Dynamic Scheduling Optimization for Component Maintenance, Repair, and Overhaul Shops

A case study for an independent component maintenance provider

T.W. Roolvink

Dynamic Scheduling Optimization for Component Maintenance, Repair, and Overhaul Shops

A case study for an independent component maintenance provider

by

T.W. Roolvink

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on July 16, 2025.

Student number:	4961382	
Project duration:	October, 2024 - July, 2025	
Thesis committee:	Dr. A. Bombelli	TU Delft, Chair
	Dr. M.J. Ribeiro	TU Delft, Supervisor
	Ir. P.C. Roling	TU Delft, Supervisor
	Dr. E.J.J. Smeur	TU Delft, Examiner
	Ir. K. Alizadeh	Company Supervisor
	L. Koningsbloem	Company Supervisor

TU Delft

Acknowledgements

It all started around nine months ago, when I first stepped into the world of component maintenance and the mysterious land of job shop scheduling. At the time, I could hardly imagine that phrases like "tardiness" or "resource allocation" would one day casually appear in my conversations. From mapping out the internal shop processes within the company, to puzzling over operational bottlenecks, to running simulations on real-life datasets, this journey has been anything but boring. There were moments of triumph, like seeing the first feasible schedule appear after hours of work, and moments of despair, such as discovering that my algorithm happily violated yet another operational constraint. If my thesis research were a game, I would say I unlocked both the "Data Detective" and "Schedule Whisperer" achievements. But every challenge served the bigger goal: finalizing my studies at TU Delft with a thesis I could be proud of. Looking back, the journey was filled with lessons, laughter, and lots of learning, and while it was not always smooth, it was undoubtedly rewarding. What a ride it has been.

I want to extend my heartfelt appreciation to my supervising team from the TU Delft, Marta Ribeiro and Paul Roling, for their invaluable guidance throughout the duration of the thesis. Without your valuable feedback, suggestions for improvement, and extended knowledge, I could not have gotten the thesis to the level it has reached. Our bi-weekly meetings over the past months provided a consistent source of motivation, clarity, and direction, helping me stay on track and deepen my understanding of the topics addressed.

Moreover, my profound appreciation is extended to Kaveh and Laurens from the independent component maintenance provider, who have helped me with their knowledge and supportive attitude from the beginning of my internship. You guided me in deepening my understanding of the world of component maintenance, and our weekly meetings were instrumental in the progress I made during my thesis. I am grateful for your constant availability and willingness to answer questions whenever they arose. I would also like to express my gratitude to the entire data science and shop departments for creating such a welcoming environment and approachable attitude. The positive team spirit made my time at the company highly enjoyable.

Lastly, I would like to thank my family and friends for their unwavering support throughout my entire journey at TU Delft. Your encouragement, patience, and belief in me have been invaluable, and I truly could not have reached this point without you.

Thijs Roolvink
Amsterdam, July 2025

Contents

List of Abbreviations

ACO	Ant Colony Optimization
B&B	Branch and Bound
B&C	Branch and Cut
BIP	Binary Integer Programming
BRKGA	Biased Random-Key Genetic Algorithm
CMRO	Component Maintenance, Repair, and Overhaul
CP	Constraint Programming
DR	Dispatching Rule
DRL	Deep Reinforcement Learning
EDD	Earliest Due Date
FCFS	First-Come-First-Serve
FFSP	Flexible Flow Shop Problem
FIFO	First-In, First-Out
FJSSP	Flexible Job Shop Scheduling Problem
FSSP	Flow Shop Scheduling Problem
GA	Genetic Algorithm
HFSP	Hybrid Flow Shop Problem
IDG	Integrated Drive Generator
ISA	International Standard Atmosphere
JSSP	Job Shop Scheduling Problem
KPI	Key Performance Indicator
LNS	Large Neighborhood Search
LP	Linear Programming
MILP	Mixed Integer Linear Programming
MIP	Mixed Integer Programming
MTWR	Most Total Work Remaining
NP	Nondeterministic Polynomial time
OEE	Overall Equipment Effectiveness
OR	Operations Research
OTD	On-Time Delivery

P&H	Pneumatics & Hydraulics
PDD	Promised Delivery Date
PSMWAP	Multi-Skilled Workforce Allocation Problem
PSO	Particle Swarm Optimization
RFE	Ready for Evaluation
RL	Reinforcement Learning
RTB	Ready to Build
SLA	Service Level Agreement
SPT	Shortest Processing Time
SWT	Shortest Waiting Time
TAT	Turn-Around Time
TIST	Technician-Intensive Shop Task
TWT	Total Weighted Tardiness
UIE	Unit in Evaluation
WIP	Work In Progress
WMDD	Weighted Modified Due Date
WO	Work Order

Introduction

Throughout my master's program at TU Delft, I developed a strong interest in operations optimization, particularly in airline planning and cargo operations. The courses I did in these fields allowed me to work on complex real-world problems and provided rewarding opportunities to apply optimization techniques in complex and challenging settings. This interest was further strengthened during the course on maintenance modeling and analysis, where I was introduced to the potential of data-driven methods within aerospace maintenance. Motivated by the desire to combine these areas in my thesis, I started searching for a graduation project that would offer such an interdisciplinary challenge.

The independent component maintenance provider, as a historically significant name in aerospace, immediately stood out to me. After discussions with several people within the company, the idea emerged to develop an optimization approach for the scheduling of component maintenance, repair, and overhaul (CMRO) shops. These environments do not have standard, plug-and-play planning solutions due to their operational complexity, making them an ideal setting for applying tailored optimization approaches.

What makes this project particularly relevant is the integration of physical operational processes with data-driven modeling. Bridging this gap required close collaboration with shop-floor leads to ensure that key features of the maintenance environment were accurately represented and that proposed solutions were practically feasible. Regular discussions with operational staff and management helped validate both the model structure and its outcomes, enhancing the practical applicability of the work. While the job shop problem has been studied extensively in the literature, environments involving human technicians rather than only machines remain underexplored. This human-centric aspect of the workshop setting contributed to the scientific and practical relevance of the project.

This thesis report consists of two parts and is structured as follows: Part I presents the scientific paper, which includes the problem statement, the methodology applied, the outcomes of the conducted experiments, and the resulting conclusions. Part II contains the supporting literature review, including a state-of-the-art analysis of relevant research, the identification of the research gap, and the research plan formulated at the start of the thesis.

I

Scientific Paper

Dynamic Scheduling Optimization for Component Maintenance, Repair, and Overhaul Shops

Thijs Roolvink,*

Delft University of Technology, Delft, The Netherlands

Abstract

Effective scheduling is challenging in Component Maintenance, Repair, and Overhaul (CMRO) operations due to the complexity of dynamically allocating resources across multiple jobs with varying priorities and technical constraints. Current industry practices typically rely on static, manual scheduling, resulting in suboptimal resource allocation and insufficient adaptability to operational disruptions. Most existing studies approach specific job shop problems by incorporating individual features, such as job prioritization or resource constraints, without considering the combined operational complexities of CMRO shops. Therefore, this research presents a scheduling model using the Flexible Job Shop Scheduling Problem (FJSSP), tailored to dynamic CMRO environments. The model uses Mixed Integer Linear Programming (MILP) to simultaneously schedule technicians and machines while accounting for specialized skill requirements, resource constraints, and job prioritization. The approach balances multiple objectives, including tardiness and earliness, to enhance shop performance metrics such as Turnaround Time (TAT) and On Time Delivery (OTD) rates. Outcomes from a case study applied to real-world data from CMRO shops demonstrate significant operational improvements, achieving a reduction in TAT of up to 34% and an improvement in OTD by approximately 23% relative to historical shop performance. Furthermore, the model incorporates schedule robustness measures, minimizing deviations from planned schedules, despite operational uncertainties. Additionally, comparative analysis with a traditional heuristic dispatching rule model confirms the superior performance of the proposed optimization framework. This framework can be broadly applied to improve scheduling efficiency and stability in CMRO shops and similar workshop environments.

1 Introduction

Component Maintenance, Repair, and Overhaul (CMRO) shops play an important role in ensuring the reliability and availability of aircraft by performing maintenance on various aviation components. These shops handle thousands of parts annually, varying from navigation instruments to power generation systems, using highly trained and experienced technicians for specific inspection, repair, and testing tasks. Efficient scheduling in these environments is important for maximizing resource utilization and maintaining service-level agreements to improve customer satisfaction, minimize operational costs, and reduce turnaround times. In current CMRO environments, scheduling processes rely on simple prioritization models and manual decision-making. The allocation of operations to technicians is mainly done by the expertise and judgment of the shop lead, rather than relying on a scheduling algorithm. While this is effective to some extent, this approach lacks the flexibility to adapt to dynamic operational changes. Additionally, this inefficient scheduling can lead to increased Work-In-Progress (WIP), delays in important repairs, and higher labor costs, making scheduling optimization an important topic in CMRO operations.

Recent advancements in scheduling optimization demonstrate the potential impact of integrating data-driven methods. For instance, new scheduling model tools have been shown to improve Overall Equipment Effectiveness (OEE) by over 3%, reduce planning-related labor hours by more than 50%, and improve sustainability and customer satisfaction (Kumar and He, 2023). Such results emphasize the importance of transitioning from manual, experience-based scheduling to automated, optimization-driven systems. There are many other examples in the aviation industry where efficient scheduling has significantly reduced maintenance efforts and improved utilization. For example, an earlier study on aircraft heavy maintenance check scheduling introduced a genetic algorithm-based approach that reduced the total number of heavy maintenance checks by 7%. Additionally, it increased aircraft utilization by 4.4%, potentially leading to significant annual maintenance cost savings (van der Weide et al., 2022). These improvements show the potential of advanced scheduling techniques in aviation maintenance operations.

Recent literature on job shop scheduling highlights advancements in exact methods, dispatching rules, and metaheuristics for solving complex scheduling problems. Mixed Integer Linear Programming (MILP) remains

*MSc Student, Sustainable Air Transport, Faculty of Aerospace Engineering, Delft University of Technology

commonly used, providing optimal solutions with flexibility to model detailed operational constraints. Several studies have explored the integration of job priorities, skilled technician constraints, robust optimization, and dynamic rescheduling separately, but few address these aspects together. Specifically within CMRO shops and similar maintenance or production settings, current models generally overlook the dual-resource allocation challenges, sequence-dependent setup times, technician assignment constraints, and rescheduling requirements together. This research addresses this gap by proposing a general framework that incorporates all of these key components and can be easily tailored to CMRO environments. To evaluate the proposed model properly, a case study was performed in collaboration with an independent component maintenance provider, where the scheduling framework was tested using real-world data from multiple CMRO shop environments.

Accordingly, this paper proposes a scheduling model developed for job shop environments commonly found in CMRO operations. The model intends to autonomously plan maintenance tasks, addressing several crucial aspects common in such environments. First, it should incorporate all relevant operational constraints to generate optimal schedules for these complex settings. Second, it should be designed to remain effective and robust in environments with dynamic disruptions such as job arrivals, operation delays, and the insertion of additional tasks. Another important consideration is the integration of multiple performance objectives, specifically tardiness and earliness, to optimally balance Turnaround Time (TAT) and On-Time Delivery (OTD).

This paper is structured as follows. Section 2 describes the problem by outlining current scheduling challenges, operational constraints, and shop-specific characteristics within the CMRO environment. Section 3 reviews the relevant literature, focusing on the research gap this paper is trying to address. The case study context, shop instances, and key performance indicators are introduced in Section 4. The methodological approach, including the exact MILP-based model, dynamic rescheduling mechanism, and robustness strategies, is presented in Section 5. Section 6 explains the simulation setup used to evaluate the scheduling models. The results, including model validation, performance, and comparative evaluations, are described and analyzed in Section 7. Section 8 discusses the operational implications of the findings. Finally, Section 9 presents the conclusions drawn from this study and outlines possible directions for future research.

2 Problem Definition

This section defines the scheduling problem in CMRO shops by describing current planning inefficiencies in Section 2.1. Next, Section 2.2 outlines routing characteristics and operational aspects of the Hydraulics & Pneumatics (H&P) and Power Generation shops. Thereafter, Section 2.3 summarizes general operational requirements that impact scheduling complexity. Finally, Section 2.4 presents the expected future scheduling model, designed to respond dynamically to disruptions.

2.1 Current Challenges and Inefficiencies

Currently, scheduling processes in many CMRO shops are primarily static and rule-based, relying heavily on manual decision-making. Production leaders decide the allocation of Work Orders (WOs) based on their experience, with almost no use of assignment tools or models (Avelino et al., 2016). While technician capacity is partially scheduled at the beginning of the week, the remaining capacity is handled reactively throughout the week, often leading to inefficiencies and delays. The causes of the current, sub-optimal way of scheduling include:

- Limited Flexibility: The current system struggles to adjust to dynamic changes such as unforeseen delays, urgent job arrivals, and variations in processing times.
- Suboptimal Resource Allocation: Technician skills and equipment capabilities are not fully taken into account, leading to lower utilization and potentially lower service levels.
- Inefficient Prioritization: The prioritization of WOs is based on a basic points-based approach, incorporating customer importance and priority levels. However, this approach does not utilize an advanced framework to dynamically schedule on these parameters.
- Semi-Manual Decision-Making: Decision-making relies on a combination of a simple prioritization model and manual decision-making, without using advanced, data-driven methods to improve and optimize the scheduling of maintenance operations. This approach does not consider the influence of current decisions on the allocation of future operations.

2.2 Shop-Specific Environments and Routing Characteristics

The two types of CMRO shops this research focuses on each have unique characteristics and constraints, which complicate the scheduling process and need customization in the scheduling model. Sections 2.2.1 and 2.2.2 describe the processes in H&P and Power Generation or Integrated Drive Generator (IDG) shops, respectively.

2.2.1 Hydraulics and Pneumatics Shop

Within the CMRO environment, the H&P shop specializes in hydraulic and pneumatic components, operating through two workgroups. An overview of the typical routing steps in the H&P shop is shown in Figure 1. When a new WO arrives, it first receives the status Ready for Evaluation (RFE). Upon assignment to a technician, the status changes to Unit in Evaluation (UIE). During this phase, an initial inspection is performed to assess the condition of the unit and determine the complexity of the repair. Based on this evaluation, a quotation is prepared and sent to the customer. If the customer approves the quote, the WO status changes to Ready to Build (RTB), and subsequently to In Progress (IP) once a technician starts working on the WO. After completion of all required tasks and a final inspection, the unit is ready to be shipped. While some WOs may follow alternative routing paths due to specific component requirements, most follow this general process.

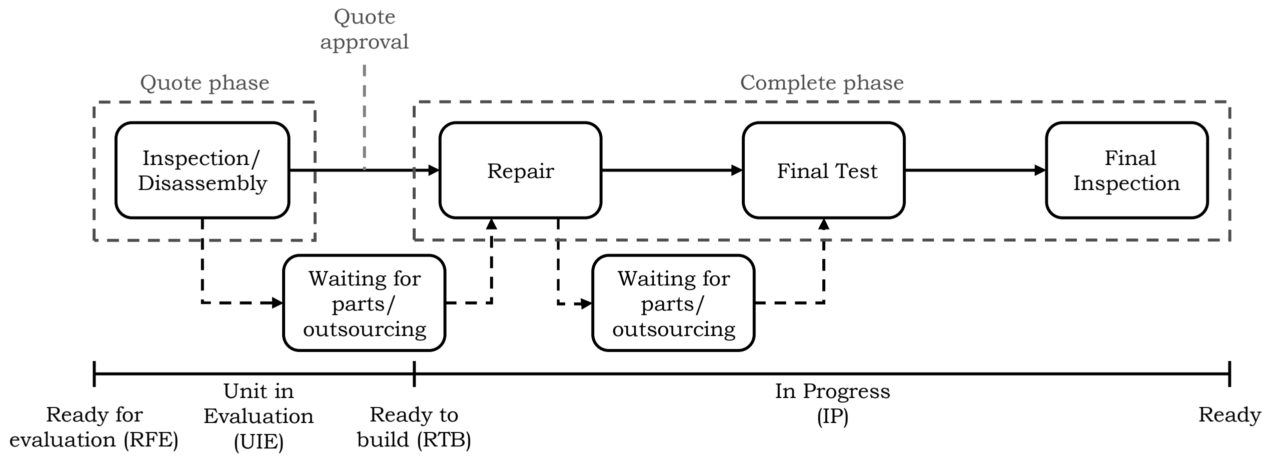

Figure 1: Routing steps and WO statuses for H&P shop.

The current way of scheduling in the H&P shop starts with the list of workable WOs, which are prioritized by the simple points-based model. This list is used by the production leads to assign technicians to the most important jobs, based on their knowledge of the individual technician skills. Based on their indication of the processing times, they choose the number of WOs to be handled by each technician. The technicians each get an overview of the WOs assigned to them on their overview screen.

2.2.2 Power Generation Shop

The Power Generation or Integrated Drive Generator (IDG) shop presents the most complex workflow, with a significantly larger number of routing steps, as shown in Figure 2. In this process, the quote approval point, which separates the quote and completion phases, occurs after the disassembly or inspection step. Furthermore, the number of operations requiring the use of the bench test machines is considerably higher than in the H&P shop, with some WOs passing through the final bench test multiple times depending on the test outcome.

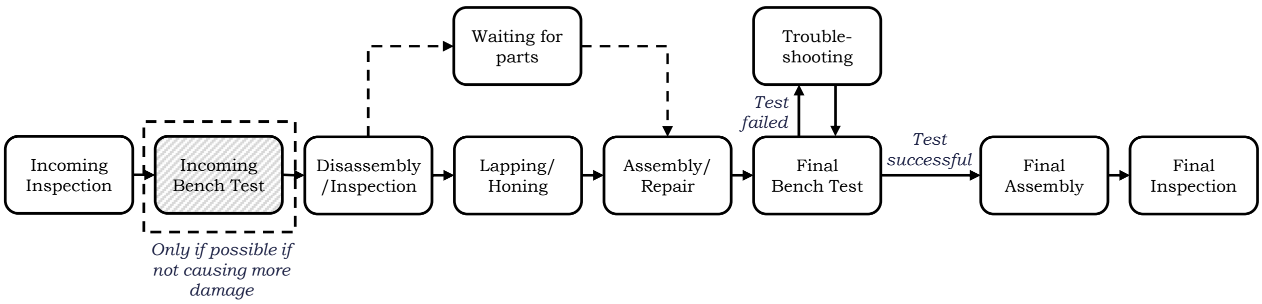

Figure 2: Routing steps for the Power Generation shop.

In Figure 3, an overview of the current way of scheduling for the power generation shop is given. Similar to the other shops, the WOs are ordered based on priority, but separately for each routing step. WOs are moved sequentially through the steps, with job assignments handled by the production leader. The jobs in each step are represented as placeholders, indicating the WOs currently in progress at that stage. With this approach, WOs are tracked and progressed correctly, creating a clear overview of where specific jobs are in the process.

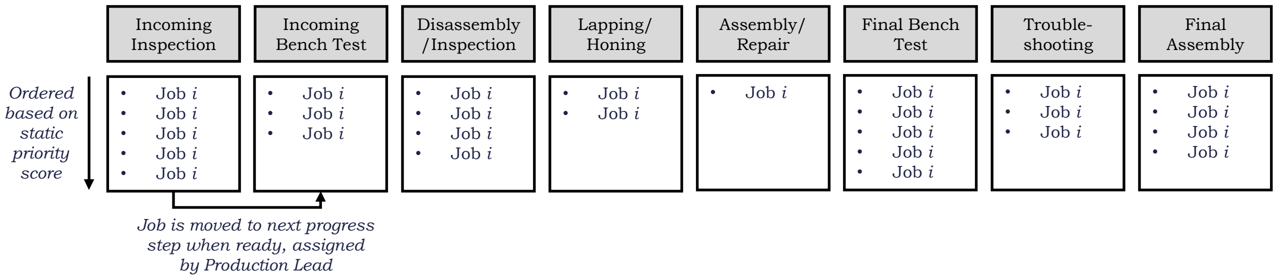

Figure 3: Current scheduling process for the Power Generation shop.

Several practical constraints and preferences further complicate the scheduling environment in the Power Generation shop. First, specialized rework operations, such as lapping and honing, may occur at different stages, sometimes during the quotation phase, and sometimes during the repair phase. A specialized technician ideally performs these operations, or when not available, another capable, right-skilled technician can step in if necessary. Notably, one technician can operate two of these machines together, increasing capacity but additionally creating further complexity in incorporating resource allocation rules in the model.

For the testing phase, failure of the final test requires that the WO must be re-inserted into the system and assigned additional routing steps, including a new test operation, to complete the repair. This reintroduction of tasks increases the variability in job flow. Furthermore, a small buffer before the final due date is preferred to avoid tasks finishing exactly at the deadline. Frequent updates to the schedule should be made at least multiple times per day to help process dynamic changes, such as unforeseen delays or urgent customer requests, while ensuring that once a job operation has been initiated, it is carried through to completion without interruption.

2.3 General Operational Aspects

Several common operational aspects impact shop performance and scheduling complexity. The final inspection step, which represents the last routing step in the workflow, must be performed exclusively by certified technicians, whose limited availability introduces additional bottlenecks in scheduling. Additionally, uncertainties in testing outcomes and processing durations complicate scheduling further, as the standard durations serve as rough indicators rather than precise estimates.

Technicians have a degree of flexibility and can operate across multiple work centers, but their allocation must respect capacity constraints. Furthermore, it is standard practice that, whenever feasible, the technician who initiates work on a particular routing step of a WO should continue handling subsequent steps. Additionally, technicians are required to perform all remaining operations of a WO back-to-back, since a WO cannot be reassigned to another technician before it is completed, except for specific cases such as test operations or final inspections. Related to resource efficiency, WOs requiring similar testing can benefit from being batched, improving the utilization and processing time on machines because no setup time is needed.

Customer prioritization further provides scheduling challenges, as high-priority customers require expedited processing, yet the current scheduling system lacks the reactivity necessary to adapt to these varying demands effectively. Furthermore, technician skills introduce another layer of scheduling complexity since specific tasks or parts often need unique technician capabilities, thus constraining flexibility in workforce allocation.

Additionally, shops communicate expected shipping dates for finalized work orders to customers, making reliable scheduling crucial. While certain unforeseen circumstances, such as delays due to parts availability, can occasionally impact reliability, maintaining accurate and dependable shipping date predictions remains an important aspect of the overall scheduling strategy. Improving scheduling accuracy enhances customer satisfaction by managing expectations effectively and ensuring clarity in communication.

2.4 Future Scheduling

Figure 4 shows the schematic representation of the future scheduling model, illustrating its key elements for practical application in CMRO environments. The model uses the required parameters, including job priorities, process durations, and due dates, to generate an initial optimized schedule. The rescheduling step is triggered automatically in response to any dynamic, disruptive events, including delays in ongoing operations, unexpected job arrivals, or the insertion of additional routing steps due to failed test operations. This feature ensures that the model can optimize schedules in reaction to real-time operational changes, enhancing overall performance and efficiency.

An important requirement for the effectiveness of this dynamic scheduling approach is a practical computational time. Due to the frequent need to re-run the model throughout the operational period, especially

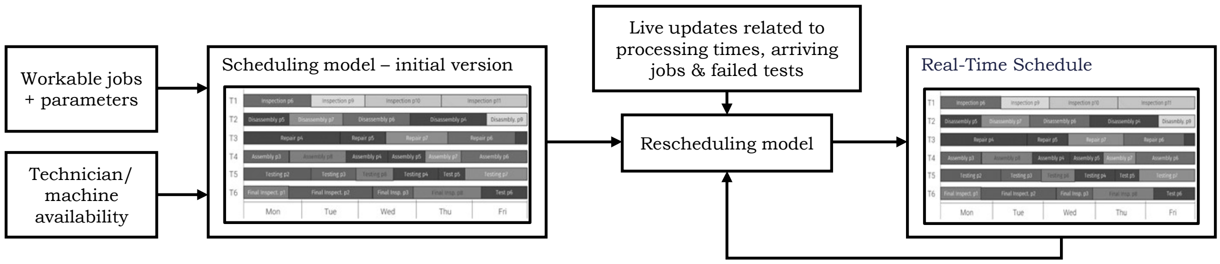

Figure 4: Overview of the future desired scheduling model for the CMRO shops.

following unforeseen disruptions, a practical computational time limit must be set to balance scheduling quality and optimality. This limit ensures timely updates and manageable model execution times, ensuring a realistic integration of the scheduling model in real-world CMRO shops.

3 Literature Review

This section reviews the literature relevant to scheduling optimization in complex operational environments such as CMRO shops. It discusses various solution methodologies in Section 3.1, including exact and heuristic approaches. Section 3.2 analyzes the incorporation of job priorities and technician skill constraints in scheduling models. Next, decomposition strategies to manage computational time are addressed in Section 3.3, followed by a review of rescheduling policies and robust scheduling methods in dynamic environments in Section 3.4. Finally, Section 3.5 identifies the research gap this thesis aims to address.

Scheduling problems are generally classified based on their machine environments, job characteristics, and optimization objectives (Pinedo, 2016). These frameworks provide structure for understanding the nature of scheduling problems and selecting the right solution approach based on their constraints, operational flows, and environmental features. The most widely studied classification is the Job Shop Scheduling Problem (JSSP), which involves multiple jobs that each follow a specific sequence of operations across machines, and is known to be nondeterministic polynomial-time hard (NP-hard) (Xiong et al., 2022). A variant of this problem is the Flexible Job Shop Scheduling Problem (FJSSP), which incorporates routing flexibility, meaning each operation of a job can be allocated to one of several alternative machines (Özgüven et al., 2010). With this routing flexibility, the problem is well-suited to apply in complex operational environments such as CMRO shops.

3.1 Solution Methods

The JSSP has been extensively studied, and many solution methods have been developed over the years. One of the earliest and most straightforward methods is Johnson's algorithm, introduced by (Johnson, 1954), which provides an optimal solution for two-machine flow shop problems. The FJSSP is solved using various solution methodologies, including exact methods, dispatching rules, mathematical programming, and metaheuristics. Dispatching rules are one of the most common methods for solving scheduling problems, often used due to their simplicity and computational efficiency. These rules prioritize jobs dynamically and are rule-based, making them useful in environments requiring real-time decision-making. Common examples are First-In-First-Out (FIFO), Most Work Remaining (MWR), and Earliest Due Date (EDD) (Meilanitasari and Shin, 2021). Despite their effectiveness, dispatching rules focus on locally optimized decisions and often fail to find a global optimum for the objective, potentially leading to suboptimal overall performance (Zahmani et al., 2021). MILP, one of the most-used solution methods for FJSSP, provides an exact optimization framework by formulating the scheduling problem using linear equations and inequalities, incorporating continuous and integer variables to represent scheduling decisions (Dauzère-Pérès et al., 2024; Hillier and Lieberman, 2015). In job shop scheduling, MILP models try to optimize objectives such as minimizing makespan, total earliness, or weighted tardiness (González-Neira et al., 2017). The strengths of MILP models are their ability to find the most optimal solutions, their flexibility with easy adaptation and implementation of various constraints and objectives, and their precision in accurately modeling complex scheduling scenarios (Wang et al., 2025), but it can require significant computational effort (Karam et al., 2017). Metaheuristics, including Genetic Algorithms (GA) (Yu et al., 2018), Ant Colony Optimization (ACO) (Qin et al., 2018), and Large Neighborhood Search (LNS) (Fathollahi-Fard et al., 2024), offer alternatives, achieving near-optimal solutions more rapidly. Recent advancements introduce reinforcement learning techniques, leveraging Markov Decision Processes and Deep Reinforcement Learning (DRL) to manage stochastic and dynamic conditions effectively (Tassel et al., 2021; Yan et al., 2022). An overview of earlier work related to the JSSP in similar environments, applying these solution methods, is provided in Table 1. As shown in this table, MILP is a commonly applied solution method for job shop scheduling problems, followed by metaheuristic approaches such as GA and ACO.

3.2 Job Priorities and Technician Skills

Including job priorities in scheduling models ensures that schedules consider strategic business objectives and customer importance. Hashimoto et al. (2011) introduced priority-based scheduling in the Technicians and Interventions Scheduling Problem (TIST), where priorities are assigned to interventions, and their objective function is designed to minimize the weighted completion time of interventions based on their priority levels. Additionally, Hsieh et al. (2024) proposes a method that minimizes Total Weighted Tardiness (TWT) by assigning weights to jobs based on factors such as customer importance, order profitability, and due dates. By incorporating these weights into the scheduling objective, the method ensures that high-priority jobs, such as urgent customer orders, are scheduled earlier, reducing overall tardiness.

The assignment of skilled technicians, necessary in CMRO environments, requires the use of qualification constraints. Ciro et al. (2015) and Damm et al. (2024) illustrate approaches to multi-skilled technician assignment, highlighting the need for binary skill parameters and adaptive assignment strategies. Additionally, the model proposed by Annear et al. (2023) dynamically assigns resources to maximize productivity while maintaining flexibility to adapt to changing and uncertain demands, possibly applicable in dynamic settings such as CMRO shops. Moreover, Aribowo et al. (2020) provides a way of integrating dedicated technicians, the concept that once a technician starts on the first operation on a given job, that technician must perform all subsequent operations on that job, aligning with specific constraints of CMRO shops.

3.3 Decomposition

Most solution methods for job shop scheduling, particularly exact approaches such as MILP, require significant computational time to compute optimal solutions when applied to large problem instances, as shown in earlier research summarized in Table 1. For example, Tighazoui et al. (2021) reports that solving an instance with 39 jobs can require several hours of computation time. This computational time can be a limiting factor in dynamic and complex environments, where quick rescheduling is useful. To address this, decomposition methods, such as time-based and machine-based decomposition techniques, manage computational complexity by breaking down large-scale problems into manageable subproblems (Pinedo, 2016). In particular, the rolling horizon approach has proven effective in dynamic environments, iteratively re-optimizing schedules within smaller time intervals (Ikli, 2022).

3.4 Rescheduling Policies and Robust Scheduling

In dynamic environments, schedules need to adapt to real-time updates such as new job arrivals and unexpected delays. It is often necessary to reschedule future events based on this new information, which significantly impacts the original scheduling decisions. Gomes et al. (2013) demonstrated this through an approach where the insertion of new jobs requires updating the initial schedule, affecting the sequencing and timing of subsequent operations. Predictive-reactive strategies represent the predominant rescheduling method in dynamic environments, where an initial schedule is adjusted after disruptions (Ouelhadj and Petrovic, 2009; Gomes et al., 2013). Rescheduling policies vary from periodic, event-driven to hybrid approaches, each offering different levels of responsiveness and stability. Kianpour et al. (2021) argues for hybrid policies combining periodic and event-driven methods, providing a balanced solution adaptable to CMRO shop environments.

Robust scheduling methodologies further improve reliability by incorporating buffer times and minimizing worst-case deviations, as demonstrated by Jamili (2016) and Fathollahi-Fard et al. (2024), significantly reducing the impact of processing uncertainties and disruptions. Other robust scheduling strategies, such as those proposed by Rahmani and Heydari (2014) and Xiong et al. (2013), focus on maintaining schedule stability during rescheduling by minimizing deviations from the initial schedule. Applying such methods in the CMRO environment can enhance operational reliability despite unforeseen disruptions.

3.5 Research Gap

Table 1 provides an overview of relevant literature on the JSSP in environments with similar properties to CMRO shops. The table compares the objective functions, solution methods, and the extent to which each work addresses elements relevant to CMRO environments. Among these studies, the approaches by Tighazoui et al. (2021) and Tliba et al. (2022) are the most promising for application in a CMRO shop. Tighazoui et al. (2021) introduces a predictive-reactive model that balances efficiency and stability, while Tliba et al. (2022) proposes an MILP model with job insertion for dynamic rescheduling. However, neither addresses all properties found in CMRO environments.

Reference paper	Sol. method	Objective	Job priorities	Technicians skills	Dynamic environment	Robust optimization	Decomposition	Computational time[a]
Johnson (1954)	Johnson's	Makespan						Low
Rahmani and Heydari (2014)	MILP	Makespan			X	X		Low
Xiong et al. (2013)	GA	Makespan			X	X		Low
Gomes et al. (2013)	MILP	Earliness/Tardiness			X			Medium
Jamili (2016)	MILP/PSO[b]	Makespan			X	X		Medium
Qin et al. (2018)	ACO	Makespan			X		X	Medium
Aribowo et al. (2020)	MILP	Makespan		X				Medium
Kianpour et al. (2021)	MILP	Earliness/Tardiness			X			Medium
Tassel et al. (2021)	RL[b]	Makespan						Medium
Tighazoui et al. (2021)	MILP	Weight. waiting time	X		X	X		High
Tliba et al. (2022)	MILP	Makespan		X	X			High
Melchiors et al. (2024)	DR	Weighted tardiness	X		X			Low
Damm et al. (2024)	GA	Multi-objective	X	X				High
Fathollahi-Fard et al. (2024)	LNS	Cost			X	X		High

[a] Computational time classifications: Low (< 3 minutes), Medium (3–30 minutes), High (> 30 minutes).
[b] PSO: Particle Swarm Optimization; RL: Reinforcement Learning.

Table 1: Overview of the relevant literature applying the JSSP.

From Table 1, it can be observed that most existing approaches focus on one or two properties needed for the scheduling model applicable in CMRO environments. This emphasizes the need to combine multiple scheduling elements, such as job prioritization, skilled technician constraints, robust optimization, and dynamic rescheduling policies, to address the complex environment of CMRO operations. Furthermore, while CMRO environments present a good example of the combination of complex, dynamic, and uncertain conditions, most existing research has not fully captured these shopfloor environments. Many state-of-the-art models address either deterministic and static industrial environments or use randomly generated instances. Additionally, the classic job shop problem has not often been applied in a workshop floor environment where there is a need to allocate jobs to both technicians and machines as separate, interdependent resources, introducing a dual-resource scheduling challenge. Moreover, what differentiates the CMRO environment is the combination of operational complexities such as sequence-dependent setup times, the incorporation of resource unavailability periods, and specialized technician assignment constraints, among other factors, which are rarely addressed together in existing scheduling literature. Overall, despite extensive research, an optimization framework simultaneously addressing all these elements within CMRO environments remains unexplored, highlighting a significant research gap that this paper aims to address.

4 Case Study

This section introduces the case study, where the proposed scheduling model is applied in two different CMRO shop environments. Section 4.1 outlines the resource configurations and additional operational aspects of the shops. Section 4.2 defines the performance metrics used for model evaluation, consistent with internal operational objectives. Lastly, Section 4.3 describes the dataset used in this study, with a focus on its completeness, accuracy, and suitability for model validation.

