

Sequence- and time-dependent maintenance scheduling in twice re-entrant flow shops

Eigbe, Eghonghon Aye; De Schutter, Bart; Nasri, Mitra; Yorke-Smith, Neil

DOI

[10.1109/ACCESS.2023.3317533](https://doi.org/10.1109/ACCESS.2023.3317533)

Publication date

2023

Document Version

Final published version

Published in

IEEE Access

Citation (APA)

Eigbe, E. A., De Schutter, B., Nasri, M., & Yorke-Smith, N. (2023). Sequence- and time-dependent maintenance scheduling in twice re-entrant flow shops. *IEEE Access*, *11*, 103461-103475. <https://doi.org/10.1109/ACCESS.2023.3317533>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Received 25 August 2023, accepted 15 September 2023, date of publication 20 September 2023, date of current version 27 September 2023.

Digital Object Identifier 10.1109/ACCESS.2023.3317533

RESEARCH ARTICLE

Sequence- and Time-Dependent Maintenance Scheduling in Twice Re-Entrant Flow Shops

EGHONGHON-AYE EIGBE¹, BART DE SCHUTTER², (Fellow, IEEE), MITRA NASRI³, (Member, IEEE), AND NEIL YORKE-SMITH¹

¹Department of Software Technology, Delft University of Technology, 2628 XE Delft, The Netherlands

²Delft Center for Systems and Control, Delft University of Technology, 2628 CN Delft, The Netherlands

³Department of Electrical Engineering, Eindhoven University of Technology, 5612 AZ Eindhoven, The Netherlands

Corresponding author: Eghonghon-Aye Eigbe (e.eigbe@tudelft.nl)

This work was supported by the Mastering Complexity (MasCot) Program, a Nederlandse Organisatie voor Wetenschappelijk Onderzoek (NWO) under the Scheduling Adaptive Modular Flexible Manufacturing Systems (SAM-FMS) project with Grant 17931.

ABSTRACT Industrial and academic interest converge on scheduling flow shops with sequence- and time-dependent maintenance. We posit that anticipatory, integrated scheduling of operational and maintenance tasks leads to superior performance to purely ‘wait-then-fix’ handling of the maintenance tasks. Motivated by an industrial problem with (sequence dependent) setup times, maximum separation constraints, and a combination of sequence- and time- dependent maintenance tasks, this paper introduces an integer programming solution, a constraint programming solution and a heuristic solution based on list scheduling. The motivating use case provides a unique combination of concerns that is to the best of our knowledge, not yet studied in the literature. We build on existing work where we can by extending models for sequence-dependent maintenance scheduling to accommodate sequence- and time-dependent maintenance scheduling and also propose other new models. We show the relative performances of our methods through empirical evaluations and also show significant improvements – up to 25% reduction in makespan – when compared to a reactive scheduling approach that does not consider maintenance in its planning. Based on our evaluations on exact methods, constraint programming models scale better than mixed integer programming models for this problem.

INDEX TERMS Flexible manufacturing systems, maintenance scheduling, makespan minimisation, re-entrant flow shops.

I. INTRODUCTION

We consider a sequence- and time-dependent maintenance scheduling problem. Our problem is motivated by an industrial use case of a large-scale printer (LSP) and is modelled as a flow shop. The operations in this problem have ordering constraints that enforce precedence and also maximum separation constraints that limit the delay between some of the operations. We also face setup time considerations. There are maintenance tasks which depend on the schedule: different sequences of operations have different deterioration effects on the machines. Additionally, the contribution of an operation to the total deterioration effect in a sequence is dependent

The associate editor coordinating the review of this manuscript and approving it for publication was Yu Liu¹.

on the timing of operations. Thus, our maintenance planning problem is both sequence- and time-dependent. The key question is how to handle both operational and maintenance tasks.

This is a challenging question because deteriorated machines produce low-quality jobs; there are thresholds beyond which deterioration of machines must be fixed by carrying out a maintenance activity. The overall objective is to find a feasible schedule that minimises the makespan.

Integrated production and maintenance planning is a challenge in many industries such as in wind farms [1], [2], in the capital goods industry [3] and the pulp and paper industry [4]. In many cases, machine deterioration is dependent on use, i.e., the maintenance required depends on how production operations have been scheduled. Sometimes, this dependence can be ignored and solutions can focus on preventive or

policy based maintenance [5], [6], [7]. Recent work in this direction has used reinforcement learning to come up with these policies [8]. In other cases, the maintenance and production planning problem can be so integrated that the effect of use patterns on maintenance cannot be ignored. Previous work along these lines has considered different ways in which maintenance planning is integrated with production planning such as time-dependent maintenance [9] and position-dependent maintenance [10], [11]. The literature also considers different models of maintenance activities with one of the most popular models being that maintenance activities affect the processing time of operations [12]. While similar problems have been tackled in the literature, our work deals with a unique combination of maximum separation constraints and a deterioration effect not on the processing times of jobs but on the quality of work produced.

We present three solutions to this problem namely, (i) a mixed integer programming solution, (ii) a constraint programming solution, and (iii) a list-scheduling based heuristic solution that extends the capabilities of existing schedulers to handle the kind of maintenance activities presented in this problem.

Through empirical evaluations, we show that in comparison to the reactive approach of scheduling only production operations and then performing maintenance activities when deterioration thresholds are crossed during a production run, our proactive approach achieves significant improvements in the makespan.

Parts of this paper were presented in a non-archival workshop [13]. In the workshop paper [13], we introduced the list-scheduling based heuristic (Section VI). The current paper presents the work in archival form, introduces two other solution methods, and also expands the scope of the evaluation to include more list-schedulers from the literature.

This paper is organised as follows: Section II discusses related work, Section III provides the background and problem definition, Sections IV, V and VI present mixed integer programming, constraint programming and heuristic solutions respectively. We perform empirical evaluations in Section VII and Section VIII concludes the paper.

II. RELATED WORK

The literature has investigated the dynamic relationship between machine deterioration and production scheduling from multiple angles ranging from ways to accurately determine the deterioration of a machine [14], [15] to actually generating schedules. We group the research themes in this field based on two categories, namely; (i) the way deterioration is modelled and (ii) the way maintenance activities are modelled.

Based on the *deterioration model*, existing research can be split into three main categories or approaches [16]. The *time-dependent* approach relates deterioration to the time at which a job is scheduled, i.e., scheduling a job later in the schedule incurs some additional deterioration which typically leads to longer processing times compared to scheduling it earlier.

Closely related to this is a *position-dependent* approach, where deterioration effect of an operation is dependent on the number of preceding completed operations. Finally, there is the *sequence-dependent* approach in which the deterioration depends on the ordering or sequence of the preceding operations on the machine. As a result of the industrial challenge addressed in this paper, we focus on the sequence-dependent case with an additional challenge that the deterioration effect of an operation on a machine is not known apriori and is itself time-dependent.

The survey of Gawiejnowicz [9] into the state of time-dependent scheduling problems has shown that the problem has been studied for single machine, parallel machine and dedicated machine use cases with a wide range of solution methods. However, situations where time-dependence of maintenance activities is coupled with sequence-dependence are unaddressed.

Yang [17] consider the position-dependent maintenance scheduling problem on a set of parallel machines assuming that machines can only be maintained once within the planning horizon and with a constant maintenance duration. References [10], [11] and [18] all consider position-dependent maintenance on a single machine with varying considerations such as the impact of time-dependent improvements in machine conditions, constraining job processing times to lie within an interval, and a combination of time and position-dependent deterioration respectively. References [12] and [19] also consider the position dependent case but both add due-window considerations for just-in-time scheduling considerations.

The *sequence-dependent* approach is a more recent addition to the literature and can be considered as a generalisation of the time- and position-dependent approaches. Notably, [20] and [21] study sequence-dependent deterioration on a set of parallel machines without and with maintenance activities respectively. Reference [22] considers iterated greedy heuristics for a similar problem and [23] considers the case where the parallel machines are not identical and processing time is based on a combination of deterioration and the speed of the assigned machine. Recently, [16] explored multiple integer programming models for solving the sequence-dependent maintenance problem on parallel machines and provided a heuristic approach for larger instances. The combination of sequence-dependent maintenance with other approaches and its effect in more complex manufacturing systems has not yet been studied.

Based on the *model of maintenance activities*, there are also different approaches in the literature. Some works such as [20] do not consider the presence of maintenance activities at all and aim to schedule in a way that deterioration is minimized. Other works such as [16] consider maintenance activities that reset the status of a machine to full health or 0% deterioration while a third category [12] considers rate-modifying maintenance activities that restore machine health by modifying the rate such that machines are able to perform work faster after maintenance. The authors of [24]

and [25], classify maintenance activities into those that completely reset the state of the machines and those that restore the machines to some better deterioration state only. Additionally, the maintenance activities can be of fixed duration or can also have varying types based on how deteriorated a machine is.

A core assumption in many scenarios is that deterioration makes machines slower, thus increasing processing times of operations. Our work differs fundamentally in this regard in that using deteriorated machines does not have an effect on processing times, but instead affects the quality of the jobs produced. Our problem defines deterioration thresholds beyond which maintenance activities must be carried out to meet the quality requirements of future jobs. We also consider the case of maintenance activities that reset the state of the machine but also consider that there exist different classes of maintenance activities each with their own deterioration thresholds and incurring different costs.

An additional complication in our problem is the presence of maximum separation constraints, which impose additional feasibility requirements on the problem. Exact solutions are able to easily model these additional requirements but heuristics run the risk of generating infeasible solutions in some cases. We therefore consider it necessary to design a solution for schedules that become infeasible due to the incorporation of maintenance activities. This concept of re-organising or repairing a changed schedule has been studied with various heuristics such as left and right shift [26], [27]. [28] combines multiple of these heuristics and a genetic schedule repair algorithm to build a solution that caters to multiple classes of schedule disturbances in a prefabrication plant.

In the context of flow shops, an example of schedule repair algorithms can be found in [29] which considers re-scheduling in a two-machine flow shop where schedules are disrupted by machine breakdowns. Additionally, [30] considers re-scheduling due to inserting new jobs in already planned schedules and [31] considers re-scheduling due to a wider range of disruptions in flow shop schedules at runtime. These cases all consider unexpected interruptions and do not have the combination of precedence and maximum separation constraints which provide an additional challenge for our problem.

