RESEARCH ARTICLE

Using maintenance options to maximize the benefits of prognostics for wind farms

G. Haddad[1], P. A. Sandborn[2], and M. G. Pecht[2]

[1] Schlumberger Technology Center, 150 Gillingham Lane, Sugar Land, TX 77478
[2] Center for Advanced Life Cycle Engineering (CALCE), University of Maryland, College Park, MD 20742

ABSTRACT

Many engineering systems incorporate prognostics and health management (PHM), which consists of technologies and methods to assess the reliability of a product in its actual life-cycle conditions to determine the advent of failure and mitigate system risks. Wind turbines are among the systems that incorporate PHM to reduce life-cycle costs and increase availability. While cost-benefit models that quantify the value of implementing prognostics within systems exist for wind energy systems, they do not specifically quantify the value of decisions after a prognostic indication. This paper introduces maintenance options as a means to quantify the value of decisions after a prognostic indication. A case study on a US land-based wind farm is discussed. An analysis of wind turbine maintenance data is presented, and the maintenance options methodology is then demonstrated to establish the value of the wait-to-maintain option. The value of waiting after a prognostic indication is determined using a model that quantifies the benefit that results from a PHM implementation that allows the decision-maker to delay maintenance actions, thereby using the remaining life of the system components rather than throwing it away.

KEYWORDS

Wind Farms, Prognostics and Health Management, Operation and Maintenance, Decision Support, Life-cycle Cost, Maintenance Options, Real Options, Condition-based Maintenance.

Correspondence
G. Haddad, Schlumberger Technology Center, 150 Gillingham Lane, Sugar Land, TX 77478.
Email: GHaddad@slb.com.

1. INTRODUCTION

Wind energy is at the forefront of alternative energy sources. The US Department of Energy (DoE) and the National Renewable Energy Laboratory (NREL), under the “20% Wind Energy by 2030” plan, announced that the US could feasibly increase wind energy’s contribution to account for 20% of the total electricity consumption in the United States by 2030 [1]. However, wind energy sources face numerous challenges that could hinder their competitiveness with traditional power sources. Managing availability—a function of reliability and maintainability—is a significant issue that needs to be addressed since it directly impacts the economic feasibility of wind farms.

Availability is the ability of a service or system to be functional when it is required for use or operation. The availability of wind turbines will significantly affect their economic viability. A key aspect of wind turbine availability is the need for nontraditional resources for maintenance. Wind farm maintenance often requires cranes for critical repair and replacement of major components (blades, gearboxes, and bearings, for example). The availability of the cranes themselves may be limited, or maintenance may have to wait until the turbine fails in order to justify the cost of bringing a crane to the site. If a turbine fails immediately after a maintenance action has been performed on it, then it may not be available until the next time a crane is on-site for maintenance (or the vessel is on-site if the wind farm is offshore). Kühn [2] studied the failure rates of 235 small wind turbines and assessed the annual frequency of failures in the turbines and the corresponding downtime for different subassemblies; some failures resulted in a downtime of as much as 25 days. It is important to note that the downtime for wind turbines is a composite of different downtimes, such as time for inspection, management decision time, time to acquire replacement parts, time to get equipment to the site, repair time (including weather delays), etc.

The *Wind Energy Operations & Maintenance Report*, published in 2010 [3], included a discussion highlighting the maintenance challenges with wind energy systems. Some of the most notable conclusions are that the operation and maintenance (O&M) costs for wind power are double or triple the figures originally projected, and they are particularly high in the United States.[1] Furthermore, gearboxes, which have been designed for a 20-year life, are failing after 6 to 8 years of operation. Similar figures were presented in a report published in 2011 [5]. These issues indicate that reliability, maintainability, and availability are key challenges to the economic viability of wind turbines and their ability to compete with traditional energy sources.

Figure 1, adapted from Arabian-Hoseynabadi et al. [6], shows the failure rates of different sub-assemblies in wind turbines. The plot shows that multiple subassemblies have a significant annual failure rate.

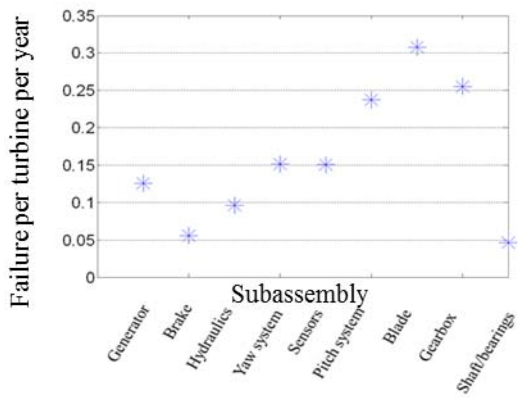

Figure 1. Reliability failure rates for wind turbine subassemblies [6].

To avoid unanticipated failures and ensure high availability, many industries, such as the aerospace industry, have begun to employ prognostics and health management (PHM) techniques that warn users (and/or maintainers) before systems fail. PHM is a discipline comprised of technologies and methods designed to assess the reliability of a product in its actual life-cycle conditions to determine the advent of failure and mitigate system risks [7][8]. The PHM system provides information on remaining useful life (RUL), which allows the decision-maker to take appropriate actions to manage the system's health prior to (or sometimes upon) failure. There are numerous applications of PHM in industries such as aircraft, wind turbines, gas turbines, and electronic systems [7].

PHM has been used on wind turbines to assess the state of health for gearboxes, bearings, oil, blades, electronics, and overall performance. A common approach is to monitor gearbox vibration using accelerometers and use vibration analysis approaches to detect faults [9]. Oil analysis is performed to safeguard the oil quality and components involved, which is generally performed offline by taking samples. The oil is analyzed for wear debris. In wear debris analysis, the quantity, size distribution, morphology, and color of the wear debris is determined, which can provide information on the wear modes, wear sources, and wear phases present in the machine [10]. The analysis of the blades is typically done by examining the physical conditions of the materials and is performed offline to monitor crack growth. Power electronics and electronic controls account for a small portion of the cost but have a potentially large effect on downtime if they fail frequently. There have been several diagnostic approaches developed for power electronic devices, especially IGBTs, but these detect faulty components rather than predict when failures will occur. Methods for PHM for wind turbine electronics include neural networks, wavelet analysis, and bond graph methods. Adams et al. [11] proposed a structural health monitoring method for wind turbines using vibration signals from the turbines. A review of condition monitoring of wind turbines can be seen in Hameed et al. [12].

As an example of the current state of wind turbine maintenance, consider a US land-based wind farm consisting of more than 100 turbines. The farm was completed in the early 2000s. The service data pertains to maintenance actions and consists of a database of all maintenance actions reported for the turbines since they were installed. The service data was obtained for all years since the turbines were installed (approximately 10 years).

Scheduled maintenance (fixed time or usage interval based) is the most common maintenance paradigm used for maintaining wind farms. However, this paradigm does not account for the actual usage and loading conditions, and,

[1] It is not uncommon for the O&M costs for "sustainment-dominated" systems [4] to be 60–70% of the total cost of ownership of the system.

if not optimized, can result in a significant amount of unscheduled maintenance actions. In order to highlight the problems with scheduled maintenance, this example presents maintenance data (also called service data) for a wind farm that is operating under a scheduled maintenance paradigm.

Figure 2 shows the cumulative maintenance cost for the wind farm under consideration for the years 2007 to 2010.[2] Each line in Figure 2 corresponds to one turbine in the farm. The plot shows abrupt jumps in cumulative cost that correspond to turbine failures (such as gearboxes) that should have been avoided by scheduled maintenance (but were not). Gearboxes are examples of components that were not supposed to fail during the first 10 years of operation, but did in fact fail.

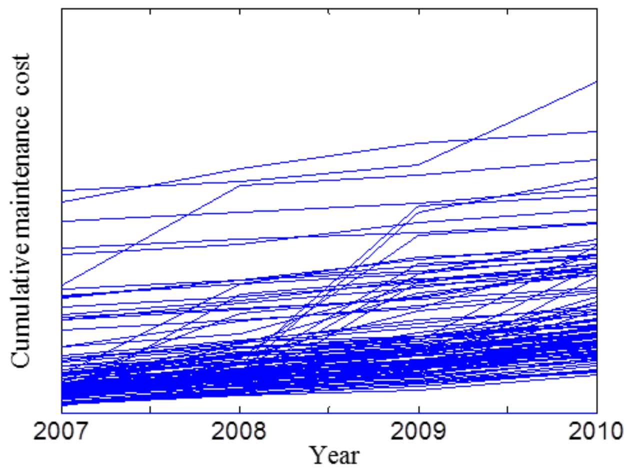

Figure 2. Cumulative maintenance cost for 3 years for the farm (more than 100 turbines).