The case study was performed with an independent component maintenance provider, a significant player in aerospace maintenance, and a former aircraft manufacturer. The current manual scheduling approach in their CMRO shops can result in suboptimal resource allocation, with only 80% of capacity scheduled at the start of the week and the remaining 20% managed reactively throughout the week. This reactive scheduling, often forced by unforeseen disruptions, can lead to increased task delays and inefficiencies because, in those cases, proactively adjusting for dynamic changes is more complex (Khoshsirat and Mousavi, 2024). These current scheduling practices show several challenges that increase the complexity of reaching the desired operational efficiency, as outlined earlier in Section 2.1. The maintenance provider operates three CMRO shops, each with specialized processes and unique characteristics. The H&P shop consists of two subgroups, one for each component category. These subgroups operate within the same shop but handle separate work orders and rely on different test benches. Secondly, the Avionics shop specializes in repairing and maintaining displays, control units, and flight data systems. Lastly, the power generation shop works on parts related to the power generation of aircraft, such as IDGs. Among the three CMRO shops at the maintenance provider, this study focuses on the Hydraulics & Pneumatics (H&P) shop and the Power Generation shop, referred to as the IDG shop.

4.1 Shop Instances

In addition to the operational differences between shops with different part specializations, as discussed in Section 2.2, the size of each shop also varies, which in turn affects the scale of the scheduling problem. Table 2 provides an overview of the available resources in each shop, including the number of technicians and test benches.

Shop	Number of technicians	Number of test benches	Different types of test benches	Total number of resources in FJSSP
Hydraulics	8	4	2	12
Pneumatics	10	8	4	18
IDG	16[b]	2	2	20[a]

[a] Includes two additional rework stations for honing and lapping operations.
[b] Excluding the test operators.

Table 2: Overview of the number of resources for each specific shop.

It is important to highlight that in the Pneumatics shop, certain test resources are rarely utilized. This limited usage reduces the complexity associated with scheduling these resources, as they do not frequently create limitations within scheduling. In contrast, the IDG shop shows different characteristics. Operations that need to be handled on test benches in this environment often result in bottlenecks in scheduling, with the need to allocate all jobs to only two test benches. Besides, a larger number of technicians need to be allocated to operations, increasing the scheduling complexity. A clear assumption can be made when comparing the three CMRO environments regarding resource intensity. The Hydraulics shop represents the smallest scale, followed by the Pneumatics shop, with the IDG shop being the largest and most resource-constrained.

Another bottleneck related to technician specialization in the CMRO shops is the limited availability of certified technicians authorized to perform final inspections. Each WO requires such an inspection before it can be signed off as completed. In the H&P shops, only 2 technicians across the technician teams are certified to conduct final inspections. Similarly, in the IDG shop, just 4 out of 16 technicians hold the required certification. This small ratio of certified technicians creates a challenging bottleneck in the workflow, especially during peak periods or when multiple inspections must be carried out simultaneously.

4.2 Key Performance Indicators

Management of the maintenance provider has identified OTD rate & TAT as the main indicators of CMRO shop effectiveness. While current performance is being closely monitored, there is clear recognition that improvements in these metrics are possible. Implementing advanced scheduling techniques is considered necessary to unlock this potential and enhance overall operational efficiency.

First, the OTD rate is calculated based on whether work orders meet their Promised Delivery Date (PDD). The contractual agreement with customers determines the PDD and is a crucial metric for evaluating scheduling reliability and customer satisfaction. Although the current OTD performance indicates room for improvement, management has set ambitious targets to achieve consistently higher rates.

Second, the TAT metric captures the duration required to complete work orders, from initiation to delivery. This metric varies notably between different shops, reflecting variations in workflow efficiency and operational constraints. TAT can be analyzed at different stages, such as the time to prepare quotes or the time required from customer approval to delivery. Reducing TAT is important, and each shop has targets to enhance efficiency. For example, the target TAT for the incoming WOs quotes was set at a maximum of 7 days for the IDG shop.

Lastly, the Service Level agreements for all customer priority groups must be consistently met to ensure operational reliability and customer trust. However, achieving high service levels for customers classified as high-priority is especially important due to their significant impact on business performance and reputation.

4.3 Dataset

The maintenance provider offered the datasets used in this research, containing historical work order data from the Hydraulics, Pneumatics, and IDG shops. This dataset was used for validating and evaluating the proposed scheduling model. These sets contain comprehensive details such as WO status updates over time, routing steps with both planned and actual hours, and technician availability.

The historical status updates for each WO include precise timestamps indicating when a part was in progress, became available for processing, or was waiting for necessary components. These detailed status logs enable a comparative analysis between currently used methods and the outcomes of the proposed model. Additionally, these timestamps state the job availability at each scheduling instance, providing essential inputs for accurate model simulations.

Data related to routing steps, representing specific operations for WOs, contains planned and actual time durations. Due to incomplete logging, it was sometimes difficult to determine the exact start and end times of technicians for each operation. In such cases, it was assumed that the technician started work from the initial routing step when specific assignments were unclear. Furthermore, planned operation durations found in the dataset are mostly accurate, though minor variations in reliability exist across different shops. Resource-related data, containing technician-specific working hours per part number, certifications, and registered weekly availability, showed high accuracy and consistency.

The dataset also includes detailed records on customer priority levels and promised delivery dates, which are contractually established for each work order. These customer-related metrics are essential for evaluating the scheduling performance in relation to service level agreements and operational efficiency targets. Overall, the quality of the dataset provided by the maintenance provider ensures that it is well-suited for testing and evaluation of the scheduling model proposed in this study.

5 Methodology

This section presents the methodological approach for this research. Section 5.1 defines the operational assumptions used to model the CMRO environment. Section 5.2 presents the exact scheduling approach, including the formulation of the MILP model and rescheduling strategies. Section 5.3 introduces the dispatching rule heuristic used as a performance benchmark.

5.1 Assumptions

To identify the most appropriate scheduling approach for the CMRO shops, the following operational assumptions about the environment were made:

- Each job has a predefined, unique, initial routing sequence.
- Processing times are estimated using standard hours, but actual durations may deviate due to uncertainty.
- Incremental job release is assumed; some jobs are available at the beginning of the scheduling horizon, while others arrive dynamically over time.
- Technician operations refer to tasks performed solely by a technician or by a technician in combination with a test machine. These operations must be performed by qualified, right-skilled technicians, respecting availability, shift schedules, and absences. Test operations can only be processed on eligible machines.
- All routing steps for a job must be executed by a single technician, with exceptions for test and rework operations in the IDG shop.
- Re-entrant flows are allowed, enabling a job to revisit machines at different stages.
- Sequence-dependent setup times are considered for test operations. A setup time of 30 minutes can be assumed.
- Operations are non-preemptive, except when technician unavailability is combined with operational delays.
- Test machines process only one job at a time and are assumed to be fully reliable.
- Buffer space between resources is assumed to be unlimited.
- Jobs have static due dates and different priority weights.

5.2 Exact Approach

This subsection describes the exact solution method used for scheduling in CMRO shops. First, the Branch and Bound and its improved form, the Branch and Cut algorithm, are described in Section 5.2.1. Next, preliminary model calibration and the evaluation of objective functions are explained in Section 5.2.2. The model implementation for the H&P shop is described in Section 5.2.3, followed by the rescheduling strategy in Section 5.2.4. Section 5.2.5 introduces the adapted MILP model for the IDG shop, followed by the dynamic resolving strategy in Section 5.2.6. Computational time analyses are discussed in Section 5.2.7, and strategies for decomposition and robustness are explained in Sections 5.2.8 and 5.2.9, respectively.

5.2.1 Branch and Bound Algorithm

The Branch and Bound (B&B) algorithm provides an effective and exact approach for solving the NP-hard problem known as the FJSSP, formulated as an MILP (Pinedo, 2016). This methodology integrates branching and bounding techniques within a tree-structured search process, enhancing computational efficiency.

The process begins with a relaxation of integrality constraints, creating a Linear Programming (LP) problem whose solution provides an initial bound (Pinedo, 2016). If the LP relaxation generates a fractional solution, branching divides the feasible region by selecting a fractional-valued decision variable and creating two subproblems: one restricting the variable below its integer floor, and the other above its integer ceiling. This recurring branching forms a search tree structure (Hillier and Lieberman, 2015).

In the bounding step, the LP relaxation at each node provides an optimal bound. Suppose this relaxation yields an infeasible solution or an objective value inferior to the best-known integer solution, the incumbent. In that case, the corresponding branch is fathomed, thereby efficiently eliminating non-promising regions of the solution space. When an LP relaxation solution is integer-feasible, it is recorded as a potential incumbent solution, updating the global bound accordingly (Hillier and Lieberman, 2015).

A more efficient variant of the B&B method is the Branch and Cut (B&C) approach. This extends the B&B with a cutting step, enhancing the bounding process by introducing additional valid inequalities, introduced as cutting planes, to tighten the LP relaxations. These cuts remove fractional solutions while retaining all feasible integer solutions. Cutting planes significantly strengthen the model formulation, thus reducing the overall search tree size (Mitchell, 2002), and generally causing improved performance relative to the exclusive use of branch-and-bound.

Figure 5 illustrates a simplified example of the B&B process, clearly demonstrating the branching decisions and bound-based fathoming. In this figure, node 1 solves the LP relaxation of the problem, generating a fractional solution, after which branching on variable x creates two child subproblems. Node 1 is fathomed because its LP bound is worse than the incumbent. Node 2 yields another fractional solution, so branching on y continues. Node 2a produces an integer feasible solution, the green node, and becomes the new incumbent, while Node 2b is infeasible. The blue nodes represent active subproblems solved by LP relaxations, the green node is the optimal integer solution found, and the gray nodes have been fathomed.

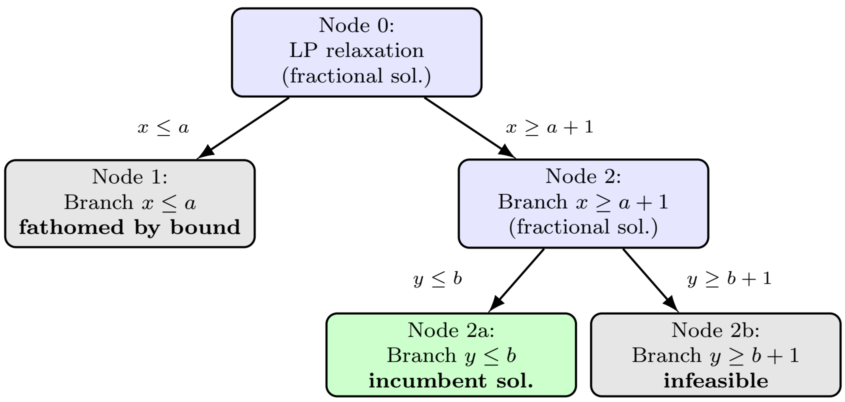

Figure 5: Illustration of a B&B search tree for a maximization MILP.

By integrating techniques available in modern MILP solvers, including branching techniques, specialized cutting-plane generation, and preprocessing routines, B&C methods efficiently navigate the extensive search space. Preprocessing procedures simplify the MILP model by eliminating redundant constraints and fixing variable bounds before search initiation. Moreover, solver-generated cuts tailored to scheduling-specific constraints, such as precedence and resource allocation constraints, provide tighter relaxations and result in exploring fewer nodes (Stecco et al., 2008).

In practice, the B&C algorithm ensures the optimality of the final schedule or accurately quantifies solution quality if prematurely terminated, offering bounds on optimality gaps. Therefore, this approach effectively balances computational efficiency with solution accuracy, making it well-suited for solving the customized JSSP for the CMRO shops. In the remainder of this paper, this algorithm is referred to as the B&B approach for consistency, as it is based on the underlying B&B framework.

5.2.2 Preliminary Model Calibration

Preliminary calibration of the proposed flexible job shop scheduling model was conducted using generated test data instances. These test instances were generated to closely resemble the actual characteristics of real-world CMRO shop data, enabling initial calibration. Specifically, the test data was used to evaluate model assumptions, modifications, and features across multiple scenarios with varying sizes and customized numbers of operations per job.

Objective choice
Multiple potential scheduling objectives were considered to ensure practical relevance and optimal performance. These included simultaneous minimization of tardiness and maximization of earliness (Equation (1)), joint minimization of tardiness and earliness (Equation (2)), and exclusive minimization of tardiness without consideration of earliness (Equation (3)). Additionally, input from consultations with shop leadership was considered to maintain alignment with operational objectives.

$$\min \sum_{i \in \mathcal{J}} w_i(\alpha T_i - \beta E_i) \quad (1) \qquad \min \sum_{i \in \mathcal{J}} w_i(\alpha T_i + \beta E_i) \quad (2) \qquad \min \sum_{i \in \mathcal{J}} w_i T_i \quad (3)$$

In these formulations, $\mathcal{J}$ represents the set of jobs that need to be scheduled, with each job $i \in \mathcal{J}$ assigned a priority weight w_i. The variable T_i represents the tardiness of job i, defined as the time by which its completion exceeds the assigned due date, while E_i indicates the earliness, or how far in advance a job is completed before its due date. The parameters α and β are coefficients used to emphasize the relative importance of tardiness and earliness, respectively.

Parameters
To evaluate model performance effectively, different parameters in the objective influencing the defined KPIs, TAT, and OTD, were given varying values. Given the operational context of the CMRO environment, multiple generated instances were input into the model to validate the performance under different parameter configurations. An essential part of this calibration phase involved optimizing the relative weighting of tardiness compared to earliness in the objective function. The weight parameters for tardiness and earliness are shown as α and β in Equation (1), respectively. While the minimization of tardiness was the primary focus, the maximization of earliness was also of significant importance. Completing jobs ahead of their due dates is preferable, particularly for assignments associated with high-priority customers.

5.2.3 H&P Shop

There exist many JSSP variants in the literature, each designed for different operational complexities and constraints. In this study, the FJSSP is selected as the foundational framework, primarily due to the variability in routing paths observed across similar jobs during different phases of maintenance processes. The flexibility in assigning operations to multiple alternative resources aligns well with the dynamic routing requirements fundamental to CMRO environments. Expanding upon the foundational framework of the FJSSP, specific features have been added or excluded to accurately reflect the characteristics and operational constraints of the CMRO shops. These decisions are driven by the analysis of the CMRO shop operations provided in Section 2.

Adjustments Traditional FJSSP
Traditional formulations of the FJSSP have predominantly focused on machines as the primary used resources, reflecting the manufacturing nature of many scheduling environments. In contrast, CMRO operations introduce a broader and more complex set of resource dependencies. In the CMRO context, operations are mostly carried out by human technicians, specialized machine operators, or with the use of specialized test benches, and can require the simultaneous use of one or two of these resource types.

A significant example is found in the scheduling of test operations within the H&P shop. These operations must be jointly assigned to a human technician and a specialized test machine. Consequently, the scheduling model must ensure that both resources, the human technician and machine, are available simultaneously, an extension beyond the standard FJSSP framework, which traditionally assumes one-to-one mappings between operations and resources. The consideration of technicians and test machines introduces an additional layer of complexity essential to reflect real-world CMRO shop environments. In addition to the multi-resource allocations, the proposed model introduces modified precedence constraints tailored to the operational logic of the shops.

Mathematical Model
The scheduling model developed for the H&P shops is formulated as an MILP model, specifically designed as a variant of the FJSSP. The foundation for this model was inspired by the dedicated technician scheduling approach proposed in Aribowo et al. (2020).

Table 3 presents an overview of all sets, indices, and parameters used in the mathematical formulation. The set $\mathcal{T}_{i,o}$, representing the qualified technicians for operation (i, o), was constructed based on historical technician assignments. Technicians with at least 25 hours of experience on the part number handled in job i were included, in agreement with shop management. For final inspection operations, this set only includes certified technicians qualified to perform the inspection tasks. The processing time of each operation, $p_{i,o}$, corresponds to the initial planned number of hours, as derived from the provided shop data. Additionally, an overview of the decision variables used to allocate technicians and test machines, sequence operations, manage technician breaks, and calculate operation timings is given in Table 3.

Symbol	**Description**
Sets and Indices	
$\mathcal{J}$	Jobs (indexed by i), including dummy job 0
$\mathcal{O}$	All operations (indexed by o)
$\mathcal{O}^{\text{tech}}$	Technician-only operations
$\mathcal{O}^{\text{test}}$	Test operations requiring test machines
$\mathcal{O}^{\text{final inspection}}$	Final inspection operations
O_i	Operations for job i
$\mathcal{T}$	Technicians
$\mathcal{M}$	Test machines
$\mathcal{T}_{i,o}$	Qualified technicians for operation (i, o)
$\mathcal{M}_{i,o}$	Compatible machines for test operation (i, o)
$\mathcal{B}_t$	Break intervals for technician t
Parameters	
$p_{i,o}$	Processing time of operation (i, o)
D_i	Due date of job i
w_i	Priority weight for job i
α, β, γ	Objective coefficients
$\text{st}_{i,j}$	Setup time when switching from job i to j on a test machine
$b_{\text{start}},\ b_{\text{end}}$	Starting and end time of break b
l_i	Index of final operation for job i
M	Large big-M constant for linearization
Decision Variables	
$X_t^{i,o} \in \{0,1\}$	1 if technician t is assigned to operation (i, o); 0 otherwise
$X_m^{i,o,j,k} \in \{0,1\}$	1 if test operation (i, o) immediately precedes (j, k) on machine m; 0 otherwise
$Y_t^{i,o,j,k} \in \{0,1\}$	1 if operation (i, o) precedes (j, k) on technician t; 0 otherwise
$W_t^{i,o,b} \in \{0,1\}$	1 if operation (i, o) finishes before break b of technician t; 0 otherwise
$S_{i,o} \in \mathbb{Z}^+$	Start time of operation (i, o)
$C_{i,o} \in \mathbb{Z}^+$	Completion time of operation (i, o)
$T_i \in \mathbb{Z}^+$	Tardiness of job i
$E_i \in \mathbb{Z}^+$	Earliness of job i

Table 3: Nomenclature for scheduling model components.

Using only the weighted tardiness and earliness objective results in the scheduling model prioritizing exclusively the final operation of each job. This approach, however, may result in scheduling inefficiencies, as earlier operations could start sooner but do not, since the optimization model solely targets the tardiness or earliness of the final operation. To handle this issue, the model introduces an additional, less weighted, objective, formulated in Equation 4, which minimizes the total completion time of all operations.

This adjustment stimulates the earlier starting and completion of intermediate operations whenever feasible. Additionally, it enhances scheduling robustness, as starting operations early can buffer against potential delays and optimizes technician utilization by ensuring that resources can start directly on their next operation when available.

$$\sum_{i \in \mathcal{J}} \sum_{o \in O_i} C_{i,o} \tag{4}$$

An illustrative example of this scenario is presented in Figure 6, where each operation is labeled as (i, j). Without the additional objective, as shown in Figure 6(a), operations 2,1; 2,2; and 4,1 could start earlier but remain delayed due to the optimization focus solely on the tardiness or earliness of the final operation. In contrast, Figure 6(b) demonstrates the result of incorporating the operation-based completion objective, where the schedule effectively advances these operations, illustrating the improvement in scheduling efficiency by ensuring the earliest feasible start times for each operation.

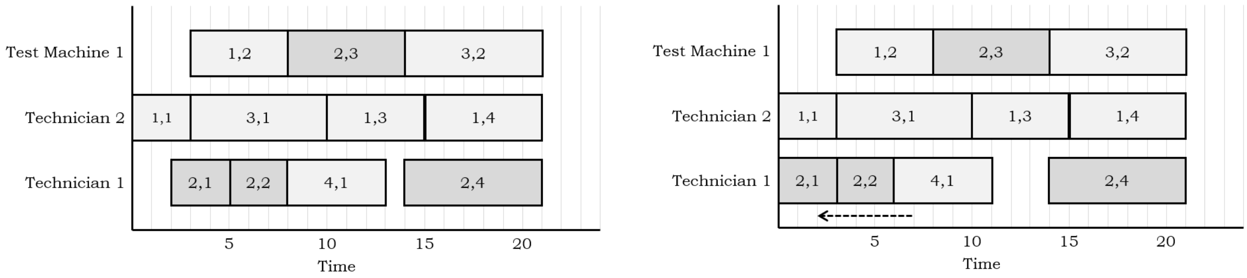

(a) Without the extra objective

(b) With the extra objective

Figure 6: Illustrative example of scheduling outcomes using the extra objective.

The objective function in Equation (5) minimizes a combination of weighted tardiness, weighted earliness, and the completion times of all operations. Hence, the model aims to achieve a minimization of tardiness while concurrently maximizing earliness. By adjusting coefficients α, β, and γ, the model balances on-time performance, early finishing, and overall schedule efficiency. In this study, the value of γ is fixed at 0.1, representing a small fraction of the total weight relative to α and β, as minimizing completion times is not a primary objective. This objective is subject to the constraints outlined in Equations (6)–(23).

$$\min \sum_{i \in \mathcal{J}} w_i(\alpha T_i - \beta E_i) + \gamma \sum_{i \in \mathcal{J}} \sum_{o \in O_i} C_{i,o} \tag{5}$$

The model addresses technician and machine assignment without relying on a traditional predefined skill matrix, wherein individual constraints ensure the feasibility of each resource-operation pair. Instead, a more efficient approach is employed by explicitly defining, for each operation, a specific set of possible technicians and, for test operations, a predefined set of suitable machines. This approach was adapted from the method proposed by Perroux et al. (2024). This direct specification of resource sets significantly reduces the computational complexity by limiting the number of constraints and decision variables, reducing the search space. Consequently, the model only includes decision variables and allocation constraints relevant to qualified resources, thereby preventing unnecessary assignment variables for resources lacking the required skills or qualifications to be allocated to an operation.

Equations (6)–(10) allocate human and machine resources. Specifically, Equation (6) ensures exactly one qualified technician performs each operation, and Equation (7) ensures that every test operation is scheduled exactly once, allocated to one of the feasible test machines. The constraint in Equation (8) forces test operations to have at most one immediate predecessor, while Equation (9) ensures the dummy job serves as the initial predecessor on every machine. Moreover, Equation (10) restricts precedences to apply only when two test operations share the same machine.

$$\sum_{t \in \mathcal{T}_{i,o}} X_t^{i,o} = 1 \quad \forall i \in \mathcal{J},\ o \in O_i \tag{6}$$

$$\sum_{j \in \mathcal{J}} \sum_{k \in O_j^{\text{test}}} \sum_{m \in \mathcal{M}_{i,o} \cap \mathcal{M}_{j,k}} X_m^{i,o,j,k} = 1 \quad \forall (i,o) \in \mathcal{O}^{\text{test}} \tag{7}$$

$$\sum_{i \in \mathcal{J}} \sum_{o \in O_i^{\text{test}}} \sum_{m \in \mathcal{M}_{i,o} \cap \mathcal{M}_{j,k}} X_m^{i,o,j,k} \leq 1 \quad \forall (j \geq 1, k) \in \mathcal{O}^{\text{test}} \tag{8}$$

$$\sum_{i \in \mathcal{J}} \sum_{o \in O_i^{\text{test}}} X_m^{i,o,0,0} \leq 1 \quad \forall m \in \mathcal{M} \tag{9}$$

$$\sum_{i \in \mathcal{J}} \sum_{o \in O_i^{\text{test}}} X_m^{i,o,j,k} \leq \sum_{i \in \mathcal{J}} \sum_{o \in O_i^{\text{test}}} X_m^{j,k,i,o} \quad \forall (j \geq 1, k) \in \mathcal{O}^{\text{test}}, m \in \mathcal{M} \tag{10}$$

Another significant adaptation from the traditional FJSSP introduced in this research is the dedicated technician constraint. Unlike standard models where technician assignments can vary between operations, the CMRO environment necessitates that once a technician initiates the first operation on a given job, the same technician must continue performing all subsequent operations on that part. This ensures that technicians are leveraging familiarity with specific job details. This requirement is explicitly integrated into the scheduling model, as introduced by Aribowo et al. (2020). It is referred to as the dedicated technician principle, enforcing the assignment of the same technician across all applicable operations of a single job. However, certain operations are deliberately excluded from this dedicated technician constraint due to their specialized nature. Specifically, the final inspection step does not fall under this constraint, as it must be performed exclusively by certified inspectors rather than general technicians. Equation (11) ensures that every job enforces the dedicated technician principle, ensuring all technician operations for a job must be performed by the same individual.

$$X_t^{i,o} = X_t^{i,o'} \quad \forall i \in \mathcal{J},\ \forall o, o' \in O_i^{\text{tech}} \setminus \mathcal{O}_i^{\text{final inspection}},\ \forall t \in \mathcal{T}_{i,o} \cap \mathcal{T}_{i,o'} \tag{11}$$

Moreover, the scheduling model introduces additional considerations for consecutive operation handling. Once a technician starts work on a specific operation, they are required to perform all subsequent operations for that job consecutively, back-to-back. However, test and final inspection operations are excluded from this back-to-back scheduling requirement. These exclusions prevent operational bottlenecks inherent to these specific steps, as queuing often occurs due to the limited availability of these resources. Requiring technicians to wait during these bottlenecks would significantly reduce technician availability and efficiency. Consequently, this application is captured within the precedence constraints of the shop scheduling models. First of all, the constraint Equation (12) ties each operation's completion time to its start time plus its processing duration. For technician-only operations, Equation (13) enforces the back-to-back scheduling, except for the final inspection step, which is covered by Equation (14).

$$C_{i,o} \geq S_{i,o} + p_{i,o} \quad \forall i \in \mathcal{J},\ o \in O_i \tag{12}$$

$$S_{i,o} = C_{i,o-1} \quad \forall\, (i, o \geq 1) \in \mathcal{O}^{\text{tech}} \setminus \mathcal{O}^{\text{final inspection}} \tag{13}$$

$$S_{i,o} \geq C_{i,o-1} \quad \forall\, (i, o \geq 1) \in \mathcal{O}^{\text{final inspection}} \tag{14}$$

Sequence-dependent setup times refer to the additional preparation time required on machines when switching operations from one job to another when part numbers differ. Within this scheduling framework, sequence-dependent setup times are particularly relevant for test machine operations. Consequently, they have been incorporated within the overlap constraints for test machine resources.

Furthermore, the H&P shop model requires allocation of test operations separately to both test machines and technicians, necessitating the simultaneous allocation of setup times to these two distinct resources. Addressing this requirement, the present model integrates sequence-dependent setup times within both overlap and precedence constraints, using a customized implementation of the methodologies described by Mousakhani (2013). To correctly allocate sequence-dependent setup times to both test machines and technicians, it is crucial to determine which job immediately precedes the current one on a given test machine. This led to the introduction of the immediate predecessor variable for test operations, indicated by $X_m^{i,o,j,k}$. This variable determines the part processed directly before the current operation on the same machine, allowing the model to check whether a part change occurs. If the part number differs, the setup time is assigned not only to the machine but also to the technician performing the test. If, instead, the same precedence variables used for technician-only operations, $Y_t^{i,o,j,k}$, were applied, which indicate all operations preceding a given operation on a resource, it would not be possible to identify which exact job is processed immediately before the current test operation on the test machine. As a result, the model would be unable to determine the preceding part number, and the overlap constraints for technicians would not be able to assess whether a setup is required.

Equation (15) ensures that if a test operation follows its preceding operations directly, its start time is shifted by the appropriate setup duration. Furthermore, Equations (16)–(17) prevent overlapping allocation to the same technician by different operations. Using a big-M formulation, if operation (i, o) precedes (j, k) on the same technician, (j, k) cannot start until (i, o) completes, including an optional setup time to prepare the test machine. The setup time is determined with the use of the immediate predecessor variable for test machine operations. Similarly, Equation (18) restricts test operations on the same machine from overlapping, ensuring each test machine handles at most one test operation at a time.

$$S_{i,o} \geq C_{i,o-1} + \sum_{j \in \mathcal{J}} \sum_{k \in O_j^{\text{test}}} \sum_{m \in \mathcal{M}_{i,o} \cap \mathcal{M}_{j,k}} \text{st}_{j,i} X_m^{i,o,j,k} \quad \forall (i, o \geq 1) \in \mathcal{O}^{\text{test}} \tag{15}$$

$$S_{j,k} \geq C_{i,o} + \sum_{h \in \mathcal{J}} \sum_{z \in O_h^{\text{test}}} \sum_{m \in \mathcal{M}_{h,z} \cap \mathcal{M}_{j,k}} \text{st}_{h,j} X_m^{j,k,h,z} - M(3 - X_t^{i,o} - X_t^{j,k} - Y_t^{i,o,j,k})$$
$$\forall i \in \mathcal{J},\ o \in O_i,\ j \in \mathcal{J},\ k \in O_j,\ t \in \mathcal{T}_{i,o} \cap \mathcal{T}_{j,k} \tag{16}$$

$$S_{i,o} \geq C_{j,k} + \sum_{h \in \mathcal{J}} \sum_{z \in O_h^{\text{test}}} \sum_{m \in \mathcal{M}_{h,z} \cap \mathcal{M}_{i,o}} \text{st}_{h,i} X_m^{i,o,h,z} - M(3 - X_t^{i,o} - X_t^{j,k} - (1 - Y_t^{i,o,j,k}))$$
$$\forall i \in \mathcal{J},\ o \in O_i,\ j \in \mathcal{J},\ k \in O_j,\ t \in \mathcal{T}_{i,o} \cap \mathcal{T}_{j,k} \tag{17}$$

$$S_{j,k} \geq C_{i,o} + \sum_{m \in \mathcal{M}_{i,o} \cap \mathcal{M}_{j,k}} \text{st}_{i,j} X_m^{i,o,j,k} - M\left(1 - \sum_{m \in \mathcal{M}_{i,o} \cap \mathcal{M}_{j,k}} X_m^{i,o,j,k}\right) \quad \forall i \in \mathcal{J},\ o \in O_i^{\text{test}},\ j \in \mathcal{J},\ k \in O_j^{\text{test}} \tag{18}$$

The model incorporates technician unavailability constraints to manage operational disruptions caused by scheduled days off or other non-working shifts. For instance, if technicians work a four-day shift schedule, the fifth day is defined as an unavailability or break period. Technician availability is represented through predefined break intervals indicating periods of unavailability. To handle these intervals, additional constraints and decision variables have been introduced, determining whether operations should be scheduled fully before or after such breaks, preventing operations from intersecting. This constraint is derived from the approach used for scheduling maintenance activities on machines as described by Perroux et al. (2023), explicitly chosen to minimize its impact on model complexity while ensuring operational feasibility.

The constraints related to unavailability intervals are given in Equations (19)–(20). In particular, (19) makes sure that if an operation is allocated to a technician with a break period, it must fully finish before the break of the technician starts. Equation (20) similarly ensures that an operation does not begin until after the break ends, preventing any partial overlap with break periods.

$$C_{i,o} \leq b_{\text{start}} + M(2 - X_t^{i,o} + W_t^{i,o,b}) \quad \forall (b, b_{\text{start}}) \in \mathcal{B},\ t \in \mathcal{T}_{i,o},\ i \in \mathcal{J},\ o \in O_i^{\text{tech}} \tag{19}$$

$$S_{i,o} \geq b_{\text{end}} - M(1 - X_t^{i,o} - W_t^{i,o,b}) \quad \forall (b, b_{\text{end}}) \in \mathcal{B},\ t \in \mathcal{T}_{i,o},\ i \in \mathcal{J},\ o \in O_i^{\text{tech}} \tag{20}$$

Lastly, Equations (21)–(23) define how earliness and tardiness are calculated relative to the due date of the job. Specifically, Equation (21) determines the earliness or tardiness of all jobs by comparing the completion time of its final operation to its predetermined due date, whilst Equation (22) ensures these earliness and tardiness metrics remain nonnegative. Additionally, (23) ensures start and completion times are nonnegative.

$$T_i - E_i = C_{i,l_i} - D_i \quad \forall i \in \mathcal{J} \tag{21}$$

$$T_i, E_i \geq 0 \quad \forall i \in \mathcal{J} \tag{22}$$

$$S_{i,o}, C_{i,o} \geq 0 \quad \forall i \in \mathcal{J},\ o \in \mathcal{O}_i \tag{23}$$

To accurately linearize overlap and unavailability constraints within this scheduling model, the value of the big-M parameter must be carefully determined. An excessively large big-M can unnecessarily enlarge the feasible solution space, negatively impacting solver performance, while a too small value may exclude feasible solutions, resulting in an infeasible model. To achieve the right balance, the parameter is computed based on estimating the maximum cumulative processing time assigned to any single resource. This estimation considers every technician and test machine individually, summing the expected processing times for all operations that can potentially be assigned to the resource, with the processing time of each operation divided by the number of feasible resource assignments available. The expected maximum workload across all resources is then identified. This maximum workload is compared to the maximum end time of the technicians' unavailability periods, presumably at the end of the time horizon. Then M will be defined either as the maximum calculated workload or as the end of the planning horizon, whichever is greater, ensuring an optimal selected M.

5.2.4 Rescheduling Approach: H&P Shop

As mentioned earlier, schedules and operations in CMRO shop environments are subjected to dynamic events, with jobs potentially experiencing delays or arriving at unpredictable times. To maintain operational feasibility and optimize technician and machine utilization throughout the week, the proposed model includes a rescheduling mechanism designed for the H&P shop. This mechanism is partly derived from prior work on dynamic scheduling under uncertainty (Fuladi and Kim, 2024; Wang et al., 2017), where reoptimization has been shown to improve responsiveness and overall performance in various environments. This mechanism is triggered in response to two primary types of dynamic events: the arrival of new jobs and updates to the processing times of existing operations. Both events necessitate reconsideration of the current schedule to ensure continued optimized schedules.

When rescheduling is required at a given time t, referred to as the current time t, the model classifies all operations based on their reference to the rescheduling point. Operations completed before t, or currently in progress at t, maintain their original start times to keep feasibility. The end times for this group can either remain fixed or be extended if actual processing times exceed the initially planned durations. This group of operations is fixed in the model and excluded from re-optimization. The remaining operations, including newly introduced jobs and jobs initially scheduled to start after t, are rescheduled within the updated optimization model. Additionally, the algorithm provides an initial solution, based on the prior allocations of future operations in the previous schedule, to the resolving B&B model, enhancing efficiency by guiding the search toward feasible, high-quality solutions.

An exception to the fixing of variables of in-progress operations occurs when the extended processing time of an ongoing operation results in a conflict with the unavailability interval of a technician. In such cases, the operation is split into two sub-operations: one representing the remaining duration up to the start of the technician's unavailability, and another capturing the remaining processing time to be allocated differently. This splitting mechanism ensures that schedules remain feasible and respect human availability constraints while minimizing disruption to the ongoing task.