In summary, there is a gap in the literature for sequence-dependent maintenance scheduling where deterioration effects of operations are not known apriori but are themselves time-dependent. The particular industrial challenge we consider has additional requirements of maximum separation that add to the complexity of the problem. Further, the schedule repair that is needed for heuristic schedulers that may produce infeasible schedules when we introduce maintenance activities, also requires new techniques.

III. PROBLEM DEFINITION

We consider a *maintenance-aware re-entrant flow shop with setup times and relative due dates* inspired by an industrial use case of a large-scale printer (LSP). The LSP prints

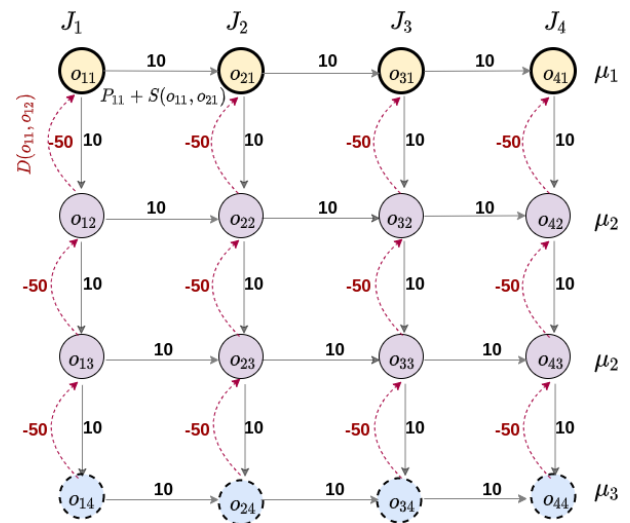


FIGURE 1. Sample re-entrant flow shop where the operations are represented by circles. Column-wise, we have operations of the same job and row-wise, we have operations on the same machine with one of these being the re-entrant machine that appears on rows 2 and 3. Operations with the same colour or boundary lines are mapped to the same machine. Setup times and maximum separation constraints are shown by solid and dashed edges respectively.

different types of duplex sheets that need to be processed twice by the same print head at a speed of 100 or more pages per minute. In this setting, jobs to be scheduled refer to sheets to be printed.

In three-field or Graham notation [32], the base problem without maintenance is defined as $F|s_i, s_{ij}, \text{limited-wait}|C_{\max}$ indicating that it is a flow shop with both sequence-dependent and independent setup times, with maximum separation constraints between operations of the same job also known as limited-wait constraints, and with an objective to minimise makespan C_{\max} . There is no preemption allowed and all jobs are released at time 0.

We represent the n -job m -machine maintenance-aware problem as the tuple $(M, J, O, P, S, D, \delta, X, O^M)$ where $M = \{\mu_1, \dots, \mu_m\}$ is the set of machines and $J = \langle J_1, \dots, J_n \rangle$ is the sequence of jobs. The set O represents the set of operations for every job $j_i \in J$ where each operation o_{ij} has a processing time P_{ij} . Each job has the same number of r operations as is in a standard flow shop. Moreover, $S : O \times O \rightarrow \mathbb{R}_{\geq 0}$ refers to setup times, which represent the required delay between the completion of an operation and the start of another operation. Setup times can exist between operations of the same job to model travelling time of a job for instance, or between operations on the same machine to model any machine preparation step that is needed between operations. Operations of the same job also have maximum separation constraints between them represented as $D : O \times O \rightarrow \mathbb{R}_{> 0}$, i.e., the maximum delay between the start times of two consecutive operations of the same job. Such constraints model the fact that operations of a job can often not be delayed indefinitely due to physical constraints in the plant like the buffer size. In a situation where such constraints do not apply,

separation constraints can simply be set to infinity and setup times to zero.

The solution to the problem is a schedule Ω , i.e., a sequence of both maintenance activities and production operations where each production operation is assigned a start time such that ω_{ij} represents the start time of operation o_{ij} .

In addition to being a flow shop with the above properties, we have a situation with re-entrancy such that the sequence of machines for each job is $\langle \mu_1, \dots, \mu_k, \mu_k, \dots, \mu_m \rangle$, i.e., there is one re-entrant machine that all jobs go through twice. Operations on the re-entrant machines are referred to as *first and second pass operations*. Re-entrancy occurs in many production processes, e.g. semi-conductor production, where wafers revisit machines at different stages of the production process and in painting processes where a job may revisit a machine for multiple coats of paint. Simple re-entrant setups have been shown to be NP-hard [33], [34]. Our motivating use case of production printing has a twice re-entrant setup arising from duplex printing.

We further have the constraints that (i) jobs are not allowed to overtake each other, (ii) the required completion order of jobs is the same as the index of the jobs, and (iii) all setup times and due date constraints are hard constraints that must be obeyed. This situation means that the only scheduling freedom is in the sequence of operations on the re-entrant machine, i.e., first and second passes of the same jobs do not necessarily have to follow each other on this machine. This means we can also think of this as a single machine scheduling problem with precedence and maximum separation constraints.

In the same vein as only needing to schedule the re-entrant machine, we limit our maintenance planning to maintaining the re-entrant machine. While other machines also require maintenance, only the re-entrant machine requires maintenance in the same time scale as the operations being carried out, creating a very tightly coupled problem compared to maintenance of other machines. Additionally, the re-entrant machine is often a key machine of concern for cost reasons – re-entrant machines are too expensive to simply duplicate and remove the re-entrancy – or for quality reasons – some products need to be handled in a delicate state (chemical products for example) and moving the product from one machine to the other would change its state.

A. DETERIORATION MODEL

In our motivating industrial problem, there is a deterioration model $\delta : \Omega \rightarrow \mathbb{R}_{\geq 0}$, that maps a scheduled sequence of operations on a machine to a deterioration state, i.e., given a sequence of operations on a machine with their corresponding start times, i.e., a schedule, $\delta : \Omega \times \rightarrow \mathbb{R}_{\geq 0}$ informs us of the machine state at the end of the sequence. Here, δ is both sequence- and time-dependent in the sense that deterioration is measured by *idle time* of a machine part, i.e., the longer a machine part has been left idle, the more deteriorated it is. These idle times follow directly not only from the sequence themselves, but also from the assigned start times

of operations in these sequences. We do not explicitly model machine parts and instead depend on the fact that different types of jobs use different machine parts and so it can be inferred which machine parts have been idle based on how long it has been since a certain job type has been scheduled. We assume that there is a set of job types $T = \{\tau_1, \dots, \tau_n\}$ that can be presented to the machine and that there is a lexicographic ordering of job types such that every set of machine parts used by a job type τ_x is contained in the set of machine parts used by a job type $\tau_{y>x}$. It then follows that at the start of an operation of type τ_x , idle time is reset to 0 for all operations of type $\tau_{y \leq x}$. Note that while there could be other kinds of problems where different jobs use completely different machine parts and such a lexicographic ordering of types is not possible, it is still a realistic assumption for many scenarios, e.g., scenarios where jobs come in different sizes and bigger sizes simply use more machine parts for production or scenarios where jobs can be customised with different add-on properties processed by additional machine parts.

Finally, we also take as input a maintenance policy X . In our problem, the policy has a set of maintenance activity classes C . For every class $c \in C$, there is a corresponding maintenance duration P_c similar to processing times of production operations. The maintenance policy further maps intervals of deterioration values $[\theta_c, \Theta_c)$ to classes of maintenance activities such that whenever the deterioration falls in $[\theta_c, \Theta_c)$, a maintenance activity of at least class c is required before further production. Thus, $[\theta_c, \Theta_c)$ defines the interval of deterioration thresholds for a maintenance activity. We assume that these intervals are non-overlapping and that maintenance activities triggered by higher thresholds, i.e., harsher deterioration, are more intense and require longer durations. The deterioration thresholds serve to capture the limits at which the quality of a job would be too low if production carries on without a maintenance activity. In this context, a low-quality job refers to a poor print typically with colours bleeding into each other, blurry prints or unintended lines running across a page.

An example problem is shown in Figure 1. The problem is represented as a constraint graph where the due dates and setup times are treated as a system of difference constraints. Operations are represented as circles and each column of operations belong to one job, while each row of operations are mapped to the same machine. Solid arrows represent the minimum separation between operations and are made of the sum of processing and setup times while dashed arrows represent the maximum separation between operations and can be thought of as relative due dates. Minimum separation edges are represented with positive values and maximum separation edges are represented with negative values as they connote *at least* and *at most* constraints on the difference between the start times of the operations they connect.

IV. INTEGER PROGRAMMING APPROACH

Mixed Integer Programming (MIP) is one of the most popular exact solving paradigms and has been applied to other

maintenance planning problems in the literature [16], [35] with some success. Due to the existence of a wide variety of commercial solvers, mixed integer formulations of a problem are valuable as solutions can be provided by these solvers. Furthermore, MIP models can give additional insight to the structure of a problem. Thus, we also consider such a solution for our problem.

In this section we present an exact MIP model for this problem. The model uses the concept of event based formulation as introduced by [16] and extends this concept to accommodate the kind of maintenance policies in this problem. The key idea here is the notion of *blocks* where a block is defined as a sequence of operations uninterrupted by a maintenance activity, i.e., a block is a sequence of operations separated from other operations in the sequence by at least one maintenance activity. For our model, we extend this idea to also include the effect of job types. A block is then defined as a sequence of operations uninterrupted by either a maintenance activity or an operation with a higher type than any other operation within the block.

We define binary variables to mark whether an operation starts a block or not. These variables are further indexed by job type and maintenance activity class, i.e., an operation can be the start of a block delineated by a class of maintenance activities and/or the start of a block delineated by a job type. Blocks of different types are allowed to overlap with additional constraints added to ensure that maintenance is not triggered more than necessary.

The model retains all variables defined in the problem definition in Section III. Indices of variables corresponding to operations are either of the form x_{ij} when both the job and operation identifier are important or of the form x_a when it is only necessary to differentiate one operation from the other. For ease of modelling, we also define binary variables Γ_{am} , $\forall o_a \in O, \mu_m \in M$ to represent machine assignment. Then, Γ_{am} is set to 1 if operation o_a is assigned to machine μ_m . Γ_{am} is not a decision variable and is part of the problem description.

