

Novel ensemble collaboration method for dynamic scheduling problems

Marko Đurasević

University of Zagreb, Faculty of Electrical Engineering and
Computing
Zagreb, Croatia
marko.durasevic@fer.hr

Francisco Javier Gil Gala

University of Oviedo, Department of Computing
Gijón, Spain
giljavier@uniovi.es

Lucija Planinić

University of Zagreb, Faculty of Electrical Engineering and
Computing
Zagreb, Croatia
lucija.planinic@fer.hr

Domagoj Jakobović

University of Zagreb, Faculty of Electrical Engineering and
Computing
Zagreb, Croatia
domagoj.jakobovic@fer.hr

ABSTRACT

Dynamic scheduling problems are important optimisation problems with many real-world applications. Since in dynamic scheduling not all information is available at the start, such problems are usually solved by dispatching rules (DRs), which create the schedule as the system executes. Recently, DRs have been successfully developed using genetic programming. However, a single DR may not efficiently solve different problem instances. Therefore, much research has focused on using DRs collaboratively by forming ensembles. In this paper, a novel ensemble collaboration method for dynamic scheduling is proposed. In this method, DRs are applied independently at each decision point to create a simulation of the schedule for all currently released jobs. Based on these simulations, it is determined which DR makes the best decision and that decision is applied. The results show that the ensembles easily outperform individual DRs for different ensemble sizes. Moreover, the results suggest that it is relatively easy to create good ensembles from a set of independently evolved DRs.

CCS CONCEPTS

• **Computing methodologies** → **Planning and scheduling.**

KEYWORDS

Genetic Programming, Scheduling, Unrelated Machines, Dispatching Rules, Ensembles

ACM Reference Format:

Marko Đurasević, Lucija Planinić, Francisco Javier Gil Gala, and Domagoj Jakobović. 2022. Novel ensemble collaboration method for dynamic scheduling problems. In *Genetic and Evolutionary Computation Conference (GECCO '22)*, July 9–13, 2022, Boston, MA, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3512290.3528807>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

GECCO '22, July 9–13, 2022, Boston, MA, USA

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9237-2/22/07...\$15.00

<https://doi.org/10.1145/3512290.3528807>

1 INTRODUCTION

Scheduling is an essential part of many real-world processes, which are directly affected by the quality of the schedules created. As a result, much of the research has focused on studying various scheduling problems and methods to achieve the best possible results. Most scheduling problems are NP-hard, which makes it difficult to obtain optimal solutions in a reasonable time [21]. As a result, scheduling problems are usually solved using various heuristic methods. However, many scheduling problems are also dynamic in nature, which means that not all problem properties are known at the beginning, but become available only during the execution of the system. This complicates the application of standard search based methods, such as genetic algorithms. Dispatching rules (DRs) have therefore become the method of choice for dynamic scheduling problems.

DRs are simple heuristics that iteratively construct the schedule [39]. At each decision point, they only determine which job should be executed next. Thus, as new jobs are released into the system or other changes occur, these rules can easily adapt to such situations. Since manual development of good DRs is difficult, research has been done on automatic development of DRs in recent years. In these studies, genetic programming (GP) [23] has emerged as the method of choice for automatic development of new DRs, as it was demonstrated that the rules generated by GP outperform the best manually developed DRs in most cases.

However, there is a limit to how good results only a single DR can achieve. This is a direct consequence of the No Free Lunch theorem [31], which states that no one method can achieve the best results on all problems. Therefore, it is also impossible for GP to develop a single DR that will produce good results on all problem instances it may encounter. Many studies have focused on improving the performance of individual DRs [5, 12], but such a rule will still perform poorly on some instances. Therefore, several studies have focused on investigating methods for multiple DRs to collaborate in solving a problem [6, 20, 27]. The motivation for such a research direction is that if a single DR cannot always make the best decisions, perhaps a collection of DRs making decisions together can produce better results.

Most studies that addressed the idea of using multiple DRs together used them in the form of ensembles. Although various ensemble learning methods have been proposed and used in the literature [3, 27], all studies showed that ensembles of DRs performed significantly better than individual DRs. This in turn motivated several follow-up studies that looked more deeply into the problem of designing ensembles.

In this study, we explore a slightly different way in which DRs forming an ensemble can work together, which differs in some aspects from the standard ensembles in the literature. More specifically, the DRs contained in the ensemble work together to solve a problem instance, but unlike standard ensembles where the independent decisions of all rules are aggregated, in these ensembles the decisions of all the rules are simulated individually. This means that each time a decision needs to be made, all the rules in the ensemble are evaluated and their decisions are assessed. The best decision is then applied to the schedule. In this way, the ensemble tries to select the best available decision and apply it to the schedule. The specific objectives of this study can be summarised as follows:

- (1) Design a novel approach in which the DRs in an ensemble work together.
- (2) Construct ensembles of DRs and apply the proposed collaboration approach to study their performance.
- (3) Analyse various aspects of the proposed ensemble collaboration method.

The remainder of the paper is organised as follows. Section 2 provides a literature review. Section 3 describes the description on the unrelated machines environment, and the application of GP to develop DRs for this environment. The proposed ensemble method is described in Section 4. Section 5 describes the experimental setup and outlines the results obtained. The conclusion and future work directions are outlined in Section 7.

2 RELATED WORK

From the first works that focused on the automatic generation of DRs for different environments such as single machines [2] and job shop [9, 11], a large number of studies in this area have been published to date [1, 13]. Several areas of automated dispatching rule design have been extensively researched, such as multi-objective optimisation [14, 15, 27, 37], application of surrogate models [16, 34], and feature selection [35, 38]. In recent years, there have also been some new research directions, which include multitask GP [32, 36], improving selection mechanisms in algorithms and operators [10, 22, 33], considering additional scheduling properties [6, 8, 19], and generating initial populations of a genetic algorithm with automatically generated DRs [30].