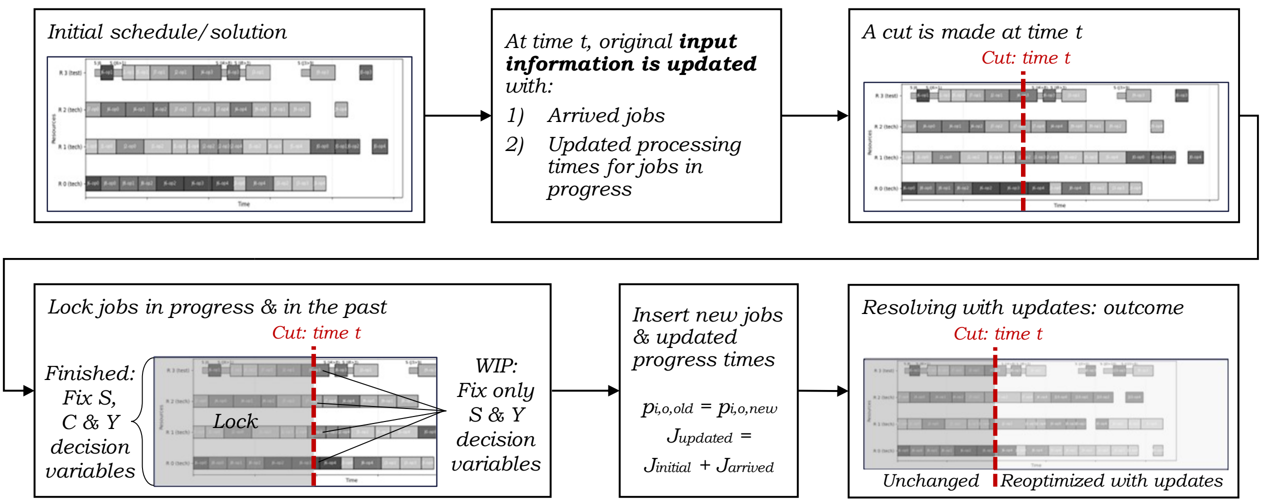

Figure 7: Overview of the proposed rescheduling approach for the dynamic FJSSP.

The rescheduling procedure proposed in this study is visualized in Figure 7. This diagram illustrates the steps to dynamically adjust the scheduling solution when operational changes occur. At each rescheduling moment, referred to as the cut, operations are classified as completed, ongoing, or not started based on their status relative to the current time. Subsequently, operations in the past are locked to maintain stability, new jobs are inserted, and updated processing times are considered to resolve the scheduling problem. Moreover, the extended version of the rescheduling algorithm for the H&P shop is outlined in Appendix A. This hybrid approach balances schedule stability with the need for dynamic responsiveness, keeping past allocations while optimizing the path forward under updated operational conditions.

5.2.5 IDG Shop

This section discusses the MILP model tailored specifically for the IDG shop. While partially similar to the previously described model for the H&P shop, the IDG shop model incorporates several substantial differences.

First, test operations within the IDG shop are exclusively assigned to test stations operated by specialized personnel. This differs from the H&P shop model, where test operations must be allocated jointly to test machines and technicians. Therefore, the immediate predecessor decision variable has been excluded, as sequence-dependent setup times apply solely to the test machine and its specialized operator, allowing technicians to undertake other jobs concurrently. Additionally, the IDG shop features two specialized rework stations for specific rework operations, namely honing and lapping, staffed by a specialized operator. These rework operators perform exclusively rework-related tasks.

Although the sets and parameters utilized in the IDG model closely resemble those of the H&P shop, except for the additional sets mentioned in Table 4, modifications have been made to the decision variables. Specifically, identical decision variables are introduced for each resource type: technicians, rework stations, and test stations, the latter two including their specialized operators, which are not regular technicians modeled as resources. An overview of the decision variables used in the IDG shop model is given in Table 5.

Symbol	**Description**
$\mathcal{O}^{\text{rework}}$	Rework operations that require handling on a rework station
$\mathcal{R}_{i,o}$	Compatible resources for operation (i,o)

Table 4: Additional sets for the IDG shop MILP model.

Symbol	**Description**
$X_r^{i,o} \in \{0,1\}$	1 if resource r performs operation (i,o); 0 otherwise
$Y_r^{i,o,j,k} \in \{0,1\}$	1 if operation (i,o) precedes (j,k) on resource r; 0 otherwise
$W_r^{i,o,b} \in \{0,1\}$	1 if operation (i,o) finishes before break b on resource r; 0 otherwise
$S_{i,o} \in \mathbb{Z}^+$	Start time of operation (i,o)
$C_{i,o} \in \mathbb{Z}^+$	Completion time of operation (i,o)
$T_i \in \mathbb{Z}^+$	Tardiness of job i
$E_i \in \mathbb{Z}^+$	Earliness of job i

Table 5: Decision variables for the IDG shop model using a resource-based notation.

The objective function for the IDG model, shown in Equation (24), is identical to that of the H&P shop model and is subject to the constraints defined in Equations (25)–(38).

$$\min \sum_{i \in \mathcal{J}} w_i(\alpha T_i - \beta E_i) + \gamma \sum_{i \in \mathcal{J}} \sum_{o \in O_i} C_{i,o} \tag{24}$$

In contrast to the H&P model, the dedicated technician constraint in the IDG shop model is further relaxed, not only for the final inspection steps but also for test and rework operations. These operations are carried out by specialized machine operators, separate from the general technician workforce. Equation (25) ensures each operation is performed by exactly one qualified technician, rework station, or test machine with an operator. Moreover, Equation (26) enforces the dedicated-technician principle.

$$\sum_{r \in \mathcal{R}_{i,o}} X_r^{i,o} = 1 \quad \forall i \in \mathcal{J},\ o \in O_i \tag{25}$$

$$X_r^{i,o} = X_r^{i,o'} \quad \forall i \in \mathcal{J},\ \forall o, o' \in O_i^{\text{tech}} \setminus \mathcal{O}_i^{\text{final inspection}},\ \forall r \in \mathcal{T}_{i,o} \cap \mathcal{T}_{i,o'} \tag{26}$$

The precedence relationships are identical to the form used in the H&P shop model, with the addition that rework operations are also excluded from back-to-back technician scheduling. Equation (27) sets each completion time at least its start time plus the processing duration. For technician-only operations, Equation (28) enforces the back-to-back scheduling, ensuring the next operation of the same job begins exactly when its predecessor completes. By contrast, excluded operations must respect Equation (29), allowing a test, rework, or final inspection to start no earlier than completion of its preceding operation, but not necessarily immediately after.

$$C_{i,o} \geq S_{i,o} + p_{i,o} \quad \forall i \in \mathcal{J},\ o \in O_i \tag{27}$$

$$S_{i,o} = C_{i,o-1} \quad \forall\, (i, o \geq 1) \in \mathcal{O}^{\text{tech}} \setminus \mathcal{O}^{\text{final inspection}} \tag{28}$$

$$S_{i,o} \geq C_{i,o-1} \quad \forall\, (i, o \geq 1) \in \mathcal{O}^{\text{test}} \cup \mathcal{O}^{\text{rework}} \cup \mathcal{O}^{\text{final inspection}} \tag{29}$$

Unlike the H&P shop, the IDG shop model treats test machine operators and test machines as a single unified resource. Therefore, setup times apply to both the operator and the machine, modeled as one resource, simultaneously, without requiring separate allocation to individual technicians. Hence, the allocation and overlap constraints for test and technician operations have been unified and identical in the IDG model, except that sequence-dependent setup times can be assigned to operations on test machines. Accordingly, the no-overlap constraints (30)–(31) prevent two technician-only and rework operations from being allocated on the same technician or rework station, whereas Equations (32)–(33) do the same for test operations sharing a test machine, including any required setup times by adding $\mathrm{st}_{i,j}$.

$$\begin{aligned} S_{j,k} \geq C_{i,o} - M\big(3 - X_r^{i,o} - X_r^{j,k} - Y_r^{i,o,j,k}\big) \\ \forall i \in \mathcal{J},\ o \in O_i^{\text{tech}} \cup \mathcal{O}_i^{\text{rework}},\ j \in \mathcal{J},\ k \in O_j^{\text{tech}} \cup \mathcal{O}_j^{\text{rework}},\ r \in \mathcal{R}_{i,o} \cap \mathcal{R}_{j,k} \end{aligned} \tag{30}$$

$$\begin{aligned} S_{i,o} \geq C_{j,k} - M\big(3 - X_r^{i,o} - X_r^{j,k} - (1 - Y_r^{i,o,j,k})\big) \\ \forall i \in \mathcal{J},\ o \in O_i^{\text{tech}} \cup \mathcal{O}_i^{\text{rework}},\ j \in \mathcal{J},\ k \in O_j^{\text{tech}} \cup \mathcal{O}_j^{\text{rework}},\ r \in \mathcal{R}_{i,o} \cap \mathcal{R}_{j,k} \end{aligned} \tag{31}$$

$$\begin{aligned} S_{j,k} \geq C_{i,o} + \mathrm{st}_{i,j} - M\big(3 - X_r^{i,o} - X_r^{j,k} - Y_r^{i,o,j,k}\big) \\ \forall i \in \mathcal{J},\ o \in O_i^{\text{test}},\ j \in \mathcal{J},\ k \in O_j^{\text{test}},\ r \in \mathcal{M}_{i,o} \cap \mathcal{M}_{j,k} \end{aligned} \tag{32}$$

$$\begin{aligned} S_{i,o} \geq C_{j,k} + \mathrm{st}_{j,i} - M\big(3 - X_r^{i,o} - X_r^{j,k} - (1 - Y_r^{i,o,j,k})\big) \\ \forall i \in \mathcal{J},\ o \in O_i^{\text{test}},\ j \in \mathcal{J},\ k \in O_j^{\text{test}},\ r \in \mathcal{M}_{i,o} \cap \mathcal{M}_{j,k} \end{aligned} \tag{33}$$

Technician breaks are handled by constraints described in Equations (34) and (35), ensuring that an operation allocated to a technician either finishes before a scheduled break or starts after that break concludes. Finally, Equations (36)–(38) define how tardiness and earliness of each job are computed relative to its due date, while also enforcing non-negativity on all relevant timing variables. These constraints remain unchanged compared to the H&P model.

$$C_{i,o} \leq b_{\text{start}} + M(2 - X_r^{i,o} + W_r^{i,o,b}) \quad \forall (b, b_{\text{start}}) \in \mathcal{B},\ r \in \mathcal{T}_{i,o},\ i \in \mathcal{J},\ o \in O_i^{\text{tech}} \tag{34}$$

$$S_{i,o} \geq b_{\text{end}} - M(1 - X_r^{i,o} - W_r^{i,o,b}) \quad \forall (b, b_{\text{end}}) \in \mathcal{B},\ r \in \mathcal{T}_{i,o},\ i \in \mathcal{J},\ o \in O_i^{\text{tech}} \tag{35}$$

$$T_i - E_i = C_{i,l_i} - D_i \quad \forall i \in \mathcal{J} \tag{36}$$

$$T_i, E_i \geq 0 \quad \forall i \in \mathcal{J} \tag{37}$$

$$S_{i,o}, C_{i,o} \geq 0 \quad \forall i \in \mathcal{J},\ o \in \mathcal{O}_i \tag{38}$$

5.2.6 Rescheduling Approach: IDG Shop

The rescheduling model proposed for the IDG shop closely resembles the algorithm previously described in Section 5.2.4. Although adjustments are implemented to align with the unique decision variables of the IDG shop MILP model, the fundamental approach remains consistent. Specifically, decision variables are fixed up to the current scheduling time, and after the current time, the schedule is re-optimized based on new information.

In the case of the H&P shop, the rescheduling mechanism was activated by two types of dynamic events: the arrival of new jobs and updates in processing times. Due to the high occurrence of test failures in the IDG shop, affecting approximately one-third of all scheduled jobs, a third dynamic event is incorporated into the rescheduling framework: the occurrence of a failed bench test. Upon encountering a test failure, the scheduling model introduces additional operations into the schedule. These include corrective repair actions, potential honing or lapping operations, and any other required routing steps, followed by a subsequent test operation to achieve a successful outcome. An illustrative example demonstrating the scheduling adjustments made in response to such a failed test event is presented in Figure 8.

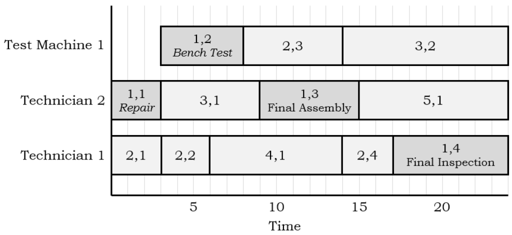

(a) Schedule before failed test

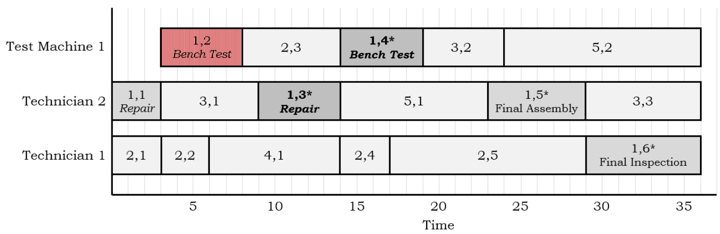

(b) Schedule after failed test

Figure 8: Illustrative example of additional routing steps inserted after a failed test in operation $(1, 2)$, with the Gantt chart shown before (a) and after (b) the test failure.

5.2.7 Computational Analysis

For practical implementation within CMRO shops, computational resources and time constraints make it infeasible to let the B&B algorithm run indefinitely until reaching full optimality. Consequently, determining a suitable computational time limit is crucial. To address this challenge, analyses were conducted for each shop to select appropriate time limits for both the initial scheduling model and subsequent rescheduling models. The analysis involves performing multiple model simulations by incrementally increasing the computational time limit and evaluating the corresponding changes in the final objective value. Additionally, operational feasibility and shop leadership input were considered for this analysis.

The results of these experiments can be visualized through convergence graphs, plotting the objective function values as a function of computational time. These graphs can demonstrate trends where substantial improvements in the objective value were initially observed, followed by a noticeable decrease in the improvement of the objective value beyond a specific computational threshold. This indicates a significant reduction in marginal gains despite constant computational effort. Based on this observed pattern, the initial time limit for the model was selected. This selection criterion effectively balances computational efficiency and solution quality, maximizing performance while minimizing resource usage. Additionally, extended analyses with increased time limits were performed to evaluate the difference in optimality, further verifying the choice of the selected time limit.

5.2.8 Time-based Decomposition

A time-based decomposition technique was implemented to address the computational complexity of solving large-scale FJSSP for CMRO environments within the Pneumatics and IDG shops. This approach reduces computational demands by removing future-scheduled jobs that fall outside a predefined scheduling window during initial planning phases, reintroducing them toward the end of the scheduling horizon. Consequently, solver efficiency is significantly improved due to the reduced solution space at each step with a smaller set of jobs, while trying to minimize the overall optimality of the final schedule (Ikli, 2022).

Initially, the scheduling model considers all available jobs at the start of the weekly simulation. Once an initial schedule is generated using the B&B algorithm, jobs whose earliest scheduled operation start times exceed a certain planning window threshold from the current time are temporarily excluded. This process reduces the number of constraints and decision variables within the mathematical model, thereby reducing the computational effort.

Throughout the simulation period, the model operates under a rolling-horizon framework, earlier explained in Section 6.2, periodically re-solving the scheduling problem in case of dynamic events such as job arrivals or deviations in job processing times. Each re-optimization applies the decomposition principle by examining job start times and selectively removing future-scheduled jobs outside the defined scheduling window.

At the end of the scheduling horizon, all previously removed jobs are reintroduced to ensure their inclusion in the schedule for the next scheduling horizon. Original job timings and operation sequences stored at their removal are restored, and a final optimization is conducted to produce an operationally feasible schedule for future time horizons. This ensures effective performance comparison with the base model in terms of objective value.

The performance of the time-based decomposition approach will be evaluated by testing different scheduling window lengths, a method also used in Stevenson et al. (2019). This approach is particularly relevant for larger-sized instances, which corresponds with the Pneumatics and IDG shops, which could have more computational challenges due to the larger number of jobs and resources involved compared to the Hydraulics shop. When demonstrating improved solver performance without negatively impacting solution quality, this decomposition approach could be expected to offer a more effective solution for managing complex and large-sized instances in scheduling scenarios within the CMRO environments.

5.2.9 Robust Scheduling

Given the operational importance of reliable schedules within CMRO environments, robustness and stability in schedules are essential. The maintenance provider prioritizes the commitment to dependable and accurate shipping dates to customers. Therefore, if a job is initially scheduled for completion on a specific date, the scheduling model must strive to maintain this original completion time, despite unforeseen disruptions. Consequently, a robust scheduling approach was proposed to improve schedule stability, reducing the impacts of dynamic events and uncertainties on promised completion times.

In alignment with the proactive-reactive framework described by Rahmani and Heydari (2014), a two-step scheduling methodology is used to ensure robustness and stability in the event of unexpected job arrivals and extended processing times of ongoing operations. Initially, a regular schedule is generated using the previously mentioned exact B&B approach, described in Section 5.2.3. Thereafter, a reactive rescheduling model is used to react to schedule disruptions. To increase stability to the re-optimized schedule, an additional objective is introduced in this reactive phase, explicitly used to minimize the deviation in job completion times relative to the initial schedule.

To incorporate robustness into the rescheduling phase, the deviation in completion times between the initial and updated schedules is quantified through an additional decision variable, representing the change in completion time for each job. This deviation is mathematically defined by Equations (39) and (40), as these constraints are included in the rescheduling model for handling disruptions.

$$\Delta_i \geq C_{i,l_i}^{new} - C_{i,l_i}^{initial} \quad \forall i \in \mathcal{J} \tag{39}$$

$$\Delta_i \geq C_{i,l_i}^{initial} - C_{i,l_i}^{new} \quad \forall i \in \mathcal{J} \tag{40}$$

In these equations, the completion times for the final operation of job i in both the initial and revised schedules are provided as inputs, respectively.

The objective function of the rescheduling model was expanded to incorporate this stability-oriented term. The revised objective function is defined in Equation 41.

$$\min \sum_{i \in \mathcal{J}} w_i(\alpha T_i - \beta E_i) + \gamma \sum_{i \in \mathcal{J}} \sum_{o \in O_i} C_{i,o}^{new} + \zeta \sum_{i \in \mathcal{J}} \Delta_i \tag{41}$$

In Equation (41), the first two terms correspond to the original weighted objective defined in Equation (5), addressing tardiness and earliness. The third term minimizes total completion time to improve schedule quality. The additional robustness term, weighted by ζ, explicitly seeks to minimize deviations from the initial schedule, enhancing stability.

The robustness weight ζ is set equal to the value of the initial scheduling objective to ensure balanced consideration between maintaining schedule stability and optimizing the primary scheduling objectives. This choice ensures that robustness is neither disproportionately favored nor neglected, allowing the rescheduling process to adaptively balance optimality and stability based on the initial schedule performance.

This robustness approach, inspired by the methodology proposed by Rahmani and Heydari (2014), has been shown to mitigate disruptions and provide stable, reliable scheduling outcomes. This integrated methodology ensures robust performance under real-world uncertainties encountered in CMRO environments by quantifying and actively minimizing completion time deviations.

5.3 Approximate Approach

To establish a comparative benchmark for the performance evaluation of the proposed exact scheduling model, an approximate heuristic approach using dispatching rules was developed and implemented. Dispatching rules are commonly used heuristic methods for solving job shop scheduling problems, mainly because of their simple logic and capacity to generate fast, effective solutions in dynamic settings (Zeiträg and Figueira, 2023).

The dispatching rule heuristic constructs the job shop schedule incrementally. At any given moment, once a machine completes its ongoing task, it selects the next operation from a set of queued operations based on a calculated priority index. This index determines the sequence in which operations are dispatched, ensuring local, immediate decision-making without requiring extensive computational resources or global optimization. Such dispatching rules, particularly dynamic variants, are advantageous for real-time adaptability to disturbances and evolving shop floor conditions.

After evaluating multiple dispatching alternatives, the Weighted Modified Due Date (WMDD) rule, derived from prior research by Kanet and Li (2004) and Melchiors et al. (2024), was selected. This dynamic rule provides the most promising results for combined tardiness and earliness objectives by incorporating the due date D_i, priority weight w_i, and current time t into its priority index calculation, as shown in Equation (42). The simulation procedure used to test the performance of the WMDD rule heuristic in a CMRO shop setting is explained in Section 6.3.

$$\text{WMDD}_i = \frac{D_i - t}{w_i} \tag{42}$$

6 Experimental Setup

This section describes the experimental framework used to assess the performance of the proposed scheduling model. Section 6.1 introduces the experimental instances, defining the generated test data and real-world datasets. Section 6.2 outlines the discrete-event simulation procedure used to simulate weekly operations and evaluate the implementation of the model for the case study. Section 6.3 describes the implementation of a dispatching rule-based heuristic used for benchmarking. The performance evaluation metrics are outlined in Section 6.4, and Section 6.5 describes the software and computational resources used.

6.1 Experimental Instances

To evaluate the performance of the scheduling model, two types of instances were used: generated datasets and realistic shop-floor data. The generated datasets, described in Section 6.1.1, were created to replicate key characteristics of CMRO shop data while performing preliminary and controlled experiments. The realistic instances, described in Section 6.1.2, were obtained from actual, historical, operational data provided by the maintenance provider to benchmark model performance under real-world conditions.

6.1.1 Generated Instances

To test, calibrate, and validate the proposed scheduling model, particularly during the initial development stages, shop-floor datasets were generated. These datasets were created to reflect the key characteristics observed in the analyzed instances of CMRO environments.

Each job within the dataset consists of a randomly determined number of operations, ranging from three to five. One of these operations was specified as a test operation, requiring processing on test machines. A randomly assigned list of feasible resources was generated for every job, representing the machines or technicians capable of performing that task. Additionally, each operation was assigned a randomly generated processing time, based on predefined intervals. A setup time was also included for test operations, within a fixed range, to simulate the realistic preparation required for such tasks compared to processing times. Generated due dates were scaled in proportion to the number of jobs and available resources, assuring that the generated instances presented a realistic and challenging level of scheduling complexity. This approach allowed for controlled experimentation with the model under conditions representative of real-world CMRO operations.

6.1.2 Realistic Instances

To test and benchmark the proposed model, realistic data from the shops, provided by the maintenance provider, was used to ensure correct implementation and measure performance.

In order to validate the performance of the scheduling model under realistic and demanding conditions, historical shop-floor data were analyzed. The historical job records were collected from datasets provided by the Enterprise Resource Planning (ERP) system of the shops, while additional data, such as technician unavailability periods, was obtained from the internal systems of the company. For each CMRO shop, all weekly instances from the past six months were explored based on the number of jobs available at the beginning of each week, the rate of jobs arriving throughout the week, and the total number of routing steps associated with those jobs. From this analysis, the most challenging weeks were identified for each shop. These weeks serve as representative high-load scenarios to test the ability of the model to generate feasible and optimal schedules under conditions of high complexity.

Table 6 provides an overview of the selected weeks, listing the number of jobs initially available on Monday and the jobs arriving during the week for each shop. Due to the relatively smaller problem sizes and lower computational load associated with the Hydraulics shop, five representative weeks were selected for this environment. In contrast, only three weeks were selected for the IDG shop, where the larger job instances and more resources experience more extensive challenges in terms of computation time.

Hydraulics shop

Week	Start date	Initial jobs	Arriving jobs
1	9/2/2024	40	31
2	9/9/2024	32	28
3	12/9/2024	42	41
4	12/16/2024	43	36
5	1/13/2025	45	45

Pneumatics shop

Week	Start date	Initial jobs	Arriving jobs
1	9/9/2024	56	46
2	9/30/2024	71	31
3	12/9/2024	52	46
4	1/13/2025	64	35

IDG shop

Week	Start date	Initial jobs	Arriving jobs
1	9/16/2024	50	55
2	9/23/2024	90	27
3	10/14/2024	73	23

Table 6: Weekly overview of initial and arriving jobs in the Hydraulics, Pneumatics, and IDG shops.

6.2 Simulation: MILP model

To evaluate the practical value of the dynamic FJSSP framework, discrete-event simulations over single working weeks were conducted for the H&P and IDG shops. The experimental simulation is set up to resemble how production control schedules and handles events on the shop floor.

Initial Week Planning

At the beginning of the simulated week, a complete job schedule is generated using the proposed B&B algorithm. This schedule includes all jobs available for planning at the start of the week, defined as those with the status RFE or RTB on Monday morning. The initial optimization is given a generous time limit, imitating the real-world setting where production control can run the algorithm to generate the schedule for the upcoming week over the weekend. The availability of technicians for the upcoming week is incorporated into the model, allowing the optimization to allocate resources effectively and take days off or other shift patterns into account.

Incorporation of Dynamic Job Arrivals

Alongside the jobs available at the beginning of the week, a second list is constructed, containing jobs that will become RTB or RFE throughout the simulation horizon. These job arrivals are pre-registered from historical data, and their release times during the week are modeled to reflect actual arrival behavior in the shop. This separation of jobs into initially available and dynamically arriving categories reproduces the operational uncertainty faced by real-world shop planning.

Execution of Simulation

Once the initial weekly schedule is computed, the simulation begins and proceeds in discrete time steps of 15 minutes. At each simulation step, the algorithm checks for two types of dynamic events: (1) the arrival of new jobs at their scheduled release time and (2) deviations in processing times of operations that complete at that time. If either of these events occurs, a rescheduling action is triggered. The model then re-optimizes the remainder of the schedule to update the most current schedule state, as described in Section 5.2.4.

This discrete rolling-horizon structure allows the simulation to reflect real-time decision-making and disruption handling on the shop floor by the proposed model. It ensures that the scheduling model needs to be adaptive, continuously responding to updated information and maintaining a feasible and efficient allocation of tasks to resources. The entire simulation process is described in Algorithm 1, which outlines the logic for integrating dynamic job arrivals and processing time deviations into the weekly scheduling cycle. The input parameters include the initial and arriving set of jobs, technician and machine data, and a simulation horizon of 40 working hours. The result is a complete, dynamically updated schedule for the selected week, providing insight into the operational performance of the proposed scheduling approach.

Algorithm 1 Weekly simulation of the H&P shop.

```
Input: initial jobs present at t=0; upcoming jobs with arrival times;
    technicians, machines; horizon H of 40 hours and time step Δt of 15 minutes
active jobs ← initial jobs
current time ← 0
schedule ← solve B&B model with (active jobs, current time)
while current time < H do
  current time ← current time + time step Δt
  rescheduling needed ← false
  // Check for arriving jobs
  for each job j in upcoming jobs do
    if arrival_time_j = current_time then
      active jobs ← active jobs ∪ {j}
      rescheduling needed ← true
    end if
  end for
  // Detect processing time overruns
  for each operation o that finishes at current time in schedule do
    if actual duration(o) > planned duration(o) then
      planned duration(o) ← actual duration(o)
      rescheduling needed ← true
    end if
  end for
  // Re-optimise if dynamic events occured
  if rescheduling needed then
    schedule ← use rescheduling B&B model with updates (schedule, active jobs, current time)
  end if
end while
Output: final schedule of the selected week
```

IDG Shop Simulation

The weekly simulation procedure implemented for the IDG shop follows the same structure as that of the H&P shop described in Algorithm 1. However, one key extension is introduced to better reflect the operational challenges observed in the IDG environment. In addition to the two dynamic event triggers used in the H&P simulation, new job arrivals and deviations in processing times, a third event type is incorporated, the occurrence of failed bench tests, as explained earlier in Section 5.2.6.

To take this into account in the simulation, the initial job data is preprocessed by removing all operations that occurred after a failed test in the historical operations, except for the final assembly and inspection. Moreover, an indication is set for the corresponding first test operation to indicate the failure of the test. During simulation, the outcome of each completed bench test operation is checked for possible failure. If a test is unsuccessful, the simulation algorithm inserts a predefined set of follow-up operations into the job routing. This mechanism is repeated for each test operation.

6.3 Simulation: Dispatching Rule Algorithm

The developed simulation to assess the performance of the dispatching heuristic under realistic operational dynamics in the H&P shop also implements a discrete-event scheduling procedure. The simulation runs iteratively in discrete time steps of 15 minutes, continuously updating the scheduling environment as jobs arrive, operations complete, or unexpected events such as overruns occur.

Initially, similar to the simulation explained in Section 6.2, a list of all jobs is computed, differentiating those that can be scheduled at the start of the simulation and those arriving later throughout the simulation.

The status of availability for each technician and machine is tracked, noting the subsequent times they become available and managing the intervals of unavailability. Jobs marked as dedicated are constrained to specific technicians throughout their operations, with exceptions made only for final inspections.

At each discrete timestep, the simulation executes the following algorithmic steps sequentially. First, in case of job arrivals, those are made available for scheduling and dispatching. After completing current operations for technicians and machines, the algorithm checks the actual hours compared to the planned hours for possible delayed operations. If differences occur, the processing and completion time of operations are adjusted accordingly, potentially triggering rescheduling of operations to avoid overlap with scheduled technician breaks.

When a technician becomes available, the algorithm considers potential operations based on the WMDD dispatching rule, explained in Section 5.3. Priority calculation occurs dynamically, considering due date, current time, and the priority weight of each operation. Operations are scheduled by selecting the highest-priority task that meets feasibility constraints, including technician compatibility, machine availability, and no intersection with break intervals.

Furthermore, the scheduling algorithm implements forced back-to-back chaining for sequential operations that require the same technician, as the operational constraints of the shop necessitate this. Operations that encounter technician unavailability intervals mid-process, due to an unexpected extended processing time, are split into two sub-operations, ensuring adherence with availability constraints and minimal operational disruption. The second operation can be separately allocated in a later stage.

The algorithm iteratively updates the job schedule, accounting for all dynamic adjustments and ensuring no overlaps between operations on shared resources or during technician breaks. Upon completing the simulation horizon, the resulting schedule is evaluated against several KPIs, including total tardiness, total earliness, and an aggregate objective function reflecting weighted tardiness and earliness, similar to the objective used in the exact approach. Furthermore, the performance of the model in processing jobs with varying priorities is evaluated.

Figure 9 illustrates the complete algorithm flow of the dispatching rule implementation used in the simulation framework. The block diagram outlines how the model progresses through time in discrete steps, repeatedly checking resource availability and dynamically allocating operations based on the WMDD priority rule.

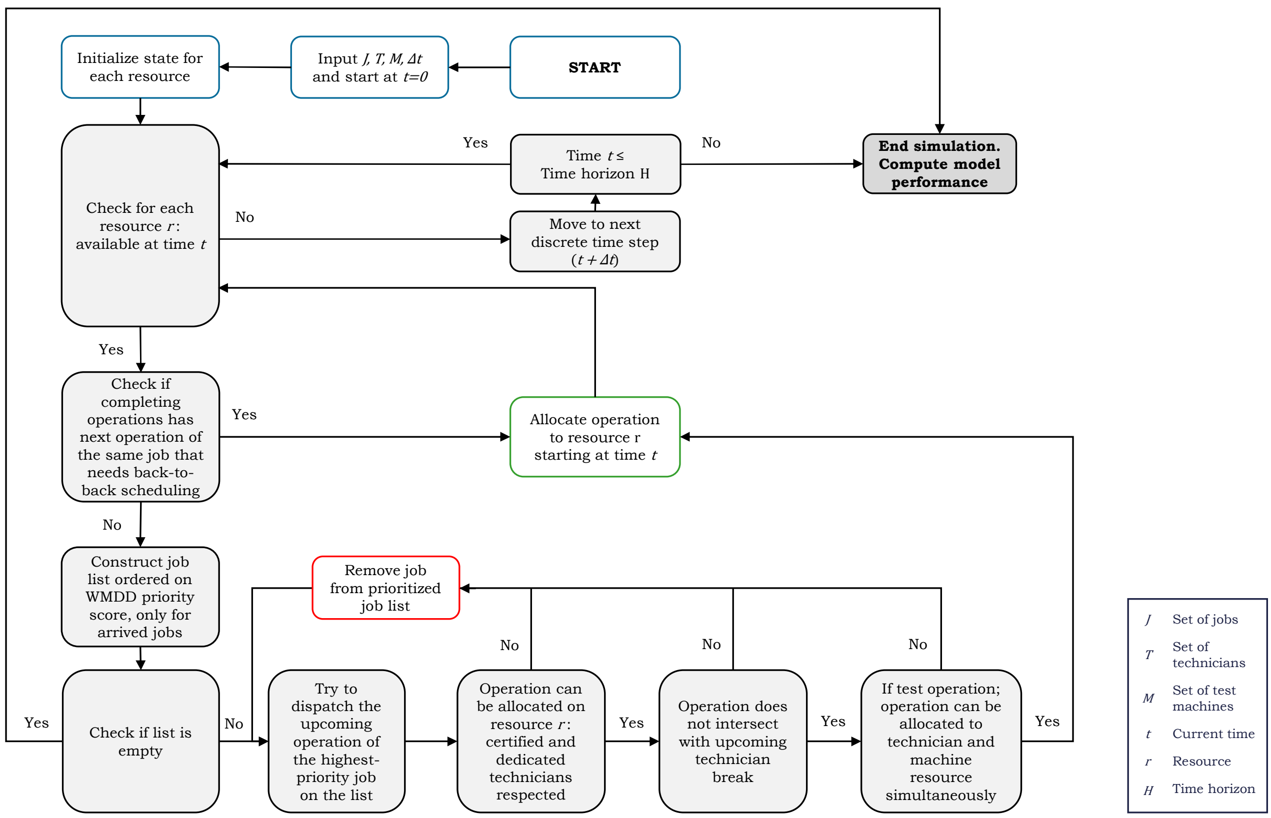

Figure 9: Block diagram of the proposed dispatching rule algorithm for the dynamic FJSSP.

This detailed simulation procedure reproduces operational practices in CMRO environments, providing a realistic benchmark to validate the effectiveness and performance of the exact scheduling model developed within this research. Moreover, this method may offer an alternative solution approach that requires less computational time compared to the exact model.

6.4 Performance Evaluation Metrics

Several key KPIs were defined to evaluate the effectiveness and practical applicability of the scheduling models developed in this study. These KPIs support in evaluation of schedule quality, operational efficiency, and comparison with real-world performance metrics. The performance indicators considered in the evaluation of schedules are described in this section.

The TAT is defined as the total elapsed time between job arrival and quotation or completion, compensated for periods of awaiting quote approval or awaiting parts, calculated with Equation (43). In this equation, C_{i,l_i} represents the completion time of the final operation of job i, A_i denotes the arrival time of job i in the shop, and T_u is the total time that the job was unworkable.

$$\text{TAT}_i = C_{i,l_i} - A_i - T_u \tag{43}$$

To indicate the difference in the proposed model performance compared to real-world shop performance, the relative decrease in TAT is calculated with Equation (44).

$$\text{TAT}_{improvement}\ [\%] = \frac{\text{TAT}_{real\ world} - \text{TAT}_{proposed\ model}}{\text{TAT}_{real\ world}} \times 100 \tag{44}$$

Next, the OTD rate is defined in Equation (45), which measures the percentage of jobs completed on or before their due date. Here, D_i is the due date of job i, C_{i,l_i} is the completion time of job i, and $|\mathcal{J}|$ represents the total number of jobs.

$$\text{OTD}\ [\%] = \frac{|\{i \in \mathcal{J} : C_{i,l_i} \leq D_i\}|}{|\mathcal{J}|} \times 100 \tag{45}$$

The formula used for defining the OTD% improvement is given in Equation (46).

$$\text{OTD}_{improvement}\ [\%] = \text{OTD}_{proposed\ model} - \text{OTD}_{real\ world} \tag{46}$$

The mean start time after arrival assesses scheduling responsiveness by taking the average duration between job arrivals and the allocation of their first operation, described in Equation 47. In this equation, $S_{i,1}$ represents the start time of the first operation of job i, and A_i denotes the arrival time in the scheduling simulation, or the time when the job becomes workable, which can occur at the start or during the scheduling horizon.

$$\text{Mean Start Time after Arrival} = \frac{1}{|\mathcal{J}|} \sum_{i \in \mathcal{J}} (S_{i,1} - A_i) \tag{47}$$

The mean waiting time, defined in Equation (48), indicates the operational efficiency by summarizing the idle durations jobs experience between arrival and processing. Here, $p_{i,o}$ represents the processing time of operation o of job i.

$$\text{Mean Waiting Time} = \frac{1}{|\mathcal{J}|} \sum_{i \in \mathcal{J}} \left((C_{i,l_i} - A_i) - \sum_{o \in O_i} p_{i,o} \right) \tag{48}$$

Additional performance metrics, such as total tardiness, earliness, mean completion time after arrival, and absolute improvement of the TAT compared to the real-world shop performance, were also included to provide extensive insights into schedule effectiveness. With these metrics, a comparison and validation against actual historical shop performance can be performed, measuring if the use of the developed scheduling model results in operational improvements in the CMRO environment explored in this case study.