For the wind farm considered here, there were four service types: 1) non-functional turbine, 2) operating with problem, 3) preventive maintenance, and 4) scheduled preventive maintenance. The service type indicates the state of the turbine when maintenance was requested. A "non-functional turbine" indicates that service was requested for a turbine that was not operational. Preventive maintenance is a request for all maintenance actions as a result of a trigger; in addition, scheduled preventive maintenance occurs twice a year. In the farm under consideration, the frequency of "non-functional turbine" request types is larger than all other types. "Operating with problem" ranks second in the service request type. If the turbine is non-operational or operating with a problem, the turbine may not produce as much power as it should, and its life-cycle cost will increase as a result.

For the service request "non-functional turbine," the number of days elapsed between the service request and the repair date of the turbine is shown in Figure 3. The histogram shows high frequency for days 0 to 20. This may be due to the fact that many faults are resolved by resetting the turbine or some other subsystem in the turbine; or the fault can be indicated during scheduled maintenance. However, there were a substantial number of service requests that took more than 20 days to resolve. The availability of maintenance resources directly impacts the number of days the turbine is not functional because it is down for maintenance.

[2] The cost values on the vertical axis in Figure 2 have been removed to preserve the confidentiality of the data.

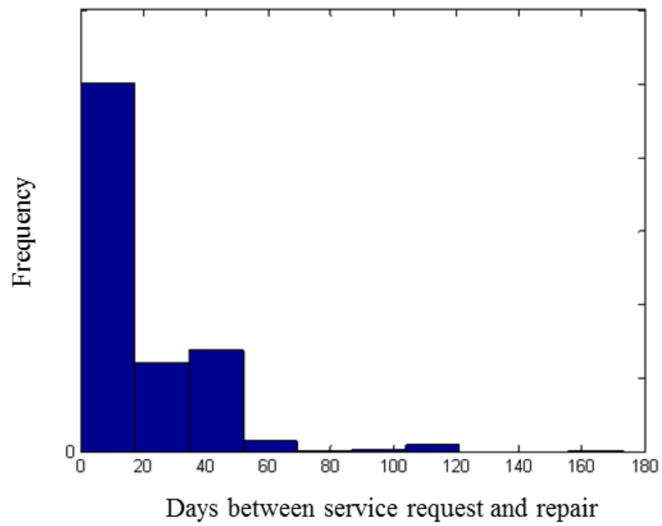

Figure 3. Time between service request and repair for non-functional turbine service requests.

For some turbines in Figure 3 more than 100 days elapsed between the service request and the time when the turbine was repaired. If the downtime lasts for 100 days, then the turbine will not generate power for almost a third of the year.

Figure 4 shows the days that elapsed between the repair of the problem and the closest scheduled maintenance cycle for the fault type: "non-functional turbine." The first bin in the histogram in Figure 4 exhibits the largest count since it can be associated with the proximity to the scheduled maintenance cycles. The count includes times before or after the scheduled maintenance. In other words, some events may have happened and were missed in the scheduled maintenance. However, there is a high frequency of events happening more than 40 days outside (before or after) the maintenance cycle. Some events happen more than 100 days outside the maintenance cycle, and as a result the turbine may have had to wait for more than 100 days until the next maintenance cycle for repair.

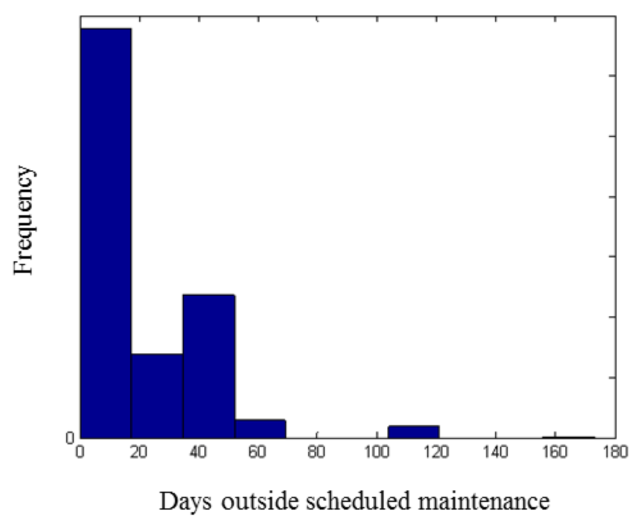

Figure 4. Time elapsed from repair to closest maintenance cycle for non-functional turbine service requests.

This discussion highlights fundamental points that may result in a decrease in the economic value of wind farms. Turbines in a farm have a large amount of downtime, and a large number of maintenance events occur outside scheduled maintenance events. The scheduled maintenance policy for a wind farm presents a number of scenarios that potentially decrease the value obtained from the wind farm. Prognostics can potentially increase the maintenance value of such systems, prevent failures, and increase the impact of renewable energy sources.

1.1 Maintenance Approaches for Wind Farms

There are different approaches to maintenance, but, fundamentally, depending on if a system has failed, when we think it will fail, and how it has failed, there are decisions that need to be made about how and when to maintain it. The goal is to optimize the maintenance of turbines by minimizing the costs of maintenance and maximizing the reliability and the availability. Corrective (unscheduled) maintenance involves maintaining a system upon failure. Corrective maintenance allows the entire life of the system or component to be used. For many systems, corrective maintenance is inefficient, as it can result in long downtimes, catastrophic failures with significant collateral damage (damage to components that were not part of the initial cause of the shutdown), and unpredictability, which can lead to low maintenance value. Scheduled maintenance can be time-based or usage-based. It has a low maintenance value if the time or amount of usage to failure is not well characterized, leading to the loss of substantial remaining useful life (RUL). Condition-based maintenance (CBM) enabled by PHM allows the reliability of a system to be monitored in real time. PHM-enabled CBM is considered to be the paradigm with the highest potential value, since it minimizes the unused RUL, avoids failures, and provides a lead time for logistics management, among other benefits. However, PHM-enabled CBM may have greater implementation and support costs than unscheduled maintenance.[3]

When a system that incorporates PHM detects an anomaly, it creates a prognostic indication that consists of an estimation of the RUL before the system fails; this creates flexibility for managing the system. RUL is the remaining useful life that a system has, and it effectively represents the lead time (subject to appropriate uncertainties) for the decision-maker or other maintenance entity to take preventive actions prior to failure. This can be described as the flexibility phenomenon, whereby entities involved with the operation, management, and maintenance of a system have the flexibility to take actions with a risk/value tradeoff.

In a system that incorporates PHM, after a prognostic indication, several different actions can be taken by the decision-maker to manage the health of the system. Examples of the actions that can be taken include fault accommodation (e.g., turning the turbine off), changing loads (e.g., operating at a lesser load), and tactical control (e.g., changing the control of the turbine). Hence, the decision-maker has a set of options among which they can choose. Bonissone [13] proposed a temporal segmentation for decisions for systems with PHM, where the tactical and operational decisions at the object level are examples of options after prognostic indication. The term *options* will be used in this paper to denote a choice or action that the decision-maker can take after a prognostic indication. The flexibility that PHM allows increases the value of the system if the decision-maker can take advantage of the options it provides. This paper develops a methodology that uses real options for the valuation of options enabled by PHM.

A real option is an alternative or choice that becomes available as a result of a business investment opportunity. It is a right, but not an obligation, to take an action (e.g., defer, expand, contract, or abandon a project) at a predetermined cost called the exercise price, for a predetermined period of time—the life of the option [14][15]. Real options analyses are decision tools for addressing the value of investments under uncertainty.

Real options are the extension of financial options to real assets. Unlike financial options, real options are not securities and they cannot be traded. A real option has an underlying asset, such as a project or a growth opportunity. For real options there is no need for a contract to specify the payoff, and the payoff can be a future cost avoidance [16]. In our case, the investment opportunity that creates maintenance options is the investment in putting PHM into the system.

Real options have been used for maintenance applications (e.g., [17]). Real options have also been applied in the maintenance, repair, and overhaul industry [18], and for scheduling joint production and preventive maintenance for the manufacturing industry when demand was uncertain [19].

