓ, transition probabilities p and the fluid horizon $\overline{\mathrm{H}}$. Given a scheduling stage T, we build and solve the problem (15) corresponding to the current state $(\mathbf{x}, \mathbf{k})$ with parameters ℓ, p and H. Let ν^*, A^*, μ^* be an optimal solution of this problem. The fluid policy selects the action with the highest probability in the fluid approximation for the next scheduling stage: $\pi_T(\mathbf{x}, \mathbf{k}) = \arg\max_{a \in \mathcal{A}} A^*_a(T)$, thus projecting back from a fluid solution to a discrete policy.

IV. Use case analysis

A. Case study

To evaluate the practical relevance of the proposed method, we apply it to a representative RTE offshore substation corresponding to the upcoming *Centre Manche 1* project found in [12]. The case study focuses on a monopole HVDC offshore substation represented with a total of seven components. Replacement times vary by component, and an annual quota of $Q = 6$ free maintenance days is assumed (equal to the maximum replacement time). We use standard MTBF values for onshore components. A detailed description of the substation model and replacement times is provided in Annex A. Accessibility scenarios are derived from 80 years of hourly wave height data from Copernicus [11]. We compute daily averages and apply a 1.0-meter wave height threshold to determine accessibility, yielding one scenario per scheduling stage each year. These scenarios vary by month but not by year. While production scenarios could theoretically be derived from wind speed data, this requires knowledge of turbine types, which we lack. Since accessibility governs maintenance feasibility on free maintenance days, it is more critical than production. For simplicity, we assume constant production within each scheduling stage, proportional to the average load factor of French offshore wind farms during the corresponding two months in 2024 [13]. Combining these assumptions, we construct a set $\mathcal{W}_T$ of 80 weather scenarios per scheduling stage T, where production depends on T but not on the weather scenario. Consider a policy π such as the fluid policy presented in the previous section. We compute the cost associated with π using Monte Carlo simulations. Several trajectories are simulated under the policy, and the average cost is returned. The procedure is as follows.

- 1. Set the current scheduling stage T to 1 and the current state to $(\mathbb{1}_{\mathcal{C}}, K)$.
- 2. Randomly select a weather scenario $w \in \mathcal{W}_T$.
- 3. Choose the action a for the current state and stage according to π.
- 4. Simulate state transitions under w and a throughout the scheduling stage and calculate the costs over daily time steps. Set the current stage to the value of the state at the end of the scheduling stage and T to $T + 1$.
- 5. Repeat steps 2–4 while $T < 180$ and sum the costs over the concession period.
- 6. Repeat steps 1–4 N times and compute the average total cost; the choice of N is discussed below.

The fluid approximation is formulated using the Gurobi 1.7.5 solver integrated with JuMP 1.29.1 in Julia 1.11.6. Solving the fluid approximation for a given state and scheduling stage takes less than 30 seconds on a computer equipped with an 11th Gen Intel® Core™ i7-11850H processor. The entire experiment requires under 5 hours of total computing time. We consider the policy given by the fluid approximation with a time horizon of one year, with the full set of 80 weather scenarios $\mathcal{W}_T$ for all scheduling stage T. We compare this to two benchmark policies using three metrics: the mean simulated cost over 30 years, the 0.05% Value at Risk ($\text{VaR}_{0.05}$), and the frequency of simulations with strictly positive costs ($\mathbb{P}(\text{cost} > 0)$). We take $N = 10,000$ for each policy, which results in standard errors below 3% of the estimated costs. The first benchmark policy is an operational rule that involves performing corrective maintenance throughout the year and scheduling preventive maintenance during the summer because accessibility conditions are better. Corrective maintenance is performed as soon as possible when capacity is below nominal capacity. Preventive maintenance is performed in July on all components that are not in state 1, if the necessary free maintenance days are available. The comparison is available in Tab. I. The operational rule is about twice as expensive as the fluid policy. It is also exposed to potentially higher costs in the tail of the cost distribution. The second benchmark policy is a fluid heuristic with surrogate subsystem costs ℓ assuming perfect accessibility. More precisely, for all T, $\mathcal{W}_T$ contains only one scenario w with h_t^w equals 1 for all $t \in \mathcal{T}_T$ and p_t^w is as before. The cost induced by this weather-ignoring policy is, on average, 34 times higher than the fluid policy. Indeed, maintenance is planned to maximize quota usage under perfect accessibility, any additional maintenance days resulting from delays experienced under actual accessibility conditions are conducted outside the established quotas. A commented simulation example with the fluid policy and a decomposition of costs by components are available in Annex B.

