The stochastic opportunistic replacement problem, part III: improved bounding procedures

(article starts on next page)

The stochastic opportunistic replacement problem, part III: improved bounding procedures

Efraim Laksman[1] · **Ann-Brith Strömberg**[1] · **Michael Patriksson**[1]

Abstract
We consider the problem to find a schedule for component replacement in a multi-component system, whose components possess stochastic lives and economic dependencies, such that the expected costs for maintenance during a pre-defined time period are minimized. The problem was considered in Patriksson et al. (Ann Oper Res 224:51–75, 2015b), in which a two-stage approximation of the problem was optimized through decomposition (denoted the optimization policy). The current paper improves the effectiveness of the decomposition approach by establishing a tighter bound on the value of the recourse function (i.e., the second stage in the approximation). A general lower bound on the expected maintenance cost is also established. Numerical experiments with 100 simulation scenarios for each of four test instances show that the tighter bound yields a decomposition generating fewer optimality cuts. They also illustrate the quality of the lower bound. Contrary to results presented earlier, an age-based policy performs on par with the optimization policy, although most simple policies perform worse than the optimization policy.

Keywords Stochastic programming · Mixed binary linear optimization · Maintenance optimization · Stochastic opportunistic replacement problem

Mathematics Subject Classification 90C11 · 90C15 · 90B25

1 Introduction

This article studies a generalization of a multi-component maintenance scheduling problem with deterministic component lives (the so-called opportunistic replacement problem, ORP;

✉ Ann-Brith Strömberg
  anstr@chalmers.se

  Efraim Laksman
  efraim.laksman@ronneby.se

  Michael Patriksson
  mipat@chalmers.se

[1] Department of Mathematical Sciences, Chalmers University of Technology and University of Gothenburg, Göteborg, Sweden

see Almgren et al. 2012), to the case when component lives are considered to be stochastic variables. The resulting *stochastic opportunistic replacement problem* (SORP) was introduced and analyzed in Patriksson et al. (2015a, b).

We consider a stochastic programming approach for the minimization of the expected cost of maintenance over the remaining planning horizon, assuming that the failure risk functions are non-decreasing with increasing component ages. Our problem formulation leads to a two-stage approximation of a multi-stage stochastic programming problem, which can be decomposed, akin to integer Benders' decomposition (Benders 1962). The decomposition requires lower bounds on the value of the *recourse function*, which represents the *expected future maintenance cost* of a specific replacement decision at the current time, given the current ages of the components. We present a means to compute such lower bounds; in particular, we improve upon the corresponding bounds from Patriksson et al. (2015b), resulting in a more efficient decomposition method.

Two identical systems with identical component ages were considered in Patriksson et al. (2015b, Prop. 2); it claims that the value of the recourse function of either system is upper bounded by that of the other, plus a non-negative value defined by the (current) replacement decisions for the two systems. Based on our improved bounds on the recourse function, we state and verify a tighter bound than the one claimed.

Beside the improved efficiency of the decomposition, we establish a lower bound on the expected cost of the solution to the SORP; it is useful since our method does not in general provide an optimal solution to the SORP.

A review over current research is given in Sect. 2. A formal definition of the problem studied is given in Sect. 3, and in Sect. 4 it is formalized how time is discretized. Section 5 contains a lower bound on the value of the SORP. Sections 6, 7, and 8 contain (the new) bounds on the value of the recourse function, a brief description of the ORPIL (a deterministic optimization problem), and a description of a decomposition method, respectively. Together, these three sections form a presentation of the optimization policy. Section 9 contains some numerical tests comparing the new optimization policy with the old one. Section 10 concludes the paper.

2 Literature survey

While the literature abounds with articles on aspects of reliability and maintenance of multi-component systems with deteriorating and/or stochastically failing components, we provide below a summary of the most interesting and (mostly) recent related articles in the area, sorted along a time-line.

van Noortwijk and Frangopol (2004) describe the difference between *proactive* and *reactive* preventive maintenance (PM) actions, applied before and after, respectively, an indication of a deterioration. As the rationality behind these principles have been questioned, an optimal PM strategy based on lifetime reliability and life-cycle costs is established. The former model was applied by the Netherlands Ministry of Transport, Public Works and Water Management (Rijkswaterstaat); the latter contributed to the further development of a bridge management methodology set up by the UK Highways Agency.

Deloux et al. (2009) develop a *predictive* maintenance policy for a continuously deteriorating system subject to stress. The system considered possesses two failure mechanisms: an excessive deterioration level and a shock. To optimize the maintenance policy, a combined *statistical process control* (SPC) and *condition-based maintenance* (CBM) approach

is proposed. The CBM is used to inspect and replace the system according to the deterioration level. SPC is used to monitor the *stress covariate*. In order to assess the performance of the proposed policy and to minimize the long-run expected maintenance cost per time unit, a mathematical model for the system cost is derived. An analysis of the numerical results highlights the properties of the proposed policy w.r.t. the different maintenance parameters.

Laggoune et al. (2009) propose a PM planning approach for a multi-component series system subjected to random failures, the cost rate being minimized under a general life distribution. The expected total cost comprises corrective and preventive costs (w.r.t. the components) as well as common costs (w.r.t. the production loss during system shutdown). When the system is down (correctively or preventively) the opportunity to preventively replace non-failed components is considered. A solution procedure combining Monte-Carlo simulations and a heuristic search is proposed and applied to component replacement in the hydrogen compressor in an oil refinery.

Wang (2012) recognizes that spare parts demands are difficult to predict based on historical data of parts usages, wherefore an optimal *inventory control policy* may be difficult to obtain. A joint optimisation of the inventory control of spare parts and the PM inspection interval is proposed. *Stochastic cost models* for spare parts inventory and maintenance are derived and enumeration is employed to find optimal solutions over a finite time horizon. A *delay-time concept* is developed for inspection modelling and used to derive the probabilities of numbers of failures and defective items identified at a PM epoch. The inventory model follows a periodic review policy, the demand being governed by the need for spare parts due to maintenance.

Carlos et al. (2013) argue that the maintenance frequency of wind farms is partly determined by wind velocity variations, which partly determine the degradation progress. Wind measurements and Monte-Carlo sampling are used to estimate the power generated, with a cost associated to production loss during maintenance. The bi-objective goal of the wind farm maintenance strategy is to minimize the total cost and maximize the energy produced, maintenance frequency being a decision variable.

Jha et al. (2013) develop a maintenance scheduling optimization model for bridge infrastructure—utilizing a prediction model to account for stochastic aspects of deterioration—as well as an experimental procedure for examining bridge deterioration. The importance of regarding structural deterioration is discussed, and the work intend to imply environmentally sustainable structures. Specifically, the use of composite materials to replace steel rebars—which are due to corrosion—within bridge decks is presented as a sustainable alternative, to reduce the need for repair and maintenance.

Stochastic control methods have a long history in *risk management* and *life-cycle cost procedures*. Papakonstantinou and Shinozuka (2014) combine stochastic control methods and Bayesian principles into *partially observable Markov decision processes* (POMDPs) that expand available policy options compared to some state-of-the-art methods. POMDPs enable optimum decisions, based on the best possible knowledge of a decision-maker at each time. The problem of finding optimal policies for the maintenance and management of aging structures through a POMDP framework with large state spaces is modelled and solved; the framework is formed using stochastic, physically based models. An example of a corroded existing structure is presented; it is based on non-stationary POMDPs, for an infinite and a finite horizon case with 332 and 14,009 states, respectively.

Comparea et al. (2015) propose and compare maintenance optimization techniques based on *genetic algorithms* (GA), the parameters of the maintenance model being affected by uncertainty and the fitness values represented by cumulative distributions. A method to rank the uncertain fitness values and a novel *Pareto dominance concept* are developed. The GA-

based methods are applied to a practical case concerning the setting of a CBM policy on the degrading nozzles of a gas turbine operating in an energy production plant.

Do et al. (2015) consider a proactive CBM technique, with perfect and/or imperfect maintenance actions, in a deteriorating system. Perfect maintenance actions restore the system to an 'as good as new' state—often at a high cost. A positive impact of imperfect maintenance actions is the relatively low cost; negative impacts are that (i) the system is restored to a state between 'as good as new' and 'as bad as old', and (ii) each action may accelerate the system's deterioration speed. An *adaptive maintenance policy* is proposed which, at each inspection time needed, selects optimal maintenance actions, the time interval between successive inspection times being determined w.r.t. a *remaining useful life* (RUL); the use of the policy is illustrated by a numerical example.

Gunn and Diallo (2015) consider systems operating in critical environments, where operators and/or regulators often specify—for major components—replacement intervals, within which failure cannot occur. A preventive replacement of a component yields an opportunity to replace other components that are within their replacement intervals—to avoid repeating the large teardown cost in short term. For a fixed time horizon, the *opportunistic indirect grouping of periodic events problem* is represented by a tree of possible combinations of replacements of major component. A *depth-first shortest path algorithm* is developed; for moderate numbers of components, out of millions of nodes created only a small portion need to be examined. For larger numbers of components and longer time horizons, the depth-first search still rapidly finds improving solutions and serves as a good heuristic approach.

Aghezzaf et al. (2016) consider the integration of production and maintenance planning in a failure-prone manufacturing system. The system's operating state (i.e., age) is stochastically predictable and it can receive PM during preplanned periods. PM is imperfect: the system is restored to a state between 'as bad as old' and 'as good as new', the latter reachable only by an overhaul. An *integrated production and PM planning model*—accounting for the system's manufacture capacity and operational reliability state—is formulated using mixed-integer non-linear optimization, and then reformulated as an extended mixed-integer linear program. Since the reformulation is computationally demanding, a 'fix-and-optimize' procedure is proposed, that utilizes some properties of the original model. The resulting procedure is tested—with good results—on instances adapted from the literature.

Atashgar and Abdollahzadeh (2016) formulate a *joint redundancy and imperfect block opportunistic maintenance optimization model*, to determine the redundancy level and a maintenance strategy that simultaneously minimizes the wind farm loss-of-load probability and life-cycle cost. Reliability thresholds are introduced for imperfect maintenance of failed and working turbines, and preventive dispatch of maintenance teams. A simulation method developed evaluates the performance measures of a wind farm system w.r.t. different types of wind turbines, maintenance activation delays and durations, and a limited number of maintenance teams. The influence on the performance of the wind farm, of the assumptions and the parameters of the simulation model, is discussed based on a sensitivity analysis. Pareto optimal solutions are derived using a multi-objective *particle swarm optimization algorithm*. A comparative study with a commonly used maintenance policy shows that the proposed strategy significantly reduces maintenance costs and loss-of-load probability.

Cherkaoui et al. (2016) assess the economic performance of CBM strategies through their long-run expected maintenance cost rate criterion as well as robustness. A cost model is developed to quantify the economic performance and robustness to assess CBM strategies. Two representative strategies—periodic and quantile-based—for inspection and replacement are compared, revealing the factors affecting their performance and robustness the most.

Rasmekomen and Parlikad (2016) describe an approach to optimize CBM of multi-component systems, where the state of certain components may affect the rate of degradation of other components. A real case is presented: an industrial cold box in a petrochemical plant; data collected on the fouling of its tubes show that the extent of fouling of one tube affects the rate of fouling of other tubes due to overloading.