6.5 Testing Set-up: Software and Hardware

All models and algorithms introduced in this study were developed in Python version 3.9.2, using both Jupyter Notebooks and Visual Studio Code as the primary development environments. Data processing and numerical operations were conducted with libraries such as NumPy, pandas, math, and datetime. For clear visualization of outcomes, both Plotly and Matplotlib were used to generate interactive and static figures, respectively. The MILP models were solved using the commercial Gurobi Optimizer, version 12.0.1. All computations were performed within a secure container environment provided by the maintenance provider, which allowed access to extended computational resources, with a maximum computation power of 2.8 GHz CPU and 27.6 GB of RAM. The data from the maintenance provider was retrieved from SQL databases connected with their ERP system and additional internal APIs, ensuring integration with operational systems.

7 Results

This section presents the results of the proposed scheduling model. Section 7.1 introduces the preliminary model calibration using generated data, where different objective functions and parameter settings are tested and validated. Section 7.2 presents the results of the case study based on historical real-world data, evaluating model performance across the three shops, with additional outcomes on robustness, decomposition, and computational efficiency.

7.1 Preliminary Model Calibration

The preliminary model calibration phase involved tuning and validating the scheduling model parameters to ensure optimal performance before testing with real-world CMRO shop data. This phase, described in Sections 7.1.1 and 7.1.2, verified the suitability of selected objective functions and parameter configurations, supporting decisions on model adjustments and practical implementation.

7.1.1 Objective Validation

The objective function used in the scheduling model was selected in collaboration with the management of the maintenance provider. However, to validate this selection and ensure the optimality of the chosen objective, an analysis was conducted using the generated CMRO shop instances explained in Section 6.1.1. The objective functions previously introduced in Section 5.2.2 were evaluated through simulation experiments. These objectives include the minimization of tardiness combined with the maximization of earliness as represented earlier by Equation (1), simultaneous minimization of both tardiness and earliness described by Equation (2), and the sole minimization of tardiness as outlined in Equation (3).

A set of 100 simulation runs was performed, each comprising 25 jobs, 8 technicians, and 4 test benches, to assess the performance differences among the three objectives. The jobs were assigned to five different priority groups, where group 5 represents the highest priority with the highest assigned weight, and group 1 represents the lowest. Table 7 summarizes the resulting comparative outcomes for each objective. These results indicate that the initially proposed objective, referred to as objective 1, consistently resulted in improved OTD percentages compared to objectives 2 and 3. Furthermore, objective 1 resulted in a significantly lower average TAT compared to the other two objectives, specifically for high-priority job groups. Conversely, the objectives minimizing or ignoring earliness, objectives 2 and 3, led to comparatively lower mean TATs for higher-priority jobs. When comparing the performance using objective 3 against objective 1, the TAT and OTD percentages for lower-priority jobs are more favorable than those for high-priority jobs. These observations emphasize the crucial role of maximizing earliness in improving service levels for high-priority customers, a major consideration within CMRO environments.

Priority	**Objective 1**		**Objective 2**		**Objective 3**	
Group	**OTD%**	**TAT**	**OTD%**	**TAT**	**OTD%**	**TAT**
1	51.29	70.77	56.83	93.27	63.71	65.81
2	64.78	58.97	66.30	92.60	70.58	66.19
3	64.27	48.80	62.80	89.21	66.83	62.99
4	66.00	43.33	65.82	90.10	67.06	64.19
5	67.03	40.15	65.66	90.44	67.84	63.33
All	**62.69**	**52.35**	**63.21**	**91.02**	**67.01**	**64.43**

Table 7: Performance comparison for the mean TAT and OTD across three objectives by priority level.

In addition, an analysis of the Gantt charts derived from simulation results identified another significant distinction among the objectives. For objectives 2 and 3, jobs delivered on schedule generally finished precisely at the deadline due to either penalties or ignoring earliness. In contrast, objective 1 encouraged early completion of tasks, with on-time jobs frequently completed significantly ahead of their due dates. To illustrate this observation, a detailed analysis of one representative simulation showed an average completion buffer of 15 time units before the due date for objective 1. Objectives 2 and 3 exhibited smaller buffers, averaging 0 and 1.5 time units, respectively. Given the operational principles of CMRO shops, which highlight early job completion as essential for reliability and customer satisfaction, objective 1 is a much better fit for organizational priorities and practical requirements. Indeed, management explicitly states that finishing exactly at the due date provides insufficient operational flexibility, making earlier completion preferable. Moreover, the aggregated performance across all priority groups indicates that objective 1 achieves a significantly better average TAT compared to the other evaluated objectives. Consequently, based on these findings, the remainder of this project proceeded with the use of objective 1, which balances the minimization of tardiness and the maximization of earliness.

7.1.2 Parameter Calibration

In preparation for testing the scheduling model with real-world CMRO data, a series of test runs were performed to determine an appropriate balance between the terms of tardiness and earliness in the objective function. Specifically, different weightings of the tardiness penalty parameter α, with a fixed earliness reward β, were tested to examine their effect on two key performance indicators: OTD% and TAT. Table 8 summarizes this parameter sensitivity analysis. The model was executed for each parameter configuration using 100 different instances, comparable in size and characteristics to the data described in Section 6.1.1.

The results indicate that the configuration $\alpha = 3$, $\beta = 1$ achieved a comparatively high mean OTD%, while on average having one of the lowest mean TAT across all tested values. This setting represents the most practical trade-off between the two main KPIs and was considered the most appropriate configuration to assess the performance of the model in a realistic CMRO context.

Furthermore, this parameter combination was selected based on consultation with CMRO shop stakeholders. It reflects their preference for a balanced priority on minimizing tardiness and maximizing earliness. Within CMRO operations, ensuring that high-priority jobs are delivered with a sufficient buffer before their due dates and maintaining high TAT are essential to maintaining expected service levels and customer satisfaction. As such, placing too much weight on reducing tardiness at the cost of early deliveries would conflict with operational priorities.

Parameters	**Mean OTD [%]**	**Mean TAT [t]**
$\alpha = 2,\ \beta = 1$	46.80	52.37
$\alpha = 3,\ \beta = 1$	47.96	52.39
$\alpha = 4,\ \beta = 1$	47.76	52.50
$\alpha = 5,\ \beta = 1$	47.84	52.50
$\alpha = 6,\ \beta = 1$	48.20	52.58
$\alpha = 7,\ \beta = 1$	48.56	52.65

Table 8: Comparison of OTD% and TAT performance across different parameter configurations.

7.2 Case Study

The case study uses historical CMRO shop data for the validation of the model. Section 7.2.1 presents the results for the H&P shops, followed by Section 7.2.2, which discusses the performance in the IDG shop. Resource utilization across all shops is analyzed in Section 7.2.3. Sections 7.2.6, 7.2.5, and 7.2.4 further evaluate robustness, decomposition, and computational performance, respectively. Lastly, Section 7.2.7 compares the exact model to a heuristic dispatching rule approach.

7.2.1 H&P Shop

To begin the real-world validation of the proposed scheduling model, the initial case studies were performed for the H&P shops. These environments provided a first setting to test the model on actual operational data. The purpose of this phase was to assess how well the model performed in replicating and improving the existing schedules in terms of TAT, OTD, and job prioritization. The section is divided into two parts, which present the performance outcomes for the Hydraulics and Pneumatics shops. Results include outcomes per priority group and comparisons with historical real-world shop performance to quantify improvements achieved by the model.

Hydraulics Shop

The scheduling model was first implemented for the Hydraulics shop, where performance was analyzed across five historically complex weeks. At the maintenance provider, customers are divided into six different priority groups based on their contractual agreements, which serve as the basis for job prioritization in daily operations. To implement these shop-floor prioritization policies, a set of weightings for these priority groups was required. Several weighting configurations were tested, including one based on the existing, points-based, prioritization approach. The configuration that simulated model behavior most consistently with current scheduling practices was selected as the most fitting configuration to test the model against real-world shop performance. An overview of the priority groups and their assigned weights is presented in Table 9.

Priority group	**Weight**
AAA	50
A	30
B	20
C	10
POOL	10
SHOP	8

Table 9: Assigned weights per priority group in scheduling objective.

Table 10 summarizes the performance of the model for the selected weeks. The results indicate that jobs in higher priority categories are generally scheduled to start earlier after arrival and achieve earlier completion times. Additionally, the observed mean idle time for high-priority groups such as AAA and A confirms the ability of the model to prioritize these jobs effectively. One minor exception is observed between the POOL and SHOP groups, likely due to the larger job volume in the POOL group rather than inefficiencies in prioritization.

Additionally, the mean tardiness and earliness values across job classes reflect the influence of the assigned priority weights. These values are normalized to adjust for due dates outside the scheduling horizon. In particular, higher-priority groups show more significant levels of earliness, reflecting the focus of the model on early completion for urgent jobs. The C and SHOP groups show relatively low mean tardiness, which may seem unexpected given the lower priority weights. This is mainly because many jobs in these categories have due dates set well beyond the scheduling horizon. As a result, only a small fraction of these jobs are classified as tardy, reducing the average tardiness observed for these groups.

Priority groups	**No. of jobs**	**Mean idle time [hrs]**	**Mean start time[a] [hrs]**	**Mean completion time[a] [hrs]**	**Mean tardiness[b] [hrs]**	**Mean earliness[b] [hrs]**
AAA	80	5.97	5.32	10.56	2.15	32.19
A	51	6.17	5.65	10.64	5.13	28.45
B	120	11.01	10.23	17.45	5.09	25.19
C	11	11.23	11.32	14.41	0.14	26.66
POOL	79	17.58	15.91	22.40	8.15	18.80
SHOP	32	13.34	12.29	16.04	4.89	24.73
All	**373**	**10.87**	**9.96**	**15.88**	**4.97**	**25.73**

[a] Mean times measured from the arrival time of the jobs.
[b] Mean tardiness and earliness are normalized for due dates that lie outside the scheduling horizon.

Table 10: Model performance per priority group for the Hydraulics shop.

Table 11 compares model-generated schedules and actual historical performance for the Hydraulics shop. The model consistently achieved improvements in both TAT and OTD, particularly for high-priority job groups. These improvements are mainly due to its ability to optimize resource allocation globally and react effectively to dynamic conditions. For example, gains are driven by more efficient scheduling of critical operations compared to existing practices. Specifically, final inspection tasks, which can only be executed by certified personnel, are normally postponed until the end of the workweek. The model, in contrast, tends to allocate such operations earlier in the week whenever feasible, as completing the final operation of a job earlier contributes significantly to improving the objective function. In several cases, this adjustment alone reduces TAT by multiple days.

Moreover, it should be noted that certain real-world factors are not incorporated in the model, including coffee break times, cleaning shifts, and other activities for which no data is available. Moreover, jobs already in progress at the start of the scheduling horizon are not explicitly represented in the model due to data limitations. Despite these constraints, the model can be expected to outperform real-world outcomes even when such externalities are considered, given the significant improvement rates. The most noteworthy indicator of performance improvement is the reduction in average TAT, as the impact on OTD is limited mainly by due dates falling outside the simulated scheduling horizon. Additionally, the jobs that belong to the POOL priority group do not improve as much as other groups. This is because, within this group, sometimes jobs are marked as critical, in which they receive additional priority, which is not considered in the model simulations done in this paper, as introducing too many different priority categories would have reduced the manageability and overview of the scheduling logic for this case study.

Priority groups	**No. of jobs**	**Mean absolute TAT improvement [days]**	**Mean relative TAT improvement [%]**	**Mean historical OTD [%]**	**Mean improved OTD [%]**
AAA	80	5.60	42.28	54.90	9.80
A	51	8.27	33.77	55.00	11.25
B	120	6.58	31.82	54.17	10.00
C	11	4.25	31.01	72.73	18.18
POOL	79	-0.01	-1.80	50.63	2.53
SHOP	32	6.50	16.72	56.25	0.00
All	**373**	**5.13**	**25.89**	**54.42**	**8.04**

Table 11: Comparison of model output with real-world Hydraulics shop performance.

Figure 10 shows an example of the Gantt chart generated based on a weekly simulation using historical data from the Hydraulics shop for a specific week. The chart confirms the successful implementation of key features, such as technician break periods, as can be noticed for technician 'EPO', and sequence-dependent setup times, as test operations assigned to both a test bench and a technician include the required setup time if the part differs from the predecessor part handled on the test bench.

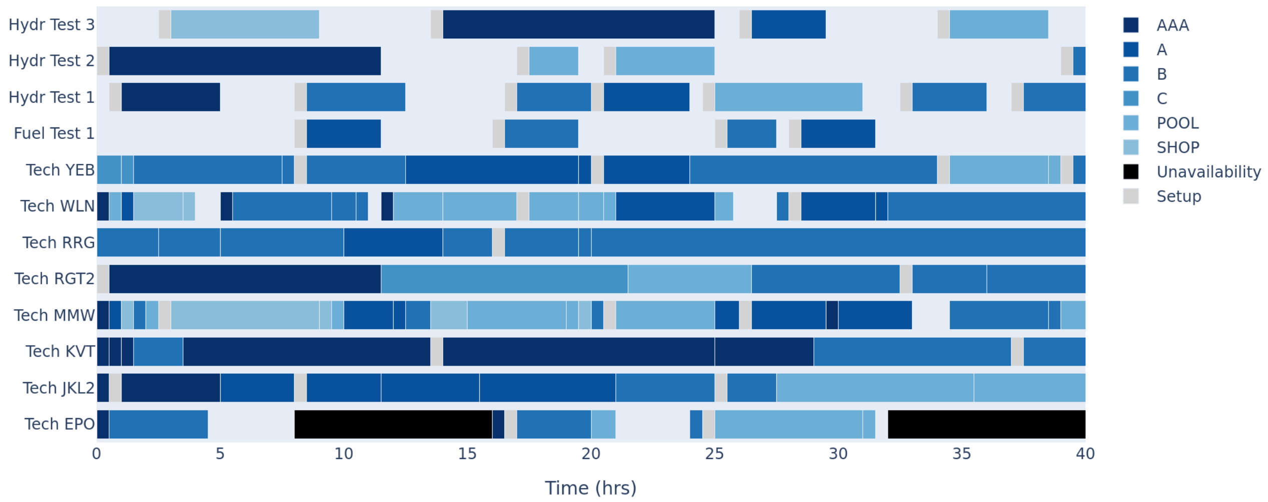

Figure 10: Gantt chart outcome of the Hydraulics shop scheduling simulation using the MILP model for the week of 9 September.

Pneumatics Shop

The model was thereafter implemented for the Pneumatics shop. Priority weights tailored to the operational performance of this shop were selected through iterative testing of multiple configurations. The final configuration used in the simulations is shown in Table 12.

Priority group	Weight
AAA	100
A	20
B	10
C	5
POOL	1
SHOP	0.5

Table 12: Assigned weights per priority group in scheduling objective.

Performance results across four representative weeks are summarized in Table 13. The model again demonstrated effective prioritization, with high-priority jobs being scheduled earlier and achieving lower completion times relative to lower-priority jobs. One significant deviation occurred in the AAA group, where performance was less favorable compared to the A group. This can be attributed to high utilization pressure on a specific single high-flow pressure test bench, which is a known bottleneck in the Pneumatics shop and is heavily used by AAA jobs. The presence of this bottleneck limited the ability of the FJSS model to improve outcomes for this group despite its high assigned weight. The normalized mean tardiness and earliness outcomes are as expected, minimizing tardiness for higher priority jobs and maximizing earliness for higher priority jobs, except for the AAA group due to operational bottlenecks.

Priority groups	No. of jobs	Mean idle time [hrs]	Mean start time[a] [hrs]	Mean completion time[a] [hrs]	Mean tardiness[b] [hrs]	Mean earliness[b] [hrs]
AAA	147	9.67	5.72	14.36	9.74	36.39
A	42	3.80	3.04	7.60	0.68	46.95
B	47	14.15	13.21	21.54	6.32	34.38
C	5	7.85	5.25	12.75	11.00	22.80
POOL	94	16.93	14.99	22.67	10.20	25.31
SHOP	50	23.83	22.38	27.59	19.05	7.48
All	**385**	**13.17**	**10.76**	**18.23**	**9.46**	**31.15**

[a] Mean times measured from the arrival time of the jobs.
[b] Mean tardiness and earliness are normalized for due dates that lie outside the scheduling horizon.

Table 13: Model performance per priority group for the Pneumatics shop.

Table 14 shows again the realized improvements in scheduling performance based on the improvement of the selected KPIs. Especially looking at the AAA group of jobs, the TAT is significantly improved because the model substantially increases the ability to handle the machine capacity problems for this group compared to manual shop planning, and even realizes a slight average increase in OTD percentage. Additionally, each priority group of jobs had decreased TATs compared to historical data, and the overall OTD percentage improved. As a result of more efficient allocation of higher priority jobs, SHOP can be handled earlier in the schedules compared to historical allocation times, which explains the significant improvement of absolute TAT.

Priority groups	No. of jobs	Mean absolute TAT improvement [days]	Mean relative TAT improvement [%]	Mean historical OTD [%]	Mean improved OTD [%]
AAA	147	3.26	28.01	58.94	1.32
A	42	3.63	59.60	85.19	0.00
B	47	5.31	26.69	63.83	6.38
C	5	3.56	23.83	50.00	0.00
POOL	94	2.06	9.80	60.47	1.16
SHOP	50	21.88	33.35	17.31	0.00
All	**385**	**5.68**	**27.49**	**57.87**	**1.52**

Table 14: Comparison of model output with real-world Pneumatics shop performance.

7.2.2 IDG Shop

The third application of the proposed FJSS model was implemented within the IDG shop. For these simulations, priority weights identical to those previously used in the Hydraulics shop were applied, as these were considered most suited based on the performance in this operational environment. These priority weights are given in Table 9.

The proposed model effectively handled the unique characteristics of the IDG shop. Especially, the single allocation of test operations due to the independent test bench operators and the capability of dynamically inserting additional operations following a failed test were successfully incorporated. This adaptability demonstrates the ability of the proposed FJSS model to be adjusted to diverse dynamic operational constraints typical for CMRO environments.

Table 15 summarizes the performance metric outcomes, derived from the three historical weekly simulations for the IDG shop. The results again confirm the expected differentiation in job handling according to priority levels. Specifically, jobs within the highest priority group, AAA, experienced the shortest start time after becoming available, averaging approximately 3 hours from arrival to allocation. Contrarily, lower-priority groups, such as A and B, had later start times after arrival, averaging roughly 11 and 19 hours, respectively. Notably, the SHOP priority group showed relatively efficient scheduling, primarily due to their smaller number of operations, shorter processing durations, and absence of test bench dependencies, which are identified as primary bottlenecks in the operational workflow. This effect is also reflected in the mean tardiness and earliness, although the low mean tardiness can be attributed partly to the small number of tardy jobs for the SHOP group, two in total.

Priority groups	No. of jobs	Mean idle time [hrs]	Mean start time[a] [hrs]	Mean completion time[a] [hrs]	Mean tardiness[b] [hrs]	Mean earliness[b] [hrs]
AAA	56	2.89	2.23	12.67	1.29	65.81
A	70	11.13	6.29	18.94	2.26	49.73
B	20	19.49	14.41	32.63	4.69	31.79
C	4	29.81	11.31	45.00	28.82	18.78
POOL	31	40.04	32.80	50.24	26.76	19.22
SHOP	115	10.83	9.63	12.53	0.46	51.77
All	**296**	**13.30**	**10.21**	**19.82**	**5.40**	**47.51**

[a] Mean times measured from the arrival time of the jobs.
[b] Mean tardiness and earliness are normalized for due dates that lie outside the scheduling horizon.

Table 15: Model performance per priority group for the IDG shop.

Furthermore, Table 16 compares the simulation results obtained using the proposed model and the actual historical performance of the IDG shop. The simulations indicated a substantial improvement in the performance metrics of the shop, particularly in terms of TAT improvements across all priority groups. Specifically, jobs in the SHOP and AAA priority categories showed impressive TAT improvements, with absolute reductions averaging over seven and nearly six days, respectively. Overall, the application of the FJSS model led to an average TAT improvement of approximately 34 percent, the highest overall TAT improvement compared to other shops. In addition, OTD rates displayed notable improvements, particularly for the A and SHOP groups, supporting the effectiveness of the proposed scheduling approach.

Priority groups	No. of jobs	Mean absolute TAT improvement [days]	Mean relative TAT improvement [%]	Mean historical OTD [%]	Mean OTD improvement [%]
AAA	56	5.78	37.68	85.71	1.79
A	70	5.26	30.53	77.14	10.00
B	20	4.41	12.23	70.00	5.00
C	4	4.68	24.84	75.00	25.00
POOL	31	3.05	13.55	67.74	0.00
SHOP	115	7.32	44.48	47.83	50.43
All	**296**	**5.86**	**34.21**	**65.88**	**22.97**

Table 16: Comparison of model output with real-world IDG shop performance.

7.2.3 Resource Utilization

To assess capacity pressure and indicate potential bottlenecks in the CMRO shops, average resource utilization was evaluated for each resource type across the Hydraulics, Pneumatics, and IDG shops. Table 17 presents the mean utilization percentages for technicians, test machines, and rework stations, computed from actual allocated and available hours over the simulated historical weeks for each shop.

Shop	Technician [%]	Test Machines [%]	Rework Stations [%]
Hydraulics	97.19	60.19	—
Pneumatics	98.00	23.17	—
IDG	94.73	97.40	72.50

Table 17: Mean resource utilization percentages by type across all shops.

The results indicate that technician resources are highly utilized in all shops, with values exceeding 94%, indicating that technician availability is a limiting factor in the scheduling process, which emphasizes the importance of optimal allocation. In comparison, test machine utilization differs significantly across the evaluated shops. The Pneumatics shop exhibits a low average test utilization of 23.17%, which could indicate that test resources are not a limiting factor in terms of capacity in this environment. Contrarily, in the IDG shop, both technicians and test machines are highly utilized, indicating tight capacity across both resource types. Test resource utilization in the Hydraulics shop is at 60.19%, falling between the low levels observed in Pneumatics and the high levels in the IDG shop. Technician allocation remains the primary challenge, as reflected by the consistently high weekly utilization rates.

7.2.4 Computational Analysis

A computational analysis evaluated the impact of different time limits on solution quality for the initial schedule and subsequent rescheduling steps. Historical weekly data instances were used to simulate and compare the performance under various computational time constraints. Specifically, the Hydraulics and Pneumatics shops were initially evaluated with a time limit of 30 minutes for generating the initial schedule and 3 minutes for each subsequent rescheduling operation. This initial time limit was based on the convergence graph, indicating that the objective value does not increase significantly after 30 minutes. The IDG shop required higher base limits due to the increased size of its instances based on the number of jobs and resources. In each case, extended time limits were tested to evaluate the extent to which additional computational effort improves the final objective value.

Table 18 presents an overview of the tested computational time limits for each shop, together with the corresponding mean improvements in the objective value and the MIP gaps indicated by the used MILP solver. The MIP gap represents the relative difference between the best lower bound found by the solver and the final objective value, where the lower bound is not necessarily optimal. The mean objective value improvements are computed as the relative difference in the objective value between the model using the base time limits and the model using extended time limits, comparing the objective at the start of the week ($t = 0$) with the final objective at the end of the simulation ($t = 40$), following multiple rescheduling operations. These minimal gains in objective value suggest that the original time limits might be appropriate and efficient when looking at the trade-off between optimality and computational effort.

Shop	Time Limit (Initial / Rescheduling)	Mean Objective Value Improvement [%]	Mean MIP Gap [%]
Hydraulics	30 min / 3 min[a]	—	1.90%
Hydraulics	1 hr / 6 min	0.30%	1.84%
Hydraulics	3 hr / 9 min	0.38%	1.77%
Pneumatics	30 min / 3 min[a]	—	1.43%
Pneumatics	1 hr / 6 min	0.58%	1.04%
Pneumatics	3 hr / 9 min	0.64%	1.10%
IDG	1 hr / 6 min[a]	—	3.06%
IDG	3 hr / 9 min	0.70%	2.54%

[a] Baseline time limits.

Table 18: Impact of computational time limits on the objective value improvement and the MIP gap

In addition to the limited improvement in objective values, the reported MIP gaps in Table 18 remain relatively small and indicate only slight reductions when increasing the time limits. Across the weekly simulations, the number of triggered rescheduling moments generally ranged between 30 and 40 times per week, depending on the number of new jobs, deviations in processing times, and failed tests.

7.2.5 Decomposition

The decomposition methodology described in Section 5.2.8 was applied to the large-scale instances of both the Pneumatics and IDG shops. The aim was to evaluate whether the proposed decomposition approach could effectively improve model performance, particularly for these larger-sized instances. To compare this approach properly, the same computational time limit for the simulations was applied to both the baseline model and the decomposition approach. This directly assesses how much the solution can improve under identical experimental conditions.

Figure 11 presents the results obtained from applying the time-based decomposition approach to two selected weeks with the highest number of starting jobs in the IDG shop, specifically the weeks of 23 September and 14 October. Figure 11(a) illustrates the number of jobs temporarily excluded from consideration due to their start times falling outside the specified decomposition window. These jobs were subsequently reintroduced into the schedule at a later stage. As expected, shorter window lengths increased the number of temporarily excluded jobs. Smaller windows reduce the size of the scheduling problem by limiting the number of jobs considered at each rescheduling step. While this reduction in complexity can potentially decrease computational effort, it also introduces the risk of suboptimal job allocation decisions due to incomplete information.

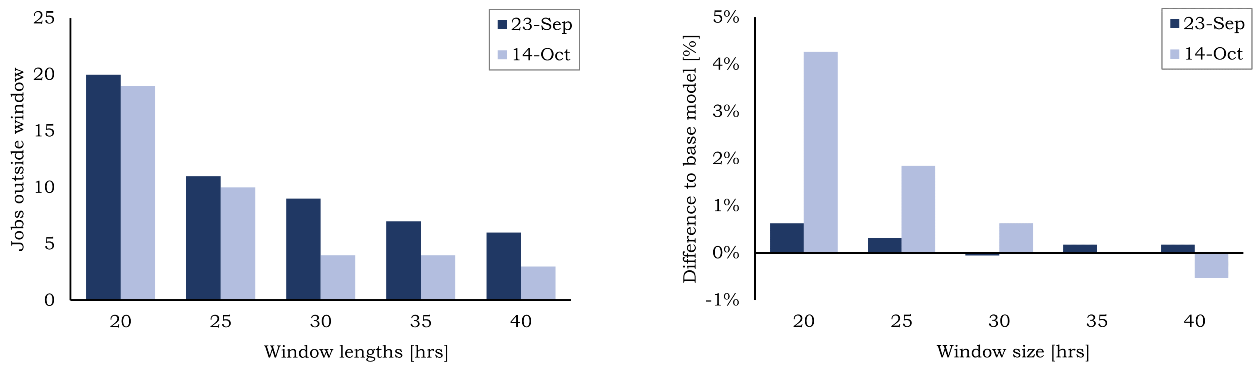

(a) Number of jobs that fell outside of the window, were removed and reinserted in the job set.

(b) Relative difference in objective value to base model without decomposition.

Figure 11: Time-based decomposition performance for the IDG shop using different window lengths

Figure 11(b) compares objective values obtained with varying window lengths relative to the baseline scenario without decomposition to estimate the optimal balance between computational efficiency and scheduling quality. In this figure, a positive difference in objective value indicates that the obtained objective, measured at the end of the simulation, is higher than in the baseline scenario, meaning the result is relatively worse. Conversely, a negative difference reflects a lower objective value and therefore an improvement over the baseline. The results show that shorter windows substantially worsen the performance measured by the objective value.

Contrarily, longer windows offer minor improvements over the baseline scenario, suggesting limited value of the time-based decomposition method in this specific scheduling context. This can be attributed to the relatively moderate workload in the shop, where few jobs are scheduled far enough into the future for decomposition to significantly reduce complexity. As a result, longer windows do not significantly unburden the problem or increase solution quality enough. Additionally, when using very short windows, jobs arriving early in the week may be excluded prematurely, which can lead to suboptimal decisions as these jobs remain unconsidered for the rest of the planning horizon. This limits the benefit of multiple window lengths, explaining the small improvements observed in this case study.

A similar pattern was observed when the decomposition approach was applied to the Pneumatics shop. The experiments generated comparable outcomes, suggesting that the benefits of the decomposition approach may be marginal under the current conditions of these specific CMRO shop environments.

7.2.6 Robustness and Stability

To evaluate the robustness and stability measures implemented for the proposed scheduling approach, an analysis was performed for both the Hydraulics and Pneumatics shops. The primary objective of implementing these measures, which penalize deviations in job completion times during reactive rescheduling, was to deliver more reliable completion estimates at the beginning of each scheduling week, thereby improving the accuracy of communicated delivery dates.

The analysis started with the Hydraulics shop. When robustness measures were not applied, the base model showed a mean deviation in job completion times of 6.31 hours between the initial schedule on Monday and the final realized schedule at the end of the scheduling horizon. This Mean Completion Time Deviation (MCTD), formulated in Equation (49), is defined as the average absolute difference between the initially scheduled and finally realized completion time for the last operation of each job. Specifically, C_{i,l_i} represents the initially scheduled completion time of the final operation of job i, while $\widehat{C}_{i,l_i}$ denotes its realized completion time at the end of the simulation. The MCTD decreased to 2.53 after incorporating robustness into the objective. The reduction is explained by the fact that the robust model explicitly minimizes these deviations in completion times during rescheduling. As a result, an improvement in schedule stability was observed. Specifically, the proportion of jobs that failed to meet their initially promised shipping dates by the end of the week decreased substantially from 15% in the base model to just 3% in the robust model. This confirms the effectiveness of incorporating robustness in improving schedule reliability.

$$\text{MCTD} = \frac{1}{|\mathcal{J}|} \sum_{i \in \mathcal{J}} \left| C_{i,l_i} - \widehat{C}_{i,l_i} \right| \tag{49}$$

However, introducing robustness had implications for job prioritization and flexibility within the scheduling process. Figure 12 illustrates these trade-offs by comparing the robust model to the base model. Incorporating robustness led to reduced flexibility in scheduling new arrivals. Mainly because the model explicitly minimizes deviations in completion times for jobs already allocated in the schedule earlier in the week, which limits the ability to reallocate these jobs to later time slots. As a result, the ability to reallocate scheduled jobs to create space for newly arriving high-priority jobs is limited. Consequently, high-priority job performance was affected, with increased mean completion times. Hence, lower-priority jobs experienced improvements in their performance metrics, as they were less frequently postponed by high-priority tasks. Figure 12(b) indicates the related decrease in TAT improvements for high-priority jobs, mainly noticeable in priority groups A and B.

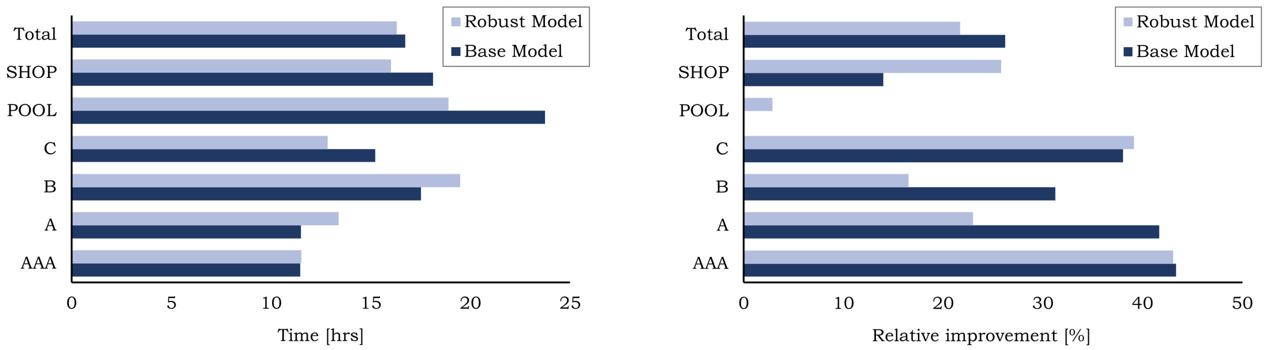

(a) Mean completion time after arrival.

(b) Mean TAT improvement

Figure 12: Hydraulics shop performance differences in: (a) mean completion time after arrival and (b) mean TAT improvement of the robust model compared to the base model.

The robustness approach was similarly implemented for the Pneumatics shop, generating comparable improvements in scheduling reliability. Specifically, the percentage of missed promised delivery dates decreased from 8% in the base model to just 2% with the robust model. Comparable with the Hydraulics shop, performance metrics for different priority groups were similarly affected. Figures 13(a) and 13(b) show that while robustness enhanced scheduling reliability, it negatively impacted mean completion times and TAT improvements for higher-priority jobs. Nevertheless, this trade-off resulted in more predictable and dependable schedules, enhancing customer satisfaction by significantly reducing uncertainty in communicated delivery dates.

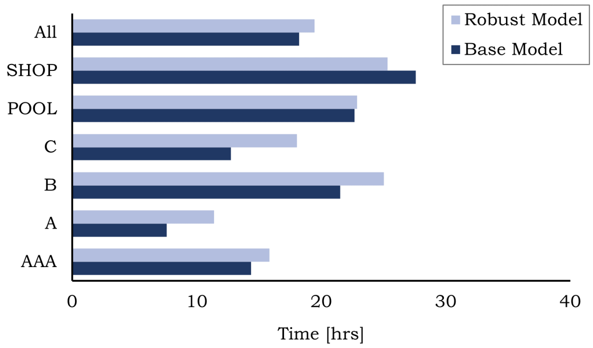

(a) Mean completion time after arrival.