Additionally, since we only plan maintenance on the re-entrant machines μ_k , many constraints only apply to operations on this machine and are denoted as R such that $R = \{o_a \in O | \Gamma_{ak} > 0\}$. Furthermore, a dummy operation o_{dummy} of processing time 0 is defined and constrained to be the first operation on each machine. We also extend the use of a job type τ to serve as a function that returns the type of an operation when written as $\tau(o)$.

The following additional variables are developed for the MIP model: ω_{ij} refers to the start time of operation $o_{ij} \in O$, B_{ab} is a binary variable relating to the precedence constraints between operations o_a and o_b . Note that B_{ab} refers only to direct precedence and not the general notion of o_a being scheduled sometime before o_b . We discretize job types such that τ_a refers to the type of o_a and assume that there is a lexicographic ordering of job types such that processing a job type with a higher value is sufficient to reset the machine

for lower job types according to the maintenance policy described in Section III.

Block starts are marked by binary variables Z_a^c and ζ_a^τ where Z_a^c determines if operation o_a starts a block of operations delineated by a maintenance activity of class c and ζ_a^τ determines if operation o_a starts a block of operations delineated by a job type τ . Idle time values are held by the variables K_a^c and L_a^τ , which correspond to the minimum time elapsed since a maintenance activity of class c preceding o_a and the minimum time elapsed since an operation of type τ preceding o_a respectively. Furthermore, deterioration values at the start of an operation o_a are held by the variable δ_a and are determined by the deterioration values K and L .

Some of the constraints are linearised using big-M variables namely, \mathcal{M}^τ and \mathcal{M}^ω . We define some bounds for these variables in Section IV-B below.

A. THE INTEGER PROGRAMMING MODEL

In this section, we define the integer programming model made up of an objective function, decision variables and constraints.

Objective

$$\min(C_{max}) \quad (1)$$

Decision variables

B_{ab}	Operation o_a directly precedes o_b	$B_{ab} \in \{0, 1\}$
ω_a	Start time of operation o_a	$\theta_a \in \mathbb{R}$
L_a^τ	Minimum time elapsed since operation of a type τ preceding o_a	$L_a \in \mathbb{R}$
K_a	Minimum time elapsed since any maintenance activity preceding o_a	$K_a \in \mathbb{R}$
δ_a	Deterioration of machine at start of o_a	$\delta_a \in \mathbb{R}$
Z_a^c	o_a starts a block delineated by a maintenance activity of class c	$Z_a^c \in \{0, 1\}$
ζ_a^τ	o_a starts a block delineated by a job of type τ	$Z_a^c \in \{0, 1\}$

Constraints

$$\omega_{(i+1)j} \geq \omega_{ij} \quad \forall o_{ij} \in O \quad (2a)$$

$$\omega_{i(j+1)} \geq \omega_{ij} + P_{ij} + S(o_{ij}, o_{i(j+1)}) \quad \forall o_{ij} \in O \quad (2b)$$

$$\omega_{i(j+1)} \leq \omega_{ij} + D(o_{ij}, o_{i(j+1)}) \quad \forall o_{ij} \in O \quad (2c)$$

$$\sum_{o_a \in O} B_{ad} = 0 \quad (2d)$$

$$\sum_{o_a \in O \cup \{o_{dummy}\}} B_{ab} = 1 \quad \forall o_b \in O \quad (2e)$$

$$\sum_{o_b \in O} B_{ab} \leq 1 \quad \forall o_a \in O \quad (2f)$$

$$B_{ab} \leq \sum_{\mu_m \in \mu} \Gamma_{am} \Gamma_{bm} \quad \forall o_a, o_b \in O \quad (2g)$$

$$\omega_b \geq B_{ab}(\omega_a + P_a + S(o_a, o_b)) \quad \forall o_a, o_b \in O \quad (2h)$$

$$\omega_b \geq B_{ab}(\omega_a + P_a + Z_b^c(P_c)) \quad \forall o_a, o_b \in O$$

$$c \in C \quad (2i)$$

$$L_b^\tau \geq B_{ab}(\omega_b - \omega_a - P_a) \quad \forall o_a, o_b \in R,$$

$$\tau \in T \quad (2j)$$

$$L_b^\tau \geq B_{ab}(L_a^\tau + \omega_b - \omega_a) - \mathcal{M}^\omega \zeta_b^\tau$$

$$\forall o_a, o_b \in R, \quad \tau \in T \quad (2k)$$

$$K_a \geq 0 \quad \forall o_a \in R \quad (2l)$$

$$K_b \geq B_{ab}(K_a + \omega_b - \omega_a) - \mathcal{M}^\omega \sum_c Z_b^c$$

$$\forall o_a, o_b \in R \quad (2m)$$

$$\delta_a \geq \min(K_a, L_a^{\tau(o_a)}) \quad \forall o_a \in R \quad (2n)$$

$$\mathcal{M}^\omega(1 - Z_a^c) + (\delta_a - \theta_c)$$

$$\geq 0 \quad \forall o_a \in R, \quad c \in C \quad (2o)$$

$$\mathcal{M}^\omega Z_a^c - (\delta_a - \theta_c)$$

$$> 0 \quad \forall o_a \in R, \quad c \in C \quad (2p)$$

$$\mathcal{M}^\tau \zeta_b^\tau - B_{ab}(\tau(o_a) - \tau)$$

$$\geq 0 \quad \forall o_a, o_b \in R, \quad \tau \in T \quad (2q)$$

$$\mathcal{M}^\tau(1 - \zeta_b^\tau) + B_{ab}(\tau(o_a) - \tau)$$

$$\geq 0 \quad \forall o_a, o_b \in R, \quad \tau \in T \quad (2r)$$

$$C_{max} = \omega_{|J|_r} + P_{|J|_r} \quad (2s)$$

The objective of the model is to minimise makespan denoted by Equation (1). The constraints in Equations (2b) to (2h) apply to all operations while the constraints in Equations (2j) to (2r) only apply to operations scheduled on the re-entrant machine. All non-binary variables are constrained to be non-negative, i.e., start times, idle times and deterioration values all have a lower bound of 0.

Equation (2a) enforces the fixed-order relationship between operations at the same level of the flow shop. Equations (2b) to (2c) enforce setup times and maximum separation constraints between operations of the same job respectively, while Equation (2d) enforces that the dummy operation has no predecessors. Equation (2e) ensures that every operation has exactly one predecessor and Equation (2f) enforces that every operation has at most one successor. Equation (2g) enforces that operations only follow each other if they are mapped to the same machine and Equation (2h) enforces that there is no overlap between operations leaving room for setup times. These make up the constraints that specify the problem without maintenance.

The maintenance constraints follow below. Equation (2i) enforces that there is no overlap between operations while leaving enough room for any maintenance activities that may have been triggered. Equation (2j) and (2k) specify constraints on the minimum time elapsed since an operation of a certain type has come through the machine. Similarly, Equations (2l) and (2m) specify the minimum time elapsed since a maintenance activity of a certain class has been scheduled. The constraints represented by Equations (2j) to (2m) are defined in a cumulative way based on predecessor operations. Equations (2k) and (2m) are activated depending on the

presence of a job type or a maintenance activity respectively. This toggle is implemented by big-M values that are activated based on the binary variables Z and ζ .

The actual deterioration value is computed by Equation (2n) which is set to the minimum of both K and L . Equation (2n) computes deterioration based on the idle time so far and is set to the minimum of K and L so that maintenance is only triggered when necessary. Equations (2o) and (2p) specify that maintenance activities are triggered whenever deterioration thresholds are crossed thereby starting a new block while Equations (2q) and (2r) similarly start a new block based on the relationship between types of operations, i.e., a new block is triggered whenever an operations predecessor has a higher type. Note that this model allows multiple maintenance classes to be triggered simultaneously if the threshold violations cross multiple thresholds. However, Equation (2h) means that the gap left for maintenance corresponds to the largest processing time of all triggered maintenance activities, thus not paying unnecessary maintenance costs.

Finally, Equation (2s) calculates the makespan which in a fixed order problem, is the finishing time of the last operation of the last job.

B. BOUNDS FOR BIG-M VALUES

1) \mathcal{M}^ω

Throughout the model, \mathcal{M}^ω is used as a big-M constraint in two instances. The first is in Equations (2m) and (2k) to sum up the minimum times since the last maintenance or the last occurrence of a type of job and in Equations (2o) and (2p) to toggle maintenance if deterioration thresholds are crossed. In each of these cases, the upper bound is the maximum possible deterioration value that can occur. Because our deterioration deals with idle times, we are then looking for a value that is larger than or equal to the maximum time the machine can be left idle.

An idea for this bound is to use an upper bound on the makespan as there always exists a solution with a better makespan than one in which the machine is left idle for the upper bound on the makespan.

This upper bound assumes the worst case which is that every operations incurs the maximum possible setup time and the maintenance activity with the longest duration occurs before every operation. Thus, the bound is

$$\mathcal{M}^\omega = \sum_{o_a \in O} \left(P_a + \max_{c \in C} (P_c) + \max_{o_b \in O} (S(o_a, o_b)) \right). \quad (3)$$

2) \mathcal{M}^τ

The tightest bound for \mathcal{M}^τ is the largest job type available in the problem. This holds because:

- \mathcal{M}^τ is an upper bound on the types of jobs,
- we assume that job types are all given integer values corresponding to their quality requirements,

- we assume there is a lexicographical ordering of these types that corresponds with the order of the integers representing each job type.

C. LINEARISING THE MODEL

Some equations are still quadratic namely Equations (2h), (2i), (2k) and (2m). They however all involve the product of a binary variable and a non-negative continuous variable and can also be linearised via the big-M method. The corresponding M is also \mathcal{M}^ω . A detailed explanation of how this linearisation is achieved can be found in [36].