Shi and Zeng (2016) describe an *opportunistic PM strategy* and analyze *multi-component systems with stochastic dependence* (similar to that by Rasmekomen and Parlikad 2016). Assuming that measures are available of the RULs of components as well as of the impact of a component's degradation on the RULs of other components, filtering theory is used for predictions. An optimization model determines timings and groupings, based on an opportunistic maintenance principle: at failure maintenance is performed also on other components if time lies in an "opportunity zone". An optimal balance is sought between the cost of performing maintenance too early on some components, and the gain from opportunistic coordination. Tests are limited to three-component systems.

Alaswad and Xiang (2017) review CBM optimization models for stochastically deteriorating single- and multi-unit systems. A CBM strategy collects and assesses real-time information, and suggests maintenance decisions based on the system's current condition. In recent decades, research on CBM has grown rapidly due to a rapid development of computer-based monitoring technologies. Research studies have shown that CBM—if planned properly—can effectively improve equipment reliability at reduced costs. The review is emphasised on mathematical modeling and optimization approaches. Focus lies on optimization criteria, inspection frequency, maintenance degree, and solution methodology. Since the modeling choice for the stochastic deterioration process greatly influences CBM strategy decisions, the CBM models are classified based on the underlying deterioration processes, namely discrete-and continuous-state deterioration, and the proportional hazard model.

3 Problem definition

Consider a system composed by a set $\mathcal{N} = \{1, \ldots, N\}$ of serially connected components, so that a component must be replaced whenever it reaches a failed state. The life of an individual component may be stochastic, but is independent of both failures and lives of other individual components in the system. (The case in which all individuals possess deterministic lives is treated in Patriksson et al. 2015a.) We will assume that at least one component possesses a stochastic life, and that the probability distributions for the lives of all the components are known.

Let $\mathcal{V}$ be the set of system states, where each state $(s, \boldsymbol{\xi}, \mathbf{a}) \in \mathcal{V}$ is composed by the time $s \in \mathbb{R}_+$, the vector $\boldsymbol{\xi} \in \mathbb{B}^N$ of component states [$\xi_n = 0$ ($\xi_n = 1$) if component $n \in \mathcal{N}$ is functioning (in a failed state)] at time s, and the vector $\mathbf{a} \in \mathbb{R}_+^N$ of ages at time s (a_n denotes the age of component $n \in \mathcal{N}$).

We denote the set of *feasible* maintenance decisions with respect to the state $(s, \boldsymbol{\xi}, \mathbf{a}) \in \mathcal{V}$, by $\mathcal{X}_{\boldsymbol{\xi}} = \{ \mathbf{x} \in \mathbb{B}^N \mid \mathbf{x} \geq \boldsymbol{\xi} \}$. A *decision policy* may be used to make a maintenance decision for a system in a given state.

Definition 1 (*Maintenance policies*) A *deterministic maintenance policy* is a function $\mathbf{p} : \mathcal{V} \to \mathbb{B}^N$ such that $\mathbf{p}(s, \boldsymbol{\xi}, \mathbf{a}) \in \mathcal{X}_{\boldsymbol{\xi}}$. Let $\mathcal{B}^N$ denote the set of random vectors with outcome space $\mathbb{B}^N$. A *stochastic maintenance policy* is a function $\mathbf{p} : \mathcal{V} \to \mathcal{B}^N$ such that $\Pr(\mathbf{p}(s, \boldsymbol{\xi}, \mathbf{a}) \in \mathcal{X}_{\boldsymbol{\xi}}) = 1$.

While, in Patriksson et al. (2015a, b), the term *maintenance policy* denotes a *deterministic maintenance policy* (stochastic maintenance policies are not employed in Patriksson et al. 2015a, b), it is here used with *either* of the attributes *deterministic* or *stochastic*. Stochastic maintenance policies are employed in the proof of Proposition 1(i), (ii).

Costs corresponding to the different decisions and/or policies must be defined. We assume that the replacement of any component $n \in \mathcal{N}$ incurs a replacement cost c_n, and that a start-up cost d is incurred each time at least one component is replaced. The *current cost* of the maintenance decision $\mathbf{x} \in \mathcal{X}_{\boldsymbol{\xi}}$ is then defined as

$$\sum_{n \in \mathcal{N}} c_n x_n + d \max_{n \in \mathcal{N}} \{x_n\}.$$

Letting $S > 0$ denote the planning horizon, the period of time during which the system must work is defined as the interval $[0, S)$.

Definition 2 (*Stochastic opportunistic replacement problem, SORP*) Given a planning horizon S, a start-up cost d, component replacement costs c_n, and probability distributions of the components' lives, for $n \in \mathcal{N}$, find a deterministic maintenance policy, such that the expected sum of current costs during the planning period $[0, S)$ is at minimum.

A deterministic maintenance policy which finds an optimal solution to the SORP is called an *optimal maintenance policy*.

Definition 3 (*Current problem*) Given a planning horizon S, a start-up cost d, component replacement costs c_n and probability distributions of the components' lives, for $n \in \mathcal{N}$, the system state $(s, \boldsymbol{\xi}, \mathbf{a}) \in \mathcal{V}$, and assuming that an optimal maintenance policy will be used for all maintenance decisions during the time span (s, S), find a maintenance decision $\mathbf{x} \in \mathcal{X}_{\boldsymbol{\xi}}$ such that the expected sum of current costs during the planning period $[s, S)$ is at minimum.

If $\mathbf{x}^*(t, \boldsymbol{\xi}, \mathbf{a})$ is an optimal solution to the current problem at state $(t, \boldsymbol{\xi}, \mathbf{a}) \in \mathcal{V}$, and $\mathbf{p}^*(s, \boldsymbol{\xi}, \mathbf{a}) = \mathbf{x}^*(s, \boldsymbol{\xi}, \mathbf{a})$ for all $(s, \boldsymbol{\xi}, \mathbf{a}) \in \mathcal{V}$, then $\mathbf{p}^*$ is an optimal maintenance policy. Thus, to maintain a system optimally, it is sufficient to solve each current problem that occurs.

4 Discretization of time in the current problem and the life distributions

The current problem at state $(s, \boldsymbol{\xi}, \mathbf{a}) \in \mathcal{V}$ is defined w.r.t. the continuous time interval $[s, S)$. We define the time unit $\delta := \frac{S-s}{T+1}$ and consider the set $\mathcal{T} := \{0, \ldots, T\}$ representing $T + 1$ discrete time points. For each $t \in \mathcal{T}$, any individual of any component which fails during the time interval $[s + t\delta, s + t\delta + \delta)$ is treated as it would fail already at the time $s + t\delta$ (or that the individual is observed to be so close to failure that a replacement is necessary before the time $s + t\delta + \delta$).[1] Hence, the discretized version of the current problem is defined for the discrete set $\{s, s + \delta, \ldots, s + T\delta\}$ of time points—hereafter represented by the set $\mathcal{T}$. Since decisions made at time $t \in \mathcal{T}$ depend on realizations of the system's states at the times $\{1, \ldots, t\}$ and of the decisions made at the times $\{0, \ldots, t-1\}$, the *discretized current problem* is a *multi-stage stochastic programming problem* (Birge and Louveaux 1997, Ch. III.7).

---

[1] Any individual with life $< \delta$ would have a zero life in the discretized model and should be discarded from the problem. Hence, we assume that all individuals possess lives $\geq \delta$.

We define the $N \times N$-matrix $\mathbf{A} := \mathrm{diag}(\mathbf{a})$ and the N-vector $\mathbf{1} := (1, \ldots, 1)^\top$. Then, $\mathbf{A}(\mathbf{1} - \mathbf{x}) \equiv \mathbf{a} - \mathbf{Ax}$ denotes the components' ages *after* the replacement decision $\mathbf{x}$. We let the function $f_t^{\mathbf{p}} : \mathbb{B}^N \times \mathbb{R}^N \to \mathbb{R}$ represent the expected cost incurred by following the stochastic maintenance policy $\mathbf{p}$ over the time steps $\{t, \ldots, T\}$, defined as

$$f_t^{\mathbf{p}}(\boldsymbol{\xi}, \mathbf{a}) := \mathrm{E}_{\mathbf{p}(t, \boldsymbol{\xi}, \mathbf{a})} \left[d \max_{n \in \mathcal{N}} \{ p_n(t, \boldsymbol{\xi}, \mathbf{a}) \} + \mathbf{c}^\top \mathbf{p}(t, \boldsymbol{\xi}, \mathbf{a}) + \mathcal{Q}_t^{\mathbf{a}}(\mathbf{p}(t, \boldsymbol{\xi}, \mathbf{a}), \mathbf{p}) \right],$$

(1a)

where the function $\mathcal{Q}_t^{\mathbf{a}} : \mathbb{B}^N \times \mathbb{B}^N \to \mathbb{R}$ is defined by

$$\mathcal{Q}_T^{\mathbf{a}}(\mathbf{x}, \mathbf{p}) := 0, \tag{1b}$$

$$\mathcal{Q}_t^{\mathbf{a}}(\mathbf{x}, \mathbf{p}) := \mathrm{E}_{\boldsymbol{\eta}} \left[f_{t+1}^{\mathbf{p}}(\boldsymbol{\eta}, \mathbf{a} - \mathbf{Ax} + 1) \right], \quad t = T - 1, \ldots, 0, \tag{1c}$$

and $\boldsymbol{\eta} \in \mathbb{B}^N$ is a random vector of component states and whose distribution is given by $\mathrm{Pr}(\eta_n = 1) = 1 - \mathrm{Pr}(\eta_n = 0) = G_n(a_n - a_n x_n + 1)$, where G_n is computed as in (3), below. The recourse function $\mathcal{Q}_t^{\mathbf{a}}$ represents the expected future costs at time t. In the special case of $\mathbf{p}$ being a deterministic maintenance policy, the right-hand-side of (1a) can be simplified to

$$d \max_{n \in \mathcal{N}} \{ p_n(t, \boldsymbol{\xi}, \mathbf{a}) \} + \mathbf{c}^\top \mathbf{p}(t, \boldsymbol{\xi}, \mathbf{a}) + \mathcal{Q}_t^{\mathbf{a}}(\mathbf{p}(t, \boldsymbol{\xi}, \mathbf{a}), \mathbf{p}).$$

For an optimal maintenance policy $\mathbf{p}^*$, the expected sum of current costs over the times $\{t, \ldots, T\}$ may be expressed as

$$f_t^{\mathbf{p}^*}(\boldsymbol{\xi}, \mathbf{a}) := \mathrm{minimum}_{\mathbf{x},z} \quad dz + \mathbf{c}^\top \mathbf{x} + \mathcal{Q}_t^{\mathbf{a}}(\mathbf{x}, \mathbf{p}^*), \tag{2a}$$

$$\text{subject to} \quad \mathbf{x} \geq \boldsymbol{\xi}, \tag{2b}$$

$$x_n \leq z, \quad n \in \mathcal{N}, \tag{2c}$$

$$\mathbf{x} \in \{0, 1\}^N, \tag{2d}$$

$$z \in \{0, 1\}. \tag{2e}$$

The *current problem* is to find $f_0^{\mathbf{p}^*}(\boldsymbol{\xi}, \mathbf{a})$ [i.e., solve the program (2) for the state $(0, \boldsymbol{\xi}, \mathbf{a})$] and the corresponding optimal solution, typically without explicit knowledge of the optimal policy $\mathbf{p}^*$ [i.e., without minimizing (2a) for each $(t, \boldsymbol{\xi}, \mathbf{a})$].