In addition to the above research directions, the creation of ensembles of DRs represents another goal that has been intensively studied in the literature. One of the first studies on the creation of ensembles of DRs for the static job shop environment was conducted in [20]. The authors use the cooperative coevolution algorithm. The algorithm contains the same number of subpopulations as there are elements in the ensemble, and each subpopulation is used to evolve a single DR for the ensemble. The DRs of the individual populations interact only when evaluated as an ensemble. In [17], the authors

apply a multilevel genetic programming method to evolve ensembles. This method did not achieve better results than the method of [20], but it evolved ensembles in less time. In [7], the authors propose NELLI-GP, an ensemble learning method for generating heuristic ensembles. The results show that the proposed method performs better than previous approaches. In [18], the authors investigate different ensemble combination schemes, including linear combination, weighted linear combination, majority voting, and weighted majority voting. The results show that the linear combination scheme generally achieves the best performance.

In [26], the authors compared 4 methods for constructing ensembles of DRs for the unrelated machines environment, including BagGP, BoostGP, cooperative coevolution, and the simple ensemble combination (SEC) method. The results showed that most of the obtained ensembles performed significantly better than individual DRs. Additional analysis revealed that the best results were obtained by ensembles with smaller size. For all tested methods, SEC mostly created the best ensembles. Since SEC is a simpler method compared to the other tested approaches, this motivated the authors to perform a more thorough analysis of the method, especially considering different ways to construct the ensembles [28]. This study also showed that, even when using different ensemble construction methods, it was possible to obtain good ensembles in a short time. The ensemble approaches from [27] were also applied to the resource constrained project scheduling problem (RCPSPP) [24]. In this work, it was shown that using ensembles it was again possible to improve the results compared to single DRs. Moreover, for this problem the SEC method also achieved the best overall results.

In [6], the authors propose a novel type of ensemble collaboration which they apply to solve the one machine problem with variable capacity. Since the considered problem is static in nature, i.e., all information about it is known in advance, the authors outline that it is possible to solve the problem instances with a selected set of DRs and then choose the best solution for each instance. The authors use a genetic algorithm to determine the best ensemble of DRs, and show that such ensembles significantly improve results over individual DRs. These ensembles are further investigated in [4], where the authors analyse the obtained ensembles and compare ensembles created from the ATC rule and from GP generated DRs. The results show that ensembles created from automatically generated DRs perform better than those obtained from the ATC rule. Finally, in [3], the authors investigate different methods for creating ensembles from DRs, namely an iterated greedy method, local search, and a memetic algorithm. The results show that all methods construct ensembles of good quality, with the memetic algorithm achieving the best ensembles overall.

Apart from their application in scheduling, ensembles of heuristics generated by GP have also been used for the capacitated arc routing problem, where they were considered to evolve smaller and more interpretable routing policies without reducing their performance.

3 BACKGROUND

3.1 Unrelated machines environment

The unrelated machines environment represents an important scheduling problem that arises in many real-world applications such as manufacturing, cloud environments, and the like [21]. This problem consists of n jobs and m machines, where each job j must be assigned to a particular machine i . Each job has the following properties:

- processing time (p_{ij}) - time needed to execute job j on machine i ,
- release time (r_j) - time when job j will be released in the system,
- due date (d_j) - time by which job j must be executed,
- weight (w_j) - importance of job j .

The aforementioned properties are taken into account when creating a schedule for a given problem, since they have a direct impact on the quality of the schedule, depending on which criterion is optimised. In this study, the total weighted tardiness (TWT) criterion is to be minimised. TWT is defined as $TWT = \sum_j w_j \max(C_j - d_j, 0)$, where C_j represents the completion time of job j . Thus, minimising this criterion reduces the time that jobs spend executing after their due date. The problem described can be classified as $R|r_j|TWT$ using the standard notation of scheduling problems [21].

In addition to the above properties of the problem, another important property of the considered problem is that it is dynamic. This means that until a job is released in the system, no information about it is known about it. Moreover, it is also not known when a job will be released into the system. Therefore, at each decision point, the decision can only be made based on the jobs that have been released so far, and no information about the future of the system can be used. This feature makes DRs a suitable choice for solving such problems.

3.2 Designing DRs with GP

DRs usually consist of two parts, a schedule generation scheme (SGS) and a priority function (PF). The SGS is tasked with constructing the entire schedule, i.e., it determines at which point a scheduling decision should be made and which jobs and machines should be considered for that decision point. The SGS used in this study is denoted in Algorithm 1 [40]. This SGS performs a decision every time at least one job and one machine are available. At that moment, it determines which job should be scheduled on which machine. For this purpose, it uses a PF that assigns a numeric value to each job-machine pair. The pair that received the lowest value is selected and the corresponding job is then scheduled on the selected machine. However, the job is scheduled only if the selected machine is free. Otherwise, this job is skipped and the next one is scheduled. The entire process is repeated as long as there are jobs left.

Since the PF determines the order in which jobs are scheduled, it is an important aspect of DR. Although several good PFs have been designed manually, such as those used in the ATC or COVERT DRs, it has been shown that automatically generated DRs can outperform them. For this reason, GP is often used to automatically design PFs that are appropriate for the problem under consideration. Thus,

Algorithm 1: SGS used by generated DRs

```

1: while unscheduled jobs are available do
2:   Wait until at least one job and machine are available
3:   for all available jobs  $j$  and each machine  $i$  in  $m$  do
4:     Get the priority  $\pi_{ij}$  of scheduling  $j$  on machine  $i$ 
5:   end for
6:   for all available jobs do
7:     Determine the machine with the best  $\pi_{ij}$  value
8:   end for
9:   while jobs whose best machine is available exist do
10:    Determine the best priority of all such jobs
11:    Schedule the job with best priority
12:   end while
13: end while

```

Table 1: Terminal set

| Terminal | Description |
|----------|---|
| pt | processing time of job j on machine i |
| $pmin$ | minimal processing time (MPT) of job j |
| avg | average processing time of job j across all machines |
| PAT | time until machine with the MPT for job j becomes available |
| MR | time until machine i becomes available |
| age | time which job j spent in the system |
| dd | time until which job j has to finish with its execution |
| w | weight of job j (w_j) |
| SL | slack of job j , $-\max(d_j - p_{ij} - t, 0)$ |

the most important part of adapting GP for such a problem is selecting the appropriate set of terminal and function nodes. The set of terminal nodes used is given in Table 1. The function set, on the other hand, consists of the addition, subtraction, multiplication, division (implemented as a protected division that returns 1 if the division by 0 occurs), and positive (unary operator that returns 0 if the argument is negative) operators. Both sets were selected based on previous studies [29].