While there have been numerous efforts to implement PHM in wind turbines, an understanding of the economic merit of implementing PHM and the value it actually provides at the system level is lacking. This paper develops a method to quantify the value of decisions after a prognostic indication. This will provide the value that the decision-maker obtains from PHM at the system-level, as well as a basis for optimizing maintenance decisions.

In this study, we use an options framework to quantify the value of the waiting-to-maintain option. This represents a cost-benefit-risk model that provides a basis to know whether waiting to maintain is beneficial. The waiting option lends itself naturally to the fundamental tradeoff in maintenance problems with prognostics. To understand options, this fundamental tradeoff can be stated as follows: the waiting option attempts to find the time to perform

[3] There are particular systems under particular circumstances for which scheduled or unscheduled (corrective) maintenance is in fact the optimum maintenance approach. It may not always be practical to implement CBM and PHM in systems (or within every component of a system). Furthermore, complex systems may use a combination of the different maintenance paradigms in order to optimize maintenance decisions.

maintenance that maximizes the combination of the revenue that can be generated by the system during its RUL and the cost of unscheduled failure that can be avoided.

The remainder of the paper is structured as follows: Section 2 presents maintenance options and their valuation approach. A case study on maintenance options for wind turbines is presented in Section 3. Finally, Section 4 concludes the work and discusses research opportunities.

2. A Methodology to Determine the Value of the Wait-to-Maintain Option

This section presents a methodology to quantify the wait-to-maintain option. Section 2.1 presents the valuation model. Section 2.2 presents the use of the methodology in a simple example, and Section 2.3 presents a generalization of the valuation of the waiting option for different time horizons. The algorithm will then be used in Section 3 in a case study on wind turbines.

2.1. Valuation Methodology

Determining the value of waiting to perform maintenance is at the heart of decision support after a prognostic indication. Decision-makers are concerned with the value of delaying an investment in maintenance given the flexibility enabled by PHM. This is essentially knowledge of the time when waiting is no longer beneficial.

We start by defining the value (V_M) as the value of the summation of cost avoidance (CA) and cumulative revenue generated from operating the system (R) up to the end of the RUL[4], which can be expressed as:

$$V_M = CA + R \quad (1)$$

The cost avoidance opportunities are expressed as the difference between the cost of resolving non-detected failures (C_{USM}) (the subscript USM stands for unscheduled maintenance) and the cost of the resolution of detected problems prior to failure (i.e., the cost of carrying out maintenance based on prognostic information, C_{PHM}). In other words, it is the difference between the cost of managing (repairing or replacing) a failed system and the cost of repairing it based on its condition.

$$CA = C_{USM} - C_{PHM} \quad (2)$$

Consider the following simple example: a system indicates an RUL of 3 time units, and V_M has an initial value of 1. A Monte Carlo simulation that follows 8 possible time histories for this example is shown in Figure 5. A prognostic indication is obtained at time 0. Each of the time steps after the prognostic indication represents a point in time when there is an opportunity for maintenance (this is a discrete time when maintenance can be performed; note, maintenance may not be possible at all future times since resources may not always be available). In our analysis, we assume that maintenance can be carried out at discrete time steps 1, 2, and 3. At every time step, we need to examine the value of continuing on to the next time step (i.e., waiting until a future time step) and compare it to the value of maintaining at the current time step. In terms of maintenance, waiting for another time step means using the system through more of its remaining useful life, thus generating more revenue, but at an increasing risk of failure. In the simplest case, where there is one time step and the value can go up or down with some probability, Figure 5 can be regarded as a simple decision tree. However, Figure 5 illustrates multiple time steps where the value can increase and/or decrease until the last time step. The actual decision process is illustrated in detail in Section 2.2.

[4] The methodology presented in this paper assumes that an RUL has been generated using a PHM system. The methodology of obtaining the RUL from the data at hand is outside the scope of this paper. There are numerous approaches for predicting a RUL for systems, e.g., [7] and [12].

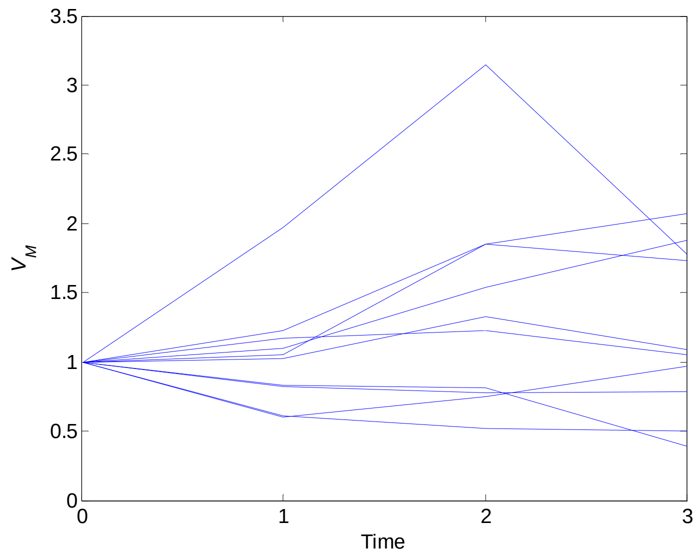

Figure 5. Simulated value (cost avoidance and revenue opportunities). Each line is a possible stochastic value path post-prognostic indication for the system.

Waiting is highly influenced by uncertainty. For example, for the paths in Figure 5, waiting for one time step shows paths with V_M higher and lower than 1. This is a result of propagating uncertainty in time. For example, consider a wind turbine that has a predicted RUL; the operator may choose to run the turbine. A factor influencing this decision is uncertainty about wind speed. The probabilities of high and low wind speeds coupled with the fact that using a wind turbine when the wind is blowing at different speeds may generate different revenues. In Figure 5, for the paths where the value is higher than 1, waiting is beneficial; this is the upside effect of uncertainty. For the paths that are lower than 1, the uncertainties will result in a recommendation of immediate maintenance (or when the first opportunity arises). If a cross-section of the data in Figure 5 is taken at a specific time step, the result will be a distribution of possible values (V_M) at that particular time step.

Time-dependent uncertain quantities can be modeled with dynamic forecasting models [20]. An example of a stochastic differential equation to propagate uncertainty with time is:

$$dX_t = \mu X_t dt + \sigma X_t dW_t \tag{3}$$

where X_t is the value of the quantity being simulated at time t (i.e., a stochastic process), μ is a drift component, σ is a variance component (or shock, also called volatility), and W_t is a Brownian motion (also known as a Weiner or continuous-time stochastic process). For example, Figure 6 shows the three uncertain quantities in equation (1) modeled using equation (3). The vertical axis represents the three uncertain quantities in equation (1): cost avoidance (CA), cumulative revenue (R), and value (V_M). The horizontal axis represents time in days; at time 0 a prognostic indication is obtained, and uncertainty is propagated for 100 days in each case. The diagram on the left is the cost avoidance (starting value of $45,941 (60% of the cost of failure of a gearbox), μ = -0.8, σ = 0.25). The diagram in the middle is the cumulative revenue of the turbine (starting value of $237, which is the daily value generated by a 600KW turbine with a capacity factor of 0.33 and a cost of energy $0.05/KWh; μ = 0.5, σ = 0.1; note that this is a cumulative quantity). The diagram on the right represents the summation of the first two quantities. The drift and volatility values are assumed for the purpose of the example; however, their actual values and the methodology to obtain them will be discussed in the case study in Section 3. The values for drift and volatility are estimated from the variation in the power curve for the revenue and from historical data for cost avoidance.