TABLE I: Cost over 30 years depending on the policy.

policy	mean	$\text{VaR}_{0.05}$	$\mathbb{P}(\text{cost} > 0)$
fluid	126	257	0.05
operational	249	1671	0.31
fluid no W	4343	8792	1.00

B. Sensitivity to model ambiguity

Up to this point, the degradation model was assumed to be known. In practice, however, multiple sources of model ambiguity may arise. Some are exogenous, such as incomplete or unavailable data from manufacturers, or the lack of failure history in offshore environments. Others stem from the modeling process itself. For some components, when representing degradation using Markov chains, the degradation states cannot be clearly delineated due to the lack of well-defined failure modes. Modeling decisions must be made regarding which components to include and how to aggregate them within the global degradation model. In this work, we assume component independence; however, this assumption may not fully reflect real-world dependencies.

We study uncertainties in the degradation model of one of the components. In this work, we rely on classical MTBF values derived from onshore component data. Offshore components are subject to harsher environmental conditions and greater mechanical stress, leading to potentially lower MTBFs than their onshore counterparts. As costs decrease with higher MTBF, a conservative approach uses MTBF values at the lower bound of plausible estimates. This reflects the higher risk associated with underestimating maintenance costs from the perspective of the TSO. Unlike the MTBF, ambiguity in transition probabilities is more difficult to manage because there is no direct relation between these probabilities and costs. In our analysis, we fix the MTBF of the pump system and explicitly account for uncertainty in transition probabilities, examining how variations in these probabilities impact costs. All other components are modeled with a fixed degradation model as in the previous section. The pump system is composed of two identical units. The system transitions from full operation (two units working) to degraded operation (one unit working), and finally to failure. More precisely, we introduce a family of transition matrices with three states parametrized by a variable θ

$$P_\theta = \begin{bmatrix} 1-\frac{1}{(1-\theta)\mathrm{MTBF}} & \frac{1}{(1-\theta)\mathrm{MTBF}} & 0 \\ 0 & 1-\frac{1}{\theta\mathrm{MTBF}} & \frac{1}{\theta\mathrm{MTBF}} \\ 0 & 0 & 1 \end{bmatrix},$$

where θ represents the average time spent in the intermediate state normalized by the MTBF. This parametrization preserves the component's MTBF. We examine three cases: $\theta = 0.1$, representing rapid degradation after redundancy loss; $\theta = 0.67$, corresponding to the previous model with identical, independently degrading, non-aging pumps; and $\theta = 0.99$, where redundancy is lost early but the component functions in degraded mode for most of its lifetime. Evaluating costs solely based on the model used for optimization is not sufficient, as uncertainties in the degradation dynamics may still persist once operations begin at the offshore sites. We consider three degradation models and compute nine cost values: each policy is optimized under one model and evaluated under all three. The columns correspond to the parameters used for optimization (i.e., to derive the fluid policy), while the rows indicate the parameters used in the simulations to estimate the policy's cost. The results, shown in Table II[1], reveal a high sensitivity of cost estimates to the choice of the degradation models. A conservative cost estimate is given by the maximum value in the column with the smallest maximum value (highlighted in bold), representing the worst-case cost of the best-performing policy.

TABLE II: Cost over 30 years depending on the MDP used for optimization and estimation.

θ **est** \ θ **opt**	0.99	0.67	0.1
0.99	**1195**	1200	1438
0.67	192	126	144
0.1	335	204	194