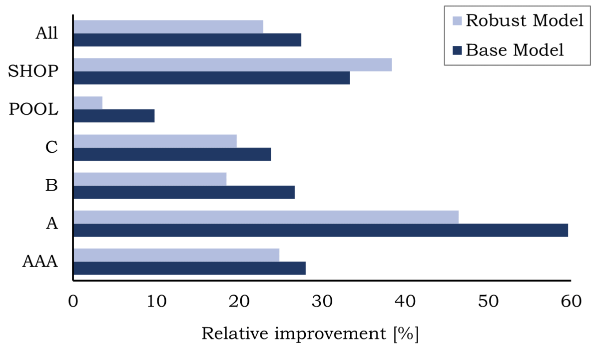

(b) Mean TAT improvement

Figure 13: Pneumatics shop performance differences in: (a) mean completion time after arrival and (b) mean TAT improvement of the robust model compared to the base model.

7.2.7 Dispatching Rule Algorithm

The effectiveness of the heuristic dispatching rule algorithm, specifically the WMDD rule described in Section 5.3, was assessed by comparing its performance against the exact scheduling approach previously presented. Figure 14 presents the relative performance differences observed for each week in the Hydraulics and Pneumatics shops. Positive differences indicate less favorable outcomes compared to the exact model, as higher values in the used metrics, such as the objective value and number of tardy jobs, reflect reduced performance. Therefore, results indicate a significantly reduced overall performance when employing the WMDD heuristic in comparison to the exact approach, which is mainly noticeable in terms of the objective value and the number of tardy jobs. Despite its lower performance, the heuristic approach functions correctly and respects all operational constraints, including those related to technician availability, precedence relations, and setup times, as can be verified in Figure 15(a), which presents a feasible weekly schedule generated for the Hydraulics shop.

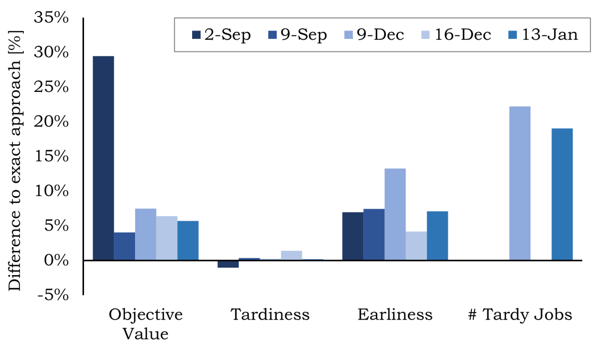

(a) Hydraulics shop.

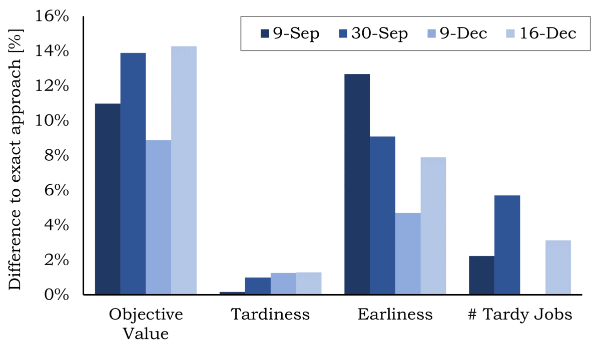

(b) Pneumatics shop.

Figure 14: Relative difference in performance of the DR approach to the exact approach for: the (a) Hydraulics and (b) Pneumatics shop.

Although the total earliness was consistently better using the exact model, the differences in total tardiness between the approximate and exact methods were comparatively smaller. This outcome is primarily due to the prioritization logic of the WMDD rule, which first allocates jobs already overdue before scheduling jobs with future due dates. An exceptional case was observed in the simulation for the week of 2 September, where total tardiness slightly improved under the heuristic approach despite a significant degradation in overall objective value measured. A detailed analysis for this specific week, presented in Table 19, indicated significant delayed starting times of high-priority jobs when using the heuristic method compared to the exact model. This is also the main reason for the considerable difference in objective value. This week, many low-priority jobs were already overdue, whereas many high-priority jobs still had remaining time before their due dates. The DR method prioritizes tardy jobs before considering earliness and allocates resources to delayed low-priority jobs before almost overdue high-priority ones, leading to a significant worsening of the objective value. In contrast, the exact model incorporates a more balanced consideration of job priority and timing, as CMRO shops prefer. This is visually confirmed in Figure 15, where high-priority jobs are scheduled earlier in the exact model's Gantt chart compared to the approximate approach. While this indicates that the heuristic may sometimes identify locally promising outcomes in terms of specific objectives such as tardiness, it also highlights its limitations in prioritizing important tasks, specifically high-priority jobs, within a globally optimized scheduling framework.

Priority Group	Exact model	Approximate model
AAA	5.65	12.91
A	5.50	21.88
B	10.90	15.11
C	13.80	26.05
SHOP	4.67	4.00
POOL	16.15	16.92

Table 19: Comparison of mean start times after job arrival between the exact and approximate model for each priority group, for the week of 2 September in the Hydraulics shop.

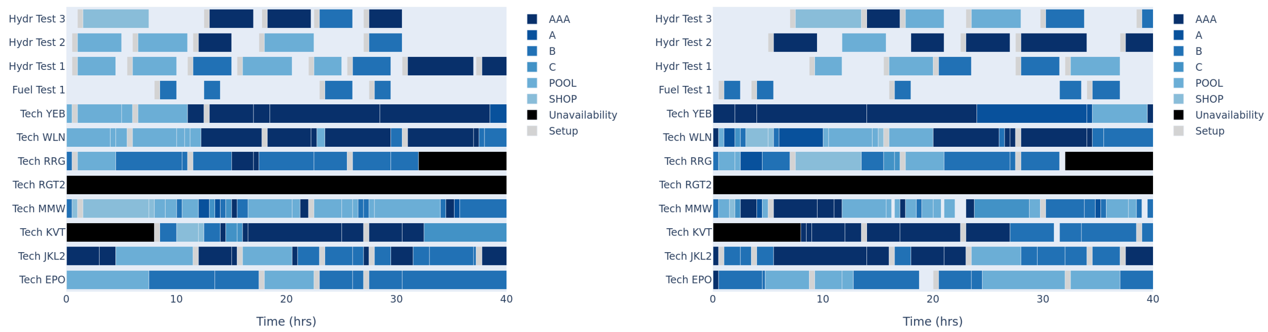

(a) Approximate approach

(b) Exact approach

Figure 15: Gantt charts of the Hydraulics shop simulation using different methods for the week of 2 September.

Despite the lower performance metrics, the heuristic dispatching approach offers substantial advantages in terms of computational efficiency. As indicated in Tables 20 and 21, the computation time needed to generate a complete weekly schedule using the simulation remained below one second for the Hydraulics shop and only a few seconds for the Pneumatics shop. Moreover, the low computational effort makes this method suitable for dynamic environments where frequent rescheduling is required. The approach can easily incorporate updated job information at each dispatching point, as shown in the simulation, where jobs became available upon arrival in the shop and delays were handled by extending the completion times when needed. However, in such settings, schedule predictability may be limited, as the dispatching list is continuously updated and the entire schedule is reconstructed from the current time at each rescheduling step. As a result, the actual schedule can deviate substantially from the initially generated plan. Therefore, this heuristic method presents a feasible alternative for scenarios where computational resources are limited, while accepting reduced scheduling optimality and predictability.

Week	Run time [s]	Total jobs
2-Sep	0.86	71
9-Sep	0.78	60
9-Dec	0.86	83
16-Dec	0.94	79
13-Jan	0.74	90

Table 20: Run time and total number of jobs per weekly instance for the Hydraulics shop.

Week	Run time [s]	Total jobs
9-Sep	2.56	102
30-Sep	2.31	102
9-Dec	3.38	98
13-Jan	2.73	99

Table 21: Run time and total number of jobs per weekly instance for the Pneumatics shop.

8 Discussion

The performance of the model received highly positive feedback from the shop leads responsible for planning at the maintenance provider. They valued the high frequency and granularity of schedule updates, finding the rescheduling time suitable for operational implementation. They also emphasized that further accuracy could be achieved through improved data input, such as better reflecting actual working hours, assistance activities in the shop, and processing time variations between technicians. The shop leads underscored the advantage of allocating final inspection tasks earlier and more balanced throughout the week, as opposed to the current practice of mostly conducting inspections at the end of the week. This improvement could reduce scheduling difficulties experienced by certified personnel, workload spikes, and ensure earlier job completions, as indicated by the improved performance of the proposed model. Additionally, by planning all available jobs from the start of the scheduling horizon, the model enables schedules to extend further into the future than is currently done, thereby improving management of customer expectations and delivery commitments, as noted by shop leads. Importantly, they also clarified that periods in the schedule where technicians are not assigned to specific tasks do not necessarily indicate inefficiency, as these gaps are often used for low-priority teardown work, cleaning duties, or supervisory responsibilities.

Within the case study, the optimized scheduling implementation resulted in a mean reduction in TAT of 34.21% within the IDG shop and absolute TAT improvements of five to seven days for high-priority repairs. Additionally, OTD performance improved by 22.97%. Similarly, comparable performance gains were achieved in the H&P shops, confirming the effectiveness of the developed exact approach. These significant improvements are mainly attributed to the global optimization approach used by the exact model, which simultaneously considers all applicable capacity constraints. By comparison, historical scheduling typically focuses on the availability of a single resource at a time, without accounting for future constraints. By strategically allocating tasks based on all upcoming routing steps and available resources, the model more effectively leverages the capabilities of the shop. Although some simplifications were made and certain real-world factors were not modeled in the case study, the observed performance improvements are primarily due to more efficient scheduling. The relatively greater improvements in TAT and OTD observed in the IDG shop, compared to the other shops, can likely be attributed to bottlenecks caused by test capacity. The limited availability of test machines creates scheduling challenges, as reflected by their high utilization rates. In those environments, the advantages of a data-driven scheduling approach over manual planning are especially noticeable, explaining the more significant performance gains in the IDG shop. In comparison, the Pneumatics shop, which only indicated 1.52% OTD improvement, had the lowest machine utilization levels, suggesting that fewer capacity bottlenecks result in lower optimization advantages.

The time-based decomposition approach did not lead to the expected performance improvements in this case study. The use of shorter windows resulted in notably worse total performance, whereas longer windows provided negligible improvements over non-decomposed models. This limited benefit is primarily due to the workload and scheduling horizon in the studied CMRO shop, where only a few jobs are initially allocated far enough into the future for decomposition to reduce complexity or improve overall scheduling quality. Furthermore, shorter scheduling windows risk excluding early-week job arrivals from subsequent rescheduling, which results in suboptimal scheduling decisions. Although time-based decomposition showed limited improvement in this case, prior studies suggest it can be more effective in longer-horizon settings.

Besides the specific case study at the maintenance provider, the developed scheduling framework has broader applicability. The flexible design, integrating dynamic rescheduling, technician specialization, and priority-based job handling, makes it applicable to other shop floors experiencing similar planning challenges. The successful application of the framework to both the H&P and IDG shops demonstrates its adaptability, showing that it can be modified to fit the operational characteristics of different environments. Therefore, the methodology could be extended to other maintenance or production environments with comparable shop floor dynamics, such as varying resource skill levels, the consideration of both human operators and machines, and handling limited test capacity. These types of environments remain unexplored in existing literature, but could benefit from data-driven scheduling optimization as demonstrated in this study.

Compared to earlier studies, the time limits of 30 to 60 minutes used in this research seem reasonable for obtaining near-optimal MILP solutions. Thörnblad et al. (2015) reported MIP gaps of 1.62% on average for 40-job instances with similar complexity and increased time limits. Our model achieved lower gaps, such as 1.43% for the pneumatics shop, despite tighter limits and additional rescheduling, indicating improved efficiency. Moreover, Elyasi and Salmasi (2013) compared a mathematical flow shop model in a dynamic setting to an SPT-based heuristic and reported an average performance gap of 23.9%, with significant variation across instances. This is in line with the differences observed between our exact model and the dispatching rule, both in the performance gap and the large deviations seen for different problem instances. Lastly, Goren et al. (2011) also observed a slight reduction in scheduling efficiency in their job shop model to enhance stability and robustness under dynamic conditions, which is consistent with our results, where minor performance losses were compensated by improved robustness during rescheduling.

However, the research identified several limitations. First, certain practical details, such as minor operational interruptions, informal breaks, and unplanned activities, were not modeled due to limited data availability. Second, the computational demands of the exact optimization model remain high, particularly compared to heuristic methods, which could pose challenges for scaling up or more frequent rescheduling scenarios, potentially changing from discrete time steps to continuous time.

Furthermore, several assumptions in this research must be considered when analyzing the results. Although agreed upon with shop leads, the use of an average setup time for all parts could lead to inaccuracies due to the significant variability in actual setup times across different parts. Moreover, the determination of technician qualifications based on historical hours worked on parts could affect result accuracy since, in practice, these qualifications are determined manually by shop leads. Additionally, machine downtime, which was not modeled in this research, could further impact real-world scheduling outcomes, though it could be integrated similarly to technician unavailability. The assignment of scores for the six priority groups kept the model manageable but ignored variability in job urgency within each group. Finally, the assumption of unlimited buffer space may affect the practical feasibility of scheduling decisions, as storage limitations within actual shop environments could limit new job allocations.

9 Conclusion and Future Research

This study successfully developed and implemented an autonomous flexible job shop scheduling model designed to optimize maintenance operations within dynamic CMRO environments. The scheduling approach uses a MILP formulation, including a multi-objective optimization that aims to minimize weighted tardiness and maximize weighted earliness. The method effectively aims to optimize performance and balance operational metrics, namely TAT and OTD. By incorporating customer-specific priority weights, the model strategically prioritizes urgent and high-value jobs, as indicated by the enhanced performance of high-priority customer groups in the simulations. Furthermore, the developed scheduling model incorporated essential real-world constraints such as technician skills and certification requirements, resource availability, including specialized equipment and workstations, and scheduled technician unavailability. This results in feasible and well-performing maintenance schedules, directly applicable on the shop floor, concluding that the proposed model effectively manages the main operational challenges in CMRO scheduling.

The practical application of this scheduling model in two CMRO shops for the case study at the maintenance provider demonstrated substantial improvements in key performance metrics. The model achieved reductions in TAT of up to 34.21% and increases in OTD of up to 22.97% compared to historical manual planning methods, underlining the effectiveness of the exact optimization approach. The highest performance gains were observed in the shop with the most constrained test capacity and the highest utilization rates. Thus, the conclusion can be drawn that the proposed model is most effective in shop environments where limited resource availability creates significant scheduling challenges.

Additionally, as demonstrated in the simulations, the dynamic rescheduling mechanism enables the model to re-optimize when new jobs arrive, processing times deviate, or additional operations are inserted, thereby maintaining feasible and efficient schedules under real-world disruptions within CMRO shops.

The implemented time-based decomposition approach, using varying scheduling window lengths, demonstrated limited practical value within the context of this case study. However, based on computational experiments, it can be concluded that the model is capable of generating well-performing solutions within the practically feasible time limits, both for the initial scheduling phase and the rescheduling throughout the planning horizon.

Incorporating robustness measures into the scheduling model was essential for improving delivery reliability, as it reduced the number of missed promised due dates despite operational disruptions. Although it affected the flexibility and performance of high-priority jobs, it can be concluded that the improvements in predictability justify the inclusion of robustness and stability measures, especially given the importance of delivery reliability in dynamic CMRO environments.

Finally, the performance gap between the exact optimization approach and the DR heuristic, which, similar to current manual practices, relies on local decision-making, emphasizes the value of global optimization in CMRO scheduling. While heuristics compute schedules within faster computational time, this comes at the cost of solution quality, making them less suitable when optimality is prioritized.

In conclusion, this paper demonstrated the potential of flexible job shop scheduling, enhanced by multi-objective optimization and dynamic rescheduling capabilities, to significantly improve operational efficiency in TAT and OTD within real-world CMRO shops. By confirming that optimization techniques, when realistically constrained, can substantially outperform traditional scheduling methods in dynamic maintenance environments, the primary research objective is reached. Furthermore, by integrating job prioritization, technician skill constraints, robustness, and dynamic rescheduling into a single scheduling framework, this study addresses the identified research gap in the existing literature.

Nevertheless, future research in this area could further enhance model accuracy by exploring variable processing times dependent on technician experience, which significantly influence task durations in CMRO environments. Moreover, decreasing computational time by using metaheuristics or hybrid optimization approaches would further enhance the practical use of the model, allowing for more frequent and faster schedule updates. Another promising area involves deep reinforcement learning techniques in the JSSP, which offer significant potential, particularly for developing adaptive scheduling policies capable of improving over time in dynamic environments. Lastly, while dynamic rescheduling effectively addresses uncertainty in the developed approach, the proactive management of processing time variability through predictive analytics and stochastic optimization represents another promising direction for future research in these environments.

References

Annear, L. M., Akhavan-Tabatabaei, R., and Schmid, V. (2023). Dynamic assignment of a multi-skilled workforce in job shops: An approximate dynamic programming approach. *European Journal of Operational Research*, 306:1109–1125.

Aribowo, W., Prastomo, O. W., and Halim, A. H. (2020). Scheduling an aircraft maintenance shop with dedicated technician and dedicated machine constraints. In *Advances in Industrial and Production Engineering*, Lecture Notes in Mechanical Engineering, pages 260–265. Springer.

Avelino, C., Bourne, D., Ferreira, F., Rasteiro, D., and Santos, J. (2016). Scheduling the repair of aircraft components – a case study. *Mathematics-in-Industry Case Studies*, 7(5):1–18.

Ciro, G. C., Dugardin, F., Yalaoui, F., and Kelly, R. (2015). A fuzzy ant colony optimization to solve an open shop scheduling problem with multi-skills resource constraints. *IFAC-PapersOnLine*, 48(3):715–720.

Damm, R. B., Chaves, A. A., Riveaux, J. A., and Ronconi, D. P. (2024). Scheduling technicians and tasks through an adaptive multi-objective biased random-key genetic algorithm. *Annals of Operations Research*.

Dauzère-Pérès, S., Ding, J., Shen, L., and Tamssaouet, K. (2024). The flexible job shop scheduling problem: A review. *European Journal of Operational Research*, 314(2):409–432.

Elyasi, A. and Salmasi, N. (2013). Stochastic flow-shop scheduling with minimizing the expected number of tardy jobs. *International Journal of Advanced Manufacturing Technology*, 66:337–346.

Fathollahi-Fard, A. M., Woodward, L., and Akhrif, O. (2024). A scenario-based robust optimization model for the sustainable distributed permutation flow-shop scheduling problem. *Annals of Operations Research*.

Fuladi, S. K. and Kim, C.-S. (2024). Dynamic events in the flexible job-shop scheduling problem: Rescheduling with a hybrid metaheuristic algorithm. *Algorithms*, 17(4):142.

Gomes, M. C., Barbosa-Póvoa, A. P., and Novais, A. Q. (2013). Reactive scheduling in a make-to-order flexible job shop with re-entrant process and assembly: a mathematical programming approach. *International Journal of Production Research*, 51(17):5120–5141.

González-Neira, E. M., Montoya-Torres, J. R., and Barrera, D. (2017). Flow-shop scheduling problem under uncertainties: Review and trends. *International Journal of Industrial Engineering Computations*, 8(3):399–426.

Goren, S., Sabuncuoglu, I., and Koc, U. (2011). Optimization of schedule stability and efficiency under processing time variability and random machine breakdowns in a job shop environment. *Naval Research Logistics*, 59(1):26–38.

Hashimoto, H., Boussier, S., Vasquez, M., and Wilbaut, C. (2011). A grasp-based approach for technicians and interventions scheduling for telecommunications. *Annals of Operations Research*, 183:143–161.

Hillier, F. S. and Lieberman, G. J. (2015). *Introduction to Operations Research*. McGraw-Hill Education, New York, NY, 10th edition.

Hsieh, F. Y., Lin, Y. K., and Wang, Y.-C. (2024). Developing algorithm and dispatching rules for scheduling a realistic flexible flow shop with setups to minimize total weighted tardiness. *The International Journal of Advanced Manufacturing Technology*.

Ikli, S. (2022). A rolling horizon approach for the dynamic scheduling of flying taxis. In *Proceedings of the 14th International Joint Conference on Computational Intelligence (IJCCI 2022)*, pages 65–73. SCITEPRESS.

Jamili, A. (2016). Robust job shop scheduling problem: Mathematical models, exact and heuristic algorithms. *Expert Systems with Applications*, 55:341–350.

Johnson, S. M. (1954). Optimal two- and three-stage production schedules with setup times included. *Naval Research Logistics Quarterly*, 1(1):61–68.

Kanet, J. J. and Li, X. (2004). A weighted modified due date rule for sequencing to minimize weighted tardiness. *Journal of Scheduling*, 7(4):261–276.

Karam, A., Attia, E.-A., and Duquenne, P. (2017). A milp model for an integrated project scheduling and multi-skilled workforce allocation with flexible working hours. *IFAC-PapersOnLine*, 50(1):13964–13969.

Khoshsirat, M. and Mousavi, S. M. (2024). A new proactive and reactive approach for resource-constrained project scheduling problem under activity and resource disruption: a scenario-based robust optimization approach. *Annals of Operations Research*, 338:597–643.

Kianpour, P., Gupta, D., Krishnan, K. K., and Gopalakrishnan, B. (2021). Automated job shop scheduling with dynamic processing times and due dates using project management and industry 4.0. *Journal of Industrial and Production Engineering*, 38(7):485–498.

Kumar, K. and He, M. (2023). How ai maintains manufacturing productivity amid reduced capex. *BCG X*.

Meilanitasari, P. and Shin, S.-J. (2021). A review of prediction and optimization for sequence-driven scheduling in job shop flexible manufacturing systems. *Processes*, 9(8):1391.

Melchiors, P., Kolisch, R., and Kanet, J. J. (2024). The performance of priority rules for the dynamic stochastic resource-constrained multi-project scheduling problem: an experimental investigation. *Annals of Operations Research*, 338:569–595.

Mitchell, J. E. (2002). Branch-and-cut algorithms for combinatorial optimization problems. In *Handbook of Applied Optimization*, pages 65–77. Oxford University Press.

Mousakhani, M. (2013). Sequence-dependent setup time flexible job shop scheduling problem to minimise total tardiness. *International Journal of Production Research*, 51(12):3476–3487.

Ouelhadj, D. and Petrovic, S. (2009). A survey of dynamic scheduling in manufacturing systems. *Journal of Scheduling*, 12:417–431.

Perroux, T., Arbaoui, T., and Merghem-Boulahia, L. (2023). A mathematical model for the flexible job shop scheduling problem with availability constraints. *IFAC-PapersOnLine*, 56(2):5388–5393.

Perroux, T., Arbaoui, T., and Merghem-Boulahia, L. (2024). The flexible job shop scheduling problem with setups and operator skills: an application in the textile industry. *IFAC-PapersOnLine*, 58(19):1090–1095.

Pinedo, M. L. (2016). *Scheduling: Theory, Algorithms, and Systems*. Springer, 5th edition.

Qin, W., Zhang, J., and Song, D. (2018). An improved ant colony algorithm for dynamic hybrid flow shop scheduling with uncertain processing time. *Journal of Intelligent Manufacturing*, 29:891–904.

Rahmani, D. and Heydari, M. (2014). Robust and stable flow shop scheduling with unexpected arrivals of new jobs and uncertain processing times. *Journal of Manufacturing Systems*, 33(1):84–92.

Stecco, G., Cordeau, J.-F., and Moretti, E. (2008). A branch-and-cut algorithm for a production scheduling problem with sequence-dependent and time-dependent setup times. *Computers & Operations Research*, 35(8):2635–2655.

Stevenson, Z., Fukasawa, R., and Ricardez-Sandoval, L. (2019). Evaluating periodic rescheduling policies using a rolling horizon framework in an industrial-scale multipurpose plant. *Journal of Scheduling*, 23:397–410.

Tassel, P., Gebser, M., and Schekotihin, K. (2021). A reinforcement learning environment for job-shop scheduling. *arXiv preprint arXiv:2104.03760*.

Thörnblad, K., Strömberg, A.-B., Patriksson, M., and Almgren, T. (2015). Scheduling optimisation of a real flexible job shop including fixture availability and preventive maintenance. *European Journal of Industrial Engineering*, 9(1):126–145.

Tighazoui, A., Sauvey, C., and Sauer, N. (2021). Predictive-reactive strategy for flowshop rescheduling problem: Minimizing the total weighted waiting times and instability. *Journal of Systems Science and Systems Engineering*, 30(3):253–275.

Tliba, K., Diallo, T. M., Penas, O., Ben Khalifa, R., Ben Yahia, N., and Choley, J.-Y. (2022). Digital twin-driven dynamic scheduling of a hybrid flow shop. *Journal of Intelligent Manufacturing*, 34:2281–2306.

van der Weide, T., Deng, Q., and Santos, B. F. (2022). Robust long-term aircraft heavy maintenance check scheduling optimization under uncertainty. *Computers & Operations Research*, 141:105667.

Wang, D., Wu, H., Zheng, W., Zhao, Y., Tian, G., Wang, W., and Chen, D. (2025). A mixed-integer linear programming model for addressing efficient flexible flow shop scheduling problem with automatic guided vehicles consideration. *Applied Sciences*, 15(6):3133.

Wang, L., Luo, C., and Cai, J. (2017). A variable interval rescheduling strategy for dynamic flexible job shop scheduling problem by improved genetic algorithm. *Journal of Advanced Transportation*, 2017:1–12.

Xiong, H., Shi, S., Ren, D., and Hu, J. (2022). A survey of job shop scheduling problem: The types and models. *Computers & Operations Research*, 142:105731.

Xiong, J., Xing, L.-n., and Chen, Y.-w. (2013). Robust scheduling for multi-objective flexible job-shop problems with random machine breakdowns. *International Journal of Production Economics*, 141(1):112–126.

Yan, Q., Wu, W., and Wang, H. (2022). Deep reinforcement learning for distributed flow shop scheduling with flexible maintenance. *Machines*, 10(3):210.

Yu, C., Semeraro, Q., and Matta, A. (2018). A genetic algorithm for the hybrid flow shop scheduling with unrelated machines and machine eligibility. *Computers and Operations Research*, 100:211–229.

Zahmani, M. H., Atmani, B., Bekrar, A., and Aissani, N. (2021). Multiple priority dispatching rules for the job shop scheduling problem. *Processes*, 9(8):1392.

Zeiträg, Y. and Figueira, J. R. (2023). Automatically evolving preference-based dispatching rules for multi-objective job shop scheduling. *Journal of Scheduling*, 26:289–314.

Özgüven, C., Özbakır, L., and Yavuz, Y. (2010). Mathematical models for job-shop scheduling problems with routing and process plan flexibility. *Applied Mathematical Modelling*, 34(6):1539–1548.

Appendices

A Dynamic Rescheduling Algorithm

Algorithm 2 Rescheduling flexible job shop problem after dynamic events

Input
  previous schedule $\mathcal{S}^{old} = \{\mathbf{S}, \mathbf{C}, \mathbf{X}, \mathbf{Y}\}$
  updated job set $\mathcal{J} \rightarrow$ including *new* jobs and *updated* processing times
  current time t
// Step 1 – Classify each operation
for each $(j, o) \in \mathcal{J}$ **do**
  if $C^{old}_{j,o}$ is before t **then**
    completed before cut: mark COMPLETED
  else if $S^{old}_{j,o}$ is before t **and** $C^{old}_{j,o}$ is equal or after t **then**
    currently in progress: mark ONGOING
  else
    mark NOT-STARTED
  end if
end for
// Step 2 – Handle resource breaks intersection for ongoing, delayed, operations
for each ONGOING (j, o) **do**
  if a technician or machine assigned to (j, o) has a break (b_s, b_e) with $t \leq b_s < C^{updated}_{j,o}$ **then**
    split (j, o) into two sub-operations:
      (j, o_a) with duration $b_s - S^{old}_{j,o}$ (fixed as ongoing),
      (j, o_b) with duration $C^{updated}_{j,o} - b_s$
  end if
end for
// Step 3 – Build a new MILP identical to the week-start model, given in Section 5.2.3
starting model $\leftarrow$ initialization with model given for the P&H shop, enforcing all operational constraints
// Step 4 – Partially fix variables so the past cannot move
for each operation classified as COMPLETED **do**
  fix $S_{j,o} = S^{old}_{j,o}$, $C_{j,o} = C^{old}_{j,o}$
  fix all resource assignment binary variables found in $\mathcal{S}^{old}$
end for
for each operation classified as ONGOING **do**
  fix start time $S_{j,o} = S^{old}_{j,o}$ and resource assignment binary variables
  if an operation has an updated processing time **then**
    fix $C_{j,o} = S^{old}_{j,o} +$ extra processing time
  end if
end for
for each operation classified as NOT-STARTED **do**
  enforce $S_{j,o} \geq \max\{t, \text{ job arrival time}\}$
end for
// Step 5 – Optimise
run starting model and added fixed variables
Output updated schedule $\mathcal{S}^{new}$

B Dispatching Rule Algorithm

Algorithm 3 Discrete-Event DR Scheduler for the Flexible Job Shop Problem

Input: jobs $\mathcal{J}$; technicians $\mathcal{T}$; machines $\mathcal{M}$; horizon H; step Δt
initialise resource states and empty maps $S, C, X^{\text{tech}}, X^{\text{mach}}$
for $t \leftarrow 0$ **to** H **step** Δt **do**
  // release newly arrived jobs
  for $j \in \mathcal{J}$ **and** $arrival_j \leq t$ **and not released do**
    mark j as released
  end for
  // finish running operations
  for all resource $r \in \mathcal{T} \cup \mathcal{M}$ **finishing at** t **do**
    correct overrun; split at breaks if needed
    update next-free time and last finished op of the technician
  end for
  // dispatch on each idle technician
  for all technician $k \in \mathcal{T}$ **and** k idle at t **do**
    try chained next-op of k's last job
    if no op scheduled **then**
      build ready list $\mathcal{O}$; compute with WMDD priority weight $\pi(i)$
      sort $\mathcal{O}$ by descending π
      for $o \in \mathcal{O}$ **do**
        if o is technician-only operation **and** feasible on k **and** no intersection with unavailability **then**
          schedule o on k
        else if o is test operation **and** free machine m exists **then**
          schedule o on k and m
        end if
      end for
    end if
  end for
end for
Output: start times S, finish times C, resource assignments $X^{\text{tech}}, X^{\text{mach}}$

C Simulation outcomes

C.1 Hydraulics shop

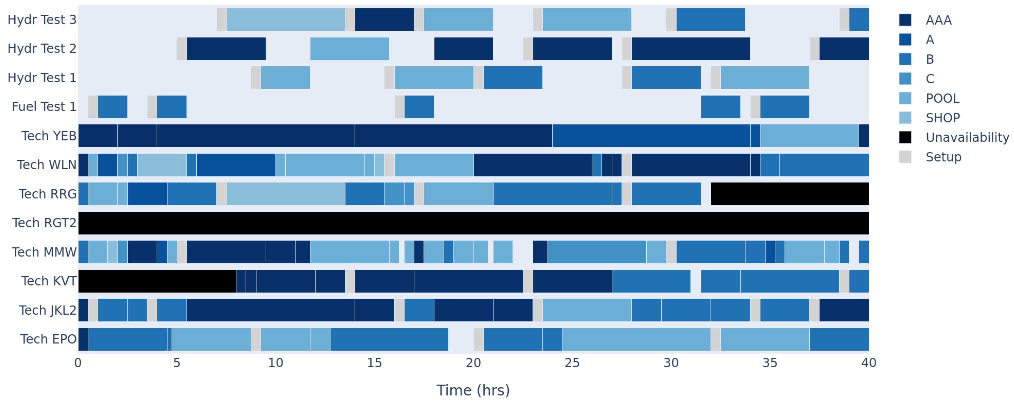

Figure 16: Gantt chart outcome of the Hydraulics shop weekly scheduling simulation using the exact approach: week of 2 September.

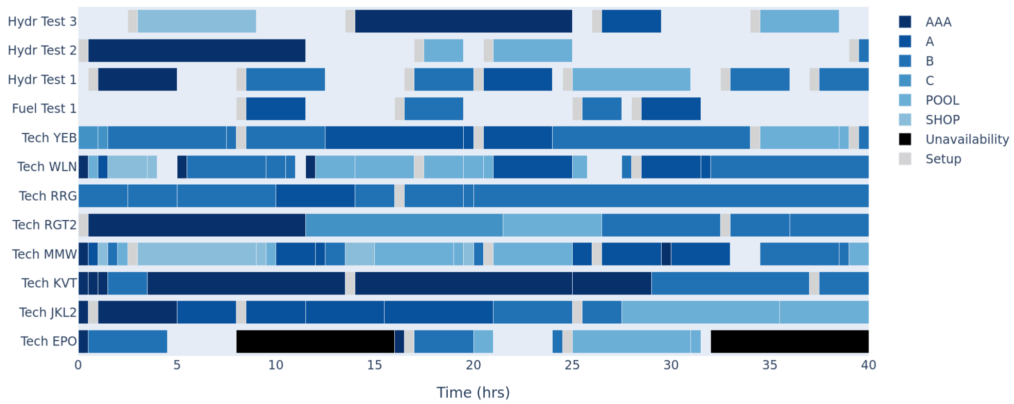

Figure 17: Gantt chart outcome of the Hydraulics shop weekly scheduling simulation using the exact approach: week of 9 September.

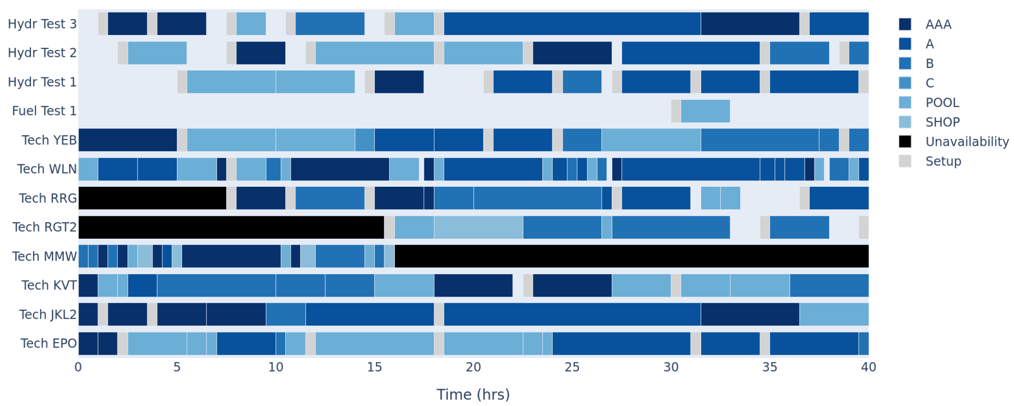

Figure 18: Gantt chart outcome of the Hydraulics shop weekly scheduling simulation using the exact approach: week of 9 December.