The lives of the individuals $i \in \mathbb{Z}_+$ of component $n \in \mathcal{N}$ are represented by the continuous[2] and independent identically distributed (i.i.d.) stochastic variables $\hat{T}_{ni}$. Let $\tau_n : \mathbb{R}_+ \to [0, 1]$ denote the probability density function of $\hat{T}_{n1}$, so that its cumulative probability function is given by

$$F_n(t) := \mathrm{Pr}(\hat{T}_{n1} \leq t) = \int_0^T \tau_n(s) \, \mathrm{d}s, \quad n \in \mathcal{N}.$$

The *failure risk function* $G_n : \mathbb{R}_+ \to [0, 1]$, defined as $G_n(a) := \mathrm{Pr}(\hat{T}_{n1} < a + 1 : \hat{T}_{n1} \geq a)$, denotes the probability that component n fails during the next time step, provided that it is in a non-failed state and of age $a \geq 0$; it is calculated as

$$G_n(a) = \frac{F_n(a + 1) - F_n(a)}{1 - F_n(a)}, \quad n \in \mathcal{N}. \tag{3}$$

---

[2] Analogous analyses for discrete (or mixed) random variables can be performed, although with a more complicated notation.

5 A lower bound on the expected costs for maintenance

We propose a lower bound on the optimal value of the SORP, as given by Definition 2. The bound is restricted to instances of the SORP for which each component $n \in \mathcal{N}$ has a stochastic life following the distribution F_n, such that the failure risk function G_n, defined in (3), is non-decreasing.[3] The bound is determined by minimizing the expected replacement costs for each component individually, as well as minimizing the expected start-up costs for replacements individually, followed by a summation of those costs.

In a one component system (i.e., for $n = N = 1$), component n is replaced only when it fails; we denote by ϕ_n the expected number of failures of such a system. In a system for which it holds that $N > 1$, the expected number of times that component $n \in \mathcal{N}$ must be replaced is at least ϕ_n. Hence, the expected cost for replacing component n (excluding start-up costs, d) is at least $c_n \phi_n$. To calculate the value of ϕ_n, let the stochastic variable $Y_n^{(m)} := \sum_{i=1}^m \hat{T}_{ni}$ denote the sum of the lives of the m first individuals of component n, and let Z_n denote the number of times that component n must be replaced, given that it is replaced only when it fails, i.e.,

$$Z_n := \max \left\{ m \in \mathbb{Z} \, \middle| \, \sum_{i=1}^m \hat{T}_{ni} < S \right\} = \max \left\{ m \in \mathbb{Z} \, \middle| \, Y_n^{(m)} < S \right\}.$$

The expected value of the number of replacements of component n, i.e., Z_n, may be found by determining the probability function for Z_n, or as the sum over m of the probability that the mth individual of component n must be replaced;[4]

$$\phi_n = \mathrm{E}[Z_n] = \sum_{m=1}^{\infty} \Pr\left(Y_n^{(m)} < S\right). \tag{4}$$

Finding an analytic expression for ϕ_n is possible only for special cases of the function F_n. Unless $\hat{T}_{n1}$ lack expected value or variance, normal approximations may be used for large enough values of m, since $Y_n^{(m)}$ is a sum of m i.i.d. stochastic variables. Another approach is to simulate outcomes of $Y_n^{(m)}$, and use the resulting relative frequencies to estimate the probabilities $\Pr(Y_n^{(m)} < S)$ for different values of m.

---

[3] When time is treated as continuous rather than discrete, the failure risk function G_n is replaced by the hazard rate function $\lambda_n : \mathbb{R}_+ \to \mathbb{R}_+$; it is for our lower bound assumed to be non-decreasing and defined by $\lambda_n(a) := \lim_{h \to 0^+} \left\{ h^{-1} \Pr(\hat{T}_{n1} < a + h : \hat{T}_{n1} \geq a) \right\} = F_n'(a)(1 - F_n(a))^{-1}$.

[4] The equivalences in (4) hold due to the relations

$$\begin{aligned}
\phi_n = \mathrm{E}[Z_n] &= \sum_{m=1}^{\infty} m \Pr(Z_n = m) = \sum_{m=1}^{\infty} m \Pr(Y_n^{(m)} < S \leq Y_n^{(m+1)}) \\
&= \sum_{m=1}^{\infty} \Pr(Y_n^{(m)} < S \leq Y_n^{(m+1)}) + \sum_{m=1}^{\infty} (m-1) \Pr(Y_n^{(m)} < S \leq Y_n^{(m+1)}) \\
&= \Pr(Y_n^{(1)} < S) + \sum_{m=2}^{\infty} (m-1) \Pr(Y_n^{(m)} < S \leq Y_n^{(m+1)}) \\
&= \sum_{m=1}^{2} \Pr(Y_n^{(m)} < S) + \sum_{m=3}^{\infty} (m-2) \Pr(Y_n^{(m)} < S \leq Y_n^{(m+1)}) = \sum_{m=1}^{\infty} \Pr(Y_n^{(m)} < S).
\end{aligned}$$

We denote by Φ the least expected number of times when at least one component must be replaced; it is given by the case when all components are replaced whenever any component fails; this follows since the failure risk functions (or, hazard rates) are non-decreasing, implying that the expected time until the next system failure may only decrease if not all components are replaced. Hence, the expected start-up cost cannot be less than $d\Phi$. To calculate the value of Φ, let $\hat{T}_i$ denote the time between system failures $i-1$ and i, and assume that all components are replaced at each system failure. Then, $\hat{T}_i$, $i = 1, 2, \ldots$, are i.i.d. stochastic variables with

$$\hat{T}_i = \min_{n \in \mathcal{N}} \left\{ \hat{T}_{ni} \right\}.$$

The distribution function of $\hat{T}_1$ is given by

$$F = 1 - \prod_{n \in \mathcal{N}} (1 - F_n).$$

Analogously to the derivations for individual components, let $Y^{(m)} := \sum_{i=1}^{m} \hat{T}_i$ denote the time for the mth system failure, and $Z := \max \{ m \in \mathbb{Z} \mid Y^{(m)} < S \}$ the number of times replacements must be made. It then holds that [cf. (4)]

$$\Phi = \mathrm{E}[Z] = \sum_{m=1}^{\infty} \Pr(Y^{(m)} < S).$$

We conclude that the sum of the least possible expected start-up costs and the least possible expected replacement costs for all components forms a lower bound on the optimal value of the SORP, namely

$$d\Phi + \sum_{n \in \mathcal{N}} c_n \phi_n.$$

The bound is based on the minimization of the replacement costs for each component, without regards for the start-up costs then incurred, and on minimization of the start-up costs without regards for the replacement costs for the components then incurred. It is only in special cases that both component replacement costs and start-up costs can be minimized simultaneously, so the bound is not tight in the general case.

6 Bounds on the value of the recourse function

Consider the expected minimum total cost resulting from the employment of an optimal maintenance policy as expressed in (2), along with the definition (1b)–(1c) of the recourse function. We will in this section establish bounds on the value of the recourse function which will prove useful in a decomposition method used for solving the current problem. In particular, we reconsider (Patriksson et al. 2015b, Prop. 1), for which we present a new and clearer proof of part (a) [i.e., Proposition 1(i)], and an improved version of part (b) [i.e., Proposition 1(ii)].[5]

---

[5] Patriksson et al. (2015b, Conjecture 1) proposes a possible improvement of Patriksson et al. (2015b, Prop. 1(b)); our improvement [i.e., Proposition 1(ii)] provides an upper bound on the recourse function which is at least as tight as the one conjectured.

We let $G(\mathbf{a})$ denote the probability that a functioning system with component ages $\mathbf{a} \in \mathbb{R}_+^N$ fails within the next time step, i.e.,

$$G(\mathbf{a}) := 1 - \prod_{n \in \mathcal{N}} \left[1 - G_n(a_n) \right].$$

For failure distributions with non-decreasing failure risk functions (see Sect. 4), the next proposition states that the expected future cost is non-decreasing with (non-decreasing) component age and provides an upper bound on the increase of the expected future cost with increasing age.

For any two vectors $\mathbf{a}, \hat{\mathbf{a}} \in \mathbb{R}_+^N$ we define the set $\bar{\mathcal{N}}(\hat{\mathbf{a}}, \mathbf{a}) := \{ n \in \mathcal{N} \mid \hat{a}_n > a_n \}$. Then, for each $t \in \mathcal{T}$, we define the function $K_t : \mathbb{R}_+^N \times \mathbb{R}_+^N \to \mathbb{R}$ as[6]

$$K_t(\hat{\mathbf{a}}, \mathbf{a}) := \sum_{k=1}^{T-t} \left[\left(\prod_{j=1}^{k-1} \left[1 - G(\hat{\mathbf{a}} + j\mathbf{1}) \right] \right) \right.$$
$$\left. \cdot \left(\sum_{n \in \bar{\mathcal{N}}(\hat{\mathbf{a}}, \mathbf{a})} c_n \left[G(\hat{\mathbf{a}} + k\mathbf{1}) - G_n(a_n + k) \right] + d \left[G(\hat{\mathbf{a}} + k\mathbf{1}) - G(\mathbf{a} + k\mathbf{1}) \right] \right) \right].$$

$$(5)$$

Proposition 1 *Assume that each component* $n \in \mathcal{N}$ *has a failure distribution with a non-decreasing failure risk function* G_n, *let* $\mathbf{p}^*$ *be an optimal maintenance policy, let the age vectors* $\mathbf{a}, \hat{\mathbf{a}} \in \mathbb{R}_+^N$ *be such that* $\mathbf{a} \le \hat{\mathbf{a}}$, *let the functions* $\mathcal{Q}_t^{\mathbf{a}}$ *and* $K_t, t \in \mathcal{T}$, *be defined by* (1b)–(1c) *and* (5), *respectively. Then, for all* $t \in \mathcal{T}$ *the following hold:*

(i) $\mathcal{Q}_t^{\mathbf{a}}(\mathbf{0}, \mathbf{p}^*) \le \mathcal{Q}_t^{\hat{\mathbf{a}}}(\mathbf{0}, \mathbf{p}^*)$, *and*
(ii) $\mathcal{Q}_t^{\hat{\mathbf{a}}}(\mathbf{0}, \mathbf{p}^*) \le \mathcal{Q}_t^{\mathbf{a}}(\mathbf{0}, \mathbf{p}^*) + K_t(\hat{\mathbf{a}}, \mathbf{a})$.

Proof (i) The inequality $\mathcal{Q}_t^{\mathbf{a}}(\mathbf{0}, \mathbf{p}^*) \le \mathcal{Q}_t^{\mathbf{a}}(\mathbf{0}, \mathbf{p})$ holds for all policies $\mathbf{p}$, since by definition, an optimal policy $\mathbf{p}^*$ provides the lowest possible expected future cost. Our idea is to find a stochastic policy $\mathbf{p}$ which, as applied to a system with component ages $\hat{\mathbf{a}}$ at time t, simulates an optimal policy $\mathbf{p}^*$ applied to a system with component ages $\mathbf{a}$ at time t, i.e., such that $\mathcal{Q}_t^{\mathbf{a}}(\mathbf{0}, \mathbf{p}) = \mathcal{Q}_t^{\hat{\mathbf{a}}}(\mathbf{0}, \mathbf{p}^*)$ holds. This means that in the policy $\mathbf{p}$, each component $n \in \bar{\mathcal{N}}(\hat{\mathbf{a}}, \mathbf{a})$ fails artificially with some probability, at each time step until it is replaced. The probability that such a component n fails artificially at time $t + k$ should equal the probability that component n—for component ages $\hat{\mathbf{a}}$—fails at time $t + k$, even though—for component ages $\mathbf{a}$—it does *not* fail at time $t + k$, conditioned on that it is not replaced at any of the times $t + 1, \ldots, t + k - 1$. The sought probability is given by

$$H_n(a_n, \hat{a}_n, k) \begin{cases} := \frac{G_n(\hat{a}_n + k) - G_n(a_n + k)}{1 - G_n(a_n + k)}, & \text{if } G_n(a_n + k) < 1, \\ \in [0, 1], & \text{if } G_n(a_n + k) = 1. \end{cases} \qquad (6)$$

Since the failure risk functions are non-decreasing, i.e., for any $0 \le a \le \hat{a}$, the inequalities $0 \le G_n(a) \le G_n(\hat{a}) \le 1$ hold, the relations in (6) imply that

$$G_n(\hat{a}_n + k) = G_n(a_n + k) + \left[1 - G_n(a_n + k) \right] H_n(a_n, \hat{a}_n, k).$$

---

[6] Whenever $u > r$, a product over indices ranging from u to r is assumed to equal 1, while a sum ranging from u to r, as well as over an empty set, is assumed to equal 0.