Additionally, Equation (2n) requires us to compute the minimum which is also a non-linear equation. We define one more auxiliary binary variable γ_a that is set to 1 if K_a is less than $L_a^{\tau(o_a)}$ and linearise the minimum constraint by replacing it with the following set of equations:

$$\delta_a \leq K_a \quad \forall o_a \in R, \quad (4a)$$

$$\delta_a \leq L_a^{\tau(o_a)} \quad \forall o_a \in R, \quad (4b)$$

$$\delta_a \geq K_a - \mathcal{M}^\omega(1 - \gamma_a) \quad \forall o_a \in R, \quad (4c)$$

$$\delta_a \geq L_a^{\tau(o_a)} - \mathcal{M}^\omega(\gamma_a) \quad \forall o_a \in R, \quad (4d)$$

$$K_a - L_a^{\tau(o_a)} \leq \mathcal{M}^\omega(1 - \gamma_a) \quad \forall o_a \in R, \quad (4e)$$

$$L_a^{\tau(o_a)} - K_a \leq \mathcal{M}^\omega(\gamma_a) \quad \forall o_a \in R. \quad (4f)$$

Equations (4a) and (4b) set the deterioration value δ_a to be upper bounded by the minimum of K_a and $L_a^{\tau(o_a)}$. However, this is not enough as δ_a is still free to take any values less than this and can lead to violations of maintenance constraints. We further use Equations (4c) and (4d) which set δ_a to be lower bounded by the minimum of K_a and $L_a^{\tau(o_a)}$. This lower bound is also achieved via big-M constraints which activate either equation based on γ_a . The combination of the lower and upper bounds ensure that δ_a is exactly set to the minimum of the two values K_a and $L_a^{\tau(o_a)}$. Finally, Equations (4e) and (4f) set γ_a to 0 if K_a is less than $L_a^{\tau(o_a)}$ and 1 otherwise.

V. CONSTRAINT PROGRAMMING APPROACH

Constraint programming (CP) has recently been shown to perform well for scheduling problems [37]. This motivates us to also explore a constraint programming solution. In this section, we present a CP model.

Our CP model uses the idea of *interval variables* and *sequence variables*. These are known constraint programming concepts [37] with the following definitions. Interval variables refer to operations to be scheduled and are declared with a length equal to the processing time of the operation. The goal of the solver is to assign a start time to each of these variables. An additional characteristic of interval variables are that they have the option to either be compulsory, i.e., they must exist in any schedule produced by the solver, or be optional. Sequence variables on the other hand, represent orderings of interval variables. The solver receives these as a set of interval variables with its goal being to decide on a sequencing of these interval variables.

Apart from constraints and variables, constraint programming also provides some auxiliary functions such as **startOf** and **typeOf** which help us access variable properties – in this case, their assigned start times and types respectively.

We define two classes of interval variables, (i) operations which are always present and each retain the representation of o_a and (ii) maintenance activities which are optional and referred to as m_a^c where m_a^c is a maintenance activity of class c that precedes an operation o_a . The variables K , L and R retain their definitions from the MIP model in Section IV.

Given that we have $|A|$ maintenance classes and $|R|$ operations in total on the re-entrant machine, we define $|A||R|$ maintenance activities since the worst case is that there is one maintenance activity of a class before every regular operation. We add constraints such that the maintenance activities are included in the sequence only when deterioration thresholds are violated.

Sequence variables are defined per machine and referenced as $Sequence_m$ for corresponding machine μ_m where $Sequence_m$ contains all operations mapped to μ_m including the optional maintenance activities. For the re-entrant machine, we define an additional sequence variable $SequencePlain_m$ as a sequence of only production operations – excluding maintenance activities – and constrain this sequence to follow the same ordering as $Sequence_m$. The purpose of this duplicate sequence is to ensure that sequence-dependent setup times are respected. The details of how we achieve this follow in the constraint definitions below.

Objective

$$\min(C_{max}) \quad (5)$$

Decision variables

$Sequence_m$	Sequence of production and maintenance operations on machine μ_m	
$SequencePlain_m$	Sequence of production operations on machine μ_m	
L_a^τ	Minimum time elapsed since operation of a type τ preceding o_a	$L_a \in \mathbb{R}$
K_a	Minimum time elapsed since any maintenance activity preceding o_a	$K_a \in \mathbb{R}$

Constraints

$$C1 : \text{before}(Sequence_1, o_{i1}, o_{(i+1)1}) \forall o_{i1} \in O \quad (C1)$$

$$C2 : \text{sameSubsequence}(Sequence_1, Sequence_m) \quad (C2)$$

$$\forall \mu_m \in \mu$$

$$C3 : \text{sameSubsequence}(Sequence_m, SequencePlain_m) \quad (C3)$$

for re-entrant machine μ_k

$$C4 : \text{startOf}(o_{i(j+1)}) \geq \text{startOf}(o_{ij}) + P_{ij} + S(o_{ij}, o_{i(j+1)}) \quad (C4)$$

$$\forall o_{ij} \in O$$

$$C5 : \text{startOf}(o_{i(j+1)}) \leq \text{startOf}(o_{ij}) + D(o_{ij}, o_{i(j+1)}) \\ \forall o_{ij} \in O \quad (C5)$$

$$C6 : \text{noOverlapDirect}(Sequence_m, P, S) \\ \forall \mu_m \in \mu \quad (C6)$$

$$C7 : \text{noOverlapDirect}(SequencePlain_k, P, S) \\ \text{for re-entrant machine } \mu_k \quad (C7)$$

$$C8 : \text{if}(\min(K_a, L_a^{\tau(o_a)}) \geq \theta^c \wedge \min(K_a, L_a^{\tau(o_a)}) < \Theta^c) \Rightarrow \\ \text{presenceOf}(m_a^c) = 1 \\ \forall o_a \in R, c \in C \quad (C8)$$

$$C9 : \text{if}(\text{presenceOf}(m_a^c) = 1) \Rightarrow K_a = 0 \\ \forall o_a \in R, c \in C \quad (C9)$$

$$C10 : \text{if}(\text{presenceOf}(m_a^c) = 0) \Rightarrow \\ K_a = K_{\text{indexOfPrev}(O_a)} + \text{startOf}(O_a) \\ - \text{startOfPrev}(O_a) \forall o_a \in R, c \in C \quad (C10)$$

$$C11 : \text{if}(\text{typeOfPrev}(o_a) \geq \text{typeOf}(o_a)) \Rightarrow \\ L_a^{\tau} = \text{startOf}(o_a) - \text{endOfPrev}(O_a) \\ \forall o_a \in R, \tau \in T \quad (C11)$$

$$C12 : \text{if}(\text{typeOfPrev}(o_a) < \text{typeOf}(o_a)) \Rightarrow \\ L_a^{\tau} = L_{\text{indexOfPrev}(o_a)}^{\tau} + \text{startOf}(o_a) \\ - \text{startOfPrev}(o_a) \forall o_a \in R, \tau \in T \quad (C12)$$

$$C13 : C_{max} = \text{endOf}(o_{|j|_r}) \quad (C13)$$

In this model, Constraint C1 enforces an ordering between the first operations of each job. The **before** constraint enforces precedence relationships between two operations in a sequence. We enforce the order of the first operations of each job which we know will be on the first machine (as we have a flow shop). Constraint C2 builds on C1 to then enforce that this ordering is respected across all other sequences using the **sameSubsequence** constraint. Note that we only enforce a subsequence because the re-entrant machine has operations on multiple levels of the flow shop. The **sameSubsequence** is set up such that only operations at the same level are constrained with the fixed ordering, which is in line with the requirements of our problem. Constraint C3 uses the **sameSubsequence** in a similar way to constrain the duplicate sequences – with and without maintenance activities included – to have the same ordering.

Next, Constraints C4 and C5 enforce the sequence-independent setup times and maximum separation constraints respectively. Both of these apply to operations of the same job as is seen with the index of operations in the constraints.

Sequence-dependent setup time and no overlap constraints are handled by Constraints C6 and C7, which ensure that both the separations required by processing times and sequence-dependent setup times are obeyed. Since maintenance activities are also included in our sequences, we ensure correctness of Constraints C6 by extending the processing and setup times accordingly with setup times set to 0 for operations before or after maintenance activities. The **noOverlapDirect** constraint works such that the separation

denoted by sequence-dependent setup times applies only between direct successors, i.e., say an operation o_a is followed by o_b with a maintenance activity m_b^c in-between, the setup time between o_a and o_b will not be enforced. Thus, setting the maintenance setup time to 0 can lead to constraint violations as the problem is now underconstrained. It is worthy of note that there exists a **noOverlapIndirect** constraint, which applies sequence dependent setup time constraints to all successors; however, this over-constrains the problem.¹ We use the **noOverlapDirect** constraint and circumvent under-constraining the problem by using the supporting Constraint C7 on a duplicate sequence without maintenance.

Constraint C8 enforces the presence of a maintenance activity whenever the minimum deterioration is within the limits of threshold violations. We do not explicitly calculate a deterioration variable δ in this model but this is essentially the left hand side of Constraint C8. We depend on the fact that our problem defines non-overlapping maintenance threshold intervals to ensure that at most one maintenance activity is triggered before an operation.

Constraints C9 and C10 deal with the computation of the minimum time elapsed since a maintenance activity has occurred. Since we are guaranteed to trigger at most one maintenance activity per operation, we do not maintain different minimum elapsed times per maintenance class as was done in the MIP model. Similarly, Constraints C11 and C12 compute the minimum time elapsed since a job of a certain type has been through the machine.

Finally, C13 calculates the makespan, which we again know to be the finishing time of the last operation of the last job.

Worthy of note is that Constraints C10 and C12 are cumulative constraints that could be expressed using the **cumulFunction** constraint, which keeps track of each interval variables contribution to a function [37]. However, many implementations of this function within available solvers require that the contribution of each interval variable be known apriori whereas, in our case, the contribution of each interval variable is itself based on decision variables [37] due to maintenance also being time-dependent.²

VI. HEURISTIC SOLUTION APPROACH

While exact approaches such as those presented in Sections IV and V have lots of advantages, they often do not scale well. In this section, we present an alternate heuristic solution approach to handle larger problem instances. The work presented in this section has appeared earlier in a workshop paper [13].

