Exact and heuristic approaches for the PDPTW-SE and the SMCTSP *

Vítor A. Barbosa[1] , **Rafael A. Melo**[2]

[1] Mobility and Industrial Management – KU Leuven
Leuven, 3001 – Belgium

[2]Instituto de Computação – Universidade Federal da Bahia (UFBA)
Salvador, 40170-115 – BA – Brazil

vitor.alvesbarbosa@kuleuven.be, rafael.melo@ufba.br

Abstract. *This thesis studies exact and heuristic optimization methods for highly constrained routing and scheduling problems. We focus on problems where operational decisions tightly couple routing and temporal coordination, making them particularly challenging from a combinatorial optimization perspective. The first contribution introduces the pickup and delivery problem with time windows and scheduling on the edges (PDPTW-SE), a new problem that integrates vehicle routing and machine scheduling decisions in scenarios where vehicles require scheduled machine assistance to overcome inter-region obstacles, such as the straits between islands or the vertical separation in multi-floor buildings. We propose an arc-based mixed-integer programming (MIP) strengthened by preprocessing and valid inequalities, together with a multi-start heuristic with an LP-based improvement procedure (MSLP). We also introduce a benchmark set representing different application scenarios. Computational experiments show that the strengthened MIP significantly improves solution quality, while the heuristic scales to larger tested instances for which the exact approach did not find feasible solutions within the time limit. The second contribution addresses the single-machine coupled-task scheduling problem (SMCTSP) to minimize the makespan. We propose a constraint programming (CP) model and a biased random-key genetic algorithm (BRKGA) enhanced with restarts, perturbation strategies, and local search. Experiments on benchmark instances demonstrate state-of-the-art performance: the CP model obtained the best-known solutions for 163 out of the 180 hardest instances, while the BRKGA achieved superior solution quality under short time limits. The results of this research have already led to publications in the Brazilian Symposium of Operational Research (SBPO) and the European Journal of Operational Research (EJOR), reinforcing the scientific relevance of the thesis.*

1. Introduction

Many real-world operations require the coordination of routing decisions with tightly constrained temporal processes, such as scheduling shared resources or machines. These interactions create complex combinatorial optimization problems that are difficult to solve

*The present document summarizes the master's thesis of Vítor Alves Barbosa [Barbosa 2025], whose original title is *"Exact and heuristic approaches for the pickup and delivery problem with time windows and scheduling on the edges, and for the single-machine coupled task scheduling problem with exact delays"*.

using traditional routing or scheduling methods in isolation. This thesis studies exact and heuristic approaches for such highly constrained routing and scheduling problems, focusing on models that integrate routing decisions with scheduling constraints.

The first problem introduces a generalization of the well-known pickup and delivery problem with time windows (PDPTW), which we refer to as the pickup and delivery problem with time windows and scheduling on the edges (PDPTW-SE). The PDPTW-SE could be applied, for instance, in the logistical planning of multiple islands or multiple-floor buildings or in post-disaster relief operations.

The second problem considers the single-machine coupled-task scheduling problem (SMCTSP) to minimize the makespan. The SMCTSP is part of the coupled-task scheduling problems (CTSPs), which encompasses several applications such as pulsed radar systems [Shapiro 1980], chemistry manufacturing [Ageev and Baburin 2007], and patient appointment scheduling in medical services [Pérez et al. 2011, Condotta and Shakhlevich 2014, Liu et al. 2019].

The main goal of this thesis is to develop and evaluate exact and heuristic methods for the PDPTW-SE and the SMCTSP problems. For the PDPTW-SE, we aim to formally introduce the problem and provide initial methods for evaluating its applicability in different scenarios by conducting computational experiments on a benchmark set of instances. For the SMCTSP, we aim to propose and develop robust methods to achieve high-quality solutions that outperform or at least match the current state-of-the-art approaches.