It follows that the inequalities $0 \leq H_n(a_n, \hat{a}_n, k) \leq 1$ hold. Therefore, $H_n(a_n, \hat{a}_n, k)$ is indeed a probability whenever $k \geq 0$ and $0 \leq a_n \leq \hat{a}_n$.

To construct the stochastic policy $\mathbf{p}$ needed, let t, $\mathbf{a}$, and $\hat{\mathbf{a}}$ be fixed, and define the mutually independent stochastic variables $X_n^{(k)} \in \text{Bin}(1, H_n(a_n, \hat{a}_n, k))$ for $n \in \mathcal{N}$ and $k \in \mathbb{Z}_{>0}$. Further, let $\boldsymbol{\alpha} \in \mathbb{Z}_{>0}^N$ and define, for all $n \in \mathcal{N}$ and $k \in \mathbb{Z}_{>0}$,

$$
\tilde{\xi}_n^{(k)} := \begin{cases} \max\left\{ \xi_n, X_n^{(k)} \right\}, & \text{if } \alpha_n = a_n + k \text{ and } a_n < \hat{a}_n, \\ \xi_n, & \text{otherwise}, \end{cases}
$$

and

$$
\tilde{\alpha}_n^{(k)} := \begin{cases} \hat{a}_n + k, & \text{if } \alpha_n = a_n + k \text{ and } a_n < \hat{a}_n, \\ \alpha_n, & \text{otherwise}. \end{cases}
$$

The stochastic policy $\mathbf{p}$ is then defined as

$$
\mathbf{p}(t+k, \boldsymbol{\xi}, \boldsymbol{\alpha}) := \begin{cases} \mathbf{p}^*(t+k, \tilde{\boldsymbol{\xi}}^{(k)}, \tilde{\boldsymbol{\alpha}}^{(k)}), & k \in \{1, 2, \dots\}, \\ \mathbf{p}^*(t+k, \boldsymbol{\xi}, \boldsymbol{\alpha}), & k \in \{0, -1, -2, \dots\}. \end{cases}
$$

Consider two systems which, at time t, function and possess component ages $\mathbf{a}$ and $\hat{\mathbf{a}}$, respectively. The policy $\mathbf{p}$ applied to the system with ages $\mathbf{a}$, acts as the policy $\mathbf{p}^*$ applied to the system with ages $\hat{\mathbf{a}}$.[7] Each component n is treated according to its actual age and state, unless it was too young at the start (i.e., if $a_n < \hat{a}_n$) and has not been replaced since then (i.e., if $\alpha_n = a_n + k$). It follows from the construction of $\mathbf{p}$ that the correct artificial ages and failure probabilities are given otherwise.

Hence, the equality $\mathcal{Q}_t^{\mathbf{a}}(0, \mathbf{p}) = \mathcal{Q}_t^{\hat{\mathbf{a}}}(0, \mathbf{p}^*)$ holds and the proposition follows.

(ii) To construct the stochastic policy $\mathbf{p}$ needed, let t, $\mathbf{a}$, and $\hat{\mathbf{a}}$ be fixed and define the mutually independent stochastic variables[8]

$$
Y_n^{(k)} \in \begin{cases} \text{Bin}\left(1, 1 - \frac{1 - G_n(\hat{a}_n + k)}{1 - G_n(a_n + k)}\right), & \text{if } G(\hat{\mathbf{a}} + k\mathbf{1}) > 0, \quad k \in \{1, 2, \dots\}, \\ \text{Bin}(1, 0), & \text{if } G(\hat{\mathbf{a}} + k\mathbf{1}) = 0, \quad n \in \bar{\mathcal{N}}(\hat{\mathbf{a}}, \mathbf{a}), \end{cases} \tag{7}
$$

and define

$$
X^{(k)} := \max_{n \in \bar{\mathcal{N}}(\hat{\mathbf{a}}, \mathbf{a})} \left\{ Y_n^{(k)} \right\}, \quad k \in \{1, 2, \dots\}. \tag{8}
$$

Let

$$
p_n(t+k, \boldsymbol{\xi}, \boldsymbol{\alpha}) := \max\left\{ p_n^*(t+k, \tilde{\boldsymbol{\xi}}^{(k)}, \tilde{\boldsymbol{\alpha}}^{(k)}), u_n^{(k)} \right\}, \quad k \in \{1, 2, \dots\}, \quad n \in \mathcal{N},
$$

where p_n and p_n^* denotes the nth component of the policy $\mathbf{p}$ and $\mathbf{p}^*$, respectively,

$$
\tilde{\boldsymbol{\alpha}}^{(k)} := \begin{cases} \hat{\mathbf{a}} + k\mathbf{1}, & \text{if } \boldsymbol{\alpha} = \mathbf{a} + k\mathbf{1}, \\ \boldsymbol{\alpha}, & \text{otherwise}, \end{cases} \quad k \in \{1, 2, \dots\}, \tag{9}
$$

$$
\tilde{\xi}_n^{(k)} := \begin{cases} \max\left\{ Y_n^{(k)}, \xi_n \right\}, & \text{if } \boldsymbol{\alpha} = \mathbf{a} + k\mathbf{1} \text{ and } n \in \bar{\mathcal{N}}(\hat{\mathbf{a}}, \mathbf{a}), \quad k \in \{1, 2, \dots\}, \\ \xi_n, & \text{otherwise}, \quad n \in \mathcal{N}, \end{cases}
$$

---

[7] No restriction is imposed by assuming that the system functions at time t, as components may possess age 0, which causes identical future costs to having older—and possibly broken—components replaced.

[8] That all failure risk functions considered are non-decreasing guarantees that the involved probabilities are well-defined.

and

$$
u_n^{(k)} := \begin{cases}
\max\left\{ X^{(k)},\ \max_{m\in\mathcal{N}}\{\xi_m\} \right\}, & \text{if } \boldsymbol{\alpha} = \mathbf{a}+k\mathbf{1} \text{ and } n\in\bar{\mathcal{N}}(\hat{\mathbf{a}},\mathbf{a}), \\
\max_{m\in\mathcal{N}}\{\xi_m\}, & \text{if } \boldsymbol{\alpha} = \hat{\mathbf{a}}+k\mathbf{1} \text{ and } n\in\bar{\mathcal{N}}(\hat{\mathbf{a}},\mathbf{a}), \\
0, & \text{otherwise,}
\end{cases}
\qquad
\begin{aligned}
& k\in\{1,2,\dots\}, \\
& n\in\mathcal{N}.
\end{aligned}
$$

For $k \le 0$, without loss of generality we let $\mathbf{p}(t+k, \boldsymbol{\xi}, \boldsymbol{\alpha}) = \mathbf{p}^*(t+k, \boldsymbol{\xi}, \boldsymbol{\alpha})$.[9]

Let $t + K$ be the first replacement time after time t, where $K \ge 1$ is a stochastic variable, and let "case $\mathbf{a}$" and "case $\hat{\mathbf{a}}$" denote the cases with component ages at time t being $\mathbf{a}$ and $\hat{\mathbf{a}}$, respectively. Due to the definition (8), the distributions of K are equivalent in "case $\mathbf{a}$" and in "case $\hat{\mathbf{a}}$". In particular, it holds that

$$
\Pr(K = k : \text{case } \hat{\mathbf{a}}) = G(\hat{\mathbf{a}} + k\mathbf{1})\prod_{j=1}^{k-1}\left[1 - G(\hat{\mathbf{a}} + j\mathbf{1})\right] = \Pr(K = k : \text{case } \mathbf{a}),
$$

where the second equality can be established by induction over k. The independence between failure times of different components means that given $K = k$, at time $t + k$ the distributions of component states for components $n \notin \bar{\mathcal{N}}(\hat{\mathbf{a}}, \mathbf{a})$ are equivalent for the two cases. Further, given $K = k$ and $\{\xi_n\}_{n\in\mathcal{N}\setminus\bar{\mathcal{N}}(\hat{\mathbf{a}},\mathbf{a})}$, due to the definitions (7) and (9), the probabilities of the possible maintenance decisions—and hence also the corresponding expected costs—at time $t + k$ are equivalent in the two cases. Combined with the fact that all components in $\bar{\mathcal{N}}(\hat{\mathbf{a}}, \mathbf{a})$ are replaced at the first replacement time after t, the distribution of component states after the first replacement after time t will be equivalent in the two cases. It follows that $\mathcal{Q}_t^{\hat{\mathbf{a}}}(\mathbf{0}, \mathbf{p}) = \mathcal{Q}_t^{\mathbf{a}}(\mathbf{0}, \mathbf{p})$. We now have that $\mathcal{Q}_t^{\hat{\mathbf{a}}}(\mathbf{0}, \mathbf{p}^*) \le \mathcal{Q}_t^{\hat{\mathbf{a}}}(\mathbf{0}, \mathbf{p}) = \mathcal{Q}_t^{\mathbf{a}}(\mathbf{0}, \mathbf{p})$.

The time for the first replacement after time t is $t + K$, whence, in case $\mathbf{a}$, policy $\mathbf{p}$ may perform replacements not performed by policy $\mathbf{p}^*$, causing extra costs. After these extra replacements, the policy $\mathbf{p}$ leaves a system in which no component individual is older than its equivalent in the system left by the policy $\mathbf{p}^*$. Hence, according to Proposition 1(i), $\mathcal{Q}_t^{\mathbf{a}}(\mathbf{0}, \mathbf{p}) \le \mathcal{Q}_t^{\mathbf{a}}(\mathbf{0}, \mathbf{p}^*) + R_t(\mathbf{a}, \bar{\mathcal{N}}(\hat{\mathbf{a}}, \mathbf{a}))$ holds, where $R_t(\mathbf{a}, \bar{\mathcal{N}}(\hat{\mathbf{a}}, \mathbf{a}))$ denotes the extra cost incurred by replacing the components $n \in \bar{\mathcal{N}}(\hat{\mathbf{a}}, \mathbf{a})$ at time $t + K$ given case $\mathbf{a}$. To calculate the extra costs in case $\mathbf{a}$, we partition the event "$K = k$" in case $\mathbf{a}$ into the two events "$\max_{n\in\mathcal{N}}\{\xi_n\} = 1$" (i.e., the system fails at time $t + k$) and "$\max_{n\in\mathcal{N}}\{\xi_n\} = 0$ and $X^{(k)} = 1$" (i.e. the system does not fail at time $t + k$):

$$
\Pr\left(K = k,\ \max_{n\in\mathcal{N}}\{\xi_n\} = 1\right) = G(\mathbf{a} + k\mathbf{1})\cdot\prod_{j=1}^{k-1}\left[1 - G(\hat{\mathbf{a}} + j\mathbf{1})\right];
$$

$$
\Pr\left(K = k,\ \max_{n\in\mathcal{N}}\{\xi_n\} = 0,\ X^{(k)} = 1\right) = \left[G(\hat{\mathbf{a}} + k\mathbf{1}) - G(\mathbf{a} + k\mathbf{1})\right]\cdot\prod_{j=1}^{k-1}\left[1 - G(\hat{\mathbf{a}} + j\mathbf{1})\right].
$$

If the system fails at time $t + k$, no extra start-up cost must be paid, but for each component $n \in \bar{\mathcal{N}}(\hat{\mathbf{a}}, \mathbf{a})$ an extra cost c_n must be paid, unless it is to be replaced according to policy $\mathbf{p}^*$.[10] The probability that component n has failed at time $t + k$, given a system failure at time $t + k$, is $G_n(a_n + k)/G(\mathbf{a} + k\mathbf{1})$ (the event of system failure is implied by the failure of

---

[9] It does not matter how $\mathbf{p}(t+k, \boldsymbol{\xi}, \boldsymbol{\alpha})$ is defined when $k \le 0$.