Our heuristic approach is based on extending list schedulers to integrate maintenance activities in the schedule. Heuristic list schedulers have been developed for the

¹Given a sequence of operations $o_a \rightarrow o_b \rightarrow o_c$, sequence-dependent setup times will be considered from $o_a \rightarrow o_b$, $o_b \rightarrow o_c$ and $o_a \rightarrow o_c$ whereas the only sequence-dependent setup times that should be considered are from $o_a \rightarrow o_b$ and $o_b \rightarrow o_c$.

²Start times of operations are decision variables.

Algorithm 1 Maintenance Aware List Scheduling (MALS)

```

1: function MALS(flow shop  $f$ , operation ordering  $order$ ,
   ranking  $rank$ )
    $\Omega \leftarrow \langle \rangle$   $\triangleright$  returns schedule  $\Omega$ 
2:    $\Omega \leftarrow \langle \rangle$   $\triangleright$  empty schedule
3:    $\Omega' \leftarrow \emptyset$   $\triangleright$  empty set of schedules
4:    $\Omega'' \leftarrow \emptyset$   $\triangleright$  empty set of schedules
5:    $o_p \leftarrow dummy$   $\triangleright$  operation initialised to dummy
   operation
6:   for  $o_c$  in  $order$  do
7:      $\Omega' \leftarrow generateOptions(o_c, f, \Omega)$ 
8:     for  $\omega \in \Omega'$  do
9:        $\omega \leftarrow triggerMaintenance(o_c, o_p, f, \omega)$ 
10:       $\Omega'' \leftarrow \Omega'' \cup \{\omega\}$ 
11:      $\Omega \leftarrow selectHighestRanked(\Omega'', rank)$ 
12:      $o_p \leftarrow o_c$ 
13:      $\Omega'' \leftarrow \emptyset$ 
return  $\Omega$ 

```

industrial problem we consider [38], [39], [40] and are also suitable for online scheduling. Thus, we look into extending them to handle integrated production and maintenance scheduling. The typical flow of a list scheduler is to *order* operations according to some metric and insert them in a schedule one after the other until all operations are scheduled [40].

A. MAINTENANCE-AWARE LIST SCHEDULING

To make a list scheduling approach maintenance-aware, we propose to evaluate the effect of any operation placement on maintenance triggering before making a decision. This leads to a schedule with the necessary maintenance activities triggered by the operation sequence already included. This is shown in Algorithm 1. In Line 1, the scheduler takes as input the flow shop to be scheduled, the chosen ordering of the operations $order$, and the ranking of decisions $rank$. Lines 2–6 initialise the variables used in the algorithm, i.e., an empty schedule Ω that is filled with operations by the algorithm, empty sets of schedules Ω' and Ω'' used to keep track of scheduling options, and an operation o_p to track the last operation that was inserted in the schedule. Specifically, o_p is initialised to a dummy operation for the first run where no insertions have occurred yet. In Line 7, the scheduler loops through each operation o_c in the chosen order and Line 8 finds positions to place the operation in the schedule being built with each possible option resulting in a different schedule stored in the set Ω' . For every one of these schedules, we trigger predicted maintenance in Line 10, which updates the schedules with predicted maintenance activities included. We keep track of the last regular operation placed in the schedule o_p to reduce the amount of work it takes to trigger maintenance as the schedule is already evaluated up to that operation o_p . Eventually, we pick the best option in Line 12 where the ‘best’ is as determined by the supplied ranking $rank$.

Algorithm 2 Trigger Maintenance

```

1: function triggerMaintenance(current operation  $o_c$ ,
   previous operation  $o_p$ , flow shop  $f$ , schedule  $\Omega$ )  $\triangleright$  returns
   schedule  $\Omega$ 
2:   for  $o_i \in \langle o_p, \dots, o_c \rangle$  do
3:      $\Delta \leftarrow \delta(\langle o_1, \dots, o_i \rangle, \mu_k)$   $\triangleright$  predict deterioration
   state
4:     if  $X(\Delta)$  is defined then  $\triangleright$  deterioration triggers
   maint.
5:        $a^c \leftarrow X(\Delta)$   $\triangleright$  insert maint. activity
6:        $\Omega \leftarrow insertMaintenanceOperation(a^c, \Omega)$ 
7:        $\Omega \leftarrow updateStartTimes(f, \Omega)$ 
8:        $feasible \leftarrow checkFeasibility(f, \Omega)$ 
9:       if  $\neg feasible$  then
10:         $\Omega \leftarrow repairSchedule(f, \Omega)$ 
return  $\Omega$ 

```




FIGURE 2. Edge update after inserting a maintenance activity. The original constraints between operations o_{22} and o_{32} are now between operation o_{22} and the maintenance activity with new edges added to connect the maintenance activity to operation o_{32} . This ensures the original constraints of the problem are present and the maintenance activity is scheduled before operation o_{32} .

The steps shown in Algorithm 1 are generic and can be customised to any list scheduler of choice. However, evaluating maintenance is performed according to the steps described in Algorithm 2. For a given schedule, we first go through the operations in the schedule from the last inserted operation o_p to the current operation being inserted o_c in Line 2. For each operation, we evaluate the deterioration state in Line 3. If a maintenance activity is triggered at any point in the schedule, the action is then inserted and the schedule is re-evaluated in Lines 5–9. We approach this by creating an operation a^c to represent the maintenance activity and adjusting the edges in the graph such that the constraints of the original problem remain intact after the insertion of the new operation. This is illustrated in Figure 2 where we show the edges added after inserting a maintenance activity. Since we have hard timing constraints between operations, inserting a maintenance activity can lead to a previously feasible schedule becoming infeasible. In such a case, a schedule repair action is triggered to return the schedule to a feasible state in Line 11. Algorithm 2 assumes that a schedule is always repairable and below in Section VI-B2, we show what the necessary conditions are for this to be true.

B. SCHEDULE REPAIR

Flow shop schedules generally need to obey a certain ordering of operations to be valid. However, re-entrant flow shops with

Algorithm 3 Schedule Repair Strategy

```

1: function repairSchedule(flow shop  $f$ , position  $n$ , schedule  $\Omega$ )
    $\Omega$   $\triangleright$  returns schedule
2:    $feasible \leftarrow false$ 
3:    $end \leftarrow false$ 
4:   while  $\neg feasible \wedge \neg end$  do
5:      $(fp', o_{fp,k}) \leftarrow penultimateFirstPass(n, \Omega)$ 
6:      $(ffp', o_{ffp,k}) \leftarrow ultimateFirstPass(n, \Omega)$ 
7:      $(sp', o_{sp,k+1}) \leftarrow lastSecondPass(n, \Omega)$ 
8:     if  $o_{ffp} = o_{1,k}$  then  $\triangleright$  first operation on machine
9:        $end \leftarrow true$ 
10:     $i \leftarrow sp' + 1$ 
11:    while  $i \leq ffp'$  do
12:       $\Omega \leftarrow removeSecondPassOp(o_{i,k+1}, \Omega)$ 
13:       $\Omega \leftarrow insertSecondPassOp(fp', o_{i,k+1}, \Omega)$ 
14:       $fp' \leftarrow fp' + 1$ 
15:       $i \leftarrow i + 1$ 
16:     $n \leftarrow fp'$ 
17:     $\Omega \leftarrow updateStartTimes(f, \Omega)$ 
18:     $feasible \leftarrow checkFeasibility(f, \Omega)$ 
19:     $\Omega \leftarrow triggerMaintenance(o_{sp}, o_{1,k}, f, \Omega)$ 
20:  return  $\Omega$ 

```

due dates have an additional validity criterion, which is the due date between operations. In a case where operations that are not completely part of the set of input operations – such as maintenance activities – have to be scheduled, due date violations become even more likely. Since these operations are only known when schedules are evaluated, we always have the possibility that a schedule becomes infeasible as a result of these insertions. Furthermore, it is still combinatorial to decide on the repaired version of the schedule that minimizes the makespan after an event that causes infeasibility occurs. We therefore need to develop a schedule repair strategy for this problem.

1) OUR STRATEGY

Schedule repair entails reorganising a schedule to obtain a state where the schedule is valid again [41]. Since we start from a valid schedule that is rendered infeasible by inserting new operations, the infeasibility is due to a due date violation, i.e., an operation has been delayed too long after its preceding operation. Therefore, the fix is to systematically bring operations closer to their predecessors. However, it is not immediately obvious which operations need to be brought forward and how far this needs to go. As such we define a recursive strategy where we take small steps forward and reevaluate the fix until the schedule is feasible again. Additionally, moving operations around can violate the maintenance policy so after re-organisation, it is necessary to re-evaluate the schedule. This solution falls under the class of proactive-reactive dynamic scheduling [42].




FIGURE 3. Schedule repair strategy showing progressive steps in the algorithm. In the first step, the schedule is infeasible because of the maintenance activity (highlighted in green). From this point on, the future steps re-organise the schedule until we achieve a feasible schedule in Step 4. In Step 5, a last step is taken to trigger maintenance again as re-ordering operations could have invalidated existing or triggered new maintenance activities. Operations encircled in dotted lines are the ultimate first pass from the point of failure, the ones circled in a thin line are the penultimate first pass, and the ones circled in a thick line are the last higher pass operation.

As shown in Algorithm 3, every time we reorganise the operations in the schedule, we first identify three key operations, namely, the *penultimate first pass* operation from the point where the schedule was broken, the *last second pass operation* from the point where the schedule was broken, and finally the *last second pass operation* that has been included in the schedule. This is shown in Lines 4–6 where we identify these key operations and their positions in the schedule. We then move all scheduled second pass operations belonging to jobs ranging from the *last second pass* to the *ultimate first pass* in the schedule – this occurs in the remove and insert calls on Lines 13–17. This way, the schedule has been reorganised such that second pass operations from the point of failure are at least a step closer to their first pass operations. We repeat this process until the schedule becomes feasible,³ moving the point of failure a step backward each iteration – this is as seen on Line 18 where the point of failure is updated ahead of the next iteration. After the schedule is deemed feasible, a last step is taken to trigger maintenance again in Line 20 as re-ordering operations could have invalidated or triggered maintenance activities. This re-ordering works because due dates exist only between consecutive operations of the same job.