1.1. Main contributions

The main contributions of this thesis are:

- **Introduction of a new optimization problem**: the PDPTW-SE, which integrates vehicle routing and machine scheduling decisions in constrained scenarios.
- **Exact optimization for PDPTW-SE**: an arc-based MIP formulation strengthened with preprocessing and valid inequalities.
- **Scalable heuristic for PDPTW-SE**: a multi-start heuristic with LP-based improvement capable of solving instances significantly larger than those tractable by exact methods.
- **Benchmark dataset and reproducibility infrastructure**: a publicly available repository containing instances, algorithms, scripts, and experimental results.
- **State-of-the-art algorithms for SMCTSP**: a constraint programming model and an enhanced BRKGA achieve the best-known results on the majority of hard benchmark instances.

2. The PDPTW-SE

The PDPTW-SE extends the PDPTW by allowing pickup and delivery nodes to be located in different regions inaccessible to regular vehicles. To fulfill the requests, a specific type of machine must be employed to transport a vehicle between regions. However, due to limitations on the availability of these machines, they must be scheduled. In this work, we aim to minimize the total completion time of all vehicles. A vehicle's completion time is defined as the difference between the arrival and departure times at the depot. The PDPTW-SE is a computationally challenging problem that integrates routing and scheduling decisions as it generalizes the PDPTW, which is known to be NP-hard [Lenstra and Kan 1981, Dumas et al. 1991].

One application of the PDPTW-SE is in the logistical planning of pickups and deliveries involving multiple islands. In this scenario, requests are placed on islands, and cargo ships are required to transport the vehicles between islands. Another application of the PDPTW-SE is in the logistical planning of pickups and deliveries by automated guided vehicles (AGVs) in multi-floor buildings, such as hospitals or hotels. In these scenarios, requests are placed on different floors, and it may be necessary to use elevators to transport the vehicles between the floors. A third possible application of the PDPTW-SE is in the logistical planning of pickup and delivery of goods in areas affected by man-made or natural disasters. In a scenario where all roads between regions are closed for regular vehicles, a possible alternative solution would be to use machines that can traverse between regions carrying regular vehicles.

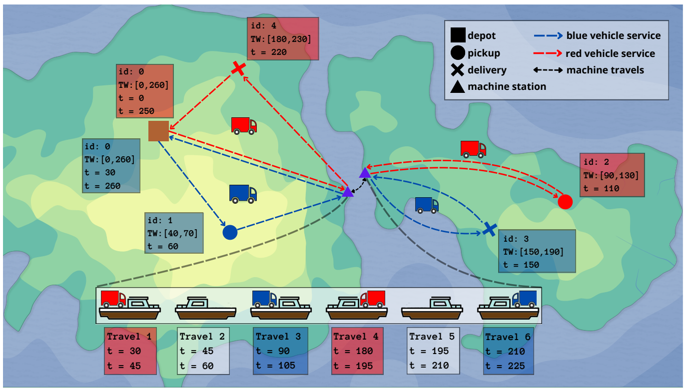

Figure 1. PDPTW-SE solution example.

Figure 1 illustrates a solution for an instance with two requests, two vehicles, two regions, and one machine. The red vehicle starts its journey at time 0, moving toward pickup node 2, which is located in another region. Hence, it arrives at the machine station at time 30, and the machine transports the vehicle to the other region, arriving at time 45. Then, the red vehicle continues its journey to node 2, while the machine returns to its previous station to wait for the blue vehicle. The blue vehicle leaves the depot at time 30, picks up the goods in node 1 at time 60, and meets the machine at its station at time 90. Next, the machine transports the blue vehicle to the other region, arriving at time 105. While the blue vehicle moves to node 3, the red vehicle picks up the goods at node 2 at time 110 and returns to the machine station. The blue vehicle arrives before the TW, but it waits to deliver the goods at node 3 at time 150, and while it returns to the machine station, the machine transports the red vehicle to the other region. Afterward, the machine returns to its previous station to transport the blue vehicle that was already waiting at the machine station, while the red vehicle continues its journey to deliver the goods at node 4. Finally, both vehicles return to the depot. The total completion time of this solution is 480, as the blue and red vehicles complete their journey in 230 and 250, respectively.