[10] In policy $\mathbf{p}^*$, component n must be replaced if it has failed; it may be replaced if it has not failed.

component n). The extra cost incurred by the event "$K = k$" due to system failure at time $t + k$, multiplied by its probability, is at least

$$\left(\prod_{j=1}^{k-1}\left[1 - G(\hat{\mathbf{a}} + j\mathbf{1})\right]\right) \cdot \left(\sum_{n \in \bar{\mathcal{N}}(\hat{\mathbf{a}}, \mathbf{a})} c_n\left[G(\mathbf{a} + k\mathbf{1}) - G_n(a_n + k)\right]\right). \tag{10}$$

If $K = k$ even though there is no system failure at time $t + k$, then an extra cost of d must be paid, as well as the cost c_n for each $n \in \bar{\mathcal{N}}(\hat{\mathbf{a}}, \mathbf{a})$. The extra cost incurred for $K = k$ and a system failure at time $t + k$, multiplied by the probability of this event, is

$$\left(\prod_{j=1}^{k-1}\left[1 - G(\hat{\mathbf{a}} + j\mathbf{1})\right]\right) \cdot \left[G(\hat{\mathbf{a}} + k\mathbf{1}) - G(\mathbf{a} + k\mathbf{1})\right] \cdot \left(d + \sum_{n \in \bar{\mathcal{N}}(\hat{\mathbf{a}}, \mathbf{a})} c_n\right). \tag{11}$$

Hence, the expected value of the extra cost (from the component ages being $\hat{\mathbf{a}}$ instead of $\mathbf{a}$) incurred by system failures over the times $t + 1, \ldots, T$, is not less than the sum, over $k = 1, \ldots, T - t$, of (10) and (11), which equals $K_t(\hat{\mathbf{a}}, \mathbf{a})$, defined in (5). The proposition follows since the inequality $K_t(\hat{\mathbf{a}}, \mathbf{a}) \geq R_t(\mathbf{a}, \bar{\mathcal{N}}(\hat{\mathbf{a}}, \mathbf{a}))$ holds.   □

The result in Patriksson et al. (2015b, Prop. 2) relates the values of the recourse function for two systems without restrictions on component ages; its proof contains, however, an error (a missing superscript). Based on our improvement of the bounds in Proposition 1, we next present a (corrected and) stronger version of Patriksson et al. (2015b, Prop. 2).

We define the component-wise maximum as $\max\{\mathbf{x}, \mathbf{y}\} := [\max\{x_n, y_n\}]_{n \in \mathcal{N}}$.

Proposition 2 *Assume that each component $n \in \mathcal{N}$ possesses a failure distribution with non-decreasing failure risk function G_n, let $\mathbf{p}^*$ be an optimal maintenance policy, and let K_t be defined as in (5). Then, for all $\mathbf{x}, \mathbf{y} \in \mathbb{B}^N$ and all $t \in \mathcal{T}$ it holds that*

$$\mathcal{Q}_t^{\mathbf{a}}(\mathbf{x}, \mathbf{p}^*) \geq \mathcal{Q}_t^{\mathbf{a}}(\mathbf{y}, \mathbf{p}^*) - K_t(\mathbf{a} - \mathbf{A}\mathbf{y}, \mathbf{a} - \mathbf{A}\max\{\mathbf{x}, \mathbf{y}\}).$$

Proof Since the equivalence $G_n(a_n - a_n \max\{x_n, y_n\} + k) = G_n(k)$ holds for any $k \geq 0$ and all $n \in \bar{\mathcal{N}}(\mathbf{a} - \mathbf{A}\mathbf{y}, \mathbf{a} - \mathbf{A}\max\{\mathbf{x}, \mathbf{y}\})$, it follows that

$$\begin{aligned}
\mathcal{Q}_t^{\mathbf{a}}(\mathbf{y}, \mathbf{p}^*) &= \mathcal{Q}_t^{\mathbf{a}-\mathbf{A}\mathbf{y}}(\mathbf{0}, \mathbf{p}^*) \\
&\leq \mathcal{Q}_t^{\mathbf{a}-\mathbf{A}\max\{\mathbf{x},\mathbf{y}\}}(\mathbf{0}, \mathbf{p}^*) + K_t(\mathbf{a} - \mathbf{A}\mathbf{y}, \mathbf{a} - \mathbf{A}\max\{\mathbf{x}, \mathbf{y}\}) \\
&\leq \mathcal{Q}_t^{\mathbf{a}-\mathbf{A}\mathbf{x}}(\mathbf{0}, \mathbf{p}^*) + K_t(\mathbf{a} - \mathbf{A}\mathbf{y}, \mathbf{a} - \mathbf{A}\max\{\mathbf{x}, \mathbf{y}\}) \\
&= \mathcal{Q}_t^{\mathbf{a}}(\mathbf{x}, \mathbf{p}^*) + K_t(\mathbf{a} - \mathbf{A}\mathbf{y}, \mathbf{a} - \mathbf{A}\max\{\mathbf{x}, \mathbf{y}\}),
\end{aligned}$$

where the first and second inequalities are obtained by application of Proposition 1(ii) and (i), respectively. The proposition follows.   □

7 The ORPIL

The model of the opportunistic replacement problem with individual lives (ORPIL; Patriksson et al. 2015a) will be of interest. The ORPIL, which generalizes the opportunistic replacement problem (ORP; Almgren et al. 2012), will be used in Sect. 8.1 to form a two-stage sample average approximation of the SORP. The basic structure of the ORP, the ORPIL, as well as the SORP, consists of a set $\mathcal{N}$ of serially connected components $n \in \mathcal{N}$, each with a replacement

cost c_n, and a start-up cost d to be paid each time—during the (discrete) planning period $\mathcal{T} = \{0, \ldots, T\}$—at least one component is replaced. The differences lie in the lives of the components, which in the case of the

ORP   are constant, and for each component are identical for all individuals;
ORPIL   are constant, but may differ between individuals;
SORP   are stochastic.

As shown in Patriksson et al. (2015a, Prop. 1), the ORPIL is NP-hard. To control its complexity, for each component $n \in \mathcal{N}$ only the first $q_n \in \{1, \ldots, T+1\}$, so-called *non-identical life individuals*, may be assigned individual lives;[11] these are indexed by the set $\mathcal{I}_n := \{1, \ldots, q_n\}$. The remaining individuals (whenever $q_n \leq T$) are assigned *identical lives* and are indexed by the set $\tilde{\mathcal{I}}_n := \{q_n + 1, \ldots, T+1\}$.[12]

Definition 4 (*Opportunistic replacement problem with individual lives; ORPIL*) Given is a discrete planning period $\mathcal{T} = \{0, \ldots, T\}$ and the start-up cost d. For component $n \in \mathcal{N}$, the replacement cost is c_n, $T_{ni} \in \mathbb{Z}_{>0}$ denotes the life of individual $i \in \mathcal{I}_n$, and $T_n \in \mathbb{Z}_{>0}$ denotes the life of the individuals $i \in \tilde{\mathcal{I}}_n$. Find a feasible maintenance schedule[13] which minimizes the sum of all replacement and start-up costs during the planning period.

For full details of the ORPIL, see Patriksson et al. (2015a).

8 Solving the discretized current problem

Two approaches to solve a two-stage approximation of the discretized current problem were presented in Patriksson et al. (2015b): (i) a two-stage deterministic equivalent, and (ii) a two-stage decomposition method, the latter of which was shown to be the computationally most efficient. Hence, we use approach (ii), as presented in Sect. 8.2. Both approaches require, in each iteration, the calculation of an approximate value of the recourse function, as presented in Sect. 8.1.

8.1 The two-stage sample average approximation

Let Ω denote the set of all possible scenarios ω for the discretized current problem (2). A scenario $\omega \in \Omega$ is defined by a life $T_{ni}(\omega) \in \mathbb{Z}_{>0}$ for each individual $i \in \mathcal{I}_n \cup \tilde{\mathcal{I}}_n$ of each component $n \in \mathcal{N}$, and is realized with probability $\Pr(\omega)$.

Each scenario, ω, defines an instance of ORPIL (see Sect. 7) with $q_n = |\mathcal{I}_n| = T + 1$, $n \in \mathcal{N}$; its solution is a maintenance schedule, which is optimal provided that scenario ω is realized (see Def. 4). Such large values of q_n make ORPIL computationally demanding, which is problematic when a large number of scenarios are to be realized. Since it is unlikely that $T + 1$ non-identical life individuals are needed for all components, it may suffice to let

---

[11] In Patriksson et al. (2015a), all components were given the same number of non-identical life individuals, i.e., $q_n = q_m$ for all $n, m \in \mathcal{N}$. Such a restriction is not necessary, and may be computationally unfavorable when the ORPIL is used in a solution scheme for the SORP.

[12] Since the life of each individual is at least δ, corresponding to one time step, at most $T + 1$ individuals of each component are needed.

[13] A *maintenance schedule* is defined as a set of decisions regarding which components to replace at what time steps. The schedule is feasible if and only if each individual $i \in \mathcal{I}_n$ ($i \in \tilde{\mathcal{I}}_n$) of each component $n \in \mathcal{N}$ spends no more than T_{ni} (T_n) time units in the system.

q_n equal the number of individuals of component $n \in \mathcal{N}$ that are *likely* to be used prior to the planning horizon; we anticipate a computational advantage from this construction.

Let $\mathbf{q} = (q_1, \ldots, q_N)$, and let the sampled multiset $\hat{\Omega}_{\mathbf{q}}$ approximate the set Ω, with

$$\hat{\Omega}_{\mathbf{q}} \subseteq \big\{ \omega \in \Omega \,\big|\, T_{ni}(\omega) = \mathrm{E}[T_n], \ i \in \tilde{\mathcal{I}}_n, \ n \in \mathcal{N} \big\},$$

where $\mathrm{E}[T_n]$ denotes a discretization of the expected life (calculated based on continuous life) of an individual of component n. A scenario $\omega \in \hat{\Omega}_{\mathbf{q}}$ is sampled by sampling $T_{n1}(\omega)$ from the life distribution of component n, given $T_{n1}(\omega) \geq a_n + 1$; sampling $T_{ni}(\omega)$ for $i \in \mathcal{I}_n \setminus \{1\}$ from the life distribution of component n, given $T_{ni}(\omega) \geq 1$; and letting $T_{ni}(\omega) := \mathrm{E}[T_n]$ for $i \in \tilde{\mathcal{I}}_n$. We let $\Pr(\omega) = |\hat{\Omega}_{\mathbf{q}}|^{-1}$ for all $\omega \in \hat{\Omega}_{\mathbf{q}}$.