4 ENSEMBLE COLLABORATION METHOD

From the literature review, it appears that DRs in the ensemble were used to collaborate in two different ways to solve different scheduling problems. In the first way, the ensembles are used to jointly make decisions during the scheduling process [20, 27, 28]. This means that all DRs included in the ensemble are evaluated at each decision point and then a single decision is made based on their individual decisions. To combine these decisions, different ensemble combination methods, such as linear combination or majority voting, can be used. On the other hand, in the second type of ensembles [3, 4, 6], each rule is applied individually to the problem instance and in the end the best solution among the rules included in the ensemble is selected as the final solution. However, such a method is only applicable in the case of static scheduling conditions, since all information about the system must be known in advance.

From the preceding description, it is clear that both types of ensembles are motivated by the fact that no single DR can perform

well on all problems, but the way the rules of the ensembles work together is different. While the first type of ensembles has been used under both static and dynamic conditions, the second type of ensembles has only been used for static problems due to their unique characteristics. However, the concept of executing rules in parallel and selecting the best result would also be interesting to consider in dynamic environments as well.

In dynamic scheduling environments, the DR does not have the information about which jobs will arrive in the future. However, when jobs are released, they wait to be scheduled until a machine becomes free. Therefore, it is to be expected that as the system runs, a certain number of jobs will accumulate in the queue, waiting for the moment when they are selected for scheduling. Traditional DRs operate in such a way that, based on the current status, they schedule a job only for the available machine and wait again when a machine becomes free. Therefore, an ensemble could be defined so that all DRs in it suggest which job should be scheduled next. However, instead of combining these decisions into one, a method can be used to determine how good a decision each rule made. The simplest method would be to determine how each decision affects the optimised criterion and select the one that results in the smallest increase in the optimised criterion. However, using only a single job to determine how good a decision is can lead to poor decisions. Therefore, each rule can be used to simulate not only the scheduling of the next job, but rather all currently released jobs under the assumption that no new jobs will be released in the future. In this way, it would be possible to get a more accurate approximation of how the next decision might affect the quality of the schedule in the future and how well each DR performs under current system conditions. It must be emphasised that no job is actually scheduled during this process, but rather the behaviour of each DR is merely simulated with the current system condition. Only when the DR with the best decision has been determined on the basis of the simulation is this decision implemented in the real schedule.

Based on the previous description, we propose an ensemble collaboration method that uses multiple DRs in parallel to build the schedule. Algorithm 2 gives an overview of the proposed ensemble collaboration method, which is similar to the standard SGS. The main difference is that at each decision point, each DR in the ensemble is used to create the schedule considering only the currently released jobs. It should be noted that the DRs here only simulate the schedule and that no jobs are actually scheduled in the real system. After running the simulations of all the rules, the rule that gave the best objective value in the simulation is selected. This rule is then used to make the scheduling decision at the current time. It should be emphasised that the rule schedules only a single job, and then moves to the next decision point where the entire procedure is repeated. Therefore, any new jobs released between these two decision points are also considered in the next decision point, which means that the procedure is suitable for dynamic environments.

Although not explicitly stated in the algorithm, there is a parameter that can be used to change the behaviour of the proposed method. As mentioned in the previous descriptions, the approximation how good a DR is at the current time can be done by considering a different number of currently released jobs. Although this parameter can be set to any value, in this study we will examine

Algorithm 2: SGS for the ensemble collaboration method

```

1: while unscheduled jobs are available do
2:   Wait until at least one job and machine are available
3:   for each rule  $k$  in the ensemble  $E$  do
4:     Build the schedule with only the released jobs at the
       current moment
5:     Store the quality of the simulated schedule in  $F[k]$ 
6:   end for
7:   Determine the rule  $k_{best}$  which has the lowest value
       objective value of the simulated schedule
8:   Apply rule  $k_{best}$  to determine next scheduling decision
9: end while

```

only the two extreme values, namely by scheduling only the next job and evaluating its impact on the schedule, and scheduling all released jobs and evaluating the performance of such a schedule. The first method is referred to as EDR-S, while the second method is referred to as EDR-M.

Based on the previous descriptions, it can be seen that the proposed ensemble collaboration method has certain similarities with the rollout heuristic [25]. At each decision point, the rollout heuristic considers all possible decisions, but to evaluate how good each of these decisions is, it uses a DR, to approximate the rest of the schedule. For this reason, the method is only applicable in static scheduling environments. The proposed ensemble collaboration method can be considered as a particular way of combining ensembles and the rollout heuristic, since it uses multiple DRs from ensembles at once to approximate the rest of the schedule similarly to the rollout heuristic. By combining these two concepts, it was possible to develop a method that is now applicable to dynamic environments.

An open question is still how to construct the ensembles. The advantage of the proposed approach is that most of the methods proposed in previous studies can be used to construct such ensembles. Since in this study we mainly focus on demonstrating the effectiveness of using the proposed ensemble collaboration method, the SEC method is used because of its simplicity, but also because of its ability to construct better ensembles than other more sophisticated methods [27]. SEC is applied using the random construction method. This method constructs p ensembles, where p is a user-specified parameter, where each of the ensembles is constructed by randomly selecting e DRs from the pool of available rules [28]. Here e represents the size of the ensemble to be constructed and is also a user-specified parameter.

5 EXPERIMENTAL STUDY

5.1 Setup

To develop rules, create ensembles, and evaluate their efficiency, a problem set of 180 instances was used. These instances contain between 12 and 100 jobs, and between 3 and 10 machines. In addition, the due dates for each instance were generated with different parameters, resulting in instances with different levels of difficulty. The instances are divided into 3 independent sets, the test set, the validation set and the training set. The training set is used by GP to develop new DRs. More specifically, GP is executed 50 times, and

Table 2: Parameters for GP

| Parameter name | Parameter value |
|-----------------------|--|
| Population size | 1000 individuals |
| Termination criterion | 80 000 function evaluations |
| Initialisation | Ramped half-and-half |
| Maximum tree depth | 5 |
| Crossover operators | Subtree, uniform, context-preserving, size-fair |
| Mutation operators | Subtree, hoist, node complement, node replacement, permutation, shrink |
| Mutation probability | 0.3 |

the best individual from each execution is stored. These DRs are used to create ensembles with the SEC method using the validation set. Finally, the test set is used to measure the quality of the created DRs and ensembles. All values mentioned in this section were obtained using this set.