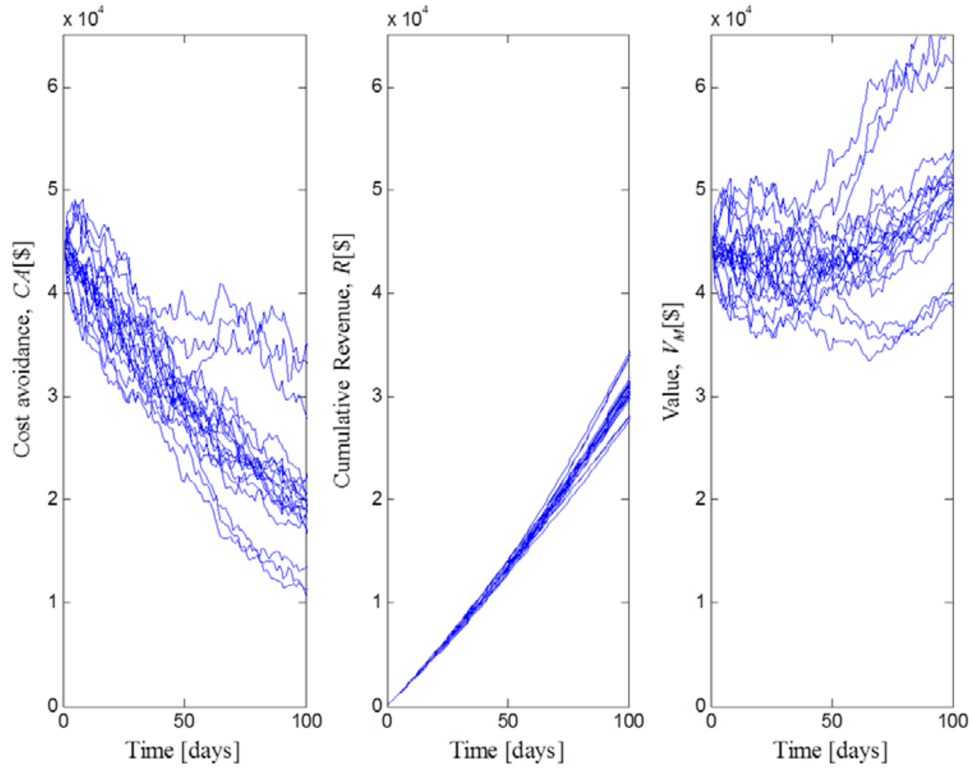

Figure 6. Uncertain quantities propagated in time.

We note that the time-dependent uncertainties are considered to account for a representation of the uncertainty in the remaining useful life. The propagation of uncertain quantities over time presents a distribution at each time step. For example, if a cross-section of the value V_M is taken at the last time step (100 days) in Figure 6, the spread in the distribution of the quantity represents all the possible outcomes if the decision-maker is to wait for the whole 100 days. The problem is, however, not static. A new prediction can be obtained one day after the first prediction and some uncertainty may be resolved and the uncertain quantities are propagated again. The following discussion presents the value of having the option to wait for up to the predicted RUL (see Section 2.3 for the generalization of waiting to other points in time). This will give a value of waiting for all possible outcomes in uncertain quantities. The valuation can be repeated whenever new information about prediction or uncertainties is obtained.

At every time step, the value V_M is compared with the cost of maintenance, C_M, which accounts for the cost of failure, $C_{Failure}$, cost of downtime $C_{Downtime}$, and penalty $C_{Penalty}$:

$$C_M = C_{Failure} + C_{Downtime} + C_{Penalty} \tag{4}$$

Comparing V_M and C_M is at the heart of the option valuation. If V_M (the summation of the revenue generated from running the turbine and the cost avoidance opportunities) is larger than C_M, then the decision-maker can harness the upside effect of uncertainties and wait to perform maintenance.

In order to valuate the wait-to-maintain option, we start by defining the following quantities: X_i is the value of V_M from operating the system at the current time; X_{i+1} is the value of running the system until the next time step when maintenance can be performed; $V(X_{i+1})$ is the value of waiting for an additional time step to maintain the system. In other words, this is the value of the cost avoidance opportunities and the revenue generated from the system, and is derived from waiting to perform (invest in) maintenance.

In order to find the optimal time for maintenance, we define a stopping rule to exercise the option. The stopping rule defines the optimal time for maintenance. For example, if the value of continuation is smaller than the value of exercising the option at the current time, it is optimal to stop and maintain since waiting does not add additional value. The stopping rule is based on finding the expectation of the option's value at time t_{i+1} conditional on the value of revenue at time t_i, given by $E[V(X_{i+1})|X_i]$. The value of continuation and the stopping rule are impacted by

C_M, which is the baseline against which all the values are compared.

The expectation function is a representation of the expected value obtained from waiting (delaying maintenance actions/investments) conditioned on the value obtained from the system and cost avoidance at the current time step. For example, if the expected value—in terms of cost avoidance and additional revenue derived from waiting (delaying a maintenance action)—is positive, then the decision-maker is better off postponing the maintenance actions and using the system through to the remaining useful life. If, on the other hand, the cost of failure is higher than the revenue generated, it will have a value of zero[5] in the expectation function, indicating that waiting is not beneficial. This can be related to the fundamental maintenance problem: how to maximize the use of RUL while minimizing the risk of failure.

Longstaff and Schwartz [21] proposed an algorithm for pricing American options (these are options that can be exercised at any time, which is the case for the wait-to-maintain option) with path dependency by approximating the conditional expectation function with a set of basis functions. This algorithm is called the least squares Monte Carlo algorithm (referred to as LSM). This algorithm is expressed as

$$E[V(X_{i+1})|X_i] \cong \sum_{n=0}^{N} \beta_{i,n}\phi_n(X_i) \tag{5}$$

where β is the coefficient of a basis function, ϕ is a basis function, and N is the number of basis functions used.

Equation (5) is essentially an approximation of the values in the uncertain paths with a set of polynomials (Laguerre polynomials are used in [21]). In this paper we have used the LSM algorithm to valuate maintenance options.

In summary, the process for valuating the waiting option is as follows:

1) Estimate the time-dependent uncertainties in the model. For the waiting option defined in this paper, the uncertain quantities that need to be propagated with time are CA and R. The parameters for the uncertain quantities can be estimated from the historical data.
2) Propagate the uncertain quantities over a time frame corresponding to the RUL. Time-dependent uncertain quantities can be modeled with dynamic forecasting models [20]. An example of such models, Geometric Brownian motion, is given in equation (3). Equation (3) represents a time-varying uncertain quantity. When simulated with a number of time histories, it represents the uncertain quantity over time.
3) Define the value of C_M; this serves as a basis of comparison for the value of waiting.
4) Start from the last time step (t_n) and move backwards. At time step t_{n-1}, for each path, consider the value that is "in the money"[6]; as a path where the value of V_M is higher than C_M. Approximate the expectation function in equation (5) with a set of basis functions using Least Squares.
5) Compare the value obtained from the expectation function of waiting from t_{n-1} to t_n and the actual value at t_{n-1}. If the value of waiting is higher, the decision-maker is better off waiting for the path considered. If the value of waiting is lower, the decision would be to maintain at t_{n-1}.
6) Repeat steps 4 and 5 until time step 0.
7) Obtain the value of waiting at the optimal decision points for each path, discount the value to time 0, and average to obtain the value of waiting until time t_n.

The value obtained from these steps represents the additional value that the decision-maker can obtain from having the option to run the system up to the end of the RUL. It is a cost-benefit-risk model that provides a basis for quantitatively determining if waiting is beneficial or if the risk of failure is high enough that the revenue generated from the system cannot be compensated for by running the system. In other words, the value of waiting is a representation of the benefit obtained from using the RUL versus throwing it away. If the decision-maker decides to maintain immediately after the prognostic indication, then the RUL is thrown away and there is no value in waiting. The steps in the valuation of the waiting option are described using an example in Section 2.2, and the generalization for calculating the value of the wait-to-maintain option for any time before and up to the end of the RUL is presented in Section 2.3

[5] The value cannot be negative because we are representing the opportunity of the upside effect generated from waiting.
[6] An option is referred to as "in the money" if its value at a particular time step is higher than C_M.

2.2. Waiting Option Valuation Example

To illustrate the valuation of the wait-to-maintain option, we present a simple example. The example considered in this section corresponds to the data presented in Figure 5. We consider a system with prognostic capabilities that indicates an RUL of 3 time units. The goal is to obtain the value of having the option to wait for up to 3 time units (the end of the RUL); in this case (where C_M is the same for all time steps), we implicitly assume the system will not fail before the end of the RUL (Section 2.3 describes lifting this restriction). In this example, the value of waiting is an additional benefit that can be obtained from the knowledge of an RUL of 3 time units. The numerical data is given in Table 1.