2.1. Proposed approaches

To solve the PDPTW-SE, we proposed an exact method and a heuristic. The exact method is an arc-based MIP formulation. The arcs represent the trips between the two nodes. A node can be the depot, a pickup, or a delivery location. Machine stations are not considered nodes of the network. Thus, arcs are subdivided into intra- and inter-region arcs.

The first part of the MIP formulation models the routing constraints for the PD vehicles. The second part formulates the scheduling constraints with MTZ subtour elimination constraints. Machine routes are constructed by establishing the precedence between inter-region arcs. The objective function minimizes the total completion time. Furthermore, we define a preprocessing step to eliminate several inadmissible arcs. We also propose several families of valid inequalities to cut off undesirable fractional solutions and potentially improve the performance of a branch-and-bound system when solving the formulation. Details concerning the MIP formulation (decision variables, constraints, objective function, preprocessing, and formulation tightening) can be found in Section 2.4 of the master's thesis [Barbosa 2025].

We also propose a multi-start heuristic with an LP-based improvement procedure (MSLP). It consists of executing a greedy insertion heuristic, followed by multiple executions of the randomized version of the greedy heuristic, where the schedule of each feasible solution obtained is optimized by solving a linear programming (LP) model. The LP model is inspired by the scheduling constraints of the MIP formulation, removing integer variables, and adding constraints based on the sequence of visits of each vehicle and the sequence of travels of each machine. Details concerning the MSLP heuristic (algorithms and LP formulation) can be found in Section 2.5 of the master's thesis [Barbosa 2025].

2.2. Computational experiments

We created a benchmark set of instances for the PDPTW-SE based on the randomized instances of [Lim and Li 2001] for the PDPTW [SINTEF 2008]. This set is composed of two families of instances, denoted *multi-island* and *multi-floor*, that represent two different potential applications of the PDPTW-SE. These differ in how the nodes and machine stations are placed into the environment and how the environment is partitioned into regions. For the multi-island instances, the request nodes are clustered into spatial regions (islands) using k-means, and each island's shape is approximated by its convex hull. A backtracking procedure selects one vertex from each hull to minimize the total pairwise distance (minimum weighted clique). Machine locations are then generated on edges near these vertices, ensuring valid spacing and integer coordinates. For the multi-floor instances, in contrast, the request nodes are assigned to different floors by randomly selecting z-coordinate from the set of available floors, forming a three-dimensional structure. Machine stations share the same (x, y) coordinates and are stacked vertically with one station per floor, positioned near the center of the spatial range while avoiding occupied points. Figure 2 illustrates the two instance families. Details concerning the computational experiments (settings, benchmark instances, results, and analysis) can be found in Section 2.6 of the master's thesis [Barbosa 2025].

For the MIP formulation, instances with up to 12 requests were considered. The solver achieved optimal solutions for up to 48.8% of the 320 instances and found feasible solutions for 95.0%, including all short-horizon instances. Performance was notably better on short-horizon instances, with smaller optimality gaps. Strengthening the formulation with valid inequalities increased both the number of optimal and feasible solutions by three. Besides, it produced equal or smaller gaps in at least 75% of the instances that the solver found a feasible solution but reached the time limit.

The MSLP heuristic was tested on 480 instances: the previous 320 plus 160 in-

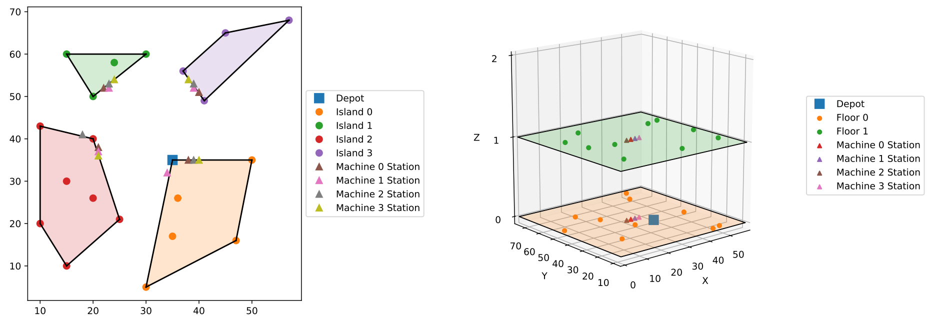

(a) Multi-island instance (b) Multi-floor instance

Figure 2. Illustration map for the two instance families

stances with up to 60 requests (120 pickup-delivery nodes). For instances with up to 12 requests, it found feasible solutions in all runs, often matching or outperforming the strengthened MIP, particularly on long-horizon instances. For larger instances (40 and 60 requests), the heuristic obtained feasible solutions for 93.8% of cases, whereas the MIP formulations found none.