For the development of DRs, the steady-state algorithm GP with a 3-tournament selection is used. The parameters of the algorithm are listed in Table 2. Since multiple operators are used for crossover and mutation, one is randomly selected from the set of defined operators each time such an operator needs to be applied. After 50 DRs have been developed, the SEC method uses them as building blocks for constructing the ensembles. The parameters that must be specified for the SEC method are the number of DRs constructed per execution and the size of the constructed ensembles. For the number of constructed ensembles, the values 100, 500, 1000, 5000, 10000, and 20000 are used. For the ensemble size, the initial tests are performed with sizes 3, 5, and 7. The SEC method is executed 30 times to obtain statistically significant results. From each execution, the best ensemble obtained on the validation set is saved and evaluated on the test set. The tables then show the minimum, median and maximum values of these 30 runs.

To test whether significant differences exist between the different parameter values, the Kruskal-Wallis test is used. The Bonferroni correction method is in post-hoc analysis. In addition, the Mann-Whitney test is used for individual pairwise comparisons. Results are considered significantly different if a p-value of less than 0.05 is obtained.

5.2 Results

Table 3 shows the results for three selected ensemble sizes using the EDR-S method. The results show that there is little difference between ensembles created using the SEC method with different numbers of constructed ensembles. This shows that even a smaller number of ensembles is sufficient to find those that perform better than individual DRs. The best overall results were obtained with the construction of 1000 ensembles of size 5. The statistical test showed that, in general, there was no significant difference between the results obtained using the SEC method with different numbers of constructed ensembles. However, in all cases significantly better results were obtained compared to individual DRs. Thus, it can be concluded that the SEC method can produce ensembles that are better than individual rules, even with a small number of samples.

Table 4 shows the results obtained with the EDR-M method. The table shows that the results obtained with the different number of constructed ensembles for the SEC method are quite similar. As in the previous method, the best overall median value is again obtained for the ensemble of size five, but this time constructed with SEC-500. However, the statistical test showed that for all three ensemble sizes there is no significant difference between the results obtained for the different SEC parameter values. Nevertheless, in all cases, the SEC method produced ensembles that were significantly better than the individual DRs.

To get a better idea of how the two tested methods compare, Figure 1 shows the results for them grouped by the different ensemble sizes. In this figure, S- SEC denotes the variant of SEC that uses the EDR-S method to evaluate ensembles, and M- SEC denotes the variant that uses the EDR-M method. For ensembles of size 3, it can be seen that there is not much difference between the methods for the number of ensembles constructed. However, the situation is different for the two larger ensembles. In these cases, we can see a significant difference between the results of the two methods. For ensemble size 5, EDR-M achieves significantly better results for all experiments except when 1000 ensembles were constructed. For ensemble size 7, EDR-M achieves significantly better results in all cases.

Based on the previous observations, the following conclusions can be drawn about the two tested. In both cases, the number of sampled ensembles in the SEC method has no significant influence on the results, which means that in both cases a smaller number of ensembles can be sampled, which is advantageous from the point of view of the execution time of the method. As for the ensemble sizes, it can be observed that there are no significant differences between the two methods at the smallest ensemble size. Thus, it seems that when the number of rules in the ensemble is small, the choice of which rule to select from the ensemble can be made even when considering a single job. However, this is not the case for the two larger ensemble sizes, as the results obtained with EDR-M are always significantly better than those obtained with EDR-S. This seems to indicate that it is more difficult to make the right decision about which DR to select from the ensemble once there are more rules in the ensemble. This makes sense because, again, it is difficult to evaluate multiple rules based on a single decision they make.

The obtained results are comparable or even better in several cases when compared to ensembles proposed in [26] and [28]. However, we leave further comparisons between the different ensemble collaboration methods and deeper analyses for future work, as the main objective of this paper was to focus on the proposed collaboration method.

6 FURTHER ANALYSIS

6.1 Influence of the ensemble sizes

To investigate the influence of ensemble size, ensembles with sizes between 2 and 20 DRs are tested. Using the SEC method, 500 ensembles are constructed for the EDR-M method and 1000 ensembles are constructed for the EDR-S method, as the methods produced their best median results at these values.

The results for different ensemble sizes are shown in Table 5. Size 1 in the table denotes the results obtained with each DR. The

Table 3: Results for the EDR-S collaboration method

| Method | Ensemble size | | | | | | | | |
|----------------|---------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| | 3 | | | 5 | | | 7 | | |
| | Min | Med | Max | Min | Med | Max | Min | Med | Max |
| Individual DRs | 15.23 | 15.94 | 17.59 | 15.23 | 15.94 | 17.59 | 15.23 | 15.94 | 17.59 |
| SEC-100 | 14.97 | 15.39 | 15.96 | 14.96 | 15.39 | 16.17 | 15.03 | 15.52 | 15.95 |
| SEC-500 | 15.00 | 15.40 | 16.04 | 15.08 | 15.49 | 15.80 | 14.99 | 15.42 | 16.04 |
| SEC-1000 | 14.86 | 15.34 | 15.88 | 15.01 | 15.29 | 15.63 | 14.79 | 15.44 | 15.88 |
| SEC-5000 | 14.93 | 15.35 | 15.89 | 15.00 | 15.49 | 16.09 | 15.07 | 15.43 | 16.00 |
| SEC-10000 | 14.80 | 15.27 | 15.93 | 15.01 | 15.31 | 15.82 | 15.07 | 15.44 | 15.91 |
| SEC-20000 | 14.96 | 15.41 | 15.91 | 14.97 | 15.38 | 15.84 | 15.14 | 15.42 | 16.01 |