Table 1. Value of V_M for 10 possible time histories (paths)

Path	t = 0	t = 1	t = 2	t = 3
1	1	1.05	1.85	1.73
2	1	1.97	3.15	1.78
3	1	0.61	0.52	0.50
4	1	1.02	1.33	1.09
5	1	0.82	0.78	0.79
6	1	1.10	1.54	1.88
7	1	0.60	0.75	0.97
8	1	1.23	1.85	2.07
9	1	0.83	0.81	0.39
10	1	1.17	1.23	1.05

Ten time histories (paths) for the value of maintenance are simulated in this example. Some paths have values higher than the initial starting point of 1, while some are lower; this is a consequence of uncertainty in the problem. Cost avoidance opportunities will decrease with time as there is a higher risk of failure of the system, but the system will generate more revenue as it is used through to the end of the RUL. For this example, we make the following assumptions: $X_0 = 1$ is the value of V_M at time 0; $C_M = 1.1$; and $t = 1, t = 2, t = 3$ are the times at which the decision-maker can perform maintenance.

At the final exercise date, the optimal strategy is to exercise the option if it is "in the money" (i.e., if its value is greater than C_M). Prior to the final date, the optimal strategy is to compare the immediate exercise value with the expected cash flows from continuing, and then exercise the option if immediate exercise is more valuable. Hence, the key here is to identify the conditional expected value of continuation. The objective is to determine the stopping rule that maximizes the value of the wait-to-maintain option at each point along each path.

Conditional on not exercising the option before the final expiration date at t = 3, the cash flows realized by the option holder from following the optimal strategy are given in Table 2. The values for t = 3 in Table 2 are obtained by subtracting 1.1 (the value of C_M) from the value at t = 3 if the option is in the money.

Table 2. Cash flow at t = 3 if the option is not exercised

Path	t = 0	t = 1	t = 2	t = 3
1	-	-	-	0.63
2	-	-	-	0.68
3	-	-	-	0
4	-	-	-	0
5	-	-	-	0
6	-	-	-	0.78
7	-	-	-	0
8	-	-	-	0.97
9	-	-	-	0
10	-	-	-	0

If the option is in the money at t = 2, the option holder must then decide whether to exercise the option immediately or continue until the next time that maintenance can be performed. As seen in Table 1, there are 6 paths where the option is in the money at t = 2 (where the value is more than 1.1). Note that only the paths that are in the money are considered for the analysis. We denote the value of V_M at t = 2 for those paths as X, and we denote the corresponding discounted cash flow received at t = 3 if the option is not exercised at t = 2 as Y. For instance, if the option is not exercised at t = 2 for path 1, then Y is calculated using $Y = (1.73 - 1.1)0.9418 = 0.59$ where e^{-rp} = 0.9418 is used to discount for p = 1 time period with an assumed discount rate of r = 0.06.

Table 3. X and Y at t = 2

Path	X	Y
1	1.85	0.59
2	3.15	0.64
3	-	-
4	1.33	0
5	-	-
6	1.54	0.73
7	-	-
8	1.85	0.91
9	-	-
10	1.23	0

To estimate the expected cash flow from waiting (continuing the option's life) conditional on the value at t = 2, we regress Y on the basis functions. For the purpose of the example, we choose a polynomial function consisting of an intercept, X, and X^2 (this choice is made for ease of representation). In the algorithm, we use Laguerre polynomials; however, other polynomials can be used as well. This will result in an expectation function that approximates the value of continuation:

$$E[Y|X] = -0.7679\, X^2 + 3.7020X - 3.4082 \quad (6)$$

With the conditional expectation function we can now compare the value of immediate exercise at t = 2 with the value from continuation, as shown in Table 4.

Table 4. Value of exercise ($\boldsymbol{X - C_M}$) and continuation $\boldsymbol{E[Y|X]}$ at t = 2

Path	Exercise	Continuation
1	0.75	0.81
2	2.05	0.63
3	-	-
4	0.23	0.16
5	-	-
6	0.44	0.47
7	-	-
8	0.75	0.81
9	-	-
10	0.13	0.00

The exercise value is obtained by subtracting C_M from the value of V_M at the current time step (step 5 in the algorithm in the previous section). In this example, we can represent it as (X – 1.1). The continuation value is obtained by substituting X into the conditional expectation function. This comparison implies that it is optimal to continue for all the paths considered since the value of continuation is higher than the exercise value at t = 2 (i.e., the conditional expectation function evaluated is higher than the current exercise value). The same steps are repeated iteratively until time t = 0 in order to obtain the stopping rule matrix (Table 5), which is a matrix consisting of binary numbers where 1 represents the optimal decision for a specific path. This optimal decision is defined as the stopping rule. For example, for Path 1 in Table 5, the optimal decision to exercise the option is at time t = 3, while for Path 8 the optimal decision is at time t = 1. Paths 5, 7, and 9 are all zeros in Table 5; these paths were not considered initially, as the value V_M is smaller than 1.1 for all time steps. This means that there is no value in waiting for Paths 5, 7, and 9.

Table 5. Stopping rule matrix

Path	t = 1	t = 2	t = 3
1	0	0	1
2	0	1	0
3	0	0	0
4	0	1	0
5	0	0	0
6	0	0	1
7	0	0	0
8	1	0	0
9	0	0	0
10	1	0	0

In order to obtain the cash flow realized, we consider the stopping rule matrix (Table 5) and consider the instances where there is a 1 in the matrix. This will lead to the cash flow matrix shown in Table 6.

Table 6. Wait-to-maintain option cash flow matrix

Path	t = 1	t = 2	t = 3
1	0	0	0.63
2	0	2.05	0
3	0	0	0
4	0	0.23	0
5	0	0	0
6	0	0	0.78
7	0	0	0
8	0.13	0	0
9	0	0	0
10	0.07	0	0

To calculate the value of the option, we discount each cash flow in the option to time 0 and average all the paths to obtain:

$$Value = [0.63(0.8353) + 2.05(0.8869) + 0.23(0.8869) + 0.78(0.8353) + 0.13(0.9418) + 0.07(0.9418)]/10 = 0.3388$$

where 0.8353 is used to discount 3 periods, 0.8869 to discount 2 time periods, and 0.9418 to discount 1 period.

In this example, if the decision-maker waits to maintain, an expected additional 0.3388 monetary units (MU) is obtained. If the decision-maker decides not to wait but instead to maintain immediately after obtaining the prognostic indication, a value of 0.3388 may be missed. This example demonstrates the use of the LSM algorithm for obtaining the value of the waiting option. Without accounting for the option to wait for three time steps to maintain, one would compare the value at t = 3 with 1.1 and then discount to time 0. This example would result in a value of 0.104 = ([(1.73–1.1)+(1.78–1.1)+(0.5–1.1)+(1.09–1.1)+(0.79–1.1)+(1.88–1.1)+(0.97–1.1)+(2.07–1.1)+(0.39–1.1)+(1.05–1.1)]/[(10)(0.8353)]). Hence, when accounting for flexibility in the decision-making process (waiting when it is favorable), the results show that the value of waiting is 224% ((0.3388–0.104)/0.104) higher than the value obtained without accounting for flexibility. The result from this example represents the additional value obtained from PHM. Waiting, an option that is enabled by PHM, is a representation of the benefit obtained from having the option to use the system through the end of the RUL.

2.3. Determining the Optimum Time to Wait—Generalization of the Waiting Option

The algorithm described in Section 2.1 and demonstrated in Section 2.2 explains how to obtain the value of having the option to wait to the end of the RUL. However, in general, the system may not survive to the end of the predicted RUL and decision-maker must be able to determine the value of waiting when the system may fail (jump to ruin) before the end of the predicted RUL, i.e., obtaining the optimal time to wait to maintain a system based on maximizing the value obtained from running the system while minimizing the negative effect of the risk of system failure. In order to obtain the value at each time, the following steps are performed:

1) Perform steps 1 to 7 of the algorithm in Section 2.1 to obtain the value of waiting until the end of the remaining useful life.
2) After obtaining the value of waiting until the end of the remaining useful life, the decision-maker is then interested in knowing the value of waiting for different times between the prognostic indication and the end of the remaining useful life. In steps 1 and 2 in the algorithm in Section 2.1, the uncertainties are estimated and the uncertain quantities are propagated in time. The summation of the cost avoidance and the revenue results in a matrix of the values of V_M. The rows of the matrix correspond to the paths being simulated, and the columns correspond to the values of V_M at different time steps. In order to obtain the value of waiting for a time less than the end of the remaining useful life, we use the same matrix for V_M but consider a smaller number of columns. For example, if the value of waiting for 2 time steps is to be calculated for the example in Section 2.2, we consider the matrix in Table 1 until $t = 2$.
3) After obtaining the matrix for V_M corresponding to the time frame under consideration, the time steps are discretized in order to perform subsequent calculations. The time steps correspond to times at which maintenance can be performed. This will give the basis for deciding whether waiting is beneficial or not; it is an integral part of analyzing the flexibility enabled by PHM.
4) Consider the new matrix for V_M and the corresponding time steps, and repeat steps 4 and 5 from the algorithm in Section 2.1 until reaching time $t = 0$.
5) At this point, the value of the waiting option can be calculated for the new time frame considered using step 7 from the algorithm in Section 2.1.
6) Consider a new time frame and truncate the matrix of V_M at a time corresponding to this new time frame and repeat steps 3 to 5 above to obtain the value of waiting.
7) Repeat steps 1 to 7 to obtain the value of waiting for different times ranging from the time of prognostic indication to the time corresponding to the end of the remaining useful life.