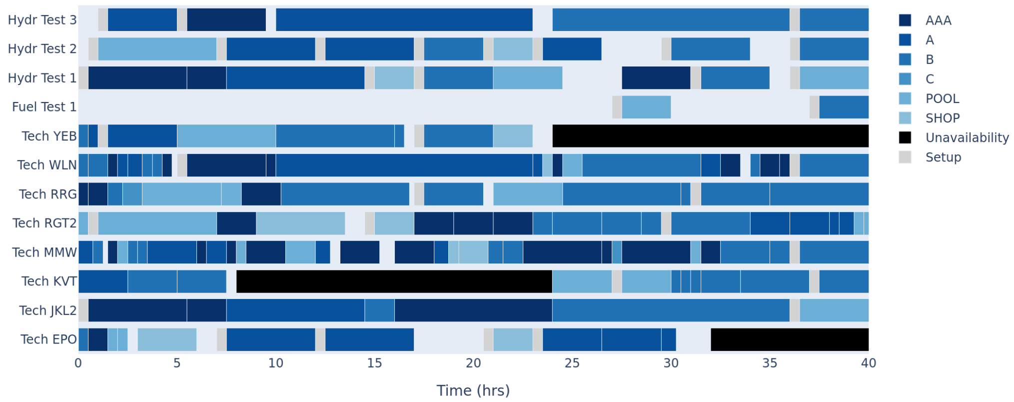

Figure 19: Gantt chart outcome of the Hydraulics shop weekly scheduling simulation using the exact approach: week of 16 December.

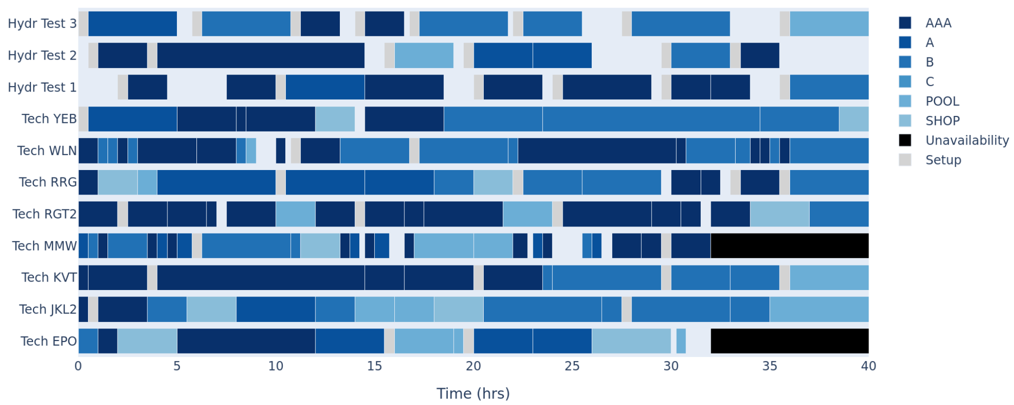

Figure 20: Gantt chart outcome of the Hydraulics shop weekly scheduling simulation using the exact approach: week of 13 January.

C.2 Pneumatics shop

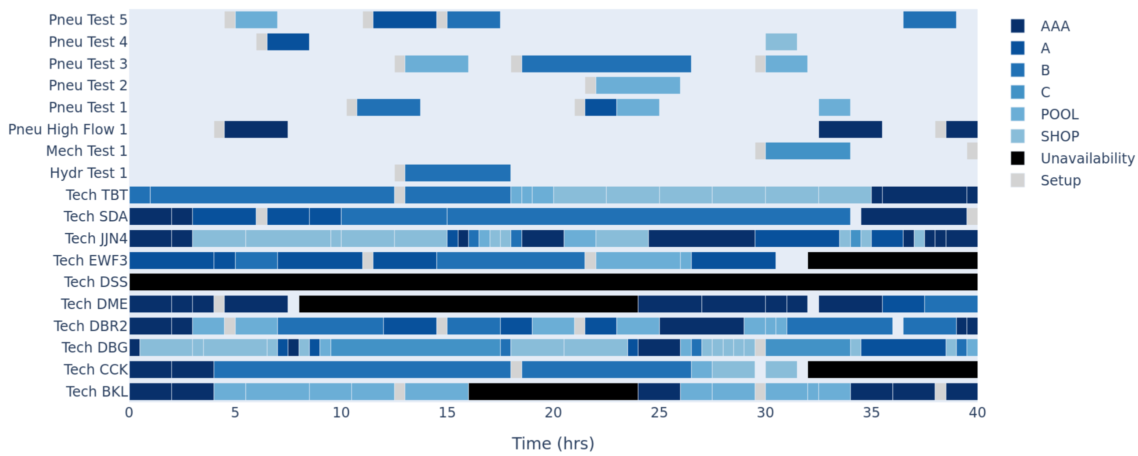

Figure 21: Gantt chart outcome of the Pneumatics shop weekly scheduling simulation using the exact approach: week of 9 September.

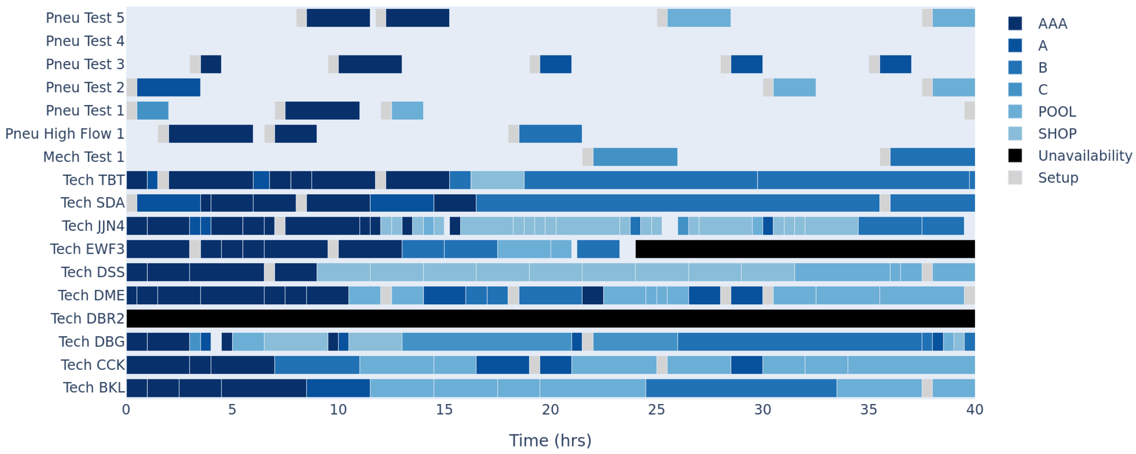

Figure 22: Gantt chart outcome of the Pneumatics shop weekly scheduling simulation using the exact approach: week of 30 September.

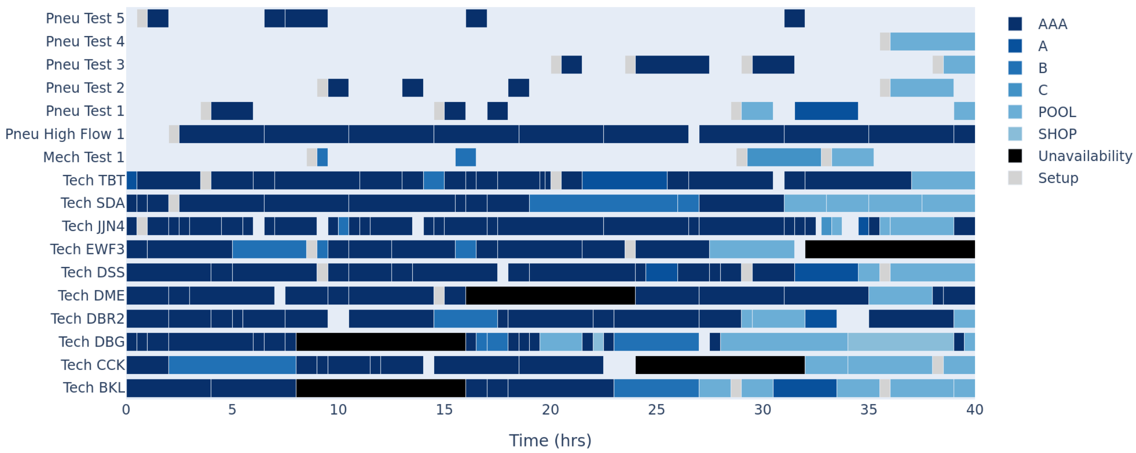

Figure 23: Gantt chart outcome of the Pneumatics shop weekly scheduling simulation using the exact approach: week of 9 December.

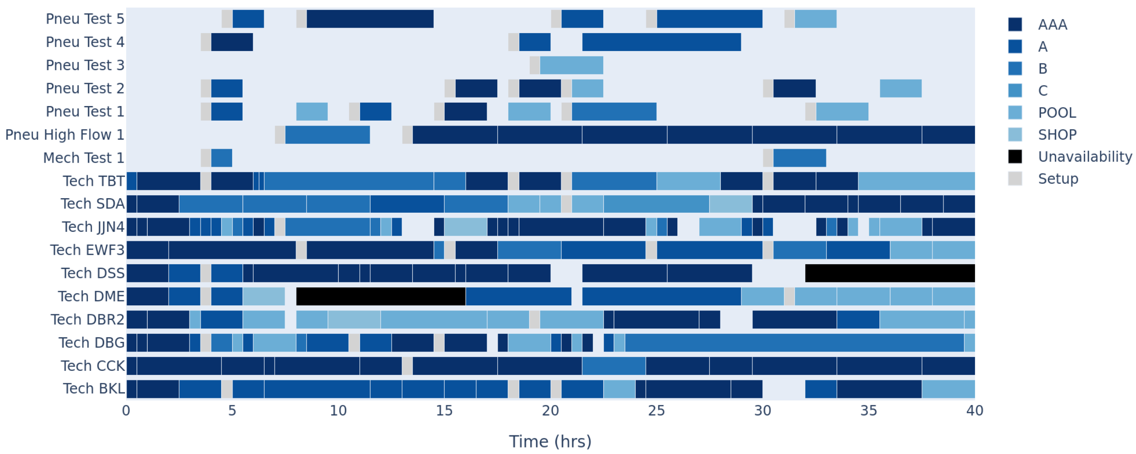

Figure 24: Gantt chart outcome of the Pneumatics shop weekly scheduling simulation using the exact approach: week of 13 January.

C.3 IDG shop

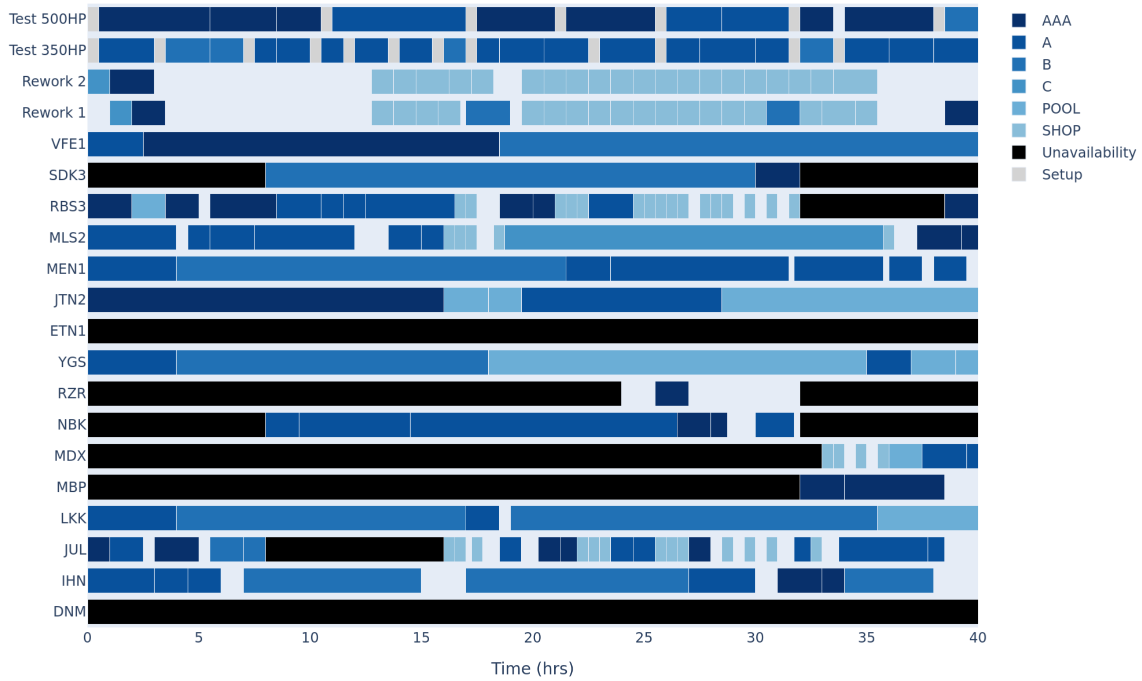

Figure 25: Gantt chart outcome of the IDG shop weekly scheduling simulation using the exact approach: week of 16 September.

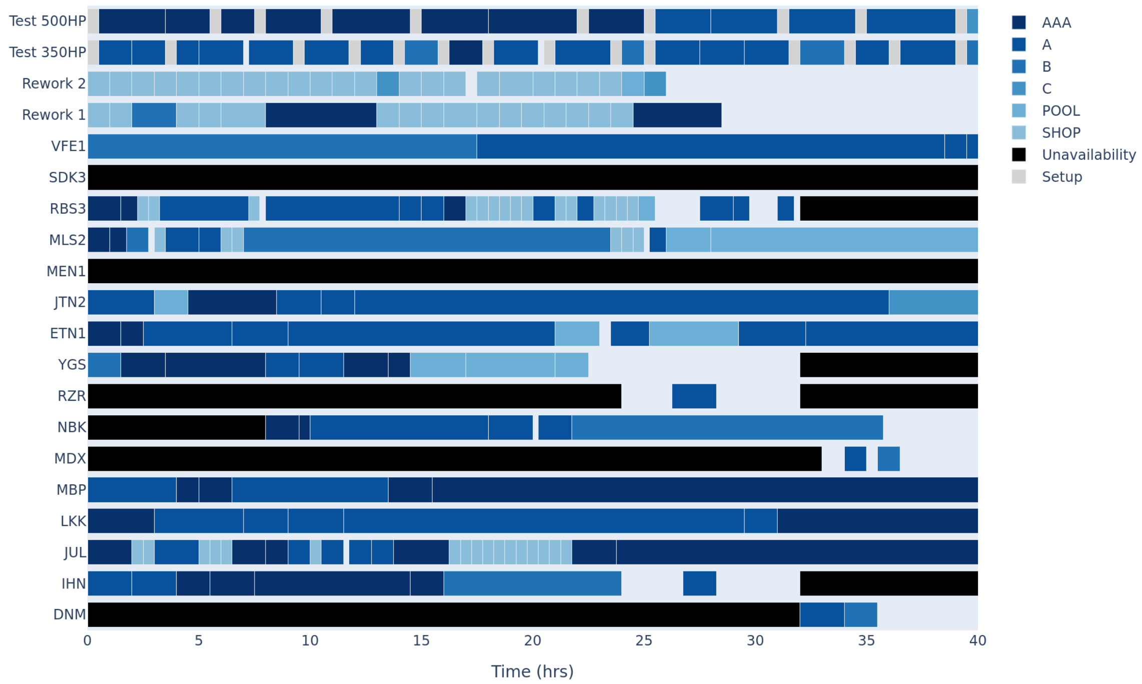

Figure 26: Gantt chart outcome of the IDG shop weekly scheduling simulation using the exact approach: week of 23 September.

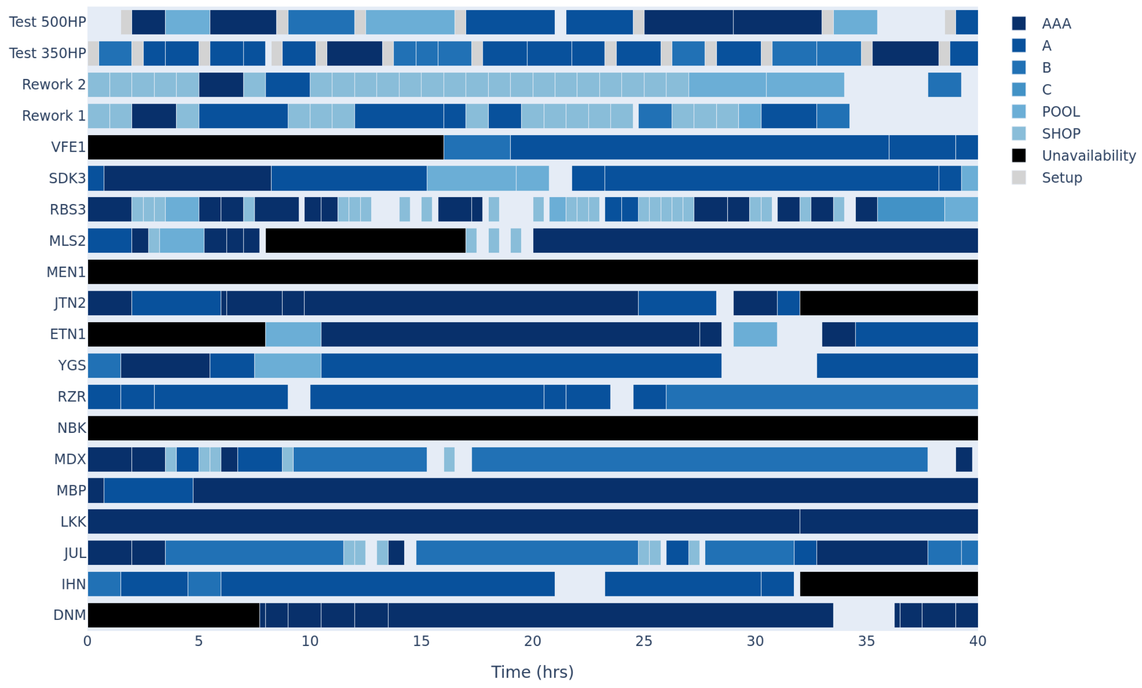

Figure 27: Gantt chart outcome of the IDG shop weekly scheduling simulation using the exact approach: week of 14 October.

II

Literature Study

State-of-the-Art Review & Research Proposal

Scheduling Optimization for Component, Maintenance, Repair and Overhaul Shops: a Case Study for an Independent Component Maintenance Provider

AE5322: Thesis Control & Operations

Thijs Roolvink

TU Delft

State-of-the-Art Review & Research Proposal

Scheduling Optimization for Component, Maintenance, Repair and Overhaul Shops: a Case Study for an Independent Component Maintenance Provider

by

Thijs Roolvink

T.W. Roolvink
4961382

Supervisors TU Delft:	Marta Ribeiro, Paul Roling
Company Supervisors:	Kaveh Alizadeh, Laurens Koningsbloem
Project Duration:	October, 2024 - July, 2025
Faculty:	Faculty of Aerospace Engineering, Delft
Date:	December 18, 2024

Contents

List of Figures

List of Tables

Nomenclature

Abbreviations

Abbreviation	Definition
BIP	Binary Integer Programming
BRKGA	Multi-Objective Biased Random-Key Genetic Algorithm
CMRO	Component Maintenance, Repair, and Overhaul
EDD	Earliest Due Date
FCFS	First-Come-First-Serve
FFSP	Flexible Flow Shop Problem
FIFO	First-In-First-Out
FSSP	Flow Shop Scheduling Problem
HFSP	Hybrid Flow Shop Problem
IDG	Integrated Drive Generators
ISA	International Standard Atmosphere
JSSP	Job Shop Scheduling Problem
KPI	Key Performance Indicator
MIP	Mixed Integer Programming
MILP	Mixed-Integer Linear Programming
MTWR	Most Total Work Remaining
OEE	Overall Equipment Effectiveness
OTD	On-Time Delivery
OR	Operations Research
PSMWAP	Multi-Skilled Workforce Allocation Problem
SLA	Service Level Agreements
SPT	Shortest Processing Time
SWT	Shortest Waiting Time
TAT	Turnaround Time
TIST	Technicians and Interventions Scheduling Problem
TWT	Total Weighted Tardiness
WIP	Work-In-Progress
WO	Work Orders

1 Introduction

Component Maintenance, Repair, and Overhaul (CMRO) shops play an important role in ensuring the reliability and availability of aircraft by performing maintenance on various aviation components. These shops handle thousands of parts annually, varying from navigation instruments to power generation systems, using highly trained and experienced technicians for specific inspection, repair, and testing tasks. Efficient scheduling in these environments is crucial for maximizing resource utilization and maintaining service-level agreements to improve customer satisfaction, minimize operational costs, and reduce turnaround times. In current CMRO environments, scheduling processes rely on simple prioritization models and manual decision-making. The allocation of operations to technicians is mainly done by the expertise and judgment of the shop lead, rather than relying on a scheduling algorithm. While this is effective to some extent, this approach lacks the flexibility to adapt to dynamic operational changes. Additionally, this inefficient scheduling can lead to increased Work-In-Progress (WIP), delays in critical repairs, and higher labor and resource costs, making optimization of scheduling an important topic in CMRO operations.

Recent advancements in scheduling optimization demonstrate the potential impact of integrating data-driven methods. For instance, new scheduling model tools have been shown to improve overall equipment effectiveness (OEE) by over 3%, reduce planning-related labor hours by more than 50%, and improve sustainability and customer satisfaction (Kumar and He, 2023). Such results emphasize the importance of transitioning from manual, experience-based scheduling to automated, optimization-driven systems. There are many other examples in the aviation industry where efficient scheduling has significantly reduced maintenance efforts and improved utilization. For example, a study on aircraft heavy maintenance check scheduling introduced a genetic algorithm-based approach that reduced the total number of heavy maintenance checks by 7%. Additionally, it increased aircraft utilization by 4.4%, potentially leading to significant annual maintenance cost savings (van der Weide et al., 2022). These improvements show the potential of advanced scheduling techniques in aviation maintenance operations.

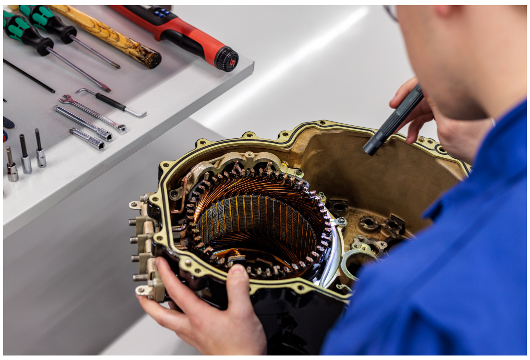

Figure 1.1: Technician working on an IDG unit in the Power Generation shop of the component maintenance provider.

This project aims to develop an advanced, automated, dynamic scheduling model tailored to diverse and complex CMRO environments. The model will integrate stochastic elements related to uncertain processing times and new job arrivals. Additionally, by incorporating dynamic properties, the proposed model will adapt to real-time updates, improving decision-making and operational efficiency on the shop floors. Another key aspect of this project is technician assignment, which could be handled with machine assignment in the Flow Shop Scheduling Problem (FSSP). Due to the high specialization of technicians in these shop environments, certain tasks can only be performed by individuals with specific skills and qualifications. Constraints such as limited equipment availability and the need to batch certain testing tasks for efficiency are also important considerations.

The model developed in this research will be applied to a case study at an independent component maintenance provider, a significant player in aerospace maintenance. The maintenance provider operates three different CMRO shops, each with specialized processes and unique characteristics. The Hydraulics and Pneumatics (P&H) shop focuses on both hydraulic and pneumatic components. This shop contains two subgroups. The Avionics shop specializes in repairing and maintaining displays, control units, and flight data systems. It is divided into multiple workgroups, each focused on a specific domain. The Power Generation shop works on parts related to the power generation of aircraft, for instance, Integrated Drive Generators (IDG). This shop is characterized by its resource allocation flexibility and a significantly more complex workflow. Due to its involvement in CMRO operations, the company offers a great opportunity to test and build the proposed scheduling model.

2
State-of-the-Art

Effective scheduling is important in optimizing operational performance within Component Maintenance, Repair, and Overhaul (CMRO) environments. Complex processes, diverse constraints, and dynamic elements such as varying job priorities, uncertain processing times, and resource limitations characterize these operations. Overcoming these challenges requires advanced scheduling techniques that outperform traditional static and rule-based approaches, using modern optimization methods and applied concepts.

The field of scheduling optimization has seen significant advancements in recent years, with applications across manufacturing, maintenance, and production industries. Research has focused on improving scheduling efficiency under real-life complications and constraints like resource availability and uncertainty. This chapter provides the state-of-the-art techniques and approaches most relevant to the scheduling challenges in CMRO shop environments, mainly related to the job and flow shop scheduling problems, but the influence of efficient resource allocation of technicians is also important.

This chapter begins by exploring the foundational problem of flow and job shop scheduling in section 2.1, explaining how these models have evolved, how they are classified, and recent developments in incorporating different operational characteristics. Next, section 2.2 focuses on research about assigning technicians with specific skills, which is important in environments where skill constraints significantly affect how tasks can be allocated. After that, section 2.3 looks at human factor constraints in scheduling, such as working-hour regulations and technician availability. Section 2.4 explains how different job priorities are included in scheduling models, while section 2.5 goes into deterministic and stochastic scheduling approaches. Section 2.6 dives into dynamic scheduling methods that respond to changes in real time, including handling uncertain job arrivals and variable processing times. Section 2.7 looks at robust optimization techniques, focusing on making schedules more reliable with high variability in parameters. Section 2.8 introduces rolling horizon and decomposition methods to solve large scheduling problems in smaller parts. In section 2.9, both exact optimization methods and other techniques such as metaheuristics are discussed. Section 2.10 provides a general look at existing research that has integrated topics mentioned in past sections of the chapter. Finally, section 2.11 identifies important gaps in the current methods and prepares the groundwork for creating a scheduling model tailored to the needs of a CMRO environment.

2.1. Classification of Scheduling Problems

Scheduling problems can be classified using well-known frameworks based on their objectives, constraints, and environmental characteristics. These classifications help in understanding the nature of scheduling problems and finding the right framework to optimize them. According to Pinedo (2016), scheduling problems can generally be categorized by their machine environments, job characteristics, and the objectives to be optimized. Additionally, scheduling problems include deterministic and stochastic variants, and they can range from single-machine to complex multi-machine setups such as flow shops and job shops.

Flow shop scheduling problems involve a set of jobs processed in the same order across multiple machines. In this environment, the sequencing of jobs significantly impacts objectives that can be used such as makespan or tardiness. In contrast, job shop scheduling allows for unique job routing across machines, adding slightly more complexity to the problem due to the varying sequences and dependencies (Pinedo, 2016). These differences are relevant to CMRO operations, where different shops show mostly characteristics of the flow shop scheduling problem (FSSP).

In examples of CMRO shops, with the focus on, for instance, Avionics, Hydraulics or Pneumatics, tasks typically follow a linear flow, similar to a flow shop environment. Similarly, a shop focused on Power Generation parts also follows a pre-determined sequence of routing steps, similar to the characteristics of a flow shop model. Because in the routing steps that are followed in all three shops, there are multiple processors, which are the multiple technicians, to whom jobs can be assigned, the shops are showing characteristics of a hybrid flow shop problem (HFSP) (Hillier and Lieberman, 2015). Additionally, the shops show some characteristics of a flexible flow shop problem (FFSP), because they involve multiple phases where each phase may have one or more technicians, referred to as machines, that can perform various tasks. According to Hillier and Lieberman, 2015, a flexible flow shop is a classification that combines elements of flow shops and job shops, creating and incorporating more flexibility in processing jobs. This will create routing flexibility, where jobs can be allocated to multiple machines within a phase, machine flexibility, where machines can handle different tasks, and resource or labor flexibility, where the labor can move between machines and phases. The routing steps of the process are referred to as phases in these flow shop models. The machines or technicians in each phase can be identical or unrelated, and the latter is applicable in the CMRO environment because of the unique skillset of each technician.

Another critical dimension of classification is the nature of the scheduling environment: static or dynamic. Static scheduling means that all jobs and resources are known in advance, making it suitable for predictable and controlled environments. Dynamic scheduling, in contrast, implements and takes into account variability in job arrivals and processing times, incorporating the real-world conditions comparable to CMRO operations (Z. Wang et al., 2020). In CMRO shop operations, the scheduling problem is mostly dynamic, driven by unpredictable job arrivals and stochastic processing times.

The inclusion of stochastic variables, such as uncertain processing times and unexpected job arrivals, introduces additional complexity to scheduling problems. Stochastic scheduling methods address these complexities by using probabilistic and robust optimization techniques, which have the goal of creating schedules resilient to variability and disruptions (Xiao et al., 2017). When considering a better representation of real-life operations, both uncertain processing durations and job arrivals must be incorporated into the model.

2.2. Skilled Technician Assignment

In scheduling environments where human resources play a crucial role, technician assignment is an important aspect to consider, especially in dynamic and skill-constrained environments. Assigning technicians to tasks requires consideration of their individual skill sets, the tasks' specific requirements, and additional constraints such as the availability of fitting work.

Technician assignment problems often involve multi-skilled personnel, where workers have diverse skill sets that must be matched to job requirements. For instance, Ciro et al. (2015) addresses this challenge in an open shop scheduling problem with multi-skills resource constraints. Their approach implemented constraints that only technicians with the required skill levels perform a given job. These constraints enforce that each operation is assigned to a qualified technician, who remains assigned to the operation throughout its duration. The authors further emphasize the complexity added by multi-skilled personnel assignment, which transforms the problem into an NP-hard optimization challenge.

Similarly, Damm et al. (2024) propose an adaptive multi-objective biased random-key genetic algorithm (BRKGA) to handle technician scheduling. Their model incorporates technician-specific skills as binary parameters, ensuring that a task can only be executed by a qualified individual. The BRKGA approach optimizes objectives such as maximizing task priority fulfillment and minimizing resource inefficiencies, showing its effectiveness in managing real-world scheduling problems with diverse constraints.

In more dynamic contexts, such as job shops, Annear et al. (2023) use an approximate dynamic programming approach to assign a multi-skilled workforce in the presence of stochastic processing times and demand uncertainty. Their model captures the real-time changes in technician availability and skill requirements, dynamically assigning resources to maximize productivity while maintaining flexibility to adapt to changing and uncertain demands.

A comprehensive review by Afshar-Nadjafi (2021) highlights the increasing importance of multi-skilling in scheduling problems, especially in environments where workers must perform tasks requiring diverse skill sets. The study indicates that Mixed-Integer Linear Programming (MILP) approaches are most general, accounting for 54.2% of the reviewed research, due to their ability to accurately incorporate constraints such as skill requirements. Additionally, metaheuristic methods, such as Genetic Algorithms and Ant Colony Optimization, represent 28.7% of the solving methodologies, offering better solutions for large and complex problems.

Another example of skilled technician scheduling is Project Scheduling with Multi-Skilled Workforce Allocation Problem (PSMWAP), which involves task scheduling and workforce allocation, addressing real-world complexities such as skill-dependent task durations and flexible working hours (Karam et al., 2017). In this model, each task requires a set of skills, with task durations influenced by the skill levels and the number of workers assigned. Multi-skilled workers, whose efficiencies vary across skills, are assigned based on skill requirements constraints, ensuring a minimum skill level. The model also considers team-building aspects by pairing expert workers with trainees to encourage knowledge transfer, enhancing workforce flexibility over time. Additionally, penalties for project tardiness, unbalanced workloads, and excessive overtime are incorporated into the objective function, reflecting the trade-offs between efficiency and cost.

By using methodologies from Ciro et al. (2015), Damm et al. (2024), and Annear et al. (2023), the proposed model for the CMRO shops integrates multi-skilled technician assignment with the scheduling of jobs. This hybrid approach tries to optimize task allocations and make sure that skill constraints are satisfied.

2.3. Human Factor Constraints

Integrating human factors such as working-hour regulations, shift patterns, and technician availability into scheduling models is essential to maintaining realistic workloads and implementing factors such as the absence of technicians on certain days, which need to be accounted for. Hashimoto et al. (2011) highlights the importance of ensuring team stability and compliance with workday limits to prevent overburdening technicians. Constraints such as non-working days and maximum daily work hours are explicitly modeled in the algorithm to comply with regulatory requirements.

Damm et al. (2024) further incorporates lunch breaks and daily work schedules into their framework, ensuring compliance with labor laws and improving the practicality of the generated schedules. Their model accounts for both individual and team-level constraints, providing an extensive approach to managing technician workloads. These researches are clear examples of incorporating human factors in an algorithm, which need to be taken into account in a hybrid flow shop when modeling human operators or technicians, as machines in a production environment.

2.4. Job Priorities

The different priorities of jobs that need to be processed will play an important role in the model proposed in this research. There are different ways of incorporating priority rules into a scheduling model. In the Technicians and Interventions Scheduling Problem (TIST) proposed by Hashimoto et al. (2011), priorities are assigned to interventions, and their objective function is designed to minimize the weighted completion time of interventions based on their priority levels. Specifically, the function assigns higher weights to tasks of greater importance. This weighting system ensures that interventions with higher priorities are scheduled as early as possible while balancing resource availability and precedence constraints.

Schworm et al. (2023) proposes a method that assigns penalties to jobs based on their completion time relative to their priority. Higher-priority jobs are weighted more heavily, resulting in earlier scheduling and minimization of delays for critical tasks. This approach implements the importance of considering job urgency within a multi-objective framework, balancing other operational goals such as workload distribution and resource utilization.

The concept of Total Weighted Tardiness (TWT) provides an effective framework for integrating job priorities into scheduling models. Hsieh et al. (2024) proposes a method that minimizes TWT by assigning weights to jobs based on factors such as customer importance, order profitability, and due dates. By incorporating these weights into the scheduling objective, the method ensures that high-priority jobs, such as urgent customer orders, are scheduled earlier, reducing overall tardiness. This approach effectively balances operational efficiency with customer satisfaction by aligning the scheduling process with organizational priorities.

It is important to state the differences between the concept of job priorities and the use of priority rules in scheduling. Job priorities, as will be incorporated in this model, reflect the relative importance of certain tasks based on external factors, such as customer requirements or the criticality of the job to operational goals. These priorities ensure that high-priority jobs, such as those for key customers or time-sensitive repairs, are given precedence in the overall scheduling objective. In contrast, priority rules are heuristics used to determine the sequencing of jobs during the scheduling process, such as First-Come-First-Serve (FCFS). While priority rules are instrumental in solving some of the scheduling problems, job priorities define the strategic importance of tasks, making sure that the schedule aligns with business objectives.

2.5. Deterministic & Stochastic scheduling

Scheduling problems are researched extensively within the operations research (OR) domain, where most of the problems and models focus on deterministic scheduling problems where parameters are non-dynamic and known beforehand. However, as Elyasi and Salmasi (2013) highlights, real-world manufacturing systems or maintenance shops often have to deal with uncertainties during the planning of production, such as resource unavailability, varying due dates, sudden job arrivals, and uncertain processing durations. These challenges are the limitations of deterministic models and the need for stochastic scheduling approaches that account for such uncertainties.

To handle these uncertainties, different approaches have been developed and discussed. One method is to consider processing times as random variables with known probability distributions, leading to stochastic scheduling models (Emmons and Vairaktarakis, 2013). In stochastic scheduling, different ways to compare and rank schedules under uncertainty include expectation ordering, stochastic ordering, and almost sure ordering. Another approach is robust scheduling, which aims to create schedules that are less sensitive to variations in processing times and other random events, which will also be explained in a later section of this review.