Table 4: Results for the EDR-M collaboration method

| Method | Ensemble size | | | | | | | | |
|----------------|---------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| | 3 | | | 5 | | | 7 | | |
| | Min | Med | Max | Min | Med | Max | Min | Med | Max |
| Individual DRs | 15.23 | 15.94 | 17.59 | 15.23 | 15.94 | 17.59 | 15.23 | 15.94 | 17.59 |
| SEC-100 | 14.88 | 15.38 | 15.78 | 14.72 | 15.21 | 15.96 | 14.84 | 15.22 | 15.62 |
| SEC-500 | 14.72 | 15.33 | 15.81 | 14.72 | 15.04 | 15.71 | 14.66 | 15.17 | 16.03 |
| SEC-1000 | 14.69 | 15.32 | 15.70 | 14.80 | 15.23 | 15.85 | 14.85 | 15.10 | 15.64 |
| SEC-5000 | 14.78 | 15.36 | 15.61 | 14.72 | 15.18 | 15.65 | 14.84 | 15.18 | 15.53 |
| SEC-10000 | 14.78 | 15.25 | 15.61 | 14.71 | 15.11 | 15.75 | 14.81 | 15.17 | 15.75 |
| SEC-20000 | 14.78 | 15.17 | 15.61 | 14.75 | 15.15 | 15.67 | 14.88 | 15.10 | 15.58 |

results show that the constructed ensembles always achieve a better median value compared to the individual DRs. However, the statistical tests show that the difference is not always significant. For the EDR-M method, the constructed ensembles always perform better than the individual DRs. On the other hand, the EDR-S method achieves significantly better results than the individual DRs for all ensemble sizes except for sizes 14, 17, 18, 19, and 20. The best median value is obtained for ensemble size 5 in both cases. This shows that already with smaller ensemble sizes quite good results can be obtained and that the use of larger ensemble sizes does not bring any real advantage.

To better illustrate the differences between the different ensemble sizes, Figure 2 shows the results in the form of boxplots. For the EDR-S method, it is interesting to see how the results slowly deteriorate as the ensemble size is increased beyond 5. Thus, in general, this method does not seem to work well for larger ensemble sizes. This is because the decision of which DR to apply is based on only a single job. This leads to poor approximations and there is a greater chance that a suboptimal decision will be made. For example, a DR might schedule a job on a machine where it will run for a very long time but meet its due date. Although this may seem like a good choice from the perspective of that job, it may have a negative impact in the long run because the machine is blocked longer and some other jobs may be late. However, by considering only a single job, it is not possible to detect and prevent such a situation.

Table 5: Results for different ensemble sizes

| | EDR-S | | | EDR-M | | |
|----|-------|--------------|-------|-------|--------------|-------|
| | Min | Med | Max | Min | Med | Max |
| 1 | 15.23 | 15.94 | 17.59 | 15.23 | 15.94 | 17.59 |
| 2 | 14.95 | 15.58 | 15.93 | 14.91 | 15.43 | 15.71 |
| 3 | 14.86 | 15.34 | 15.88 | 14.72 | 15.33 | 15.81 |
| 4 | 14.99 | 15.45 | 15.83 | 14.84 | 15.27 | 15.64 |
| 5 | 15.01 | 15.29 | 15.63 | 14.72 | 15.04 | 15.71 |
| 6 | 14.81 | 15.36 | 15.92 | 14.84 | 15.13 | 16.03 |
| 7 | 14.79 | 15.44 | 15.88 | 14.66 | 15.17 | 16.03 |
| 8 | 15.03 | 15.42 | 15.91 | 14.89 | 15.18 | 15.72 |
| 9 | 15.12 | 15.46 | 15.96 | 14.86 | 15.13 | 15.68 |
| 10 | 15.16 | 15.46 | 16.19 | 14.94 | 15.15 | 15.74 |
| 11 | 15.06 | 15.54 | 16.19 | 14.84 | 15.26 | 15.58 |
| 12 | 15.02 | 15.58 | 16.16 | 14.82 | 15.19 | 15.37 |
| 13 | 15.20 | 15.51 | 16.07 | 14.85 | 15.20 | 15.49 |
| 14 | 15.26 | 15.70 | 16.12 | 14.97 | 15.24 | 15.51 |
| 15 | 14.99 | 15.63 | 16.15 | 14.90 | 15.31 | 15.57 |
| 16 | 15.18 | 15.62 | 16.06 | 14.77 | 15.23 | 15.55 |
| 17 | 15.21 | 15.81 | 16.72 | 14.87 | 15.21 | 15.67 |
| 18 | 15.17 | 15.73 | 16.20 | 14.94 | 15.17 | 15.48 |
| 19 | 15.17 | 15.70 | 16.29 | 14.86 | 15.22 | 15.56 |
| 20 | 15.34 | 15.76 | 16.17 | 14.91 | 15.26 | 15.60 |