3. Case Study: Waiting Options for Maintaining Wind Turbines

In this section, we demonstrate the wait-to-maintain option on a US land-based wind turbine. The data considered in this section corresponds to the same farm described in Section 1. In this section, we consider a subset of 7 wind turbines from the farm and use the power data provided by the manufacturer of the turbines from a one-year period. The case study first considers the value of waiting for one subsystem indicating remaining useful life, and then two systems with prediction of remaining useful life, using the waiting option to define a dynamic maintenance threshold.

Figure 7 is an actual power curve of a wind turbine (a single turbine over one year of operation) that shows the power produced as a function of wind speed.[7] Each data point is the average over 10 minutes. The data was then scaled to 600KW (the turbine size considered in the case study). The scaling does not affect the power–wind relationship.

[7] In an ideal scenario, the power should increase after the wind speed reaches 3.5 m/s, keep increasing until reaching the cut-off speed of 13 m/s, and then stay constant, thus exhibiting an S-curve. The noise in the actual curve is due to multiple causes: noise in the data from sensors, malfunctions of the turbine, lag between the control signal and the actual blades turning (which will result in a delay in reaching the maximum power for a particular wind speed), and downtime.

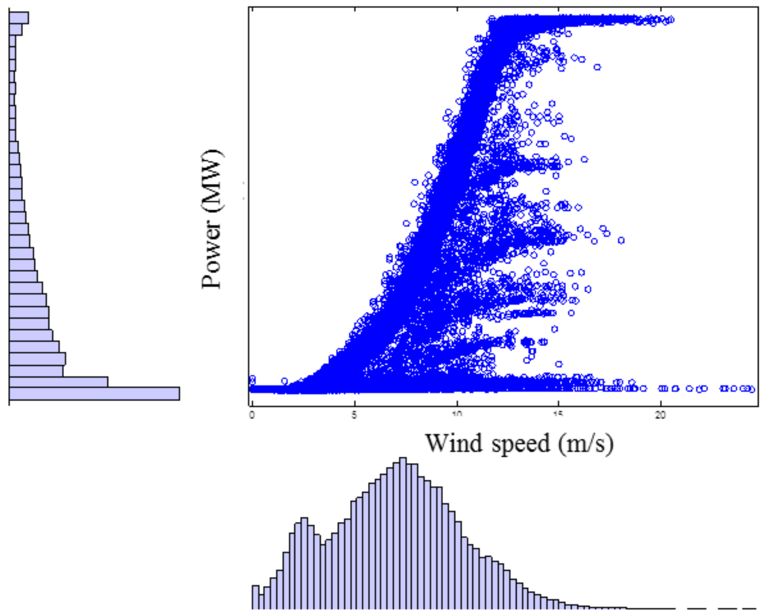

Figure 7. Actual wind turbine power curve over one year of operation.

3.1. The Value of Waiting

Uncertainties are at the heart of the maintenance options valuation. In fact, if there was no uncertainty, it would be trivial to determine the value of the option. Including these uncertainties is a central element of the analysis that causes the multiple stochastic paths after prognostic indication as discussed in Section 2. The revenue from a turbine (equation (7)), depends on the capacity factor, C_f. The capacity factor is the ratio of the actual power output of a turbine over a period of time to the maximum theoretical power that the turbine can generate:

$$R = (N_{dy})(24)(WT_{PR})(C_{EH})(C_f) \quad (7)$$

where N_{dy} is the number of days, WT_{PR} is the power rating of the turbine (600KW in this example), and C_{EH} is the cost of energy (assumed to be \$0.05/KWh). To highlight the uncertainty in the power generated by the turbine (which affects revenue), the power is averaged every day and plotted in Figure 8. Dividing this power by 600KW gives the capacity factor.

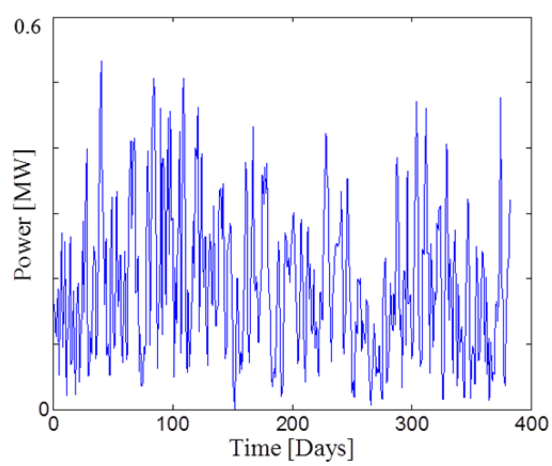

Figure 8. Daily power for a healthy turbine.

After obtaining the time-series for the capacity factor, we estimate the drift and shock parameters for the revenue from data using simulated maximum likelihood estimation and obtain a drift of 0.18 and a shock of 0.5 (the capacity factor is modeled in this manner since wind speeds change over the course of the year). Using the drift and shock parameters, we can propagate the uncertainty with time (assuming an initial capacity factor of 0.33, the average capacity factor for similar turbines [22]), and simulate 100 paths for the revenue in equation (7). In the case study, we use 3 Laguerre polynomials, 100 paths, and assume C_M to be the same as the cost of unscheduled maintenance, which is assumed to be $11,640 (the cost of a pitch mechanism in 2006 dollars [23]).

We first consider one failure mode with a remaining useful life of 37 days. The cost avoidance is assumed to have an initial value of $6,984, and a drift component of -0.6 and shock of 0.25 for the cost avoidance (the drift and the shock are assumptions; they can be estimated when historical maintenance data is available). In this example, we assume that at 80% of the RUL (80% of 37 days), there is an increase in the risk of collateral damage if the turbine fails. This will result in a cost of failure that is larger than the cost of the failure of the subsystem under consideration (the one that indicated an RUL). An example of collateral where a failure of one subsystem can lead to a cost that is larger than the subsystem that failed is discussed by Ragheb and Ragheb [24]. In [24], the failure of a gearbox can induce damage in other subsystems or even damage the entire turbine; this will eventually lead to a cost of failure that is higher than the cost of failure of the gearbox alone. While we make an assumption of an increase in the cost of failure in this example, this cost can be estimated from historical data when failure data is recorded and used to draw conclusions on the cost and risk of inducing collateral damage in other subsystems in the turbine. In this example, we assume that the risk of collateral damage increases at 80% of the RUL, and we model an increase in the risk of failure by inducing a jump in the cost avoidance that is simulated by a sharp drop in the mean of the cost avoidance. In other words, as we get closer to the end of the RUL, the cost avoidance decreases (because of the risk of collateral damage increases). We simulate this change with a jump process, which is a stochastic process that has a jump at a specific time instant:

$$dX_t \ = \ \mu X_t dt + \sigma X_t dt W_t + X_t dq_t \tag{8}$$

Equation (8) is similar to equation (3) but with the addition of the last term, which is a Poisson process, to model a jump in the value. We assume μ = –0.6 and σ = 0.25. The values are again assumed, but they can be estimated when historical data is available.

In option pricing, jump diffusion models are used to account for discontinuities in the behavior in asset pricing. Discontinuities can be caused by sudden fluctuations in the market such as a market crash or a sudden deviation in the value of assets. The jump diffusion process is based on the Poisson process, which can be used for modeling systematic jumps caused by surprise effect. An example with the embedded assumptions on the jump diffusion process can be seen in [21].

Once uncertainties in the revenue and cost avoidance are propagated in time, we obtain the paths for V_M (from equation (1)). At this stage, we apply the LSM algorithm and compare V_M to C_M (C_M is assumed to be $11,640 in this example) to obtain the value of the waiting option.