V. Conclusions and Perspectives

This paper presents a method for evaluating the penalties paid by the TSO to energy producers in the event of unexpected failures or unplanned maintenance of future offshore substations, using an MDP framework. The cost is evaluated through maintenance scheduling that incorporates degradation models and weather conditions. Strategic investment decisions are treated as parameters within the MDP, allowing for direct cost comparisons, provided that the degradation models can be represented as Markov chains. Numerical results on the case study show the potential high costs incurred due to the lack of weather forecasts at the time when maintenance is scheduled. This could motivate improvements in renewable energy production and accessibility forecasts. Although this was not addressed in the numerical applications of this paper, it would be possible to incorporate more operational constraints within the same framework. For example, we could include limited workforce capacity to better reflect operational limitations in maintenance planning. The objective function can be adapted to reflect the specific priorities of the TSO, such as incorporating maintenance costs alongside the penalties. The absence of historical offshore failure data introduces ambiguity in degradation modeling, which in turn impacts the accuracy of cost estimates. This point is critical for the TSO, which must mitigate the risk. A common initial approach is to select a single degradation model from the ambiguity set—the collection of plausible degradation models and optimize accordingly. However, given the critical nature of cost estimation and the numerical results presented, this method appears to be insufficient. Future work will focus on integrating model ambiguity directly into the policy computation process.

[1]Table II shows that costs do not vary monotonically with θ. When θ is too large, costs increase because there are insufficient free maintenance days to perform preventive maintenance at each scheduling stage when pumps are in state 2. Conversely, when θ is too small, costs are also high, as transitions from 1 to 2 and from 2 to 3 can occur within the same scheduling stage, before the next maintenance opportunity.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process
During the preparation of this work, the authors used GPT-5 mini in order to assist in writing code for generating plots and in drafting the manuscript. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

References

[1] RTE, *Energy Pathways to 2050*, 2021.
[2] Y. Wu and H. Zhao, *Optimization maintenance of wind turbines using Markov decision processes*, 2010 International Conference on Power System Technology, 2010.
[3] F. Besnard and L. Bertling, *An Approach for Condition-Based Maintenance Optimization Applied to Wind Turbine Blades*, IEEE Transactions on Sustainable Energy, vol.3, no.2, pp. 77-83, Jul 2010.
[4] F. Besnard, K. Fischer and L. Bertling Tjernberg, *A Model for the Optimization of the Maintenance Support Organization for Offshore Wind Farms*, IEEE Transactions on Sustainable Energy, vol. 3, no. 2, pp. 443-450, Apr 2013.
[5] M. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, Wiley Series in Probability and Statistics, 1994.
[6] W. Wang and H. Pham, *Reliability and Optimal Maintenance*, Springer Series in Reliability Engineering, 2006.
[7] B. Powell, *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, Wiley Series in Probability and Statistics, 2011.
[8] D. Adelman and A. Mersereau, *Relaxations of Weakly Coupled Stochastic Dynamic Programs*, Operations Research, vol. 56, no. 3, pp. 212-227, May 2008.
[9] J. Hawkins, *A Lagrangian decomposition approach to weakly coupled dynamic opti mization problems and its applications.* , PhD thesis, Massachusetts Institute of Technology, 2006.
[10] D. Bertsimas and M. Mišić, *Decomposable Markov Decision Processes: A Fluid Optimization Approach*, Operations Research, vol. 64, no. 6, pp. 1537-1555, Oct 2016.
[11] H. Hersbach, B. Bell, P. Berrisford et al *ERA5 hourly data on pressure levels from 1940 to present. Copernicus Climate Change Service (C3S) Climate Data Store (CDS)* [Online]. Available: https://cds.climate.copernicus.eu.
[12] RTE, *Strategic development plan for the French transmission grid (SDDR)*, 2025.
[13] RTE, "Generation", Electricity analysis and data. [Online]. Available: https://analysesetdonnees.rte-france.com/en/generation/wind.

ANNEX A
SUBSTATION MODEL

The structure of the substation is summarized in Table III. We consider a substation with one converter, one water circuit, 12 fans, 2 pumps, 2 transformers, one converter cooling system, and one busbar system. All elements of the same type are grouped together as a single component.

TABLE III: Number of states n, MTBF and replacement time d for each component.

c	Component	n	MTBF (years)	d (days)
1	Converter	12	588	6
2	Water Circuit	3	43	1
3	Fans	8	43	1
4	Pumps	3	43	2
5	Transformers	3	3333	1
6	Converter Cooling	3	122	1
7	Busbar	3	4762	1

For components composed of multiple identical subcomponents, the discrete states are associated with the number of subcomponents still working. We compute transition probabilities assuming independent failures and no aging, which yields binomial transitions whose parameters are calibrated to match the component MTBF. More details on the modeling of each component follow.