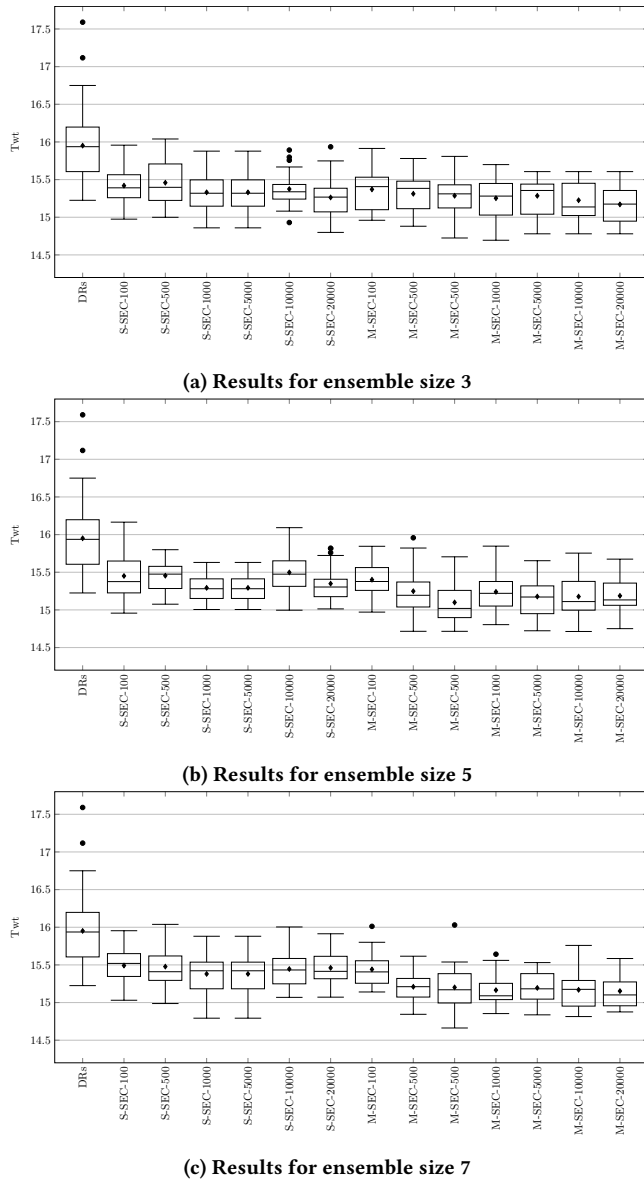


Figure 1: Boxplot representation of the results

For the EDR-M method, it can also be seen that the results improve up to size 5, after which they slowly deteriorate. However, the differences are not as significant as for the EDR-S method. This is a consequence of the fact that in this method the approximation is based on all released jobs, which gives a better idea of how the decision of each DR might affect the immediate future. Although this method is more stable with larger ensembles, we again find that the results deteriorate slightly. Thus, it appears that larger ensembles introduce more noise into the planning process, i.e., with more choices, it seems more difficult for the method to select the best one.

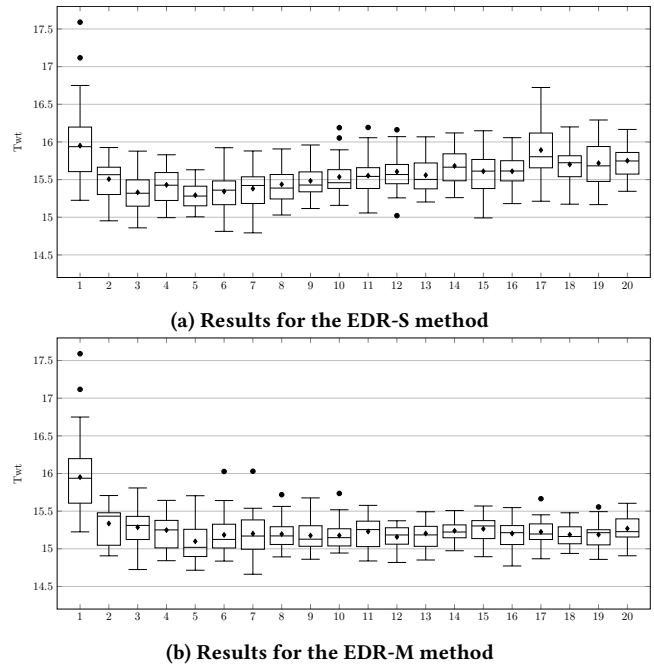


Figure 2: Results for different ensembles sizes


6.2 DR occurrence in ensembles

Figure 3 shows the number of occurrences for each DR in the 30 generated ensembles using S-SEC-1000 and M-SEC-500. The figures show the frequency for all three tested ensemble sizes. The figure shows that the methods clearly have a preference for selecting certain DRs. This is most evident for ensemble size 3. In this case, both methods are more biased towards selecting specific DRs. Even more, it can be seen that the EDR-S method even selects a specific DR in almost all ensembles. As the ensemble size increases, the tendency to select certain DRs decreases, but it can still be seen that the method favours certain DRs more than others.


It is also interesting to observe that in many cases both ensemble collaboration methods select the same DRs more frequently. The same is also true for certain DRs that are rarely selected by both methods. Thus, it appears that certain DRs are inherently less likely to be part of an ensemble. Such an analysis could be used to reduce the set of eligible DRs used to form the ensembles, and thus also reduce the search space. In additional analyses, it would also be interesting to investigate whether the rules that are more or less suitable for constructing ensembles have certain properties that could be used to detect such rules.

6.3 Ensemble execution time


Table 6 presents the average execution time of each DR and the ensemble on the entire test set. As expected, the execution time of the ensemble is higher than that of the individual DRs. For the EDR-S method, it can be observed that the increase in execution time is not significant. This is due to the fact that each rule makes only a single decision, resulting in only a small overhead. It is interesting to observe that the increase in execution time is linearly dependent on



(a) Results for ensemble size 3



(b) Results for ensemble size 5



(c) Results for ensemble size 7

Figure 3: Frequency of occurrence of DRs in the ensembles

Table 6: Average execution time of DRs and ensembles on the test set

| Method | Ensemble size | | |
|--------|---------------|-------|-------|
| | 3 | 5 | 7 |
| DRs | 0.273 | | |
| EDR-S | 0.352 | 0.491 | 0.633 |
| EDR-M | 1.197 | 2.192 | 2.991 |

the number of DRs in the ensemble. However, the EDR-M method incurs a much larger overhead since it must schedule multiple orders at each decision point. However, it can still be seen that all 60 instances are solved in less than three seconds for all ensemble sizes. This shows that the method is applicable in dynamic environments despite the additional computational overhead. The increase in execution time is again linear and depends on the number of DRs in the ensemble.

7 CONCLUSION

In this paper, we propose a novel ensemble collaboration method for solving dynamic scheduling problems. This method is inspired by existing ensemble learning methods used for static scheduling environments and the rollout heuristic. Experiments show that the proposed method can achieve superior results compared to individual DRs, even when using a simple method based only on random selection for ensemble construction. This shows that it is easy to find combinations of DRs that work well for the considered ensemble method, but also that such ensembles can be constructed in a small amount of time. Additional analysis of different ensemble sizes has shown that the EDR-M variant of the method is more resilient and more consistently produces good results than the EDR-S variant.

Several directions are planned to extend this work. First, different methods for creating ensembles will be investigated to determine if they can increase the quality of the results or at least create ensembles of similar quality but in less time. Second, a more in-depth analysis of the method will be conducted. This includes an analysis of the frequency with which DRs appear in the ensembles, as well as a deeper analysis of the parameter that determines how many jobs are scheduled when creating the simulated schedule and its impact on runtime. Finally, a detailed comparison between the proposed and existing ensemble methods is also performed to analyse and outline the strengths and weaknesses of both approaches.

ACKNOWLEDGMENTS

This work has been supported in part by Croatian Science Foundation under the project IP-2019-04-4333 and by the Spanish Government under research projects PID2019-106263RB-I00.