Figure 9 shows the value of waiting for the subsystem in the turbine under consideration where a remaining useful life of 37 days is assumed. The horizontal axis on Figure 9 is the waiting time; at time 0 a prognostic indication is obtained and the system will fail at 37 days. The vertical axis on Figure 9 is the value of the waiting option; this is the additional value that the decision-maker obtains from running the turbine through its remaining useful life. Note that this is not the value of PHM (such a value is better captured by an ROI model, e.g., [25]), but rather it is the additional benefit from capturing the upside effect of uncertainties (such as the high probability of high wind speeds). One way to calculate the value of PHM using the waiting options is by compounding the analysis of multiple options across multiple subsystems over the life of the turbine and comparing it to the cost of implementing and sustaining PHM. Figure 9 shows the value of waiting from 5 days to 37 days.

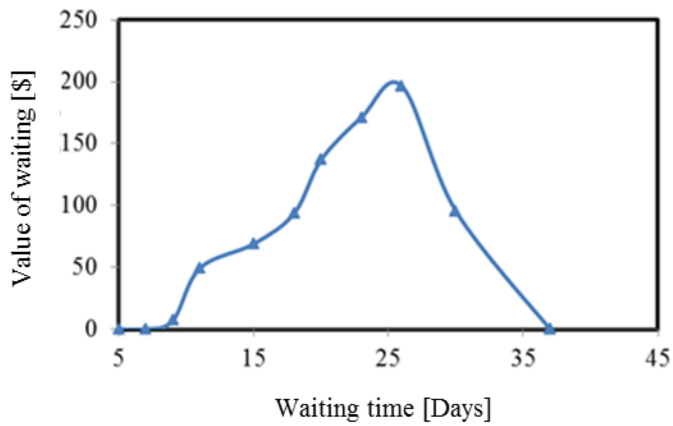

Figure 9. Value of waiting for a 37 day predicted RUL.

The value of waiting in Figure 9 is 0 for the first 7 days. This means that waiting for 7 days will not generate additional value from PHM. This is due to the fact that the cumulative revenue generated in 7 days will not compensate for the cost of failure. However, after 7 days, the cumulative revenue generated from running the turbine compensates for the risk of failure; this is when the value of waiting starts to increase. The value of waiting has a peak at approximately 25 days, with a value of waiting slightly larger than $200. The value from waiting is the additional benefit that the decision-maker obtains from PHM (knowledge of the remaining life and running the system through the RUL). The magnitude of the values in Figure 9 may seem small compared to the system's cost. However, this example calculates the value of the waiting option for one particular subsystem experiencing one failure. In reality, the PHM system will generate the waiting option over the lifetime of the turbine (many possible failures). Moreover, there are more maintenance options beyond waiting that can potentially create more value. Finally this example is for one turbine only; in reality, the turbine is part of a fleet of turbines (i.e., a wind farm) and the value would be aggregated for the lifetime of the fleet.

We note that the value of waiting starts at 0, since the valuation starts when the prognostic indication is obtained. At this time, there is no potential to account for the cost of failure from running the turbine since there is no generation at time 0. The more we run the turbine through the RUL, the more revenue we will generate (it is cumulative). Hence, at time 0, the value of waiting cannot be larger than 0. Furthermore, the value cannot be negative because we are representing the opportunity of the upside effect generated from waiting. If the cost of failure is larger than the revenue, then waiting has no value (0). This is the asymmetry in options. In other words, we are not representing the difference in cost of failure and revenue, but instead the additional value we get from the option (i.e., waiting). Note that the value in this example decreases after 25 days (and eventually goes to 0) because of the assumptions made that the cost of failure increases at 80% of the RUL. Had this assumption not been made (i.e., if there was no effect of collateral damage on cost), the value of waiting would have increasing until the end of the RUL. For example, the value of waiting for 3 time units in the example in Section 2.2 is 0.3388 (and not 0). When the cost of maintenance and failure is available from historical data, this assumption can be relaxed and estimates from actual data can be used to define the distributions that are assumed in this example.

3.2. Placing the Dynamic Maintenance Threshold

Now we demonstrate the value of waiting by considering two turbines exhibiting prognostic indications for different components at different times. Considering a timeline from 0 to 53 days, the first turbine has a failure on day 37, and the second one on day 53. The prognostic indication is assumed to be obtained 37 days prior to failure for both turbines. When the prognostic indication of the first turbine is obtained, the decision-maker does not have any information about the time when the second turbine will fail. The value of waiting is first obtained from the least squares Monte Carlo algorithm for the data from turbine 1. The uncertainty parameters in the revenue are estimated from the power data (similar to the example in Section 3.1) for 37 days prior to failure for turbine 1 (drift of 0.198 and shock of 0.58). The value of waiting starts increasing on day 9. On day 16 (37 days before the end of life of turbine 2), a prognostic indication for turbine 2 is obtained. After the failure of turbine 1, a penalty equivalent to the cost of generation lost by 1 turbine is imposed. The decision-maker is interested in knowing when to maintain given the prognostic information and the uncertainties associated with the operation of the system. The time for maintenance can be formulated as an integer programming optimization problem:

$$arg\,\max_{t}\left(\sum_{1}^{n} c_{t_n} X_0\right) \tag{9}$$

where t is the time at which maintenance can be supported; c_t is the value of the wait-to-maintain option at time t; X_0 is the maintenance decision (X_0 = 0 is the decision to wait and X_0 = 1 is the decision to maintain).

Figure 10 shows the value of waiting with the annotation of different events.

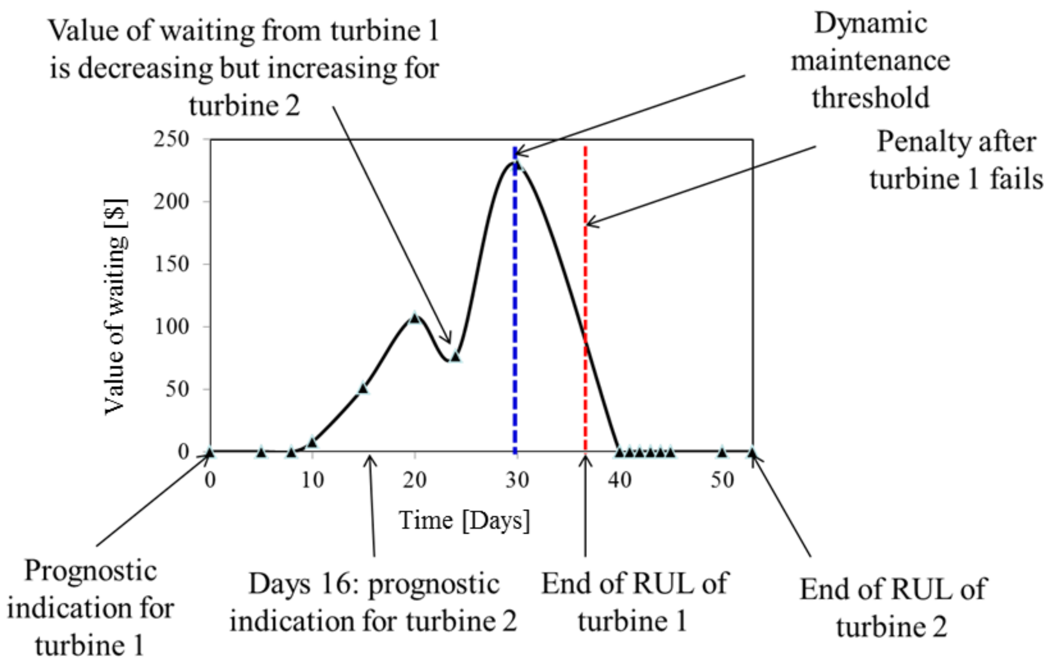

Figure 10. Dynamic threshold based on the value of the option to wait for two turbines.

The value of waiting exhibits two local maxima in Figure 10. The first one is influenced by the maximum waiting value for turbine 1. But when turbine 2 indicates an RUL, the results show that the dynamic threshold that maximizes the value of waiting corresponds to day 30. This is 7 days before the failure of turbine 1, when a penalty starts to accrue. In other words, if turbine 1 had been considered individually, it would have been maintained on day 20; but when new information about a prediction of failure in turbine 2 becomes available, the value of waiting is aggregated since the revenue generated by both turbines is now compared to the cost of failure. The combined value of waiting dictates the optimal maintenance threshold for the assumed conditions, which we call the dynamic maintenance threshold. The decrease after day 30 is due to the proximity of the end of the RUL of turbine 1 and the failure imposed if a turbine is down. One result from this example is that the optimal value obtained from using the PHM system is found in not waiting until the end of the RUL for either turbine.