- ***Converter.*** The converter is modeled as 1680 independent submodules. The substation is assumed to deliver full power if at least 1670 submodules are operational; otherwise no power is transmitted.
- ***Fans.*** Of the 12 fans, two are redundant. Capacity decreases proportionally to the number of broken fans once failures exceed the redundancy threshold.
- ***Transformers and pumps.*** Of the two transformers, one is redundant. The same applies to the pumps.
- ***Other components.*** For the water circuit, the converter's cooling system and busbar we assume three degradation states. The substation delivers full power in states 1 and 2 and no power is transmitted in state 3. Transitions occur only to the next worse degradation level; transition probabilities are chosen so that the expected times in states 1 and 2 are equal and calibrated to match the component's MTBF.

The capacity function is defined as follows:

$$C(x,q) = \prod_{c\in\mathcal{C}} \mathbb{1}_{\{q_c=0\}} \times \prod_{\substack{c\in\mathcal{C}\\ c\neq 3}} \mathbb{1}_{\{x_c<n_c\}} \times \big(100 - 20(x_3-3)^+\big).$$

The first factor ensures that the capacity is zero whenever at least one component is undergoing maintenance (i.e., some $q_c > 0$). The remaining factors express how capacity depends on component degradation states. The maximum capacity is 100.

ANNEX B
SIMULATION RESULTS

Fig. 2 illustrates a simulation under the fluid policy. Component 4 is broken so its maintenance is scheduled at the start of the next planning stage. Due to poor accessibility, the maintenance requires more free days than available, and the TSO pays a penalty. To take advantage of free maintenance days, all other maintenances are scheduled at the same time except for component 1, which has the longest maintenance, to avoid extra costs. Then a minor failure occurs on component 6. No action is taken that summer as maintenance days are exhausted; quotas are renewed the following year, and maintenance is performed then. There is no maintenance on 1 because it is in state 2 (out of 12), which is not critical.

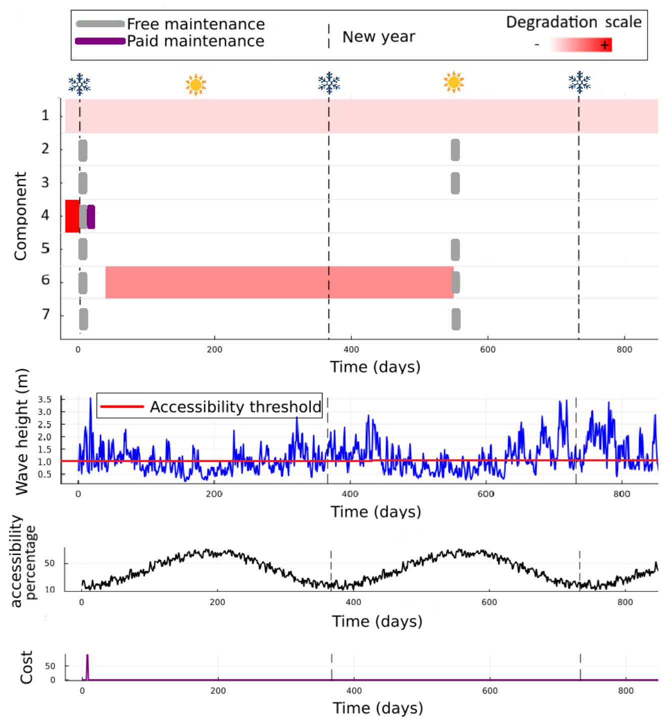

Fig. 2: Simulation results over 14 planning stages for a single weather and degradation trajectory with the fluid policy.

Top plot: degradation levels and maintenance by component. Second plot: wave height for the considered scenario. Third plot: accessibility percentage across 80 weather scenarios. Bottom plot: incurred cost.

Table 3 shows the mean simulated cost by component under the fluid policy, whose total cost (126) is shown in Fig. I. The pump is the costliest component due to its low MTBF and long maintenance time.

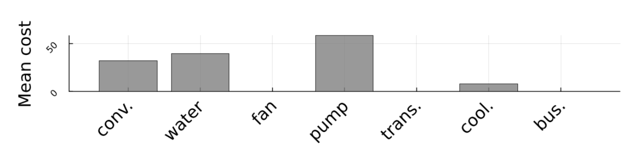

Fig. 3: Mean costs per simulation associated with failures and maintenance for each component.