Figure 3 shows an example of the schedule repair process. In Step 1, the schedule is infeasible after the insertion of a maintenance activity highlighted in green. The ultimate first pass is identified as o_{42} , the penultimate first pass as o_{32} and the last second pass as o_{13} . The operations after the maintenance activity are then brought forward as can be seen in the new placement of o_{23} in Step 2. This continues in Steps 3 and 4 until the schedule is evaluated to be feasible.

It is valuable to point out that the overall algorithm proposed is flexible enough to adopt other repair strategies depending on the use case. An alternate example could be

³It is always possible to find a feasible solution as long as the maintenance policy in use is safe. This is as shown in Theorem 1 below.




FIGURE 4. Slack between operations o_{22} and o_{23} .

the strategy of reducing the rate of production to prevent or delay maintenance activities. A host of possible rescheduling and repair strategies are surveyed in [41].

2) SAFE MAINTENANCE POLICIES

A maintenance policy X maps a deterioration state of the machine to an appropriate maintenance activity. The policy in use determines when and where maintenance activities are necessary. As discussed above, inserting a maintenance activity in a schedule may make the schedule infeasible. We define a *safe* maintenance policy as a policy that ensures that there exists at least one maintenance-aware solution to the flow shop provided there is a feasible schedule for the flow shop alone without considering maintenance activities. Since a schedule becoming infeasible after a maintenance insertion is a result of a violated due date, there should be enough room between consecutive first and second passes of the same job to fit a particular maintenance activity unless the policy is such that the maintenance activity cannot be triggered between first and second passes of the same job. Concretely, this means that the processing time of any maintenance activity a^c that can be triggered between passes of the same job o_{ik} and $o_{i(k+1)}$ should fit in the available time between them, i.e.,

$$P_c \leq D(o_{ik}, o_{i(k+1)}) - P_{ik} - S(o_{ik}, o_{i(k+1)}) \quad \forall o_{ik}, o_{i(k+1)}. \quad (6)$$

Theorem 1: Given an infeasible schedule, the schedule repair strategy defined in Algorithm 3 is always able to return it to a state of feasibility in at most $|J|$ iterations, where $|J|$ is the number of jobs in the schedule, provided that a solution exists for the problem and the maintenance policy in use is safe.

Proof: For an insertion of a maintenance activity a^c between operations o_{ik} and $o_{i(k+1)}$ to become infeasible due to a due date violation, it means that $o_{i(k+1)}$ has been delayed too long, i.e., $\omega_{i(k+1)} - \omega_{ik} > D(o_{ik}, o_{i(k+1)})$. To avert this, the maintenance activity must be able to fit in the slack between both operations. Bearing in mind that other operations could be placed between o_{ik} and $o_{i(k+1)}$, the slack $\Psi(o_{ik}, o_{i(k+1)})$ left between o_{ik} and $o_{i(k+1)}$ is

$$\begin{aligned} \Psi(o_{ik}, o_{i(k+1)}) &= D(o_{ik}, o_{i(k+1)}) - P_{ik} \\ &\quad - \max((S(o_{ik}, o_a) + P_a + \dots + S(o_a, o_b) + P_b), \\ &\quad S(o_{ik}, o_{i(k+1)})), \end{aligned} \quad (7)$$

where operations o_a, \dots, o_b represent operations possibly placed between o_{ik} and $o_{i(k+1)}$. Figure 4 shows an example

where the slack between operations o_{22} and o_{23} is

$$\begin{aligned} \Psi(o_{22}, o_{23}) &= D(o_{22}, o_{23}) - P_{22} - \max(S(o_{22}, o_{13}) - P_{13} \\ &\quad - S(o_{13}, o_{23}), S(o_{22}, o_{23})). \end{aligned} \quad (8)$$

The repair algorithm progressively brings operations closer to their direct predecessors by at least one step per iteration. In the last possible iteration of the schedule repair, each operation $o_{i(k+1)}$ follows its direct predecessor o_{ik} . It follows that this occurs in at most $|J|$ iterations of the schedule repair as the re-entrant machine can only have $|J|$ higher pass operations to be re-ordered. At this point, Equation (7) becomes

$$\Psi(o_{ik}, o_{i(k+1)}) = D(o_{ik}, o_{i(k+1)}) - P_{ik} - S(o_{ik}, o_{i(k+1)}). \quad (9)$$

For this to be infeasible, it means that a^c cannot fit in $\Psi(o_{ik}, o_{i(k+1)})$, i.e., $P_c > D(o_{ik}, o_{i(k+1)}) - P_{ik} - S(o_{ik}, o_{i(k+1)})$, which violates the rules of a safe maintenance policy shown in Equation (6). ■

VII. EXPERIMENTAL RESULTS

This section evaluates the empirical performance of the three solution approaches we propose. We apply the heuristic approach as an add-on to three existing list schedulers in the literature to evaluate the applicability of this approach to list-scheduling. We compare our heuristics against the two exact approaches (integer and constraint programming) to evaluate their accuracy and scalability.

A. EXPERIMENTAL SETUP

All experiments are performed on a 16-core 1.9GHz AMD machine running Ubuntu 20.04 with 32GB RAM. Algorithms are implemented in C++ and the MIP and CP models are solved by CPLEX version 22.1 and CP Optimizer version 22.1, respectively. The exact approaches are all given a 30 minute timeout.

We generate benchmarks according to the types of jobs typically presented in our industrial use case as described in Table 1. We generate benchmarks with patterned arrivals of job types such that jobs of a type appear in repeated blocks, e.g., a set of 50 jobs can be made of 20 type 1 jobs followed by 10 type 2 jobs and then 20 type 3 jobs. We randomise the length of the blocks and number of times these blocks repeat to mimic arrival patterns of jobs in practice. We generate 50 instances for each job size in $\{5, 10, 50, 100, 150, 200, 300, 500, 1000\}$.

Our heuristic approach is implemented as an extension to three schedulers from the literature – Bounded Heuristic Constraint Scheduler (BHCS) [39], As Soon As Possible (ASAP) Scheduler, and Modified Nawaz-Enscore-Ham (MNEH) Heuristic [43]. BHCS is a list scheduler developed specifically for our use case, while the ASAP scheduler is also a list scheduler that uses the same ordering requirements as BHCS but places operations as soon as possible (ASAP). MNEH is a modification of the popular NEH heuristic [44], [45] that is suitable for re-entrancy. Maintenance-incorporated versions of these schedulers are referred to as

TABLE 1. Properties of jobs in use case. All timings are in seconds and job travelling times are treated as setup times between operations of the same job.

Type	$P(o_{i1})$	$P(o_{i2})$	$P(o_{i3})$	$P(o_{i4})$	$D(o_{i1}, o_{i2})$	$D(o_{i2}, o_{i3})$	$D(o_{i3}, o_{i4})$
0	0.25	0.30	0.30	0.21	0.85	12.30	1.00
1	0.35	0.42	0.42	0.30	0.95	12.42	1.12
2	0.50	0.59	0.59	0.42	1.10	12.59	1.29
3	0.70	0.84	0.84	0.60	1.30	12.84	1.54
4	0.99	1.19	1.19	0.85	1.59	13.19	1.89

(a) Job processing times and due dates

Machine	Setup Time
μ_1	0.20
μ_2	0.05
μ_3	1.00

Path	Travelling Time
μ_1 to μ_1	0.60
μ_2 to μ_2	10.00
μ_2 to μ_3	0.70

Activity Class	Duration	Deterioration States
1	0.5s	10 – 15
2	10s	15 – 30
3	20s	30 – ∞

(b) Machine setup times

(c) Job travelling times

(d) Maintenance policy




FIGURE 5. Makespan improvement of maintenance-included versions over base versions. Instances where the solver timed out without providing any solution are marked with *.

MIBHCS, MIASAP, and MINEH respectively where the MI prefix refers to “maintenance incorporated”. In each of these experiments, we tune the heuristic approach to include a maintenance activity if a deterioration threshold is crossed or if 90%⁴ of the upper bound of a threshold that affects the quality of an operation further down the line is crossed. Since we insert maintenance between two operations, we always have complete information about the next operation. We can also reliably infer what operations are further down the line for the entire planning window based on which operations have already been scheduled.

⁴This value can be tuned. We chose 90% after performing a parameter sweep that showed this value performed best.

In the basic schedulers – BHCS, ASAP, and MNEH – maintenance is reactive and interrupts the schedule during production runs. We simulate the behaviour of reactive maintenance in these schedulers by evaluating the completed schedules they produce for maintenance and compare these with versions of the scheduler that incorporate our proactive maintenance heuristic.

B. PERFORMANCE EVALUATION

Figure 5 compares the makespan of the schedules produced by MIBHCS, MIASAP, and MINEH to the makespan of schedules produced by BHCS, ASAP, and MNEH respectively. We also compare the exact solutions CP and MIP with the best solutions provided by MIBHCS, MIASAP, and