REFERENCES

- [1] Jürgen Branke, Su Nguyen, Christoph W. Pickardt, and Mengjie Zhang. 2016. Automated Design of Production Scheduling Heuristics: A Review. *IEEE Transactions on Evolutionary Computation* 20, 1 (2016), 110–124. <https://doi.org/10.1109/TEVC.2015.2429314>

- [2] C. Dimopoulos and A.M.S. Zalzalá. 2001. Investigating the use of genetic programming for a classic one-machine scheduling problem. *Advances in Engineering Software* 32, 6 (2001), 489–498. [https://doi.org/10.1016/S0965-9978\(00\)00109-5](https://doi.org/10.1016/S0965-9978(00)00109-5)
- [3] Francisco J. Gil-Gala, Carlos Mencía, María R. Sierra, and Ramiro Varela. 2020. Learning ensembles of priority rules for online scheduling by hybrid evolutionary algorithms. *Integrated Computer-Aided Engineering* 28, 1 (Dec. 2020), 65–80. <https://doi.org/10.3233/ICA-200634>
- [4] Francisco J. Gil-Gala, María R. Sierra, Carlos Mencía, and Ramiro Varela. 2020. Combining hyper-heuristics to evolve ensembles of priority rules for on-line scheduling. *Natural Computing* (June 2020). <https://doi.org/10.1007/s11047-020-09793-4>
- [5] Francisco J. Gil-Gala, María R. Sierra, Carlos Mencía, and Ramiro Varela. 2021. Genetic programming with local search to evolve priority rules for scheduling jobs on a machine with time-varying capacity. *Swarm and Evolutionary Computation* 66 (2021), 100944. <https://doi.org/10.1016/j.swevo.2021.100944>
- [6] Francisco J. Gil-Gala and Ramiro Varela. 2019. Genetic Algorithm to Evolve Ensembles of Rules for On-Line Scheduling on Single Machine with Variable Capacity. In *From Bioinspired Systems and Biomedical Applications to Machine Learning*. Springer International Publishing, 223–233. https://doi.org/10.1007/978-3-030-19651-6_22
- [7] Emma Hart and Kevin Sim. 2016. A Hyper-Heuristic Ensemble Method for Static Job-Shop Scheduling. *Evolutionary Computation* 24, 4 (2016), 609–635. https://doi.org/10.1162/EVCO_a_00183
- [8] Kristijan Jaklinović, Marko Đurasević, and Domagoj Jakobović. 2021. Designing dispatching rules with genetic programming for the unrelated machines environment with constraints. *Expert Systems with Applications* 172 (2021), 114548. <https://doi.org/10.1016/j.eswa.2020.114548>
- [9] Domagoj Jakobovic and Leo Budin. 2006. Dynamic Scheduling with Genetic Programming. 73–84. https://doi.org/10.1007/11729976_7
- [10] Atiya Masood, Gang Chen, Yi Mei, Harith Al-Sahaf, and Mengjie Zhang. 2020. A Fitness-based Selection Method for Pareto Local Search for Many-Objective Job Shop Scheduling. In *2020 IEEE Congress on Evolutionary Computation (CEC)*. 1–8. <https://doi.org/10.1109/CEC48606.2020.9185881>
- [11] Kazuo Miyashita. 2000. Job-Shop Scheduling with Genetic Programming. In *Proceedings of the 2nd Annual Conference on Genetic and Evolutionary Computation (Las Vegas, Nevada) (GECCO'00)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 505–512.
- [12] Su Nguyen, Yi Mei, Bing Xue, and Mengjie Zhang. 2019. A Hybrid Genetic Programming Algorithm for Automated Design of Dispatching Rules. *Evolutionary Computation* 27, 3 (09 2019), 467–496. https://doi.org/10.1162/evco_a_00230
- [13] Su Nguyen, Yi Mei, and Mengjie Zhang. 2017. Genetic programming for production scheduling: a survey with a unified framework. *Complex & Intelligent Systems* 3, 1 (Feb. 2017), 41–66. <https://doi.org/10.1007/s40747-017-0036-x>
- [14] Su Nguyen, Mengjie Zhang, Mark Johnston, and Kay Chen Tan. 2013. Dynamic Multi-objective Job Shop Scheduling: A Genetic Programming Approach. In *Studies in Computational Intelligence*. Springer Berlin Heidelberg, 251–282. https://doi.org/10.1007/978-3-642-39304-4_10
- [15] Su Nguyen, Mengjie Zhang, and Kay Chen Tan. 2015. Enhancing genetic programming based hyper-heuristics for dynamic multi-objective job shop scheduling problems. In *2015 IEEE Congress on Evolutionary Computation (CEC)*. 2781–2788. <https://doi.org/10.1109/CEC.2015.7257234>
- [16] Su Nguyen, Mengjie Zhang, and Kay Chen Tan. 2017. Surrogate-Assisted Genetic Programming With Simplified Models for Automated Design of Dispatching Rules. *IEEE Transactions on Cybernetics* 47, 9 (2017), 2951–2965. <https://doi.org/10.1109/TCYB.2016.2562674>
- [17] John Park, Yi Mei, Su Nguyen, Gang Chen, Mark Johnston, and Mengjie Zhang. 2016. Genetic Programming Based Hyper-heuristics for Dynamic Job Shop Scheduling: Cooperative Coevolutionary Approaches. In *Lecture Notes in Computer Science*. Springer International Publishing, 115–132. https://doi.org/10.1007/978-3-319-30668-1_8
- [18] John Park, Yi Mei, Su Nguyen, Gang Chen, and Mengjie Zhang. 2017. An Investigation of Ensemble Combination Schemes for Genetic Programming based Hyper-heuristic Approaches to Dynamic Job Shop Scheduling. *Applied Soft Computing* 63 (11 2017). <https://doi.org/10.1016/j.asoc.2017.11.020>
- [19] John Park, Yi Mei, Su Nguyen, Gang Chen, and Mengjie Zhang. 2018. Investigating a Machine Breakdown Genetic Programming Approach for Dynamic Job Shop Scheduling. In *Lecture Notes in Computer Science*. Springer International Publishing, 253–270. https://doi.org/10.1007/978-3-319-77553-1_16