As mentioned by Pinedo (2016), a good scheduling model should address these forms of randomness. Uncertainties like machine breakdowns can be modeled as part of the processing times or as separate stochastic processes that determine machine availability. By including these uncertainties in scheduling models, more effective scheduling policies can be developed that are more suitable to reflect the complexities of real-world manufacturing environments.

2.6. Dynamic Scheduling

Dynamic scheduling addresses the challenges of unforeseen events and real-time changes in manufacturing environments, such as job releases, delays in process durations, or shifts in due dates. Unlike static scheduling, which relies on predetermined parameters and information, dynamic scheduling continuously adapts to real-time information to maintain efficiency and continuity in operations.

Three primary strategies define dynamic scheduling approaches. The first, completely reactive scheduling, generates no pre-schedule and relies on real-time decision-making using dispatching rules or heuristics. This approach is quick to implement but often results in suboptimal system performance due to poor quality outcomes (Z. Wang et al., 2020). The second strategy, predictive-reactive scheduling, calculates an initial schedule for the production, and when dynamic events occur, the planning is adjusted. Predictive–reactive scheduling is the most commonly used strategy in existing literature

(Gomes et al., 2013; Ouelhadj and Petrovic, 2009). Finally, robust proactive scheduling anticipates potential disturbances by incorporating buffers or slack into schedules, improving their ability to handle disruptions. Bevelander (2022) highlights methods such as slack allocation and real-time data integration to enhance robustness, ensuring schedules remain effective in dynamic environments. While effective in minimizing deviations, this method requires extensive robustness modeling, which can be computationally intensive and difficult in some environments.

In addition to these strategies, policies such as periodic rescheduling, event-driven rescheduling, and hybrid approaches are employed to determine when adjustments should be made. Periodic rescheduling updates schedules at regular intervals, offering stability but potentially overlooking urgent disruptions. Event-driven rescheduling reacts to specific changes, such as urgent job arrivals, while hybrid approaches combine both methods for greater flexibility (Kianpour et al., 2021).

For CMRO environments, dynamic scheduling is particularly valuable due to the unpredictable nature of repair and maintenance workflows. The integration of predictive-reactive strategies, along with hybrid rescheduling policies, ensures that schedules remain both efficient and resilient, ensuring that the dynamic priorities and resource constraints are respected in real-time operations.

2.7. Robust Optimization

In stochastic scheduling, handling environments with uncertainty in job processing times is a significant challenge when trying to create reliable and optimized schedules. One effective approach to handling this uncertainty is robust scheduling, which focuses on developing solutions that are resilient to variations in processing times and other unpredictable factors. A method of robust scheduling that the book written by Emmons and Vairaktarakis (2013) uses is the minimax regret criterion to minimize the worst-case deviation from the optimal makespan. Instead of specifying exact processing times, this method involves setting ranges for uncertain task durations, resulting in schedules that are particularly resistant to fluctuations. Although the minimax regret approach is more computationally demanding than deterministic methods, experiments show that schedules remain relatively stable and effective even when processing time ranges are broadened up to 50

Additionally, Fathollahi-Fard et al. (2024) demonstrates the potential benefits of robust optimization models in manufacturing operations. By addressing uncertainties like machine breakdowns, variable processing times, and random arrivals of new jobs, their scenario-based robust optimization model improves reliability. It measures the performance of the model under different robustness coefficients, and the optimal value is determined with a sensitivity analysis. Using such robust models leads to more reliable schedules and can improve key performance metrics. This shows the importance of taking into account robust scheduling approaches in real-world production scenarios like CMRO shops, where uncertainties are very common and can significantly impact operational efficiency.

To measure the robustness of a hybrid flow shop schedule, several metrics can be used. For instance, Goren and Sabuncuoglu (2009) emphasizes the expected realized performance as a key robustness measure, where the idea is to focus on the expectation of actualized schedule performance, for instance, the total flow time or tardiness, under various disruption scenarios. This approach can also be extended to assess the variability of schedule performance, for instance, by incorporating terms such as the standard deviation or variance of realized outcomes. Another approach, as used by Jamili (2016), introduces explicit formulas to compute necessary buffer times that can ensure a predetermined level of robustness. In this method, one can define a maximum acceptable expected delay or a probability threshold indicating that the probability of exceeding a certain delay limit remains below a predefined level. By adjusting these coefficients, the planner can optimize and customize the degree of resilience in the schedule.

A different approach is highlighted in Rahmani and Heydari (2014), where robust optimization models incorporate uncertain parameters directly in the decision-making process. Here, coefficients that represent the level of uncertainty or adjustable parameters related to buffer sizes can be integrated into the model. These types of parameters help ensure the schedule remains relatively stable and effective across diverse scenarios. This can be realized by constraining the expected deviations to remain below a certain threshold.

To define schedule robustness more quantitatively, prior literature has proposed several metrics and approaches. One simple measure is the likelihood that a schedule remains feasible or achieves a performance target despite uncertainties in parameters. This can be expressed as a service-level metric that indicates the chances of completing all tasks by a specified deadline, even with stochastic processing times or resource disruptions, as suggested by Himmiche et al. (2023). This probability-based measure enables decision-makers to easily understand the level of robustness as the probability that the schedule will stay within an acceptable limit.

2.8. Rolling Horizon and Decomposition Techniques

To manage the complexity of large-scale scheduling problems, decomposition methods are mostly used to break down a problem into smaller subproblems, which can make the problem easier to solve. According to Pinedo (2016), there are several types of decomposition methods: machine-based decomposition, job-based decomposition, time-based decomposition, and hybrid methods that combine these approaches.

In Machine-Based Decomposition, machines are scheduled individually, often starting with the most critical ones. Techniques like the shifting bottleneck procedure, implemented by Cayo and Onal (2020), are examples of this approach. This method can be very useful for flexible flow shops and job shops, but determining machine criticality and solving the resulting subproblems can be difficult. Job-Based Decomposition focuses on scheduling jobs one at a time, prioritizing them to minimize their impact on the overall schedule. This method is useful when there are specific timing constraints between the operations of a job.

Time-Based Decomposition or Rolling Horizon Procedures divide the scheduling horizon into smaller time intervals. A schedule is created for each interval separately, not including the events outside that period. After scheduling the current interval, the process moves on to solving the next period, as explained in S. Wang et al. (2013). This approach is useful in dynamic environments where jobs arrive over time or other parameters are uncertain.

Hybrid Methods combine elements of machine-based, job-based, and time-based decomposition to leverage the advantages of each approach. For example, a hybrid method might first use time-based decomposition to divide the scheduling horizon and then apply machine-based decomposition within each time interval. Kress et al. (2019) uses such a hybrid method, where machine-based and time-based decomposition is used to solve a flexible job shop scheduling problem.

Other well-known examples of decomposition methods are Benders decomposition and Lagrangian Relaxation, which are often used to solve complex scheduling problems in earlier research. For example, Tan and Terekhov (2018) applied a Benders decomposition to a flexible flow shop scheduling problem aiming to minimize makespan. Their method outperformed traditional MIP models when looking at computational time and their ability to reach the most optimal solutions. Similarly, Bragin et al. (2021) used the Lagrangian Relaxation method to address the job-shop scheduling problem in a highly dynamic setting with a low number of jobs. With the use of this method, they reduce computational time and achieve faster convergence. Their technique achieves near-optimal solutions faster compared to commercial solvers.

Decomposition techniques, including rolling horizon procedures, are useful and important methods for solving complex and large scheduling problems. By breaking down a large problem into smaller parts, they make it more practical and easier to solve.

2.9. Solution Methods

The flow shop problem has been extensively studied, and lots of solution methods have been developed over the years. One of the earliest and simplest methods is Johnson's algorithm, introduced by Johnson (1954), which provides an optimal solution for two-machine flow shop problems. The algorithm is straightforward and efficient, sequencing jobs to minimize makespan based on specific criteria for task ordering. Despite its low complexity, Johnson's algorithm remains relevant in modern applications for solving small-scale or specific cases of the flow shop problem, showing its practicality in scheduling research.

Dispatching Rules

Using dispatching rules is a common approach for handling scheduling problems, often employed due to their simplicity and computational efficiency. These rules prioritize jobs dynamically and are rule-based, making them useful in environments requiring real-time decision-making. Common examples are Shortest Processing Time (SPT), First-In-First-Out (FIFO), and Earliest Due Date (EDD) (Meilanitasari and Shin, 2021). These static rules allocate priorities based on predefined criteria, such as the order of job arrival or the processing time required. More dynamic rules, such as Shortest Waiting Time (SWT) or Most Total Work Remaining (MTWR), adapt priorities based on current system states, resulting in better responsiveness to changing conditions (Zahmani et al., 2021).

Despite their effectiveness, dispatching rules focus on locally optimized decisions and often fail to find a global optimum for the objective, leading to suboptimal overall performance. For instance, relying only on SPT minimizes processing times at the machine level but may increase overall tardiness or imbalance workloads. As Zahmani et al. (2021) suggests, combining multiple rules or tailoring rules for individual machines can enhance performance, but this requires careful calibration and simulation to avoid local optima. While dispatching rules are suitable for small-scale or very dynamic systems, their limited scope and lack of holistic optimization often result in combining them with more advanced methods for complex scheduling environments.

Mathematical Models

Generally used methods in flow shop scheduling mathematical models, particularly Mixed Integer Linear Programming (MILP). MILP provides an exact optimization framework by formulating the scheduling problem using linear equations and inequalities, incorporating continuous and integer variables to represent scheduling decisions (Hillier and Lieberman, 2015). This approach is highly flexible and can implement a variety of strict constraints and objectives, making it suitable for complex environments like hybrid flow shops. In earlier research, other variants of integer programming approaches, such as pure Mixed Integer Programming (MIP) and Binary Integer Programming (BIP), have also been employed to tackle flow shop scheduling problems, often focusing on more narrowly defined or structurally simpler instances. Examples of this can be found in the later section, which gives an overview of the state-of-the-art research.

In flow shop scheduling, MILP models try to optimize objectives such as minimizing makespan, total earliness, or weighted tardiness (González-Neira et al., 2017). The decision variables typically include sequencing variables, which determine the order in which jobs are processed, timing variables, which define start and completion times of jobs on machines, and assignment variables, which are especially important in hybrid flow shops with multiple machines per phase, where binary variables assign jobs to specific machines.

Constraints in the MILP model establish feasible and practical schedules by enforcing machine capacity and processing order. Precedence relations can also be constrained if this applies, so a job cannot be allocated to the next machine before it completes the previous one. Additional constraints can include setup times, due dates, machine availability, and other practical considerations relevant to the specific scheduling environment (Thörnblad, 2013).

The strengths of MILP models are their ability to find the most optimal solutions, their flexibility with easy adaptation and implementation of various constraints and objectives, and their precision in accurately modeling complex scheduling scenarios. However, MILP models also have challenges like higher computational complexity, where solution times can increase exponentially with problem size, which can lead to scalability issues, as large-scale problems may become too difficult for exact methods to solve within a reasonable time (Karam et al., 2017). Additionally, these models have high data requirements, needing detailed and accurate input data for effective modeling of the problem.

To mitigate the issues experienced in MILP models, several strategies are used. Model simplification involves reducing the number of variables and constraints where possible to make the problem easier to solve, as is demonstrated in the research of Floudas and Lin (2005). Decomposition techniques break the problem into smaller, more manageable subproblems that are solved successively or simultaneously, of which an example can be found in Kunath et al. (2022). Hybrid methods combine MILP with metaheuristics like Genetic Algorithms or Tabu Search to balance solution quality and computational effort, using the best qualities of both exact and approximate optimization techniques. An example of this is given in the research of Hajji et al. (2023), where a mathematical model and a Tabu Search are combined to find a solution within a more efficient computational time.

In practical applications, MILP models have been successfully used in various industries and applications to improve scheduling efficiency and reduce operational costs. Innovations in optimization software and computing power continue to further improve their applicability, making MILP a valuable and promising method in both academic research and industrial practice.

Metaheuristics

Johnson's algorithm and dispatching rules are examples of basic heuristic solutions, more general metaheuristic approaches can be used to solve flow shop scheduling problems as well. Metaheuristics are higher-level, problem-independent optimization strategies designed to explore the solution space more extensively. They are adaptable and often used to find near-optimal solutions for complex or large-scale instances. The most-used metaheuristic methods to solve the flow shop scheduling problem, according to González-Neira et al. (2017), are given below:

- *Genetic Algorithm (GA)*: A widely used metaheuristic, inspired by natural evolution for solving flow shop scheduling problems due to its flexibility and strengths in exploring large solution spaces. For instance, in Yu et al. (2018), a GA was developed to minimize total tardiness in a hybrid flow shop environment with unrelated machines and machine eligibility constraints. The method incorporates a dynamic technique that balances workload across machines while maintaining tight schedules, showing a better performance in comparison with the state-of-the-art literature. The advantages of GA include its robustness and adaptability to different problem configurations. However, GA may require significant computational time for parameter tuning and large-scale problems.
- *Particle Swarm Optimization (PSO)*: A metaheuristic inspired by the group behavior of swarms. In Madenoğlu (2021), PSO was applied to minimize makespan in a hybrid flow shop problem with sequence-dependent setup and transportation times. The method uses forward scheduling and dynamic job allocation to improve performance, demonstrating significant advantages over other heuristics like GA and NEH in both quality and computational efficiency. The simplicity of PCO, fast convergence, and less need for parameter tuning make it highly effective, though it can sometimes give less optimal outcomes because of early convergence without the right initialization and parameter selection.
- *Simulated Annealing (SA)*: SA is another metaheuristic used for hybrid flow shop problems. In Hajji et al. (2024), a tailored SA was applied to minimize makespan in a hybrid flow shop with dedicated machines and delivery constraints. This approach demonstrated low deviations from the lower bound and strong performance on larger instances, particularly when combined with an effective cooling schedule and neighborhood operators. While SA performs well in exploring complex solution spaces, its performance heavily depends on careful parameter tuning and computational resources.
- *Ant Colony Optimization (ACO)*: ACO is a metaheuristic that copies the principles of ants searching for food, using trails to find the solution. In Qin et al. (2018), an improved ACO was developed for dynamic flow shop scheduling with varying processing times. Key improvements used in this research were a rolling rescheduling strategy to handle dynamic events and an adaptive path compression technique to improve computational efficiency. This method showed a strong performance in minimizing makespan and adapting to uncertainty, though ACO may require careful tuning of heuristic parameters to avoid too early convergence.
- *Tabu Search (TS)*: TS is a metaheuristic that uses memory to explore the solution space by avoiding previously visited solutions stored in a list. In Dodu and Ancău (2020), a TS was applied to the Permutation Flow Shop Problem, starting with an NEH heuristic-based initial solution and using adaptive distance to explore unexplored areas. This approach demonstrated strong performance in minimizing makespan, especially on medium to large problem instances. However, this method is also really dependent on the initial solution and needs detailed parameter settings.

In addition to the discussed methods, numerous other metaheuristics can be used to solve the challenging environment of flow shop scheduling. Furthermore, hybrid methods, which combine the approaches of multiple heuristics or metaheuristics, are popular for solving dynamic hybrid flow shop problems. These hybrid approaches, such as combining Genetic Algorithms with Simulated Annealing, leverage synergies between techniques to improve the solution quality and adaptability, particularly in real-world, dynamic applications.

Emerging and Alternative Approaches

Recent advancements in artificial intelligence, particularly reinforcement learning (RL), have introduced novel approaches to flow shop scheduling. Deep reinforcement learning (DRL) can model the problem as a Markov Decision Process (MDP). In this method, an agent can learn scheduling policies through trial and error. For example, Tassel et al. (2021) demonstrated that DRL can outperform traditional dispatching rules in dynamic environments, while X. Wang et al. (2024) highlighted its potential in flexible scheduling systems despite scalability challenges.

The method developed by Infantes et al. (2024), which combines Graph Neural Networks and Deep Reinforcement Learning techniques, shows that traditional OR methods still outperform their methods in finding the optimal solution, especially for the deterministic variant, but also for smaller instances of the stochastic problem. However, according to the results of Yan et al. (2022), the performance of Deep Reinforcement Learning is comparable to the metaheuristic method that is called Genetic Algorithms, being slightly better in terms of optimality within the same computational time.

To provide an indication of the wide variety of solution methods available within different domains, objectives that can be targeted, the number of dispatching rules, and the factors of uncertainty that models can account for, an overview is given in Figure 2.1. This figure illustrates the research framework, categorizing the approaches into domains, methods, and objectives. The domain differentiates between optimization and prediction focuses, the method categorizes the techniques into metaheuristics, mathematical modeling, and machine learning, and the objective domain highlights performance indicators, including single and multi-objective optimization, often concerning time, cost, or environmental factors.

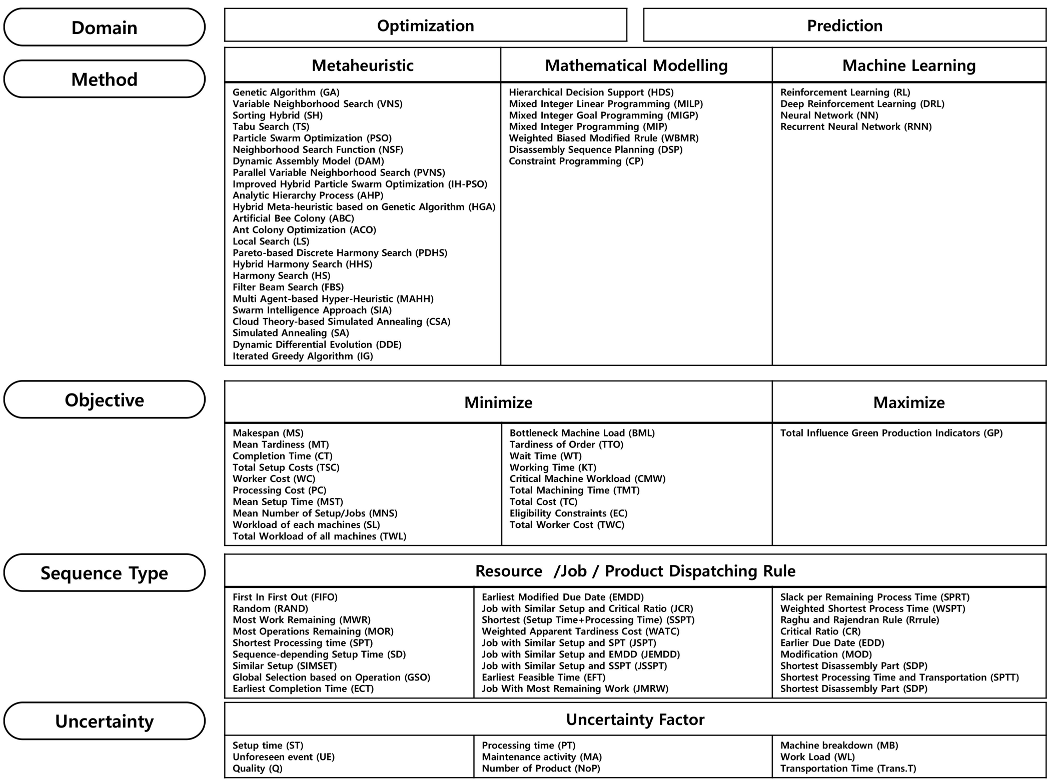

Figure 2.1: Research framework illustrating methods and objectives in sequence-driven scheduling (Meilanitasari and Shin, 2021).

2.10. Overview of State-of-the-Art Research

As previously discussed, the scheduling problem in CMRO shops can be approached using various models and methods, all with different combinations of characteristics and techniques. To understand the current landscape and identify potentially useful methodologies for our model, an extensive review of the most relevant state-of-the-art research was conducted. Table 2.1 provides an overview of various papers, comparing key characteristics and elements that align with the components of this problem or techniques potentially implemented in this research.

Each column in the table represents a critical aspect of the scheduling problem or a theoretical approach that is potentially relevant to our research. Problem classification categorizes the type of scheduling problem addressed in each paper, such as flow shop, job shop, or hybrid flow shop, indicating the operational environment modeled. The problem researched in this thesis is the hybrid flow shop problem. The solution method is the technique used to solve the scheduling problem, including MILP, heuristics, metaheuristics, and reinforcement learning approaches, which can be categorized into exact or approximate solutions. The objectives are the optimization goal or performance measure optimized in the study. The weighted job priority denotes whether the study incorporates job priorities with weights to reflect their importance in the scheduling process. Skilled technician assignment indicates if the assignment of jobs to technicians with specific skills is considered, which means that the restriction of technicians who can work on specific components is implemented. Dedicated technicians or machines indicate if the study takes into account the option of whether certain technicians or machines are dedicated to specific tasks or jobs, resulting in extra constraints for the assignment process. If the study takes into account the uncertainties of variable processing times or unexpected job arrivals should also be included in the table. The dynamic properties, as discussed in section 2.6, indicate if the scheduling approach adapts to changing variables over time, such as rescheduling in response to event-driven changes. As discussed in section 2.7, robust optimization uses optimization techniques to create solutions resilient to uncertainty and variability in the system. Rolling horizon or other decomposition techniques, explained in section 2.8, show whether the study uses methods to solve large scheduling problems by splitting them into smaller sub-problems.

The foundation of this research is the work by Tliba et al. (2023), which addresses a hybrid flow shop problem in a perfume manufacturing environment, incorporating a digital twin for dynamic scheduling. Their model integrates skilled technician constraints by defining a set of qualified machines for each job and stage, showing a lot of similarities to the CMRO problem explored in this paper. Aribowo et al. (2020) provides a unique way of integrating dedicated technicians into a mathematical scheduling model, aligning with the technician-specific constraints of CMRO environments. Tighazoui et al. (2021) proposes a predictive-reactive strategy for minimizing waiting times and instability in dynamic flow shop rescheduling, suggesting innovative methods that lay the foundation for approaches to handle dynamic updates in this research. Roslöf et al. (2002) demonstrates a combination of MILP and heuristics for solving large-scale scheduling problems, presenting effective techniques for managing job arrivals dynamically. S. Zhang and Wang (2018) introduces methods to handle sequence-dependent setups and part sharing in flexible job shops, leveraging MILP and constraint programming for dynamic job arrivals. Elyasi and Salmasi (2013) implements stochastic modeling by using expected values of uncertain parameters and machine-based decomposition, providing ways to address uncertainty. Qin et al. (2018) highlights the application of a rolling horizon in combination with an improved ant colony algorithm for dynamic scheduling in hybrid flow shops. Jamili (2016) demonstrates robust scheduling techniques using direct buffer formulations, while Fathollahi-Fard et al. (2024) explores robust optimization by incorporating stability coefficients in their model. Finally, Avelino et al. (2016) provides a CMRO-specific study that, while simple in its dispatching rule application, includes aspects that are highly relevant to real-world shop-floor challenges. These state-of-the-art papers are some examples from Table 2.1 that form the foundation for the framework and methodologies used in this research.

Moreover, in Table 2.1, one of the first aspects that stands out is that many studies focus on optimizing traditional objectives like minimizing makespan or total tardiness, and a minority of the papers used incorporated weighted job priorities. For instance, Roslöf et al. (2002), Fan et al. (2021), and Tighazoui et al. (2021) implemented job priorities in their models, indicating a potential gap in the literature where the priority weights of jobs are not extensively researched in different settings.

Table 2.1: Overview of the Bibliography on Scheduling Problems

Paper	Problem Classification	Solution Method	Exact or Approximate Solution	Objective	Weighted Job Priority	Skilled Technician Assignment	Dedicated Technicians /Machines	Uncertain Processing Times	Unexpected Job Arrivals	Dynamic Properties	Robust Optimization	Rolling Horizon/ Decomposition
Roslöf et al., 2002	Industrial Scheduling Problem (Single-machine)	MILP and SUP heuristic	Exact	min$(C_{\max} + wT)$	x				x	x		
Elyasi and Salmasi, 2013	Flow Shop with Multiple Machines	LP/NLP models	Exact	min$(\#T)$						x		x
Gomes et al., 2013	Flexible Job Shop Environment	MILP-based Predictive-Reactive Scheduling	Exact	min$(E + T)$					x	x		
Avelino et al., 2016	Assignment and Scheduling Problem	BIP	Exact	min(T)		x			x	x		
Jamili, 2016	Job Shop	MIP/B&B/ BS/PSO	Exact & Approximate	min$(C_{\max})$				x			x	
Tang et al., 2016	Flexible Flow Shop	Improved Ant Colony Algorithm	Approximate	min$(C_{\max})$		x		x	x	x		
Lv et al., 2017	Hybrid Flow Shop	Genetic Regulatory Network-Based Method	Approximate	min$(C_{\max})$				x		x		x
Qin et al., 2018	Hybrid Flow Shop	Improved Ant Colony Algorithm	Approximate	min$(C_{\max})$				x		x	x	x
Liu et al., 2018	Permutation Flow Shop	Iterated Greedy Algorithm	Approximate	min$(C_{\max})$					x	x	x	
S. Zhang and Wang, 2018	Flexible Assembly Job-Shop	MIP/CP/DR	Exact & Approximate	min$(C_{\max})$					x	x		
Aribowo et al., 2020	Flexible Flow Shop Scheduling	MILP	Exact	min$(C_{\max})$			x					
Schumacher et al., 2020	Hybrid Flow Shop	Tabu Search/ Local Search	Approximate	min$(C_{\max})$		x		x			x	
Tassel et al., 2021	Job Shop	Reinforcement Learning	Approximate	min$(C_{\max})$								
Fan et al., 2021	Job Shop	Genetic Programming-Based Hyper Heuristic	Approximate	min(Avg T)	x	x				x	x	
Tighazoui et al., 2021	Flow Shop	MILP/ Heuristics	Exact & Approximate	min(wCP)	x				x	x		
X. Zhang et al., 2021	Hybrid Flow Shop	Agent-Based Approach	Approximate	min$(C_{\max})$				x	x	x		
Tliba et al., 2023	Flow Shop with Multiple Processors	MILP - Digital Twin	Exact	min$(D_{\max} + Z)$		x		x	x	x		
Fathollahi-Fard et al., 2024	Permutation Flow Shop	Adaptive Large Neighborhood Search	Approximate	min(Cost)				x	x		x	

Secondly, the assignment of skilled technicians is not always integrated into most scheduling models. Only some studies, such as Avelino et al. (2016), Tang et al. (2016), Schumacher et al. (2020), Fan et al. (2021), and Tliba et al. (2023), are specifically considering skilled technician assignment in the right context, showing that the human resource dimension with skill constraints is often not taken into account in research done in the past. Additionally, the concept of dedicated technicians is rarely taken into account in past studies; only one study could be found that researches this application, the research done by Aribowo et al. (2020), which indicates the potential innovations that can be found regarding this technique. Moreover, when considering complex constraints like skilled technicians and dedicated technicians in past research, most of the time, exact modeling approaches are used, which could indicate that this is the best-fitting approach for such environments.

Moreover, only a limited number of studies address uncertain processing times and unexpected job arrivals together, which are essential characteristics for modeling the stochastic nature of CMRO operations. Studies like Jamili (2016), Tang et al. (2016), and Schumacher et al. (2020) integrated techniques to handle uncertain processing times but do not integrate robust optimization techniques or focus on dynamic scheduling.

Furthermore, the integration of robust optimization and decomposition techniques like rolling horizon is not used much in the existing literature. For example, Fathollahi-Fard et al. (2024) employs robust optimization in a permutation flow shop but does not consider skilled technician assignment or weighted job priority.

2.11. Research Gap

The observations made about the state-of-the-art research overview, regarding Table 2.1, show a clear gap in the current research, which is that there are no studies, to our current knowledge, that simultaneously address weighted job priorities, skilled and dedicated technician assignment, variable processing times, dynamic job arrivals, and robust optimization within an exact solution framework. This gap emphasizes the need for an extensive scheduling model that integrates these aspects, tailored to the complex and dynamic environment of CMRO shops.

Second, current literature tends to focus on one or two of these aspects applicable to CMRO environments. For example, approaches may incorporate uncertainty in processing times or consider dynamic job arrivals, but they rarely also include technician skill constraints or robust optimization methods. Moreover, the application of the dedicated technician concept, although a highly crucial factor in many real-world operations, such as the CMRO operations, remains unexplored in flow shop and hybrid flow shop scheduling models.

Furthermore, while CMRO environments present a good example of the combination of complex, dynamic, and uncertain conditions, most existing research has not fully captured these environments. Many state-of-the-art models address either deterministic and static industrial environments or use randomly generated instances, which are less complex to model and do not validate the results in real-life scenarios.

Moreover, while there have been advances in modeling uncertain and dynamic conditions, especially through approximate solution methods, few researchers have combined dynamic properties with exact models to handle variable processing times and unexpected job arrivals. This is especially relevant for highly dynamic and complex environments, such as CMRO shops, where frequent disruptions and changing task requirements create substantial operational challenges.

In summary, there is a clear need for an exact optimization model that combines weighted job priorities, multi-skilled and dedicated technicians, variable processing times, dynamic job arrivals, and robust optimization principles. Addressing this gap will contribute significantly to both academic research and practical applications in CMRO operations, providing a complete solution to the complex scheduling challenges faced in these environments.

3
Case Study

To implement and test the developed model in this research, a case study at the CMRO shops at the maintenance provider will be conducted. The CMRO shops of the maintenance provider ensure that the operational reliability of various aviation components is maintained. Current scheduling practices in the CMRO shops show several challenges that increase the complexity of achieving operational efficiency, customer satisfaction, and on-time delivery performance. These challenges are more complex to solve due to the diverse characteristics of their three CMRO shops. This chapter discusses the limitations of the current system, outlines the specific environment of each shop, and underlines key performance metrics, forming the foundation for developing an advanced scheduling model.

3.1. Current Challenges and Inefficiencies

Currently, scheduling processes in CMRO shops are mostly static and rule-based, depending heavily on manual decision-making. The allocation of work orders (WOs) is decided by production leaders based on their experience, with almost no use of assignment tools or models. While 80% of technician capacity is scheduled at the beginning of the week, the remaining 20% is handled reactively, often leading to inefficiencies and delays.

The causes of the sub-optimality of the current way of scheduling include:

- Limited Flexibility: The current system struggles to adjust to dynamic changes such as unforeseen delays, urgent job arrivals, and variations in processing times.
- Suboptimal Resource Allocation: Technician skills and equipment capabilities are not fully taken into account, leading to lower utilization and potentially lower service levels.
- Low On-Time Delivery (OTD) Rate & Turnaround Time (TAT): In 2024, the average OTD rate was only 82.4%. Meanwhile, the average TAT varied significantly across shops, at 32 days for IDG, 51 days for Avionics, and 60 days for the P&H shop.
- Inefficient Prioritization: The prioritization of WOs is based on a basic points-based approach, based on customer importance levels and priority levels such as Aircraft On Ground (AOG), critical, and routine labels, but this approach does not utilize an advanced framework to dynamically schedule on these parameters.
- Semi-Manual Decision-Making: Decision-making relies on a combination of a simple prioritization model and manual decision-making, without making use of advanced, data-driven methods to improve and optimize the scheduling of maintenance operations.

3.2. Shop-Specific Environments and Routing Characteristics

Each CMRO shop at the maintenance provider has unique characteristics and constraints, which complicate the scheduling process and need customization in the scheduling model. Figures 3.1 and 3.3

illustrate the routing steps for the Hydraulics & Pneumatics (P&H) and Avionics shops, and the Power Generation (IDG) shop, respectively.

3.2.1. Hydraulics and Pneumatics (P&H)

The P&H shop specializes in hydraulic and pneumatic components, operating through two workgroups. Technician specialization is critical, with tasks often assigned to the same technician for the entire process. However, tasks can be reassigned if necessary, depending on part availability. Next to the routing steps, of which an overview is given in Figure 3.1, the current way of scheduling is described in Figure 3.2. The list of workable WOs, that are prioritized by the simple points-based model, is used by the production leads to assign technicians to those most critical jobs, based on their knowledge of the individual technician skills. Based on their indication of the processing times, they choose the number of WOs to be handled by each technician. The technicians each get an overview of the WOs assigned to them on their overview screen.

Important aspects shop environment:

- Preferred Technician Assignment: Tasks are ideally handled by the same technician who performed the initial inspection.
- Customer Prioritization: High-priority customers demand faster turnaround, but the current system does not dynamically accommodate such needs.
- Skill Requirements: Specific parts require technicians with unique capabilities, limiting scheduling flexibility.

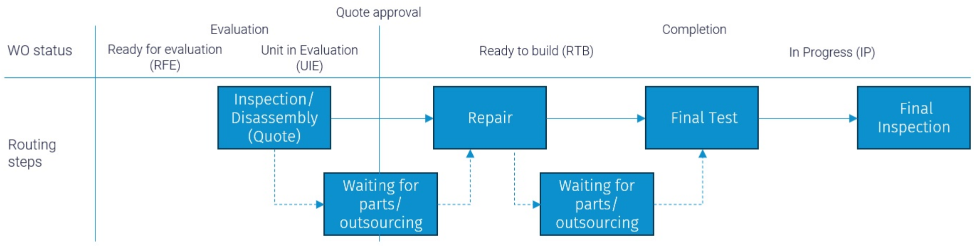

Figure 3.1: Routing steps and WO statuses for H&P and Avionics shop.

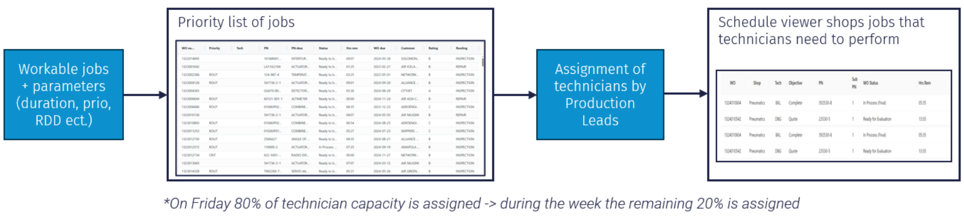

Figure 3.2: Current scheduling process for H&P and Avionics shop.

3.2.2. Avionics

The Avionics shop is divided into multiple workgroups, each specializing in distinct components such as displays and control units. Jobs requiring specialized testing equipment are often batched for efficiency. Also, the routing steps of this shop, as well as the way of scheduling, are almost identical to the P&H shop, shown in Figures 3.1 and 3.2.

Important aspects shop environment:

- Batching: Jobs requiring similar testing can benefit from being batched, improving the processing time on certain machines because no setup time is necessary.
- Workgroup Specialization: Each workgroup focuses on a specific domain, requiring precise allocation of tasks.

Additionally, each subshop has specific conditions that must be taken into account. For instance, one subshop may require that only one WO is completed per day, while another prioritizes a specific part above all others. In some cases, certain technicians are preferred for specific groups of parts within a subshop. These factors indicate examples of the complexities of the avionics shop environment.