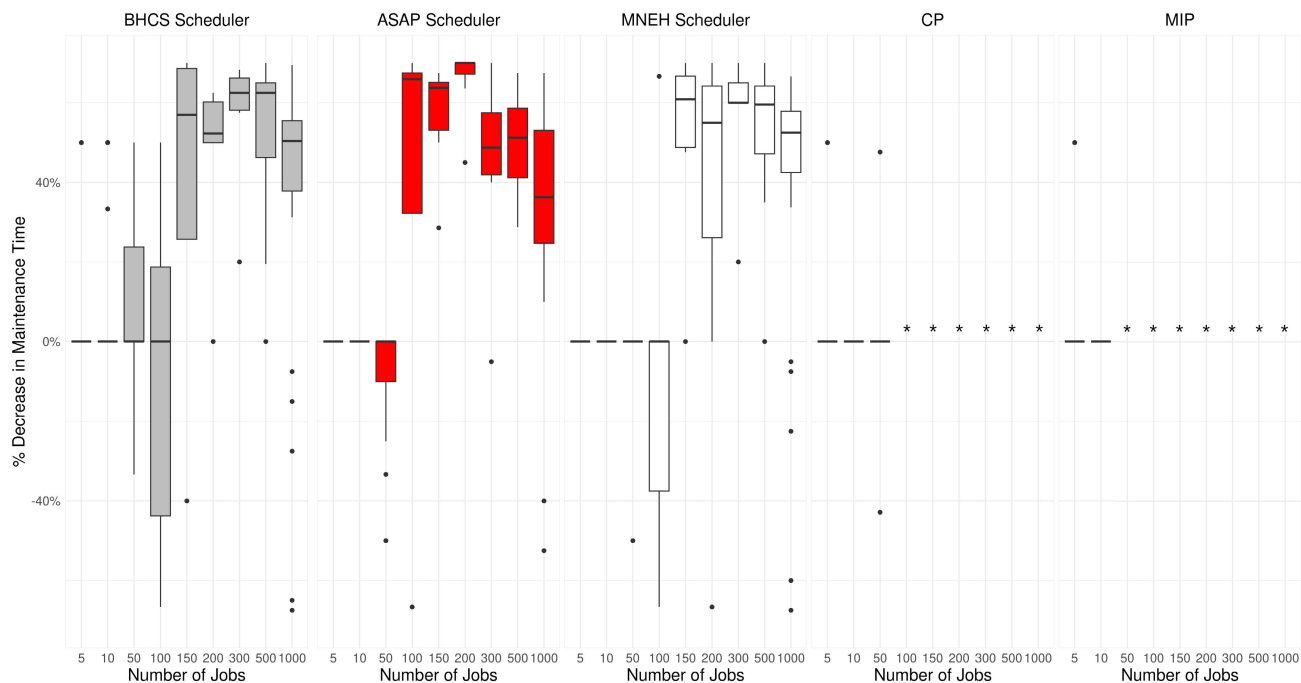


FIGURE 6. Duration of maintenance activities. Instances where the solver timed out without providing any solution are marked with *.

MINEH. MNEH has the least performance improvement due to it not being a pure list scheduler. With MNEH, only relative positions of operations are decided in each iteration and there is no partial sequence that is guaranteed to remain the same from one iteration to the next; as such the evaluation of the deterioration of a machine loses some meaning from one iteration to the next since sequences change at each iteration. The exact approaches – CP and MIP – should ideally always be better than all of the heuristic approaches but they are sometimes worse because they do not always solve till optimality within the time out.

Figure 6 shows the distribution of the time spent on maintenance. We see that with maintenance-included versions, we spend up to 70% less time on maintenance. This is because considering deterioration allows us to perform maintenance before machines deteriorate to a state where we have to pay larger maintenance costs. The difference is also this significant because there is up to one order of magnitude difference between the durations of different maintenance activities for this use case (see Table 1). This difference translates to shorter makespans for the schedulers.

In both Figures 5 and 6, there are instances where the heuristic approach worsens the results particularly for smaller job sets. The instances that are worsened by the heuristic are a result of (i) scenarios where the heuristic maintenance trigger is too conservative and performs maintenance even though the job set could be completed without it, and (ii) scenarios where the list scheduler picks a sequence that triggers shorter maintenance activities.

Neither of the exact solutions are able to scale to provide solutions for larger job sets within the 30 minute

time out – this accounts for the missing columns in Figures 5 and 6. In Table 2 we show the performance of the CP and MIP solutions. We see that the CP model is able to solve more instances than the MIP model but for the instances where the MIP model is able to provide solutions, the optimality gap is smaller.

The runtime increases with the number of jobs as expected and Table 3 shows the average runtime over the job size of the different schedulers compared in this evaluation. The exact approaches are given a 30 minute timeout and in bold are the solutions with the worst run times for a job size. Instances where no solution was provided by a method before time out are left unfilled and are the worst for that job size. The heuristic solutions are able to provide solutions in runtimes below 350ms for job sizes up to 500. Above that, the runtime grows to 1800ms. The biggest time sink for the heuristic solutions is how often the maintenance evaluation and consequently schedule repair is triggered.⁵ This is based on the operation of the base scheduler itself. MNEH evaluates whole sequences while ASAP and BHCS evaluate partial sequences at every decision point thus triggering maintenance evaluations more often, leading to higher runtimes.

In summary, we find that the heuristic approach is scalable and can produce competitive results compared to exact solvers even for small instance sizes. In general, we also find that apart from improving the actual goal of reduced makespan, integrated production and maintenance planning can also reduce the total time spent on maintenance which can result in reduced costs in some cases.

⁵Runtimes of ASAP, BHCS, and MNEH are similar.

TABLE 2. Performance of CP and MIP solutions.

Jobs	Optimality gap (%)		Time to find first solution (s)		% of instances solved	
	CP	MIP	CP	MIP	CP	MIP
5	0.60	0.00	0.82	2.44	100	100
10	1.10	4.60	2.38	419.80	100	26
50	166.01	-	84.95	-	100	0

TABLE 3. Average runtime of solution methods (s).

Jobs	CP	MIASAP	MIBHCS	MINEH	MIP
5	505.16	0.00	0.00	0.00	5.81
10	1083.43	0.01	0.01	0.00	1522.37
50	1611.23	0.67	0.56	0.07	-
100	-	2.72	1.94	0.32	-
150	-	7.06	4.84	1.12	-
200	-	14.19	8.65	2.23	-
300	-	42.10	23.11	5.94	-
500	-	164.71	84.80	30.19	-
1000	-	1756.40	854.44	303.66	-

VIII. CONCLUSION

Efficient maintenance scheduling is important for sustained productivity of industrial processes. This paper studied the problem of sequence- and time-dependent maintenance and presented three solution methods namely, mixed integer programming, constraint programming and a heuristic solution. As the problem is motivated by an industrial use case, we have evaluated all the methods on jobs in this case. We show that list scheduling heuristics can be extended to include proactive maintenance with significant performance gains over reactive approaches.

This paper considers maintenance activities that are on the same time scale as the jobs themselves. An interesting future direction is to include longer-term maintenance planning in the scope and to investigate the combined problem of production and maintenance planning over multiple time scales.

Additionally, we solve the problem from a predictive maintenance perspective, i.e., where maintenance actions are carried out based on the health status of machines. However, this requires knowledge of how machines deteriorate and this information is not always available. Many other papers consider a preventive maintenance perspective where the challenge is either scheduling around a set maintenance schedule or determining what the maintenance schedule itself should be. While we know that preventive maintenance runs the risk of either maintaining machine too little or too often compared to the needs-based approach of predictive maintenance, and both preventive and predictive maintenance have been shown to outperform reactive maintenance approaches, it is still interesting to compare both approaches and determine what problem properties make it necessary to use one or the other. This is because even when complete information on the health status of machines is available, the gains made by integrating them in the decision making process may not necessarily be worth the increased runtime.

ACKNOWLEDGMENT

The authors thank Hadi Ara, Joost van Pinxten, and Joan Marcè i Igual for their support.

REFERENCES

- [1] J. Kang and C. G. Soares, "An opportunistic maintenance policy for offshore wind farms," *Ocean Eng.*, vol. 216, Nov. 2020, Art. no. 108075.
- [2] Z. Ren, A. S. Verma, Y. Li, J. J. E. Teuwen, and Z. Jiang, "Offshore wind turbine operations and maintenance: A state-of-the-art review," *Renew. Sustain. Energy Rev.*, vol. 144, Jul. 2021, Art. no. 110886.
- [3] S. Chansombat, P. Pongcharoen, and C. Hicks, "A mixed-integer linear programming model for integrated production and preventive maintenance scheduling in the capital goods industry," *Int. J. Prod. Res.*, vol. 57, no. 1, pp. 61–82, Jan. 2019.
- [4] F. N. Avilés, R. M. Etchepare, M. M. Aguayo, and M. Valenzuela, "A mixed-integer programming model for an integrated production planning problem with preventive maintenance in the pulp and paper industry," *Eng. Optim.*, vol. 55, no. 8, pp. 1352–1369, Aug. 2023.
- [5] H. Gharoun, M. Hamid, and S. A. Torabi, "An integrated approach to joint production planning and reliability-based multi-level preventive maintenance scheduling optimisation for a deteriorating system considering due-date satisfaction," *Int. J. Syst. Sci., Oper. Logistics*, vol. 9, no. 4, pp. 489–511, Oct. 2022.
- [6] R. J. K. Netto, E. de Freitas Rocha Loures, E. A. P. Santos, and C. F. dos Santos, "Joint industrial preventive maintenance and production scheduling: A systematic literature review," in *Proc. Int. Conf. Flexible Automat. Intell. Manuf.* Cham, Switzerland: Springer, 2023, pp. 614–621.
- [7] N. Zhang, K. Cai, Y. Deng, and J. Zhang, "Determining the optimal production–maintenance policy of a parallel production system with stochastically interacted yield and deterioration," *Rel. Eng. Syst. Saf.*, vol. 237, Sep. 2023, Art. no. 109342.
- [8] A. Valet, T. Altenmüller, B. Waschneck, M. C. May, A. Kuhnle, and G. Lanza, "Opportunistic maintenance scheduling with deep reinforcement learning," *J. Manuf. Syst.*, vol. 64, pp. 518–534, Jul. 2022.
- [9] S. Gawiejnowicz, "A review of four decades of time-dependent scheduling: Main results, new topics, and open problems," *J. Scheduling*, vol. 23, no. 1, pp. 3–47, Feb. 2020.
- [10] S.-J. Yang, "Single-machine scheduling problems with both start-time dependent learning and position dependent aging effects under deteriorating maintenance consideration," *Appl. Math. Comput.*, vol. 217, no. 7, pp. 3321–3329, Dec. 2010.
- [11] W. Liu, X. Wang, L. Li, and P. Zhao, "A maintenance activity scheduling with time-and-position dependent deteriorating effects," *Math. Biosci. Eng.*, vol. 19, no. 11, pp. 11756–11767, 2022.
- [12] B. Mor and G. Mosheiov, "Scheduling a maintenance activity and due-window assignment based on common flow allowance," *Int. J. Prod. Econ.*, vol. 135, no. 1, pp. 222–230, Jan. 2012.
- [13] E.-A. Eigbe, B. De Schutter, M. Nasri, and N. Yorke-Smith, "Predictive maintenance scheduling in twice re-entrant flow shops with relative due dates," in *Proc. 15th Workshop Scheduling Planning Appl. (SPARK) 32nd Int. Conf. Automated Planning Scheduling (ICAPS)*, 2022.
- [14] Y.-C. Su, F.-T. Cheng, M.-H. Hung, and H.-C. Huang, "Intelligent prognostics system design and implementation," *IEEE Trans. Semicond. Manuf.*, vol. 19, no. 2, pp. 195–207, May 2006.
- [15] G. A. Susto, A. Schirru, S. Pampuri, S. McLoone, and A. Beghi, "Machine learning for predictive maintenance: A multiple classifier approach," *IEEE Trans. Ind. Informat.*, vol. 11, no. 3, pp. 812–820, Jun. 2015.
- [16] M. Delorme, M. Iori, and N. F. M. Mendes, "Solution methods for scheduling problems with sequence-dependent deterioration and maintenance events," *Eur. J. Oper. Res.*, vol. 295, no. 3, pp. 823–837, Dec. 2021.