For a first-stage decision $\bar{x}_n \equiv \bar{v}_{n1}^0$, $n \in \mathcal{N}$, define the constraints

$$v_{n1}^0 = \bar{x}_n, \quad n \in \mathcal{N}. \tag{12}$$

Further, define the function $F_{\mathrm{ORPIL}} : \mathbb{B}^N \times \hat{\Omega}_{\mathbf{q}} \to \mathbb{R}$ such that $F_{\mathrm{ORPIL}}(\mathbf{x}, \omega)$ equals the optimal objective value of the ORPIL for scenario ω, with the constraints (12) added to ensure that the first-stage decision is $\mathbf{x}$. A two-stage sample average approximation of $\mathcal{Q}_0^{\mathbf{a}}(\bar{\mathbf{x}}, \mathbf{p}^*)$ is given by

$$\hat{\mathcal{Q}}_s^{\mathbf{a}}(\bar{\mathbf{x}}) := \big| \hat{\Omega}_{\mathbf{q}} \big|^{-1} \sum_{\omega \in \hat{\Omega}_{\mathbf{q}}} F_{\mathrm{ORPIL}}(\bar{\mathbf{x}}, \omega) - \Big(\mathbf{c}^\top \bar{\mathbf{x}} + d \max_{n \in \mathcal{N}} \{\bar{x}_n\} \Big). \tag{13}$$

Given a decision $\bar{\mathbf{x}}$, the expected sum of current costs during the planning period[14] $[s, S)$ is approximated by computing the values $F_{\mathrm{ORPIL}}(\bar{\mathbf{x}}, \omega)$, $\omega \in \hat{\Omega}_{\mathbf{q}}$. For the optimization in (13) to be tractable, the number $|\hat{\Omega}_{\mathbf{q}}|$ of scenarios must be sufficiently small, but too few scenarios would make the two-stage sample average approximation unstable. Effects of varying $|\hat{\Omega}_{\mathbf{q}}|$ was studied by Patriksson et al. (2015b, Section 7.2), who concluded that $|\hat{\Omega}_{\mathbf{q}}| = 100$ is sufficiently large for the instances studied.

8.2 The decomposition method

We employ *integer L-shaped decomposition*; see Laporte and Louveaux (1993), to which our description below refers. The so-called *deterministic equivalent optimization problem* of our *two-stage stochastic optimization problem* is to

$$\mathrm{minimize}_{\mathbf{x} \in \mathcal{X}_{\bar{\xi}}, \, z^0 \in \{0,1\}} \ \Big\{ \mathbf{c}^\top \mathbf{x} + dz^0 + Q(\mathbf{x}, z^0) \,\Big|\, x_n \leq z^0, \, n \in \mathcal{N} \Big\}, \tag{14}$$

where the recourse function $Q : \mathbb{B}^{N+1} \to \mathbb{R}$ is defined by [cf. (13)]

$$Q(\mathbf{x}, z^0) := \mathrm{E}_\omega \big[F_{\mathrm{ORPIL}}(\mathbf{x}, \omega) \big] - \Big(\mathbf{c}^\top \mathbf{x} + dz^0 \Big), \quad (\mathbf{x}, z^0) \in \{0, 1\}^{N+1}.$$

The recourse function is iteratively approximated using a set of *optimality cuts*; for our problem setting, the L-shaped decomposition consists of solving a sequence, for $\ell = 0, 1, \ldots$, of current problems to

---

[14] The time s in the left-hand side of (13) is part of the state of the current problem (see Def. 3); it appears in ORPIL in the sense that $S - s = (T + 1)\delta$ and is to be interpreted as "$t = 0$". In our implementation, as s increases, T decreases while δ is kept approximately constant. Altering δ would affect the number of time steps in ORPIL that are spanned by a (individual or identical) life.

$$\text{minimize}_{\mathbf{x}, z^0, \theta} \quad \mathbf{c}^\top \mathbf{x} + dz^0 + \theta, \tag{15a}$$

$$\text{subject to} \quad \mathbf{x} \geq \boldsymbol{\xi}, \tag{15b}$$

$$x_n \leq z^0, \qquad n \in \mathcal{N}, \tag{15c}$$

$$K_0(\mathbf{a} - \mathbf{A}\mathbf{y}^l, \mathbf{a} - \mathbf{A}\max\{\mathbf{x}, \mathbf{y}^l\}) + \theta \geq \mathcal{Q}_0^{\mathbf{a}}(\mathbf{y}^l, \mathbf{p}^*), \qquad l = 1, \ldots, \ell, \tag{15d}$$

$$\theta \geq 0, \tag{15e}$$

$$(\mathbf{x}, z^0) \in \{0, 1\}^{N+1}. \tag{15f}$$

For $\ell = 0$ all the optimality cuts (15d) are relaxed. Let $(\bar{\mathbf{x}}, \bar{z}^0, \bar{\theta})$ denote the solution to (15); if it holds that $\bar{\theta} \geq Q(\bar{\mathbf{x}}, \bar{z}^0)$, then $(\bar{\mathbf{x}}, \bar{z}^0, \bar{\theta})$ is optimal in (14); otherwise let $\ell := \ell + 1$ and add an optimality cut to next current problem.

Feasible maintenance at time s corresponds to the constraints (15b)–(15c); if feasible replacements are made at time s, the set of feasible future replacements is non-empty (e.g., replace all non-functional components at times $s + 1, s + 2, \ldots$). The optimality cuts provide lower bounds (at least one of which is tight) on the recourse function $Q(\cdot)$, and the bounds from Prop. 2 are defined by highly non-linear functions. Further, for computational efficiency the exact evaluation of $\mathcal{Q}_t^{\mathbf{a}}(\mathbf{x}, \mathbf{p}^*)$ is replaced by the approximation (13).

We assume that $\xi_n = 1$ for at least one $n \in \mathcal{N}$ (i.e., at least one component is not functioning; recall Sect. 3), since otherwise no maintenance would be performed; hence, the start-up cost d will always be included in the current problem (Definition 3). Let $\mathcal{L}$ denote an index set corresponding to the optimality cuts generated so far, let $\mathbf{y}^\ell \in \mathcal{X}_{\boldsymbol{\xi}}$ denote a maintenance decision, $\ell \in \mathcal{L}$ (i.e., the optimal value for $\mathbf{x}$ in (15) when $\ell - 1$ optimality cuts (15d) have been generated), let the function K_0 be defined as in (5). The current problem is then iteratively approximated as

$$f_{0,\mathcal{L}}^{\mathbf{p}^*}(\boldsymbol{\xi}, \mathbf{a}) := \min_{\mathbf{x} \in \mathcal{X}_{\boldsymbol{\xi}}} \left\{ \mathbf{c}^\top \mathbf{x} + d \right.$$
$$\left. + \max\left\{ 0, \max_{\ell \in \mathcal{L}} \left\{ \mathcal{Q}_0^{\mathbf{a}}(\mathbf{y}^\ell, \mathbf{p}^*) - K_0(\mathbf{a} - \mathbf{A}\mathbf{y}^\ell, \mathbf{a} - \mathbf{A}\max\{\mathbf{x}, \mathbf{y}^\ell\}) \right\} \right\} \right\}. \tag{16}$$

Proposition 3 *Define the recourse function $\mathcal{Q}_t^{\mathbf{a}}$ as in (1), assume that the failure risk functions $G_n, n \in \mathcal{N}$, are non-decreasing, and let the function $f_{0,\mathcal{L}}^{\mathbf{p}^*}$ be defined as in (16). Then, it holds that*

$$f_{0,\mathcal{L}}^{\mathbf{p}^*}(\boldsymbol{\xi}, \mathbf{a}) \leq \mathbf{c}^\top \mathbf{x} + d + \mathcal{Q}_0^{\mathbf{a}}(\mathbf{x}, \mathbf{p}^*), \quad \mathbf{x} \in \mathcal{X}_{\boldsymbol{\xi}}.$$

Proof Let $\ell_{\mathbf{x}} \in \operatorname{argmax}_{\ell \in \mathcal{L}} \left\{ \mathcal{Q}_0^{\mathbf{a}}(\mathbf{y}^\ell, \mathbf{p}^*) - K_0(\mathbf{a} - \mathbf{A}\mathbf{y}^\ell, \mathbf{a} - \mathbf{A}\max\{\mathbf{x}, \mathbf{y}^\ell\}) \right\}$ for $\mathbf{x} \in \mathcal{X}_{\boldsymbol{\xi}}$. By Prop. 2 then follows that

$$\mathcal{Q}_0^{\mathbf{a}}(\mathbf{x}, \mathbf{p}^*) \geq \mathcal{Q}_0^{\mathbf{a}}(\mathbf{y}^{\ell_{\mathbf{x}}}, \mathbf{p}^*) - K_0(\mathbf{a} - \mathbf{A}\mathbf{y}^{\ell_{\mathbf{x}}}, \mathbf{a} - \mathbf{A}\max\{\mathbf{x}, \mathbf{y}^{\ell_{\mathbf{x}}}\}), \quad \mathbf{x} \in \mathcal{X}_{\boldsymbol{\xi}}.$$

Since the recourse function is non-negative, i.e., $\mathcal{Q}_0^{\mathbf{a}}(\mathbf{x}, \mathbf{p}^*) \geq 0$, it then follows that

$$\mathcal{Q}_0^{\mathbf{a}}(\mathbf{x}, \mathbf{p}^*) \geq \max\left\{ 0, \mathcal{Q}_0^{\mathbf{a}}(\mathbf{y}^{\ell_{\mathbf{x}}}, \mathbf{p}^*) - K_0(\mathbf{a} - \mathbf{A}\mathbf{y}^{\ell_{\mathbf{x}}}, \mathbf{a} - \mathbf{A}\max\{\mathbf{x}, \mathbf{y}^{\ell_{\mathbf{x}}}\}) \right\}, \quad \mathbf{x} \in \mathcal{X}_{\boldsymbol{\xi}}.$$

Let $\bar{\mathbf{x}} \in \mathcal{X}_{\boldsymbol{\xi}}$ be optimal in (16). The inequality

$$\mathbf{c}^\top \mathbf{x} + d + \max \left\{ 0, \ \mathcal{Q}_0^{\mathbf{a}}(\mathbf{y}^{\ell_{\mathbf{x}}}, \mathbf{p}^*) - K_0(\mathbf{a} - \mathbf{A}\mathbf{y}^{\ell_{\mathbf{x}}}, \mathbf{a} - \mathbf{A} \max \{ \mathbf{x}, \mathbf{y}^{\ell_{\mathbf{x}}} \}) \right\}$$

$$\geq \ \mathbf{c}^\top \bar{\mathbf{x}} + d + \max \left\{ 0, \ \mathcal{Q}_0^{\mathbf{a}}(\mathbf{y}^{\ell_{\bar{\mathbf{x}}}}, \mathbf{p}^*) - K_0(\mathbf{a} - \mathbf{A}\mathbf{y}^{\ell_{\bar{\mathbf{x}}}}, \mathbf{a} - \mathbf{A} \max \{ \bar{\mathbf{x}}, \mathbf{y}^{\ell_{\bar{\mathbf{x}}}} \}) \right\}$$

then holds for all $\mathbf{x} \in \mathcal{X}_{\boldsymbol{\xi}}$ and the proposition follows. $\qquad\qquad\square$

If $f_{0,\mathcal{L}}^{\mathbf{p}^*}(\boldsymbol{\xi}, \mathbf{a}) = \mathbf{c}^\top \mathbf{y}^\ell + d + \mathcal{Q}_0^{\mathbf{a}}(\mathbf{y}^\ell, \mathbf{p}^*)$ holds for some $\ell \in \mathcal{L}$, then $\mathbf{y}^\ell$ is optimal in (2). Otherwise, $f_{0,\mathcal{L}}^{\mathbf{p}^*}(\boldsymbol{\xi}, \mathbf{a}) \leq f_0^{\mathbf{p}^*}(\boldsymbol{\xi}, \mathbf{a})$ holds, i.e., (16) yields a lower bound on the optimal value of (2), and the recourse function for the best lower bound found is evaluated. The method is formalized in Algoritm 1, in which $\hat{\mathcal{Q}}_s^{\mathbf{a}}(\mathbf{y}^\ell)$ approximates $\mathcal{Q}_0^{\mathbf{a}}(\mathbf{y}^\ell, \mathbf{p}^*)$.