The methodology is applicable to multiple systems in a turbine and multiple turbines in a fleet (farm). It is able to set a threshold based on the uncertainties and prognostic information. It should be noted that the analysis in this example presents a calculation for the value of waiting that is projected for the predicted RUL. In reality, the problem is dynamic: as new information about the degradation or the uncertainties associated with the operation is obtained, the model can be updated and the threshold can be set dynamically to maximize the value obtained from the PHM system.

4. SUMMARY AND CONCLUSIONS

Engineering systems are incorporating PHM to enable high availability and lower life-cycle costs. Wind turbines have emerged as candidates for PHM, and there has been a substantial amount of work done to develop PHM methodologies to predict the onset of failures. This paper presents the concept of maintenance options and a method to quantify the waiting option for wind turbines, which is a means of quantifying the value of decisions after prognostic indication.

By formalizing maintenance options within a real options framework and developing a cost-benefit-risk model that incorporates the value of flexibility (or options), we provide a model for quantifying the system-level value of PHM after prognostic indication. The model evaluates individualized maintenance policies for different system instances and quantifies the value of PHM at all points in time from prognostic indication to the end of the remaining useful life. The model can be updated in real-time and can generate a value of the waiting to maintain option at any time. In this paper, we assume that an RUL is obtained from a PHM system; the methodologies for predicting this RUL are not discussed and are outside the scope of this work.

The approach developed in this paper is also demonstrated in an example to set a dynamic maintenance threshold. This approach can be applied to multiple systems that might not have the same prognostic distance or the same failure modes, and the maintenance threshold is based on maximizing the value of waiting across a fleet of systems. The example discussed in this paper considers two turbines with one failure each. The methodology can be applied to a fleet of turbines with multiple failures and potentially more options than the waiting option. When the value from the option is maximized, a new threshold of maintenance can be set; we define it as the dynamic maintenance threshold. The methodology can also be used to support outcome-based contracts. Furthermore, the maintenance options provide a framework for system-level decision support, whereby the decision-maker can base the management of the system after prognostic indication on a dollar value—a capability that does not yet exist and can also enable pricing availability-based contracts.

The waiting option for wind turbines is only one type of maintenance option. There are potentially numerous options. Reducing the load on a system when a prognostic indication is obtained to reduce degradation is an example of another option. The valuation of such options may be complex, but their valuation is critical for justifying the implementation of PHM economically.

The methodologies presented in this paper help to address the value of decisions after prognostic indication. There are a number of directions for future research. While the convergence of the valuation algorithm is addressed in [21], deriving another relationship for the value of maintenance or using other stochastic processes may require convergence to be reconsidered. Furthermore, the time-dependent uncertain quantities are assumed to be described by Brownian motion in this paper. While Brownian motion has been used to describe the time-dependent uncertainty in wind speed, other types of uncertainties may best be modeled by different stochastic differential equations. When historical data is available, the uncertain parameters can be estimated directly from the data and the assumptions can be relaxed. Finally, we model a sharp increase in the cost of failure using a jump process. This assumption can also be relaxed when historical maintenance data becomes available.

Finally, we note that the model presented in this paper is discussed in the context of a single prediction; we assume that remaining useful life is predicted for the system, and we compute the value of waiting for that particular time. However, the methods presented in this paper are applicable whenever new information is available: every time a new prediction is obtained or when an uncertainty is resolved, the valuation model can be run again to obtain the value of waiting for every time step until the predicted RUL.

ACKNOWLEDGMENTS

The authors would like to thank the National Science Foundation, Division of Design and Manufacturing Innovation (Grant Number: CMMI-1129697), for their support. The authors would also like to thank the more than 100 companies and organizations that support research activities at the Center for Advanced Life Cycle Engineering (CALCE) at the University of Maryland annually and, specifically, the CALCE Prognostics and Health Management Consortium members.

References

[1] U.S. Department of Energy. 20% wind energy by 2030, increasing wind energy's contribution to U.S. electricity supply. *Energy Efficiency and Renewable Energy 2008, Tech. Rep. DOE/GO-102008-2567*.

[2] Kühn P. Big experience with small wind turbines—235 small wind turbines and 15 years of operational results. *Proceedings of EWEC 2007*, Milano, Italy.

[3] Asmus P and Seitzler M. The wind energy operations and maintenance report. *Wind Energy Update*, 2010.

[4] Sandborn P and Myers J. Designing engineering systems for sustainment. *Handbook of Performability Engineering*, ed. K.B. Misra, Springer, pp. 81-103, London, 2008.

[5] The wind energy operations and maintenance report, *Wind Energy Update*, 2011.

[6] Arabian-Hoseynabadi H, Oraee H, and Tavner P. Failure modes and effects analysis (FMEA) for wind turbines. *International Journal of Electrical Power & Energy Systems* 2010*;* **32**, 817-824.

[7] Pecht M. *Prognostics and Health Management of Electronics*. Wiley, 2008, Hoboken, NJ.

[8] Cheng S, Azarian M, and Pecht M. Sensor systems for prognostics and health management," *Sensors* 2010; **10**, 5774-5797.

[9] Huang Q, Jiang D, Hong L, and Ding Y. Application of wavelet neural networks on vibration fault diagnosis for wind turbine gearbox. *Proc. 5th International Symposium on Neural Networks* 2008, 313-320.

[10] Ebersbach S, Peng Z, and Kessissoglou N. The investigation of the condition and faults of a spur gearbox using vibration and wear debris analysis techniques. *Wear* 2006; **260,** 16-24.

[11] Adams D, White J, Rumsey M, and Farrar C. Structural health monitoring of wind turbines: method and application to HAWT. *Wind Energy* 2011; **14**, 603-623.

[12] Hameed Z., Hong Y, Cho Y, Ahn S, and Song C. Condition monitoring and fault detection of wind turbines and related algorithms: A review. *Renewable and Sustainable Energy Reviews* 2009; **13**, 1–39.

[13] Bonissone P. Domain knowledge and decision time: a framework for soft computing applications. *International Symposium on Evolving Fuzzy Systems* 2006, Ambleside, UK.

[14] Dixit A and Pindyck R. *Investment Under Uncertainty*, Princeton University Press 1994, Princeton, NJ.

[15] Copeland T and Antikarov V. *Real Options: A Practitioner's Guide*, Texere 2001, New York, NY.

[16] Wallace S. Stochastic programming and the option of doing it differently. *Annals of Operations Research* 2010; **177**, 3-8.

[17] Koide Y, Kaito K, and Abe M. Life-cycle cost analysis of bridges where the real options are considered. *Proceedings of the Current and Future Trends in Bridge Design, Construction and Maintenance* 2001, 387-395.

[18] Miller L and Park C. Economic analysis in the maintenance, repair, and overhaul industry: an options approach. *The Engineering Economist* 2004; **49**, 21-41.

[19] Jin X, Li L, and Ni J. Options model for joint production and preventive maintenance system. *International Journal of Production Economics* 2009; **119**, 347-353.

[20] de Neufville R and Scholtes S. *Flexibility in Engineering Design*, MIT Press, Cambridge, Massachusetts, 2011.

[21] Longstaff F and Schwartz E. Valuing American options by simulation: a simple least-squares approach, *The Review of Financial Studies* 2001; **14**,113-147.

[22] Andrawus J, Watson J, Kishk M, and Adam A. The selection of a suitable maintenance strategy for wind turbines. *International Journal of Wind Engineering* 2006; **30**, 471-486.

[23] Kahrobaee S and Asgarpoor S. Risk-based failure mode and effect analysis for wind turbines. *North American Power Symposium (NAPS)*, Boston, MA, August 2011

[24] Ragheb A and Ragheb M. Wind turbine gearbox technologies. *Fundamentals and Advanced Topics in wind Power* 2011, **8**, 189-206.

[25] Feldman, K. Jazouli, T. and Sandborn, P. A methodology for determining the return on investment associated with prognostics and health management, *IEEE Trans. on Reliability* 2009; **58**, 305-316.