- [17] S.-J. Yang, "Parallel machines scheduling with simultaneous considerations of position-dependent deterioration effects and maintenance activities," *J. Chin. Inst. Ind. Eng.*, vol. 28, no. 4, pp. 270–280, Jun. 2011.
- [18] L. Jin, X. Yu, and Z. Dong, "Single-machine scheduling with piece-rate maintenance, interval constrained processing times and rejection penalties," in *Proc. 3rd Joint Int. Inf. Technol. Mech. Electron. Eng. Conf. (JIMEC)*, 2018, pp. 32–37.
- [19] H. Zhu, M. Li, Z. Zhou, and Y. You, "Due-window assignment and scheduling with general position-dependent processing times involving a deteriorating and compressible maintenance activity," *Int. J. Prod. Res.*, vol. 54, no. 12, pp. 3475–3490, Jun. 2016.
- [20] A. J. Ruiz-Torres, G. Paletta, and E. Pérez, "Parallel machine scheduling to minimize the makespan with sequence dependent deteriorating effects," *Comput. Oper. Res.*, vol. 40, no. 8, pp. 2051–2061, Aug. 2013.
- [21] A. J. Ruiz-Torres, G. Paletta, and R. M'Hallah, "Makespan minimisation with sequence-dependent machine deterioration and maintenance events," *Int. J. Prod. Res.*, vol. 55, no. 2, pp. 462–479, Jan. 2017.
- [22] V. L. A. Santos and J. E. C. Arroyo, "Iterated greedy with random variable neighborhood descent for scheduling jobs on parallel machines with deterioration effect," *Electron. Notes Discrete Math.*, vol. 58, pp. 55–62, Apr. 2017.
- [23] J. Ding, L. Shen, Z. Lü, and B. Peng, "Parallel machine scheduling with completion-time-based criteria and sequence-dependent deterioration," *Comput. Oper. Res.*, vol. 103, pp. 35–45, Mar. 2019.
- [24] A. S. Xanthopoulos, A. Kiatipis, D. E. Koulouriotis, and S. Stieger, "Reinforcement learning-based and parametric production-maintenance control policies for a deteriorating manufacturing system," *IEEE Access*, vol. 6, pp. 576–588, 2018.
- [25] Y. Wang, E. Elahi, and L. Xu, "Selective maintenance optimization modelling for multi-state deterioration systems considering imperfect maintenance," *IEEE Access*, vol. 7, pp. 62759–62768, 2019.
- [26] R. J. Abumaizar and J. A. Svestka, "Rescheduling job shops under random disruptions," *Int. J. Prod. Res.*, vol. 35, no. 7, pp. 2065–2082, Jul. 1997.
- [27] E. Kutanoglu and I. Sabuncuoglu, "Routing-based reactive scheduling policies for machine failures in dynamic job shops," *Int. J. Prod. Res.*, vol. 39, no. 14, pp. 3141–3158, Jan. 2001.
- [28] W. T. Chan and T. H. Wee, "A multi-heuristic GA for schedule repair in precast plant production," in *Proc. 13th Int. Conf. Automated Planning Scheduling (ICAPS)*, 2003, pp. 236–245.
- [29] A. Allahverdi, "Two-machine proportionate flowshop scheduling with breakdowns to minimize maximum lateness," *Comput. Oper. Res.*, vol. 23, no. 10, pp. 909–916, Oct. 1996.
- [30] P. Caricato and A. Grieco, "An online approach to dynamic rescheduling for production planning applications," *Int. J. Prod. Res.*, vol. 46, no. 16, pp. 4597–4617, Aug. 2008.
- [31] K. Katragjini, E. Vallada, and R. Ruiz, "Rescheduling flowshops under simultaneous disruptions," in *Proc. Int. Conf. Ind. Eng. Syst. Manage. (IESM)*, Oct. 2015, pp. 84–91.
- [32] R. L. Graham, E. L. Lawler, J. K. Lenstra, and A. H. G. R. Kan, "Optimization and approximation in deterministic sequencing and scheduling: A survey," *Ann. Math.*, vol. 5, pp. 287–326, Jan. 1979.
- [33] M. Y. Wang, S. P. Sethi, and S. L. van de Velde, "Minimizing makespan in a class of reentrant shops," *Oper. Res.*, vol. 45, no. 5, pp. 702–712, Oct. 1997.
- [34] H. Emmons, G. Vairaktarakis, H. Emmons, and G. Vairaktarakis, "Reentrant flow shops," *Flow Shop Scheduling: Theoretical Results, Algorithms, and Applications*, 2013, pp. 269–289.
- [35] F. Hnaien, F. Yaloui, A. Mhadhbi, and M. Nourelfath, "A mixed-integer programming model for integrated production and maintenance," *IFAC-PapersOnLine*, vol. 49, no. 12, pp. 556–561, 2016.
- [36] H. P. Williams, *Model Building in Mathematical Programming*. Hoboken, NJ, USA: Wiley, 2013.
- [37] P. Laborie, J. Rogerie, P. Shaw, and P. Vilfm, "IBM ILOG CP optimizer for scheduling: 20+ years of scheduling with constraints at IBM/ILOG," *Constraints*, vol. 23, no. 2, pp. 210–250, Apr. 2018.
- [38] U. Waqas, M. Geilen, J. Kandelaars, L. Somers, T. Basten, S. Stuijk, P. Vestjens, and H. Corporaal, "A re-entrant flowshop heuristic for online scheduling of the paper path in a large scale printer," in *Proc. Design. Autom. Test Eur. Conf. Exhib. (DATE)*, Mar. 2015, pp. 573–578.
- [39] J. V. Pinxten, U. Waqas, M. Geilen, T. Basten, and L. Somers, "Online scheduling of 2-Re-entrant flexible manufacturing systems," *ACM Trans. Embedded Comput. Syst.*, vol. 16, no. 5s, pp. 1–20, Oct. 2017.
- [40] R. van der Tempel, J. van Pinxten, M. Geilen, and U. Waqas, "A heuristic for variable re-entrant scheduling problems," in *Proc. 21st Euromicro Conf. Digit. Syst. Design (DSD)*, Aug. 2018, pp. 336–341.
- [41] G. E. Vieira, J. W. Herrmann, and E. Lin, "Rescheduling manufacturing systems: A framework of strategies, policies, and methods," *J. Sched.*, vol. 6, no. 1, pp. 39–62, 2003.
- [42] D. Ouelhadj and S. Petrovic, "A survey of dynamic scheduling in manufacturing systems," *J. Scheduling*, vol. 12, no. 4, pp. 417–431, Aug. 2009.
- [43] B. Jeong and Y.-D. Kim, "Minimizing total tardiness in a two-machine re-entrant flowshop with sequence-dependent setup times," *Comput. Oper. Res.*, vol. 47, pp. 72–80, Jul. 2014.
- [44] M. Nawaz, E. E. Enscore, and I. Ham, "A heuristic algorithm for the m-machine, n-job flow-shop sequencing problem," *Omega*, vol. 11, no. 1, pp. 91–95, Jan. 1983.
- [45] W. Liu, Y. Jin, and M. Price, "A new improved NEH heuristic for permutation flowshop scheduling problems," *Int. J. Prod. Econ.*, vol. 193, pp. 21–30, Nov. 2017.



EGHONGHON-AYE EIGBE received the M.Sc. degree (cum laude) in embedded systems from the Delft University of Technology, where she is currently pursuing the Ph.D. degree with the Software Technology Department. Her current research interests include robust scheduling, maintenance planning, and the intersection of machine learning and optimization.



BART DE SCHUTTER (Fellow, IEEE) received the Ph.D. degree (summa cum laude) in applied sciences from KU Leuven, Belgium, in 1996. He is currently a Full Professor and the Head of department with the Delft Center for Systems and Control, Delft University of Technology, The Netherlands. His current research interests include multi-level and multi-agent control, model predictive control, learning-based control, and control of hybrid systems, with applications in intelligent transportation systems and smart energy systems. He is a Senior Editor of the IEEE TRANSACTIONS ON INTELLIGENT TRANSPORTATION SYSTEMS and an Associate Editor of the IEEE TRANSACTIONS ON AUTOMATIC CONTROL.



MITRA NASRI (Member, IEEE) received the B.Sc. and M.Sc. degrees in computer engineering (software systems) from the Iran University of Science and Technology, Iran, in 2005 and 2008, respectively, and the Ph.D. degree from the University of Tehran, Tehran, Iran, in 2015. She is currently an Assistant Professor with the Eindhoven University of Technology (TU/e), Eindhoven, The Netherlands. Before joining TU/e, she was an Assistant Professor with the Delft University of Technology (TUDelft). Her current research interests include modeling, designing, and verifying real-time systems. She has been an ACM Member, since 2018. Since 2022, she has been an Executive Member of the IEEE Technical Committee on Real-Time Systems (TCRTS) which steers RTSS, RTAS, and ICCPS conferences, and the IEEE Benelux Chapter on Communication and Vehicular Technology (COM/VT).



NEIL YORKE-SMITH received the Ph.D. degree in artificial intelligence from Imperial College London, in 2004. He is currently an Associate Professor with the Delft University of Technology, The Netherlands, where he directs the STAR Laboratory. His current research interests include data-driven optimization, reinforcement learning, agent-based modeling, and social simulation. He is a Senior Member of AAAI and ACM.

...