Algorithm 1: Integer L-shaped decomposition

input : A system state $(t, \boldsymbol{\xi}, \mathbf{a}) \in \mathcal{V}$
output: A maintenance decision $\mathbf{x}^*$
$\ell \leftarrow 1; \mathcal{L} \leftarrow \{\ell\}; \mathbf{x}^* \leftarrow \mathbf{y}^1 \in X_{\boldsymbol{\xi}}; \psi^* \leftarrow \hat{\mathcal{Q}}_s^{\mathbf{a}}(\mathbf{y}^1); \bar{\mathbf{x}} \leftarrow$ optimum in (16); $\ell_{\bar{\mathbf{x}}} \leftarrow 1$;
while $\mathbf{c}^\top \bar{\mathbf{x}} + d + \max \left\{ 0, \ \hat{\mathcal{Q}}_s^{\mathbf{a}}(\mathbf{y}^{\ell_{\bar{\mathbf{x}}}}) - K_0(\mathbf{a} - \mathbf{A}\mathbf{y}^{\ell_{\bar{\mathbf{x}}}}, \mathbf{a} - \mathbf{A} \max \{ \bar{\mathbf{x}}, \mathbf{y}^{\ell_{\bar{\mathbf{x}}}} \}) \right\} < \mathbf{c}^\top \mathbf{x}^* + d + \psi^*$ **do**

> $\ell \leftarrow \ell + 1; \mathcal{L} \leftarrow \mathcal{L} \cup \{\ell\}; \mathbf{y}^\ell \leftarrow \bar{\mathbf{x}}$;
> **if** $\mathbf{c}^\top \mathbf{y}^\ell + \hat{\mathcal{Q}}_s^{\mathbf{a}}(\mathbf{y}^\ell) < \mathbf{c}^\top \mathbf{x}^* + \psi^*$ **then**
> > $\mathbf{x}^* \leftarrow \mathbf{y}^\ell; \psi^* \leftarrow \hat{\mathcal{Q}}_s^{\mathbf{a}}(\mathbf{y}^\ell)$;
> **end**
> $\bar{\mathbf{x}} \leftarrow$ optimum in (16); $\ell_{\bar{\mathbf{x}}} \in \text{argmax}_{l \in \mathcal{L}} \left\{ \hat{\mathcal{Q}}_s^{\mathbf{a}}(\mathbf{y}^l) - K_0(\mathbf{a} - \mathbf{A}\mathbf{y}^l, \mathbf{a} - \mathbf{A} \max \{ \bar{\mathbf{x}}, \mathbf{y}^l \}) \right\}$;

end
return $\mathbf{x}^*$;

---

9 Numerical experiments

We have performed numerical experiments to show the relevance and effects of the theoretical results developed. The test instances and the policies employed are presented in Sects. 9.1 and 9.2, respectively. Section 9.3 reports on the goodness (i.e., tightness) of the lower bound developed in Sect. 5, while Sect. 9.4 presents the effect of the improved bound on the recourse function [Proposition 1(ii)] on the optimization process.

For more elaborate numerical tests of the optimization policy, including computation times, see Patriksson et al. (2015b), which investigates the effect of the number of non-identical life individuals as well as of the number of scenarios used.

9.1 The test instances

For each of the four test instances, T1–T4, each component n has a Weibull-distributed life with scale parameter α_n and shape parameter β_n. The discrete time points for maintenance decisions are given by $\mathcal{T} := \left\{ 0, \ldots, \max \{ \lceil (S - s)/\delta \rceil, 3 \} \right\}$. The data are shown in Table 1.

- Instance T1 (Patriksson et al. 2015b) was constructed to illustrate a poor performance of (general) age-based policies (see Sect. 9.2). (Our use of a different age-based policy than that in Patriksson et al. (2015b) does, however, not verify this observation; see Sect. 9.3.)

Table 1 Data for the test instances

Instance	T1			T2			T3			T4		
$\{T, N, d, \delta\}$	$\{50, 3, 50, 1\}$			$\{50, 4, 5, 1\}$			$\{100, 7, 5, 2\}$			$\{60, 7, 5, 1\}$		
n	c_n	α_n	β_n	c_n	α_n	β_n	c_n	α_n	β_n	c_n	α_n	β_n
1	1	20	3	2	5	6	1	10	2	1	15	2
2	1	20	3	4	10	6	2	20	3	5	82	3
3	100	20	3	6	15	6	3	30	2	5	81	2
4				8	20	6	4	40	3	3	33	2
5							5	50	2	5	74	6
6							6	60	3	1	7	6
7							7	70	2	3	47	3

- Instance T2 (Patriksson et al. 2015b) was constructed to illustrate the performance of algorithms when the spread of the components' lives is small (i.e., high values of the shape parameter). Then the SORP tends to be more deterministic (like the ORP; see Almgren et al. 2012) and optimization methods are expected to outperform heuristics.
- Instance T3 contains more components and time steps than T1 and T2. The ratio between a component's replacement cost c_n and its expected life $\alpha_n \Gamma(1 + 1/\beta_n)$ is similar for all n. The start-up cost d is within the range of the component replacement costs, $c_n, n \in \mathcal{N}$, so that each component will cause a significant portion of the total replacement costs.[15] T3 is intended to force the optimization policy to create many optimality cuts; hence, our improved lower bound on the recourse function yields a more efficient implementation of the optimization policy.[16]
- In instance T4 the costs and expected lives of components are related as in T3, albeit the "regularity" is avoided according to the following: After deciding upon T, N, and d, for each $n \in \mathcal{N}$, c_n and β_n are sampled from a uniform distribution on $\{1, \ldots, 5\}$ and $\{2, \ldots, 7\}$, respectively. Then α_n is sampled from a uniform distribution on $\{5c_n, \ldots, 20c_n\}$. The values of q_n are chosen such that the total life of the non-identical life individuals for each component exceeds 40.

For each test instance, 100 different scenarios were generated and used for performance testing. These scenarios are called *simulation scenarios* to be distinguished from the scenarios used within the optimization policy (see Sect. 9.2).

9.2 The policies

For testing and illustration we have implemented and used the following four policies.

- A *run-to-failure policy*, in which only failed components are replaced. As per previous assumptions, when at least one component has failed, the time is discretized, and other

---

[15] For extreme values of the cost parameters good solutions are found by simple heuristics: (a) If $d \gg \max_{n \in \mathcal{N}} \{c_n\}$: replace either all or none of the components. (b) If $d \ll \min_{n \in \mathcal{N}} \{c_n\}$: replace only failed components. (c) If $c_m \ll \min_{n \in \mathcal{N} \setminus \{m\}} \{c_n\}$ for some $m \in \mathcal{N}$: replace component m if any other component is replaced. (d) If $c_m \gg \max_{n \in \mathcal{N} \setminus \{m\}} \{c_n\}$ for some $m \in \mathcal{N}$: replace component m only when it is failed.

[16] For a discretized current problem with N components in state $\boldsymbol{\xi} \in \mathbb{B}^N$, at most $|\mathcal{X}_{\boldsymbol{\xi}}| = 2^{N - \mathbf{1}^{\mathrm{T}} \boldsymbol{\xi}}$ optimality cuts can be generated, $\mathbf{1}^{\mathrm{T}} \boldsymbol{\xi}$ being the number of failed components. Hence, for small values of N (such as in T1 and T2), differences in numbers of generated optimality cuts may seem insignificant.

Algorithm 2: Simulated annealing

input : A set $\mathcal{N}$ of components with life distributions $F_n : \mathbb{R}_+ \to [0, 1]$ and initial guesses $\hat{\sigma}_n$ for soft life thresholds, $n \in \mathcal{N}$. A planning horizon $S > 0$, a number $0 < \epsilon \ll S$, and set $\mathcal{S}$ of simulation scenarios. An upper limit $r_{\max}$ on # iterations with no improvement.

output: Soft life thresholds σ_n, $n \in \mathcal{N}$

$C^{\text{best}} \leftarrow \infty$; $r \leftarrow 0$;

while $r < r_{\max}$ **do**

    $C \leftarrow$ average maintenance cost over all scenarios in $\mathcal{S}$, given soft life thresholds $\hat{\sigma}_n$;

    if $C < C^{\text{best}}$ **then**

        $\mid$ $C^{\text{best}} \leftarrow C$; $\sigma_n \leftarrow \hat{\sigma}_n$, $n \in \mathcal{N}$; $r \leftarrow 0$;

    end

    $\hat{\sigma}_n \leftarrow \min \big\{ \max \{ \text{N}(\sigma_n, \sigma_n/(r + 1)), \epsilon \}, S + \epsilon \big\}$, $n \in \mathcal{N}$, where $\text{N}(\cdot, \cdot)$ is a random number from the normal distribution; $r \leftarrow r + 1$;

end

return σ_n, $n \in \mathcal{N}$;

---

components which will fail within the next time step are also considered to be failed. Hence, several components may "fail" simultaneously.

- The *optimization policy* (the two-stage decomposition method; see Sect. 8). It occurs in an *old implementation*, based on the lower bound on the recourse function from Patriksson et al. (2015b), and a *new implementation*, based on the (tighter) lower bound on the recourse function given in Proposition 1(ii). Hence, the new implementation is expected to run faster than the old one. The optimization policy requires an arbitrary feasible starting solution. Both implementations are initialized at the solution provided by the run-to-failure policy, which is reasonable when the planning horizon is near. In general, however, the optimization policy may be improved through better starting solutions (obtained by some fast policy).

- The *expected value (EV) optimization policy*, which is used by Almgren et al. (2012) to solve the SORP. It approximates the current problem by—for each component—assigning to the current individual its expected life, conditioned on its age, and by assigning to all future individuals the component's expected life.[17]

- An *age-based policy*: at failure of any component, maintenance is performed and each component whose age surpasses its assigned *soft life* threshold is also replaced. This policy is akin to that of Crocker and Kumar (2000), in which maintenance is enforced whenever the age of any component reaches its assigned *hard life* threshold. Hard life thresholds are not needed in our set-up, in which failures are assumed either to be detectable only very shortly before occurring, or to cause no extra costs. The soft life thresholds are variables in an optimization problem to minimize the expected maintenance cost over the planning period; in our numerical experiments (see Sect. 9.3) it was solved using *simulated annealing* (Algorithm 2), with the components' expected lives used as initial guesses for the soft life thresholds. To prevent from getting stuck in poor local optima, the best solution from 100 separate runs of Algorithm 2 was used.