- [20] John Park, Su Nguyen, Mengjie Zhang, and Mark Johnston. 2015. Evolving Ensembles of Dispatching Rules Using Genetic Programming for Job Shop Scheduling. 92–104. https://doi.org/10.1007/978-3-319-16501-1_8
- [21] Michael L. Pinedo. 2012. *Scheduling*. Springer US. <https://doi.org/10.1007/978-1-4614-2361-4>
- [22] Lucija Planinić, Marko Đurasević, and Domagoj Jakobović. 2021. On the Application of ϵ -Lexicase Selection in the Generation of Dispatching Rules. In *2021 IEEE Congress on Evolutionary Computation (CEC)*. 2125–2132. <https://doi.org/10.1109/CEC45853.2021.9504982>
- [23] Riccardo Poli, William B. Langdon, and Nicholas Freitag McPhee. 2008. *A Field Guide to Genetic Programming*. Lulu Enterprises, UK Ltd.
- [24] Mateja Dumić and Domagoj Jakobović. 2021. Ensembles of priority rules for resource constrained project scheduling problem. *Applied Soft Computing* 110 (2021), 107606. <https://doi.org/10.1016/j.asoc.2021.107606>
- [25] Marko Đurasević and Domagoj Jakobović. 2020. Automatic design of dispatching rules for static scheduling conditions. *Neural Computing and Applications* 33, 10 (Aug. 2020), 5043–5068. <https://doi.org/10.1007/s00521-020-05292-w>
- [26] Marko Đurasević and Domagoj Jakobović. 2017. Comparison of ensemble learning methods for creating ensembles of dispatching rules for the unrelated machines environment. *Genetic Programming and Evolvable Machines* 19, 1-2 (April 2017), 53–92. <https://doi.org/10.1007/s10710-017-9302-3>
- [27] Marko Đurasević and Domagoj Jakobović. 2017. Evolving dispatching rules for optimising many-objective criteria in the unrelated machines environment. *Genetic Programming and Evolvable Machines* 19, 1-2 (Sept. 2017), 9–51. <https://doi.org/10.1007/s10710-017-9310-3>
- [28] Marko Đurasević and Domagoj Jakobović. 2019. Creating dispatching rules by simple ensemble combination. *Journal of Heuristics* 25, 6 (May 2019), 959–1013. <https://doi.org/10.1007/s10732-019-09416-x>
- [29] Marko Đurasević, Domagoj Jakobović, and Karlo Knežević. 2016. Adaptive scheduling on unrelated machines with genetic programming. *Applied Soft Computing* 48 (2016), 419–430. <https://doi.org/10.1016/j.asoc.2016.07.025>
- [30] Ivan Vlašić, Marko Đurasević, and Domagoj Jakobović. 2019. Improving genetic algorithm performance by population initialisation with dispatching rules. *Computers & Industrial Engineering* 137 (2019), 106030. <https://doi.org/10.1016/j.cie.2019.106030>
- [31] D.H. Wolpert and W.G. Macready. 1997. No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation* 1, 1 (1997), 67–82. <https://doi.org/10.1109/4235.585893>
- [32] Fangfang Zhang, Yi Mei, Su Nguyen, Kay Chen Tan, and Mengjie Zhang. 2021. Multitask Genetic Programming-Based Generative Hyperheuristics: A Case Study in Dynamic Scheduling. *IEEE Transactions on Cybernetics* (2021), 1–14. <https://doi.org/10.1109/TCYB.2021.3065340>
- [33] Fangfang Zhang, Yi Mei, Su Nguyen, and Mengjie Zhang. 2020. Guided Subtree Selection for Genetic Operators in Genetic Programming for Dynamic Flexible Job Shop Scheduling. In *Lecture Notes in Computer Science*. Springer International Publishing, 262–278. https://doi.org/10.1007/978-3-030-44094-7_17
- [34] Fangfang Zhang, Yi Mei, Su Nguyen, and Mengjie Zhang. 2021. Collaborative Multifidelity-Based Surrogate Models for Genetic Programming in Dynamic Flexible Job Shop Scheduling. *IEEE Transactions on Cybernetics* (2021), 1–15. <https://doi.org/10.1109/TCYB.2021.3050141>
- [35] Fangfang Zhang, Yi Mei, Su Nguyen, and Mengjie Zhang. 2021. Evolving Scheduling Heuristics via Genetic Programming With Feature Selection in Dynamic Flexible Job-Shop Scheduling. *IEEE Transactions on Cybernetics* 51, 4 (2021), 1797–1811. <https://doi.org/10.1109/TCYB.2020.3024849>
- [36] Fangfang Zhang, Yi Mei, Su Nguyen, Mengjie Zhang, and Kay Chen Tan. 2021. Surrogate-Assisted Evolutionary Multitask Genetic Programming for Dynamic Flexible Job Shop Scheduling. *IEEE Transactions on Evolutionary Computation* 25, 4 (2021), 651–665. <https://doi.org/10.1109/TEVC.2021.3065707>
- [37] Fangfang Zhang, Yi Mei, and Mengjie Zhang. 2019. Evolving Dispatching Rules for Multi-objective Dynamic Flexible Job Shop Scheduling via Genetic Programming Hyper-heuristics. In *2019 IEEE Congress on Evolutionary Computation (CEC)*. 1366–1373. <https://doi.org/10.1109/CEC.2019.8790112>
- [38] Fangfang Zhang, Yi Mei, and Mengjie Zhang. 2019. A Two-Stage Genetic Programming Hyper-Heuristic Approach with Feature Selection for Dynamic Flexible Job Shop Scheduling. In *Proceedings of the Genetic and Evolutionary Computation Conference (Prague, Czech Republic) (GECCO '19)*. Association for Computing Machinery, New York, NY, USA, 347–355. <https://doi.org/10.1145/3321707.3321790>
- [39] Marko Đurasević and Domagoj Jakobović. 2018. A survey of dispatching rules for the dynamic unrelated machines environment. *Expert Systems with Applications* 113 (2018), 555–569. <https://doi.org/10.1016/j.eswa.2018.06.053>
- [40] Marko Đurasević and Domagoj Jakobović. 2020. Comparison of schedule generation schemes for designing dispatching rules with genetic programming in the unrelated machines environment. *Applied Soft Computing* 96 (2020), 106637. <https://doi.org/10.1016/j.asoc.2020.106637>