Results also show that adding an extra machine significantly improves feasibility, especially in instances with four regions, and generally reduces solution values. About 75% of heuristic solutions are within 5% of the best-known value. The LP improvement procedure further enhanced solution quality, with average improvements above 5% and up to 43%, particularly for long-horizon instances.

3. The SMCTSP

The coupled-task scheduling problem (CTSP) consists of scheduling jobs, each with at least two coupled-tasks. In each job, the succeeding task must be started after the preceding task is finished, with an exact delay between them. For this work, we tackle the single-machine version of the problem, where each job has two tasks that must be separated by an exact delay. The goal consists of minimizing the makespan. Such a minimization is challenging because tasks can be scheduled within the exact delay of a job, whether by nesting jobs between other job tasks or by interleaving tasks of different jobs. Henceforth, the complexity of this problem relies on the existence of several non-singleton jobs, i.e., jobs with exact delays larger than other job processing times, which allows nesting and interleaving operations. Otherwise, these jobs can be randomly appended at the end of a schedule without loss of generality [Li and Zhao 2007].

Applications of the CTSP include pulsed radar systems in which a radar transmits a pulse and receives its reflection to keep track of targets [Shapiro 1980], chemistry manufacturing processes [Ageev and Baburin 2007], patient appointments in medical services [Pérez et al. 2011, Condotta and Shakhlevich 2014, Liu et al. 2019], among others. Another potential application involves the scheduling of tasks operated by robots in smart homes, especially in smart kitchens [Sharma and Reddy 2024]. In this scenario, kitchen robots [Mepani et al. 2022, Jiang and Zhou 2022] can take advantage of idle moments between sub-tasks to perform other sub-tasks, thus reducing the total preparation time [Bautista et al. 2023, Yi et al. 2022]. [Yi et al. 2022] studied a similar problem involving multiple tasks operated by dual-arm kitchen robots in a controlled environment.

The single-machine coupled-task scheduling problem (SMCTSP) to minimize makespan is classified as a strongly NP-hard problem [Orman and Potts 1997, Sherali and Smith 2005]. [Condotta and Shakhlevich 2012] showed that obtaining an optimal solution is NP-hard even if there is a predefined sequence for the initial (or final) tasks and all the jobs have unit processing times.

Figure 3 illustrates the single-machine coupled-task scheduling problem, along with the notations used, which are detailed at Section 3.2 of the master's thesis [Barbosa 2025].

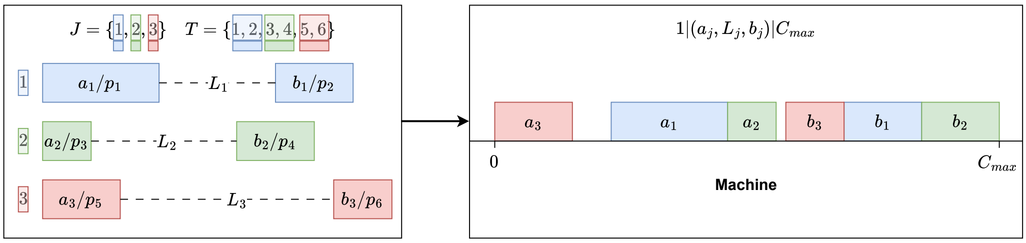

Figure 3. Illustration of the single-machine coupled-task scheduling problem

3.1. Proposed approaches

To solve the SMCTSP, we proposed an exact method and a heuristic. The exact method is a start-time-based CP model, i.e., there is essentially one variable representing the start time of a task. To ensure no overlap between tasks, we apply a global constraint, which can be modeled as a conjunction of non-overlap constraints or differently depending on the underlying solver. Details regarding the CP model are available at Section 3.2 of the master's thesis [Barbosa 2025].