9.3 The optimization policy versus the lower bound

Table 2 presents, for the instances T1–T4, lower bounds on the optimal maintenance costs and the maintenance costs resulting from the four policies described in Sect. 9.2. It also

---

[17] The EV optimization policy equals the special case of the optimization policy with $q_n = 0$, $n \in \mathcal{N}$, resulting in identical scenarios, so that one scenario is sufficient for the estimation of the recourse function.

Table 2 Lower bounds on the expected value of costs (c_x, x = lower_bound) and average costs (100 simulation scenarios/instance) resulting from the four policies (c_x, x = policy). The gaps are computed as ($c_{policy}/c_{lower\_bound} - 1) \cdot 100\%$. In the optimization policy, we used $q_n = 2$, $n \in \mathcal{N}$, for T1–T3, and **q** = (4, 2, 2, 3, 2, 6, 2) for T4

Instance	T1		T2		T3		T4	
	c_x	Gap (%)	c_x	Gap (%)	c_x	Gap (%)	c_x	Gap (%)
Lower bound	422		128		130		74	
Optimization policy	466	10.4	145	13.3	171	31.5	76	2.7
EV optimization policy	528	25.1	151	18.0	182	40.0	81	9.5
Run-to-failure policy	566	34.1	169	32.0	183	40.8	83	12.2
Age-based policy	460	9.0	146	14.1	172	32.3	77	4.1

presents the relative differences in maintenance costs (so-called gaps) between each policy and the lower bound. For each instance, there is a significant gap between the costs resulting from the optimization policy and the respective lower bound, but as concluded in Sect. 5, the bound is not expected to be tight, so it is not likely that any policy will produce solutions with costs close to the lower bound.

In our experiments, the age-based policy and the optimization policy resulted in approximately the same average cost for all four instances. A different age-based policy was employed by Patriksson et al. (2015b), and that age-based policy performed significantly worse for some instances, in particular instance T1.

9.4 The number of generated optimality cuts

While applying the old and the new implementation of the optimization policy on 100 simulation scenarios of each test instance, a vast number of discretized current problems was solved and for each of which a number of optimality cuts were generated; the corresponding frequencies are shown in Fig. 1a–d. Although the same simulation scenarios were used for both implementations of the optimization policy, they occasionally result in different maintenance decisions. One reason is that optimality cuts are typically generated in different orders, resulting in different scenarios being used for ORPIL, which leads to different approximations of the recourse functions. Another reason is that—in either implementation—the first optimal maintenance decision found is selected; hence, in case of multiple optima, the order in which the optimality cuts are generated has an impact on the optimal solution selected. Once the selected maintenance decisions within a scenario differ between the implementations, the subsequent current problems will also differ. In Fig. 1 only current problems encountered by both implementations are counted: a higher frequency of the scenarios are solved by fewer optimality cuts with the new optimization policy than with the old one, e.g. in instance T1 the new policy will most frequently require 2 optimality cuts (and slightly above 2 cuts on average), whereas the old policy will most frequently require 4 optimality cuts (and above 2.5 cuts on average). Each generated optimality cut results in one instance of ORPIL to be solved; with ORPIL being the computationally heaviest part of the algorithm, the computation time is approximately proportional to the number of generated optimality cuts.

The histograms in Fig. 1 represent tilted distributions of the numbers of optimality cuts generated: Current problems (occuring during the solution of the simulation scenarios) being
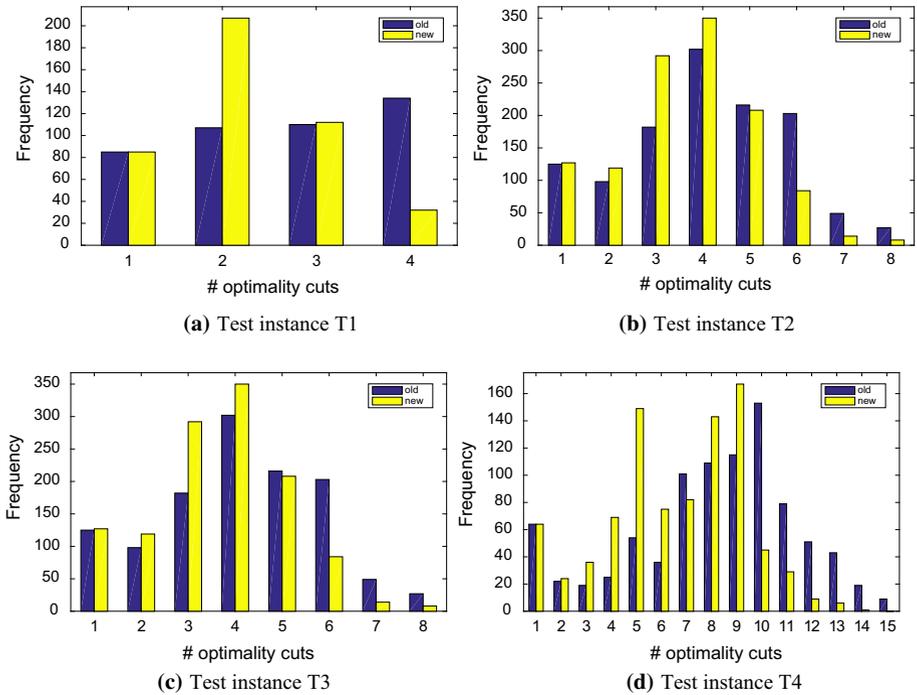
Fig. 1 Frequencies of numbers of optimality cuts generated by the old and new, respectively, implementation of the optimization policy

excluded (due to the implementations proposing different maintenance decisions at earlier stages) are more prevalent when close to than far from the planning horizon. Hence, very few optimality cuts are generated when near to the horizon—often then an obvious decision is to replace only failed components (i.e., the starting solution in Algorithm 1).

10 Conclusions and future research

This article deals with a stochastic opportunistic replacement problem for a system of components with economic dependencies. In particular, we attempt to solve the so-called *current problem*, i.e., decide "which components to replace now", at a time when at least one component must be replaced.

A simple but non-trivial lower bound on the expected maintenance cost is presented; it is based on the cost that would result from a perfect coordination of replacements (using a minimal number of expected replacement occasions) with no loss of component life (due to early replacements). The lower bound cannot be expected to be tight (i.e., equal the optimal value), except in some extreme cases. It provides, however, a restriction on the gain reachable by any optimization policy.

We improve—in terms of computing efficiency—a previously developed decomposition method for a two-stage approximation of the current problem: the improved lower bound on the recourse function value results in fewer optimality cuts being generated. Also the

two-stage approximation is generalized, allowing for different numbers of non-identical life individuals for different components.

Experiments showing that the optimization policy usually outperforms simpler policies was presented by Patriksson et al. (2015b). In contrast, our experiments indicate that a fairly simple age-based policy performs on par with the optimization policy.

References

Aghezzaf, E.-H., Khatab, A., & LeTam, P. (2016). Optimizing production and imperfect preventive maintenance planning's integration in failure-prone manufacturing systems. *Reliability Engineering & System Safety*, *145*, 190–198.

Alaswad, S., & Xiang, Y. (2017). A review on condition-based maintenance optimization models for stochastically deteriorating system. *Reliability Engineering and System Safety*, *157*, 54–63.

Almgren, T., Andréasson, N., Patriksson, M., Strömberg, A.-B., Wojciechowski, A., & Önnheim, M. (2012). The opportunistic replacement problem: Theoretical analyses and numerical tests. *Mathematical Methods of Operations Research*, *76*(3), 289–319.

Atashgar, K., & Abdollahzadeh, H. (2016). Reliability optimization of wind farms considering redundancy and opportunistic maintenance strategy. *Energy Conversion and Management*, *112*, 445–458.

Benders, J. F. (1962). Partitioning procedures for solving mixed-variables programming problems. *Numerische Mathematik*, *4*(1), 238–252.

Birge, J., & Louveaux, F. (1997). *Introduction to stochastic programming*. New York: Springer.

Carlos, S., Sánchez, A., Martorell, S., & Marton, I. (2013). Onshore wind farms maintenance optimization using a stochastic model. *Mathematical and Computer Modelling*, *57*(7–8), 1884–1890.

Cherkaoui, H., Huynh, K. T., & Grall, A. (2016). On the assessment of performance and robustness of condition-based maintenance strategies. *IFAC-PapersOnLine*, *49*(12), 809–814.

Comparea, M., Martini, F., & Zio, E. (2015). Genetic algorithms for condition-based maintenance optimization under uncertainty. *European Journal of Operational Research*, *244*(2), 611–623.

Crocker, J., & Kumar, U. D. (2000). Age-related maintenance versus reliability centred maintenance: A case study on aero-engines. *Reliability Engineering & System Safety*, *67*(2), 113–118.

Deloux, E., Castanier, B., & Bérenguer, C. (2009). Predictive maintenance policy for a gradually deteriorating system subject to stress. *Reliability Engineering and System Safety*, *94*(2), 418–431.

Do, P., Voisin, A., Levrat, E., & Iung, B. (2015). A proactive condition-based maintenance strategy with both perfect and imperfect maintenance actions. *Reliability Engineering and System Safety*, *133*, 22–32.

Gunn, E. A., & Diallo, C. (2015). Optimal opportunistic indirect grouping of preventive replacements in multicomponent systems. *Computers & Industrial Engineering*, *90*, 281–291.

Jha, M. K., Head, M., & Gar, S. P. (2013). Metaheuristic applications in bridge infrastructure maintenance scheduling considering stochastic aspects of deterioration. In A. H. Gandomi, X.-S. Yang, S. Talatahari, & A. H. Alavi (Eds.), *Metaheuristic applications in structures and infrastructures* (pp. 539–556). Amsterdam: Elsevier.

Laggoune, R., Chateauneuf, A., & Alssani, D. (2009). Opportunistic policy for optimal preventive maintenance of a multi-component system in continuous operating units. *Computers & Chemical Engineering*, *33*(9), 1499–1510.

Laporte, G., & Louveaux, F. (1993). The integer *L*-shaped method for stochastic integer programs with complete recourse. *Operations Research Letters*, *13*(3), 133–142.

Papakonstantinou, K. G., & Shinozuka, M. (2014). Optimum inspection and maintenance policies for corroded structures using partially observable Markov decision processes and stochastic, physically based models. *Probabilistic Engineering Mechanics*, *37*, 93–108.

Patriksson, M., Strömberg, A.-B., & Wojciechowski, A. (2015a). The stochastic opportunistic replacement problem, part I: models incorporating individual component lives. *Annals of Operations Research*, *224*, 25–50.

Patriksson, M., Strömberg, A.-B., & Wojciechowski, A. (2015b). The stochastic opportunistic replacement problem, part II: a two-stage solution approach. *Annals of Operations Research*, *224*, 51–75.

Rasmekomen, N., & Parlikad, A. K. (2016). Condition-based maintenance of multi-component systems with degradation state-rate interactions. *Reliability Engineering and System Safety*, *148*, 1–10.

Shi, H., & Zeng, J. (2016). Real-time prediction of remaining useful life and preventive opportunistic maintenance strategy for multi-component systems considering stochastic dependence. *Computers & Industrial Engineering*, *93*, 192–204.

van Noortwijk, J. M., & Frangopol, D. M. (2004). Two probabilistic life-cycle maintenance models for deteriorating civil infrastructures. *Probabilistic Engineering Mechanics*, *19*(4), 345–359.

Wang, W. (2012). A stochastic model for joint spare parts inventory and planned maintenance optimization. *European Journal of Operational Research*, *216*(1), 127–139.