3.2.3. Power Generation (IDG)

The Power Generation shop presents the most complex workflow, with a significantly higher number of routing steps and equipment limitations. In Figure 3.3, an overview of the current way of scheduling for the power generation shop is given. Similar to the other shops, the WO is ordered based on its priority, but separately for each routing step. WOs are moved sequentially through the steps, with job assignments handled by the Production Lead. The jobs in each step are represented as placeholders, indicating the WOs currently in progress at that stage. With this approach, WOs are tracked and progressed in the right way, and this creates a good overview of where certain jobs are in the process.

Important aspects shop environment:

- Equipment Limitations: Limited availability of critical testing machines creates bottlenecks in the process. For example, there are only two test machines available for incoming tests and final inspections.
- Unknown Durations: Testing outcomes and durations are often unpredictable, complicating scheduling. There are standard durations that are indications of the processing times of certain repairs or part numbers, but these are most of the time not accurate.
- Flexible Resource Allocation: Technicians can work across multiple work centers, but capacity limits must be respected.
- Technician Flexibility: Technicians are encouraged to handle entire workflows but may need reassignment in exceptional cases.

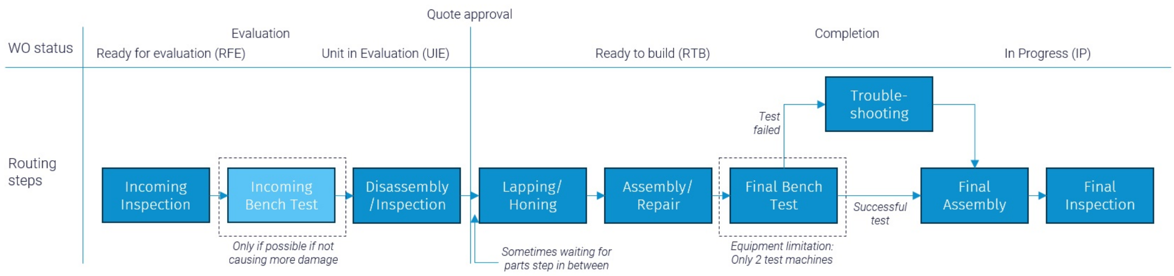

Figure 3.3: Routing steps and WO statuses for Power Generation shop.

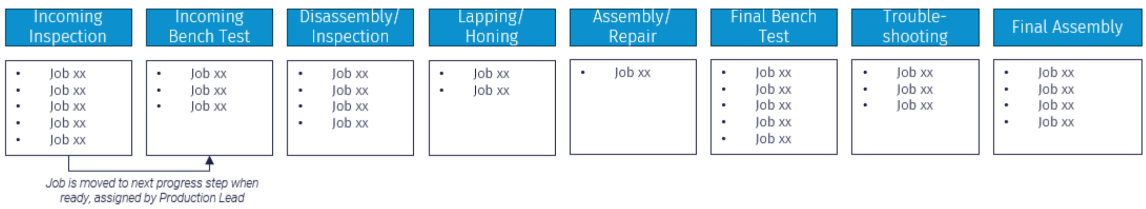

Figure 3.4: Current scheduling process for the Power Generation shop.

Several practical constraints and preferences further complicate the scheduling environment in the Power Generation shop. First, certain specialized operations, such as lapping and honing, may occur at different stages, sometimes during the quotation phase, and sometimes during the repair phase. A specialized technician ideally performs these operations, or when not available, another capable, right-skilled technician can step in if necessary. Notably, one technician can operate two of such machines together, increasing capacity but also creating further complexity in incorporating correct resource allocation algorithms in the model.

For the testing phase, failure of the final test requires that the WO be re-inserted into the system and re-sequenced for repair. This reintroduction of tasks increases the variability in job flow. Furthermore, a small buffer before the final due date is preferred to avoid tasks finishing exactly at the deadline. Frequent updates to the schedule should be limited to twice per day to help process dynamic changes, such as unforeseen delays or urgent customer requests, while making sure that once a job step has been initiated, it is carried through to completion without interruption.

Additionally, some of the shop characteristics of all three shops are identical, for instance, the rule that, if possible, a technician is assigned to each routing step of a WO where he has started working on, which will be important to incorporate in a model in the earlier stages, because of the application in all three shops. Additionally, the priority of certain customers, which was explained in subsection 3.2.1, is a characteristic that is important for all of the CMRO shops of the maintenance provider. An overview of the aspects, with the indication of importance to the shops, is given in Table 3.1.

Shop	**Uncertain Processing Times**	**Equipment Limitations**	**Preferred Technician Assignment**	**Customer Prioritization**	**Batching Jobs**	**Specialized Workgroups**	**Dynamic Routing**	**Technician Skill Requirements**	**Flexible Resource Allocation**
Hydraulics & Pneumatics (P&H)	✓	✓	✓	✓	✓	✓		✓	
Avionics	✓	✓	✓	✓	✓	✓		✓	
Power Generation (IDG)	✓	✓	✓	✓	✓		✓	✓	✓

Table 3.1: Shop aspects and their applicability across CMRO shops.

3.3. Key Performance Indicators and Performance

The effectiveness of scheduling practices is measured using several KPIs, including OTD rates, average TAT, and service levels for AAA customers. Performance metrics that will be used and the room for improvement emphasized:

- On-Time-Delivery (OTD) rate: Approximately 82.4% of WOs meet the Promised Delivery Date (PDD), which could be significantly improved. The PDD is determined based on the type of contract with the customer. The desired target for the OTD rate is 90%.
- Average TAT: The P&H shop averages 60 days, while Avionics averages 51 days. The IDG shop performs relatively better at 32 days but still fails to meet the 30-day target. This KPI can be measured in different ways, from the TAT for orders that only need to be quoted to the TAT after the quote approval by the customer. For example, the TAT for quotes of incoming parts is desired for the Power Generation shop to be 7 days.
- Service Level for High Priority customers: This KPI focuses on ensuring high-priority customers, such as AAA-customers, achieve a higher OTD rate, almost always achieving the overall target of 90%.

The proposed model aims to bring the performance of the shops closer to achieving the targets set by the maintenance provider. While no specific minimum performance improvement objectives have been set for the model, its primary goal is to achieve the targets of a 90% OTD rate and an average TAT of 30 days.

3.4. Implementation of Proposed Model

This case study implements the proposed theory to test the optimized and dynamic scheduling model tailored to the diverse needs of the CMRO shops at the maintenance provider. By integrating state-of-the-art Operations Research (OR) techniques and real-time updates, the model aims to improve resource allocation, reduce TAT, and enhance OTD rates. In Figure 3.5, the scheduling model that includes all of the features important for the shops is given schematically. In this model, all important parameters like job priority, process times, and due dates are input into the model, and an initial schedule is created. Additionally, rescheduling takes place every time a disruptive event, such as a change in the processing time of a WO in WIP, triggers this process.

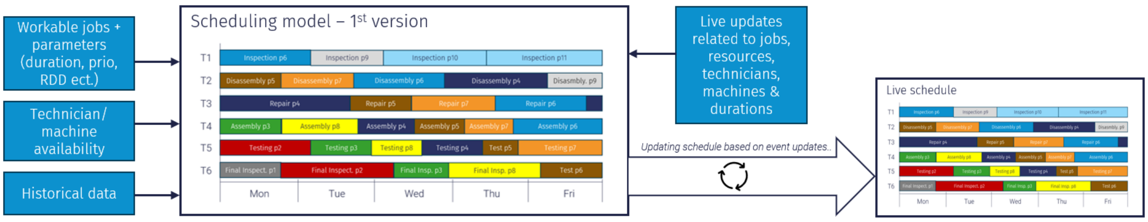

Figure 3.5: Overview of the future aspired scheduling model example for the Power Generation shop.

This project aims to bridge the gap between theoretical optimization models, such as flow shop theories, and practical applications in CMRO environments, contributing to the wider field of scheduling and operational research.

4 Research Proposal

Throughout the project, minor adjustments were made to the research question and sub-questions to reflect minor changes in the methodological approach and the research scope. These refinements have been incorporated into this chapter to describe an accurate research plan.

4.1. Research Question

The goal of this research is to develop an automated, data-driven, and flexible scheduling model that improves the operational performance of CMRO shops. Considering the complexities such as flexible job shop scheduling, unexpected job arrivals, varying technician skill requirements, the need to prioritize high-value customers, and other challenging aspects of the shop floor environment, the following main research question was formulated:

How can a flexible job shop scheduling model be developed and implemented to autonomously plan the maintenance operations for a dynamic CMRO environment, thereby improving the key performance indicators Turn-Around-Time (TAT) and On-Time Delivery (OTD)?

4.2. Sub-Questions

Several sub-questions are stated in this subsection to address the main research question thoroughly. Each sub-question aims to address a specific aspect of the main problem, ensuring that all critical aspects, constraints, and objectives are considered:

1. **Objective and Performance Evaluation:**
   (a) To what extent can a mathematical optimization model outperform a manual, experience-based scheduling approach in terms of improvements in TAT and OTD?
   (b) How can the multi-objective formulation, specifically combining weighted tardiness and weighted earliness, be integrated into the scheduling model, and how does this objective compare to other alternative objectives in terms of their impact on the defined KPIs?
2. **Customer and Job Priorities:**
   (a) How can weighted job priorities, combining customer importance and job urgency, be integrated into the hybrid flow shop scheduling model?
   (b) What is the impact of incorporating priority weights on scheduling outcomes, especially concerning service level agreements (SLA) and gains for critical customers such as AAA customers?
3. **Resource Allocation and Skill Constraints:**
   (a) How can the model ensure that technicians are assigned to tasks that match their skills, without requiring continuous intervention from production leaders?

   (b) Can the model incorporate constraints related to dedicated technicians or specialized operations?
   (c) How can unavailability periods and other human factor constraints be integrated into the model to maintain realistic technician workloads and ensure optimal allocation?

4. **Uncertainty and Dynamic Properties:**
   (a) How can job arrivals, uncertain processing times, and dynamic routing steps be effectively incorporated into a real-time rescheduling model, and what is the impact on schedule performance and computational effort under a discrete-time rescheduling framework?
   (b) How can robust scheduling improve shipping date reliability under uncertainty in job processing times and the dynamic arrival of new work orders?

5. **Computational Feasibility and Scalability:**
   (a) Can exact optimization approaches solve real-world CMRO problem instances in a reasonable computational time, and are they suitable for real-world implementation?
   (b) How does the model's computational time scale with increasing problem size and complexity, and can strategies such as rolling horizon planning or other decomposition techniques reduce computational effort without significantly compromising solution quality?

6. **Comparative Analysis Across Different Shops:**
   (a) How does the proposed model handle the operational differences between CMRO shops, and is the model easily adaptable to each unique shop environment?
   (b) To what extent are specific model adjustments needed for shops with more complex workflows or limited critical resources, and what is the impact on the scheduling model performance for each specific shop?

7. **Heuristic approach:**
   (a) Can approximate methods or heuristic approaches provide high-quality solutions with acceptable optimality gaps?
   (b) How does the computational time of these heuristic methods compare to exact optimization approaches, and are they better suited for real-world CMRO environments?

These sub-questions aim to break down the main research question into structured and measurable components. By addressing each of these sub-questions, the goals of this research will be achieved by using a holistic approach that not only improves CMRO shop floor scheduling performance but also ensures that this is resilient to real-world uncertainties and scalable to adapt to various CMRO shop environments.

4.3. Hypotheses

In line with the findings from the state-of-the-art review given in chapter 2, it is hypothesized that a well-designed flexible job shop scheduling model, incorporating weighted job priorities, skill-based technician assignments, and robust handling of uncertainty, will outperform current semi-manual approaches in improving TAT and OTD rates. By implementing well-fitting objectives into a mathematical optimization framework, the resulting schedules are expected to calculate high-quality solutions that respect operational priorities and resource constraints.

Furthermore, by referring to the techniques discussed in existing research, it is expected that integrating stochastic or robust optimization methods, as well as appropriate rescheduling policies, will lead to improved schedule stability and adaptability under uncertain and dynamic conditions. While exact optimization methods could face computational challenges in large-scale, real-world scenarios, the literature suggests that decomposition techniques can be used to manage computational times, resulting in near-optimal solutions being found more efficiently even in large-scale, real-world scenarios. Otherwise, optionally, metaheuristics can deliver near-optimal solutions within acceptable computation times.

In summary, the hypothesis is that by using and implementing these well-supported approaches found throughout the current body of science, from skill-based personnel planning and priority weighting to dynamic rescheduling and hybrid solution methods, this research can achieve an automated, data-driven scheduling model that enhances operational performance for CMRO shops by reducing TAT and increasing OTD rates.

4.4. Planning of Research Activities

The planning of the research activities is illustrated in Figure 4.1. The Gantt chart outlines the phases and corresponding activities of the project, beginning with the literature review and research definition, followed by model development, testing, and implementation. The chart also highlights important milestones such as the mid-term review, green light meeting, and final thesis defense.

A one-week vacation is planned, as indicated in the chart. Additionally, the approximate flow shop model phase, which changed to a flexible job shop model during the project, implementing a meta-heuristic approach, is conditional and will only be executed if time allows. This decision will be evaluated during the mid-term review based on the progress made on the exact model phase up to that point.

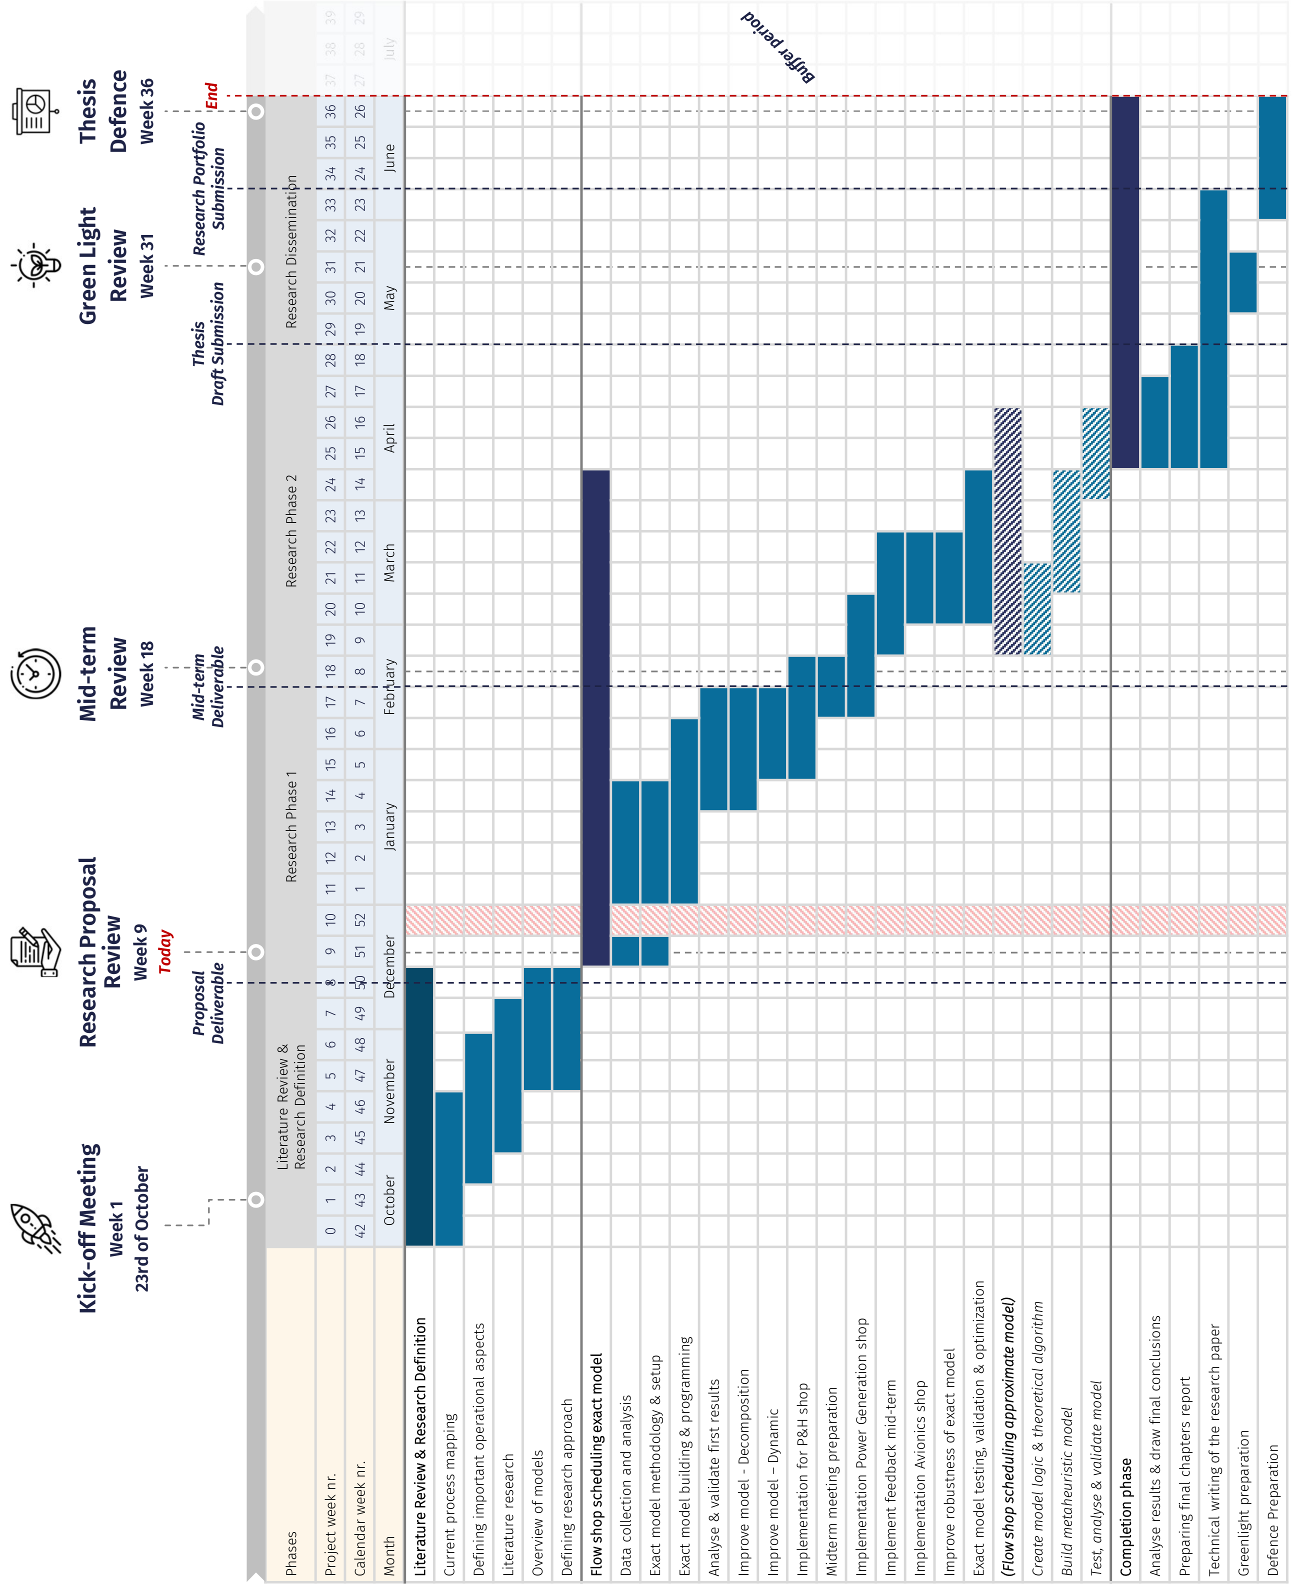

Figure 4.1: Gantt chart giving an overview of the planning of the Thesis.

References

Afshar-Nadjafi, B. (2021). Multi-skilling in scheduling problems: A review on models, methods, and applications. *Computers & Industrial Engineering*, *151*, 107004. https://doi.org/10.1016/j.cie.2020.107004

Annear, L. M., Akhavan-Tabatabaei, R., & Schmid, V. (2023). Dynamic assignment of a multi-skilled workforce in job shops: An approximate dynamic programming approach. *European Journal of Operational Research*, *306*, 1109–1125. https://doi.org/10.1016/j.ejor.2022.08.049

Aribowo, W., Prastomo, O. W., & Halim, A. H. (2020). Scheduling an aircraft maintenance shop with dedicated technician and dedicated machine constraints. In M. N. O. Zahid et al. (Eds.), *Advances in industrial and production engineering* (pp. 260–265). Springer Nature Singapore. https://doi.org/10.1007/978-981-15-0950-6_40

Avelino, C., Bourne, D., Ferreira, F., Rasteiro, D., & Santos, J. (2016). Scheduling the repair of aircraft components – a case study. *Mathematics-in-Industry Case Studies*, *7*(5), 1–18. https://doi.org/10.1186/s40929-016-0007-2

Bevelander, T. S. (2022). *Pro-active scheduling through the use of multi-agent systems in high-mix-low-volume shops* [Master's thesis, Eindhoven University of Technology]. https://research.tue.nl/en/studentTheses/3085caf3-c004-4ac3-9459-af6ce8c4c68b

Bragin, M. A., Wilhelm, M. E., Yu, N., & Stuber, M. D. (2021). Toward robust manufacturing scheduling: Stochastic job-shop scheduling. *Journal of Advanced Manufacturing and Processing*, *14*(8). https://doi.org/10.1109/TEM.2021.2785774

Cayo, P., & Onal, S. (2020). A shifting bottleneck procedure with multiple objectives in a complex manufacturing environment. *Production Engineering*, *14*, 177–190. https://doi.org/10.1007/s11740-019-00947-7

Ciro, G. C., Dugardin, F., Yalaoui, F., & Kelly, R. (2015). A fuzzy ant colony optimization to solve an open shop scheduling problem with multi-skills resource constraints. *IFAC-PapersOnLine*, *48*(3), 715–720. https://doi.org/10.1016/j.ifacol.2015.06.167

Damm, R. B., Chaves, A. A., Riveaux, J. A., & Ronconi, D. P. (2024). Scheduling technicians and tasks through an adaptive multi-objective biased random-key genetic algorithm. *Annals of Operations Research*. https://doi.org/10.1007/s10479-024-06325-6

Dodu, C. E., & Ancău, M. (2020). A tabu search approach for permutation flow shop scheduling. *Studia Universitatis Babe–Bolyai, Informatica*, *65*(1), 104–115. https://doi.org/10.24193/subbi.2020.1.08

Elyasi, A., & Salmasi, N. (2013). Stochastic flow-shop scheduling with minimizing the expected number of tardy jobs. *International Journal of Advanced Manufacturing Technology*, *66*, 337–346. https://doi.org/10.1007/s00170-012-4328-4

Emmons, H., & Vairaktarakis, G. (2013). *Flow shop scheduling: Theoretical results, algorithms, and applications* (Vol. 182). Springer Science+Business Media. https://doi.org/10.1007/978-1-4614-5152-5

Fan, H., Xiong, H., & Goh, M. (2021). Genetic programming-based hyper-heuristic approach for solving dynamic job shop scheduling problem with extended technical precedence constraints. *Computers & Operations Research*, *134*, 105401. https://doi.org/10.1016/j.cor.2021.105401

Fathollahi-Fard, A. M., Woodward, L., & Akhrif, O. (2024). A scenario-based robust optimization model for the sustainable distributed permutation flow-shop scheduling problem. *Annals of Operations Research*. https://doi.org/10.1007/s10479-024-05940-7

Floudas, C. A., & Lin, X. (2005). Mixed integer linear programming in process scheduling: Modeling, algorithms, and applications. *Annals of Operations Research*, *139*, 131–162. https://doi.org/10.1007/s10479-005-3446-0

Gomes, M. C., Barbosa-Póvoa, A. P., & Novais, A. Q. (2013). Reactive scheduling in a make-to-order flexible job shop with re-entrant process and assembly: A mathematical programming approach.

International Journal of Production Research, *51*(17), 5120–5141. https://doi.org/10.1080/00207543.2013.793428

González-Neira, E. M., Montoya-Torres, J. R., & Barrera, D. (2017). Flow-shop scheduling problem under uncertainties: Review and trends. *International Journal of Industrial Engineering Computations*, *8*(3), 399–426. https://doi.org/10.5267/j.ijiec.2017.2.001

Goren, S., & Sabuncuoglu, I. (2009). Optimization of schedule robustness and stability under random machine breakdowns and processing time variability. *IIE Transactions*, *42*(3), 203–220. https://doi.org/10.1080/07408170903171035

Hajji, M. K., Hadda, H., & Dridi, N. (2023). Makespan minimization for the two-stage hybrid flow shop problem with dedicated machines: A comprehensive study of exact and heuristic approaches. *Computation*, *11*(7), 137. https://doi.org/10.3390/computation11070137

Hajji, M. K., Hamlaoui, O., & Hadda, H. (2024). A simulated annealing metaheuristic approach to hybrid flow shop scheduling problem. *Advances in Industrial and Manufacturing Engineering*, *9*, 100144. https://doi.org/10.1016/j.aime.2024.100144

Hashimoto, H., Boussier, S., Vasquez, M., & Wilbaut, C. (2011). A grasp-based approach for technicians and interventions scheduling for telecommunications. *Annals of Operations Research*, *183*, 143–161. https://doi.org/10.1007/s10479-009-0545-0

Hillier, F. S., & Lieberman, G. J. (2015). *Introduction to operations research* (10th). McGraw-Hill Education.

Himmiche, S., Marangé, P., Aubry, A., & Pétin, J.-F. (2023). Robustness evaluation process for scheduling under uncertainties. *Processes*, *11*(2), 371. https://doi.org/10.3390/pr11020371

Hsieh, F. Y., Lin, Y. K., & Wang, Y.-C. (2024). Developing algorithm and dispatching rules for scheduling a realistic flexible flow shop with setups to minimize total weighted tardiness. *The International Journal of Advanced Manufacturing Technology*. https://doi.org/10.1007/s00170-024-14509-4

Infantes, G., Roussel, S., Pereira, P., Jacquet, A., & Benazera, E. (2024). Learning to solve job shop scheduling under uncertainty [Accepted for CPAIOR 2024]. *Proceedings of the 2024 CPAIOR Conference*. https://arxiv.org/abs/2404.01308

Jamili, A. (2016). Robust job shop scheduling problem: Mathematical models, exact and heuristic algorithms. *Expert Systems with Applications*, *55*, 341–350. https://doi.org/10.1016/j.eswa.2016.01.054

Johnson, S. M. (1954). Optimal two- and three-stage production schedules with setup times included. *Naval Research Logistics Quarterly*, *1*(1), 61–68. https://doi.org/10.1002/nav.3800010110

Karam, A., Attia, E.-A., & Duquenne, P. (2017). A milp model for an integrated project scheduling and multi-skilled workforce allocation with flexible working hours. *IFAC-PapersOnLine*, *50*(1), 13964–13969. https://doi.org/10.1016/j.ifacol.2017.08.2221

Kianpour, P., Gupta, D., Krishnan, K. K., & Gopalakrishnan, B. (2021). Automated job shop scheduling with dynamic processing times and due dates using project management and industry 4.0. *Journal of Industrial and Production Engineering*, *38*(7), 485–498. https://doi.org/10.1080/21681015.2021.1937725

Kress, D., Müller, D., & Nossack, J. (2019). A worker constrained flexible job shop scheduling problem with sequence-dependent setup times. *OR Spectrum*, *41*, 179–217. https://doi.org/10.1007/s00291-018-0537-z

Kumar, K., & He, M. (2023). How ai maintains manufacturing productivity amid reduced capex. *BCG X*. https://www.bcg.com/x/the-multiplier/how-ai-maintains-manufacturing-productivity-amid-reduced-capex

Kunath, S., Kühn, M., Völker, M., Schmidt, T., Rühl, P., & Heidel, G. (2022). Milp performance improvement strategies for short-term batch production scheduling: A chemical industry use case. *SN Applied Sciences*, *4*(87), 1–28. https://doi.org/10.1007/s42452-022-04969-2

Liu, W., Jin, Y., & Price, M. (2018). New meta-heuristic for dynamic scheduling in permutation flowshop with new order arrival. *International Journal of Advanced Manufacturing Technology*, *98*, 1817–1830. https://doi.org/10.1007/s00170-018-2171-y

Lv, Y., Zhang, J., & Qin, W. (2017). A genetic regulatory network-based method for dynamic hybrid flow shop scheduling with uncertain processing times. *Applied Sciences*, *7*(1), 23. https://doi.org/10.3390/app7010023

Madenoğlu, F. S. (2021). Solving optimization problem with particle swarm optimization: Solving hybrid flow shop scheduling problem with particle swarm optimization algorithm. In B. A. Mercangöz

(Ed.), *Applying particle swarm optimization* (pp. 264–273). Springer Nature Switzerland. https://doi.org/10.1007/978-3-030-70281-6_14

Meilanitasari, P., & Shin, S.-J. (2021). A review of prediction and optimization for sequence-driven scheduling in job shop flexible manufacturing systems. *Processes*, *9*(8), 1391. https://doi.org/10.3390/pr9081391

Ouelhadj, D., & Petrovic, S. (2009). A survey of dynamic scheduling in manufacturing systems. *Journal of Scheduling*, *12*, 417–431. https://doi.org/10.1007/s10951-008-0090-8

Pinedo, M. L. (2016). *Scheduling: Theory, algorithms, and systems* (5th). Springer. https://doi.org/10.1007/978-3-319-26580-3

Qin, W., Zhang, J., & Song, D. (2018). An improved ant colony algorithm for dynamic hybrid flow shop scheduling with uncertain processing time. *Journal of Intelligent Manufacturing*, *29*, 891–904. https://doi.org/10.1007/s10845-015-1144-3

Rahmani, D., & Heydari, M. (2014). Robust and stable flow shop scheduling with unexpected arrivals of new jobs and uncertain processing times. *Journal of Manufacturing Systems*, *33*(1), 84–92.

Roslöf, J., Harjunkoski, I., Westerlund, T., & Isaksson, J. (2002). Solving a large-scale industrial scheduling problem using milp combined with a heuristic procedure. *European Journal of Operational Research*, *138*(1), 29–42. https://doi.org/10.1016/S0377-2217(01)00140-0

Schumacher, C., Buchholz, P., Fiedler, K., & Gorecki, N. (2020). Local search and tabu search algorithms for machine scheduling of a hybrid flow shop under uncertainty. In K.-H. Bae, B. Feng, S. Kim, S. Lazarova-Molnar, Z. Zheng, T. Roeder, & R. Thiesing (Eds.), *Proceedings of the 2020 winter simulation conference* (pp. 1456–1467). IEEE. https://doi.org/10.1109/WSC48552.2020.9383982

Schworm, P., Wu, X., Klar, M., Glatt, M., & Aurich, J. C. (2023). Multi-objective quantum annealing approach for solving flexible job shop scheduling in manufacturing. *Journal of Intelligent Manufacturing*. https://doi.org/10.1007/s11740-023-01234-7

Tan, Y., & Terekhov, D. (2018). Logic-based benders decomposition for two-stage flexible flow shop scheduling with unrelated parallel machines. *Canadian Conference on Artificial Intelligence*, 56–68. https://link.springer.com/chapter/10.1007/978-3-319-89656-4_5

Tang, D., Dai, M., Salido, M. A., & Giret, A. (2016). Energy-efficient dynamic scheduling for a flexible flow shop using an improved particle swarm optimization. *Computers in Industry*, *81*, 82–95. https://doi.org/10.1016/j.compind.2015.10.001

Tassel, P., Gebser, M., & Schekotihin, K. (2021). A reinforcement learning environment for job-shop scheduling. *arXiv preprint arXiv:2104.03760*. https://arxiv.org/abs/2104.03760

Thörnblad, K. (2013). *Mathematical optimization in flexible job shop scheduling: Modelling, analysis, and case studies* [Doctoral dissertation]. https://doi.org/10.13140/RG.2.2.24918.86080

Tighazoui, A., Sauvey, C., & Sauer, N. (2021). Predictive-reactive strategy for flowshop rescheduling problem: Minimizing the total weighted waiting times and instability. *Journal of Systems Science and Systems Engineering*, *30*(3), 253–275. https://doi.org/10.1007/s11518-021-5490-8

Tliba, K., Diallo, T. M., Penas, O., Ben Khalifa, R., Ben Yahia, N., & Choley, J.-Y. (2023). Digital twin-driven dynamic scheduling of a hybrid flow shop. *Journal of Intelligent Manufacturing*, *34*, 2281–2306. https://doi.org/10.1007/s10845-022-01922-3

van der Weide, T., Deng, Q., & Santos, B. F. (2022). Robust long-term aircraft heavy maintenance check scheduling optimization under uncertainty. *Computers Operations Research*, *141*, 105667. https://doi.org/10.1016/j.cor.2021.105667

Wang, S., Liu, C.-a., Wang, H.-y., & Li, Y. (2013). Dynamic job-shop scheduling research based on rolling horizon technology. *Proceedings of the 20th International Conference on Industrial Engineering and Engineering Management*, 1045–1053. https://doi.org/10.1007/978-3-642-40063-6_102

Wang, X., Zhong, P., Liu, M., Zhang, C., & Yang, S. (2024). A novel method-based reinforcement learning with deep temporal difference network for flexible double shop scheduling problem. *Scientific Reports*, *14*, 9047. https://doi.org/10.1038/s41598-024-59414-8

Wang, Z., Zhang, J., & Si, J. (2020). Dynamic job shop scheduling problem with new job arrivals: A survey. In Z. Deng (Ed.), *Advances in artificial intelligence and computation* (pp. 664–671). Springer. https://doi.org/10.1007/978-981-32-9050-1_75

Xiao, S., Sun, S., & Jin, J. (2017). Surrogate measures for the robust scheduling of stochastic job shop scheduling problems. *Energies*, *10*(4), 543. https://doi.org/10.3390/en10040543

Yan, Q., Wu, W., & Wang, H. (2022). Deep reinforcement learning for distributed flow shop scheduling with flexible maintenance. *Machines*, *10*(3), 210. https://doi.org/10.3390/machines10030210

Yu, C., Semeraro, Q., & Matta, A. (2018). A genetic algorithm for the hybrid flow shop scheduling with unrelated machines and machine eligibility. *Computers and Operations Research*, *100*, 211–229. https://www.sciencedirect.com/science/article/pii/S030505481830217X

Zahmani, M. H., Atmani, B., Bekrar, A., & Aissani, N. (2021). Multiple priority dispatching rules for the job shop scheduling problem. *Processes*, *9*(8), 1392. https://doi.org/10.3390/pr9081392

Zhang, S., & Wang, S. (2018). Flexible assembly job-shop scheduling with sequence-dependent setup times and part sharing in a dynamic environment: Constraint programming model, mixed-integer programming model, and dispatching rules. *IEEE Transactions on Engineering Management*, *65*(3), 487–500. https://doi.org/10.1109/TEM.2017.2785774

Zhang, X., Sun, X., Liu, Y., Zhang, C., Zhao, L., & Shen, W. (2021). An agent-based approach for dynamic scheduling in hybrid flow shops. *Proceedings of the 2021 IEEE 24th International Conference on Computer Supported Cooperative Work in Design (CSCWD)*, 24–30. https://doi.org/10.1109/CSCWD49262.2021.9437879