The heuristic is a BRKGA. Each solution is encoded as a vector X composed of $|J|$ random keys, in which the j-th key corresponds to the job $j \in J$. Our decoder is a polynomial-time heuristic in which the keys are used to define a sorted sequence of the jobs $\sigma^{jobs} = (\sigma^{jobs}_1, \sigma^{jobs}_2, \ldots, \sigma^{jobs}_n)$, where σ^{jobs}_k, $k \in \{1, \ldots, n\}$, is the k-th job in the sequence. The job σ^{jobs}_k is thus inserted into the solution after the finish of the initial task of σ^{jobs}_{k-1} in the earliest possible time without overlapping with those jobs already scheduled. Despite the simplicity of our BRKGA's decoder, we show that it can provide reasonably good approximations to the optimal schedules, and the BRKGA embedded with such a decoder will be able to explore the set of possible solutions more efficiently. Figure 4 illustrates a decoding.

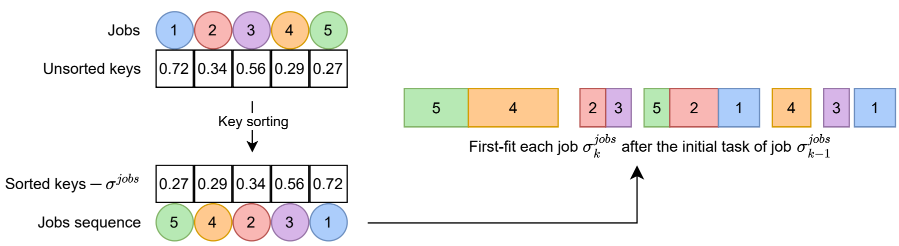

Figure 4. Decoder example.

We also integrate an initial solution generator, a first-improvement local search, and a perturbation method with the BRKGA. The initial solution generator is a multi-start procedure involving an adaptive version of the decoder. The local search works on the encoded vectors of the BRKGA, where the neighborhood is defined by moving a job to a different place. Besides, the local search is executed periodically or after an improvement on the eligible individuals of the elite set. Finally, the perturbation method comprises different intensity shakings and resets over different scenarios to prevent the BRKGA from being stuck at a same solution for several iterations. We also inject a predefined type of solution after each perturbation to avoid the BRKGA from losing its convergence structure. Details regarding the BRKGA heuristic (encoding and decoder, initial solution generator, local search, perturbation, and overall framework) can be found in Section 3.3 of the master's thesis [Barbosa 2025].

3.2. Computational experiments

We considered the general benchmark set proposed in [Khatami et al. 2019, Khatami et al. 2020]. There are eight different instance sizes, with up to 100 jobs (200 tasks), each group with 30 randomly generated instances divided into three categories, totaling 240 instances. We evaluated in total three approaches: the constraint programming model, the baseline BRKGA with restarts (BRKGA-R), and the BRKGA with restarts, shakings, and local search (BRKGA-R-S-LS). Details concerning the computational experiments (settings, benchmark instances, parameter settings, and analysis) can be found in Section 3.4 of the master's thesis [Barbosa 2025].

Concerning the 180 instances with at least 15 jobs that were not fully solved to optimality in the literature, the computational results show that the proposed approaches are effective. Among the exact methods, the CP model outperformed state-of-the-art MIP formulations, achieving the best objective values, the smallest gaps, and the highest number of optimal solutions, while also being the only method to find feasible solutions for all instances. With 3600 seconds and multiple threads, the CP model obtained the best-known solutions for 163 out of 180 instances (90.56%).

Under identical computational settings (single thread and short time limit), the BRKGA produced superior solution quality compared to CP, delivering high-quality solutions within short runtimes. The shaking and local search components proved essential to its performance, particularly local search.

To address premature population homogeneity, we introduced a controlled shaking strategy that preserves useful elite information while resetting only the non-elite portion of the population. Combined with solution injection and selective local search, this mechanism substantially improved the standard BRKGA with restarts.

4. Impact and dissemination

Papers The results of this thesis have already been disseminated in relevant scientific venues. Chapters 2 and 3 were submitted to EJOR. The paper for Chapter 2 has already been published, while the second paper has been made available as a preprint and is currently in its second round of reviews. Preliminary versions of both research lines were presented at SBPO. The list of currently published papers follows:

- [Barbosa et al. 2026] *Journal* EJOR: presents Chapter 2 of the master's thesis [Barbosa 2025].
- [Barbosa and Melo 2025] *Pre-print*: presents Chapter 3 of the master's thesis [Barbosa 2025].
- [Barbosa et al. 2024] *Conference* LVI SBPO 2024: presents a preliminary work for the SMCTSP.
- [Barbosa et al. 2023] *Conference* LV SBPO 2023: presents a preliminary work for the PDPTW-SE.

Impact This thesis advances exact and heuristic optimization for highly constrained routing and scheduling problems through contributions in problem modeling, algorithm design, and computational evaluation. It introduces the PDPTW-SE, a new routing and scheduling problem with relevant practical applications, and proposes for it an exact MIP formulation, a scalable heuristic with an LP improvement procedure, and a benchmark set, all made publicly available with code, scripts, and results[1]. It also proposes a CP model and an enhanced BRKGA for the SMCTSP, achieving state-of-the-art computational performance. In particular, the CP approach obtained the best-known solutions for 163 of the 180 hardest benchmark instances, while the BRKGA delivered superior solution quality under short time limits. These results have already generated concrete scientific impact, including an article published in the European Journal of Operational Research, a second journal submission under advanced review, a public preprint, and prior recognition at SBPO. Together, these outcomes demonstrate the scientific maturity, reproducibility, and relevance of the thesis contributions.

Acknowledgments

This study was financed in part by the Brazilian Federal Agency for Support and Evaluation of Graduate Education (CAPES) – Finance Code 001. This work was partially supported by the Brazilian National Council for Scientific and Technological Development (CNPq), grants 314718/2023-0, 445324/2024-4, and 421005/2025-4. This work was also partially supported by the FAPESB INCITE project PIE0002/2022.

References

Ageev, A. A. and Baburin, A. E. (2007). Approximation algorithms for UET scheduling problems with exact delays. *Operations Research Letters*, 35(4):533–540.

Barbosa, V. A. (2025). Exact and heuristic approaches for the pickup and delivery problem with time windows and scheduling on the edges, and for the single-machine coupled task scheduling problem with exact delays. Dissertação de mestrado, Universidade Federal da Bahia, Salvador.

Barbosa, V. A., Melo, R., and Januario, T. (2024). Programação por restrições e algoritmo genético de chaves aleatórias enviesadas para o problema de escalonamento de tarefas acopladas em uma única máquina minimizando o makespan. In *Proceedings of LVI Brazilian Symposium on Operations Research*, volume Vol 56, 2024 – 309602 of *SBPO 2024*, Fortaleza, CE, Brazil. Galoá.

[1]Repository available at: https://github.com/ab-vitor/pdptw-se

Barbosa, V. A. and Melo, R. A. (2025). Constraint programming model and biased random-key genetic algorithm for the single-machine coupled task scheduling problem with exact delays to minimize the makespan. https://doi.org/10.48550/arXiv.2512.23150.

Barbosa, V. A., Tiwari, S., and Melo, R. A. (2023). O problema de coleta e entrega com janelas de tempo e escalonamento nas arestas. In *Proceedings of LV Brazilian Symposium on Operations Research*, volume Vol 55, 2023 – 161147 of *SBPO 2024*, São José dos Campos, SP, Brazil. Galoá.

Barbosa, V. A., Tiwari, S., and Melo, R. A. (2026). The pickup and delivery problem with time windows and scheduling on the edges. *European Journal of Operational Research*. https://doi.org/10.1016/j.ejor.2026.01.036.

Bautista, M. G. A., Suganob, N. J., Concepcion, R., Bandala, A., and Dadios, E. (2023). Automated cooking systems: Benefits, challenges, and future directions. In *2023 IEEE 15th International Conference on Humanoid, Nanotechnology, Information Technology, Communication and Control, Environment, and Management (HNICEM)*, pages 1–6, Coron, Palawan, Philippines. IEEE.

Condotta, A. and Shakhlevich, N. (2014). Scheduling patient appointments via multilevel template: A case study in chemotherapy. *Operations Research for Health Care*, 3(3):129–144.

Condotta, A. and Shakhlevich, N. V. (2012). Scheduling coupled-operation jobs with exact time-lags. *Discrete Applied Mathematics*, 160(16-17):2370–2388.

Dumas, Y., Desrosiers, J., and Soumis, F. (1991). The pickup and delivery problem with time windows. *European Journal of Operational Research*, 54(1):7–22.

Jiang, A. Z. and Zhou, M. (2022). Design of Affordable Self-learning Home Cooking Robots. In *2022 IEEE International Conference on Networking, Sensing and Control (ICNSC)*, pages 1–6, Shanghai, China. IEEE.

Khatami, M., Salehipour, A., and Cheng, T. (2019). Data for "Coupled task scheduling with exact delays: Literature review and models". V1. 10.17632/dd7ht5k5pn.1.

Khatami, M., Salehipour, A., and Cheng, T. (2020). Coupled task scheduling with exact delays: Literature review and models. *European Journal of Operational Research*, 282(1):19–39.

Lenstra, J. K. and Kan, A. R. (1981). Complexity of vehicle routing and scheduling problems. *Networks*, 11(2):221–227.

Li, H. and Zhao, H. (2007). Scheduling coupled-tasks on a single machine. In *2007 IEEE Symposium on Computational Intelligence in Scheduling*, pages 137–142. IEEE.

Lim, A. and Li, H. (2001). A Metaheuristic for the Pickup and Delivery Problem with Time Windows . In *Proceedings 13th IEEE International Conference on Tools with Artificial Intelligence. ICTAI 2001*, page 160, Los Alamitos, CA, USA. IEEE Computer Society.

Liu, Z., Lu, J., Liu, Z., Liao, G., Zhang, H. H., and Dong, J. (2019). Patient scheduling in hemodialysis service. *Journal of Combinatorial Optimization*, 37(1):337–362.

Mepani, M. M., Gala, K. B., Mishra, T. A., Suresh Bhole, K., Gholave, J., and Daingade, S. (2022). Design of robot arm for domestic culinary assistance. *Materials Today: Proceedings*, 68:1930–1945. 4th International Conference on Advances in Mechanical Engineering.

Orman, A. J. and Potts, C. N. (1997). On the complexity of coupled-task scheduling. *Discrete Applied Mathematics*, 72(1-2):141–154.

Pérez, E., Ntaimo, L., Wilhelm, W. E., Bailey, C., and McCormack, P. (2011). Patient and resource scheduling of multi-step medical procedures in nuclear medicine. *IIE Transactions on Healthcare Systems Engineering*, 1(3):168–184.

Shapiro, R. D. (1980). Scheduling coupled tasks. *Naval Research Logistics Quarterly*, 27(3):489–498.

Sharma, K. and Reddy, S. R. N. (2024). Design of smart automated cooker_a survey and feasibility study. *Multimedia Tools and Applications*.

Sherali, H. D. and Smith, J. C. (2005). Interleaving two-phased jobs on a single machine. *Discrete Optimization*, 2(4):348–361.

SINTEF (2008). Instances and best solutions for the PDPTW. Online reference at `http://www.sintef.no/pdptw`, last access on August 09, 2024.

Yi, J.-s., Luong, T. A., Chae, H., Ahn, M. S., Noh, D., Tran, H. N., Doh, M., Auh, E., Pico, N., Yumbla, F., Hong, D., and Moon, H. (2022). An Online Task-Planning Framework Using Mixed Integer Programming for Multiple Cooking Tasks Using a Dual-Arm Robot. *Applied Sciences*, 12(8):4018.