Determining the Characteristic of Difficult Job Shop Scheduling Instances for a Heuristic Solution Method

Helga Ingimundardottir and Thomas Philip Runarsson

School of Engineering and Natural Sciences, University of Iceland hei2@hi.is and tpr@hi.is

Abstract. Many heuristic methods have been proposed for the jobshop scheduling problem. Different solution methodologies outperform other depending on the particular problem instance under consideration. Therefore, one is interested in knowing how the instances differ in structure and determine when a particular heuristic solution is likely to fail and explore in further detail the causes. In order to achieve this, we seek to characterise features for different difficulties. Preliminary experiments show there are different significant features that distinguish between easy and hard JSSP problem, and that they vary throughout the scheduling process. The insight attained by investigating the relationship between problem structure and heuristic performance can undoubtedly lead to better heuristic design that is tailored to the data distribution under consideration.

1 Introduction

Hand crafting heuristics for NP-hard problems is a time-consuming trial and error process, requiring inductive reasoning or problem specific insights from their human designers. Furthermore, within a problems class, such as job-shop scheduling, it is possible to construct problem instances where one heuristic would outperform another. Depending on the underlying data distribution, different heuristics perform differently, commonly known as the *no free lunch* theorem [1]. The success of a heuristic is how it manages to deal with and manipulate the characteristics of its given problem instance. So in order to understand more fully how a heuristic will eventually perform, one needs to look into what kind of problem instances are being introduced to the system. What defines a problem instance, e.g. what are its key features? And how can they help with designing better heuristics?

In investigating the relationship between problem structure and heuristic effectiveness one can research what [2] calls *footprints* in instance space, which is an indicator how an algorithm generalises over the instance space. This sort of investigation has also been referred to as *landmarking* [3]. It is evident from experiments performed in [2] that one-algorithm-for-all problem instances is not ideal. An algorithm may be favoured for its best overall performance, however

it was rarely the best algorithm available over various subspaces of the instance space. Thus when comparing different algorithms one needs to explore how they perform w.r.t. the instance space, i.e. their footprint.

In this study, the same problem generator is used to create 1,500 problem instances, however the experimental study in section 3 shows that MWRM works well/poorly on a subset of the instances. Since the problem instances are only defined by processing times and its permutation, the interaction between the two is important, because it introduces hidden properties in the data structure making it easy or hard to schedule with for the given algorithm. These underlying characteristics or features define its data structure. So a sophisticated way of discretising the instance space is grouping together problem instances that show the same kind of feature behaviour, in order to infer what is the feature behaviour between *good* and *bad* schedules.

It is interesting to know if the difference in the structure of the schedule is time dependent, is there a clear time of divergence within the scheduling process? Moreover, investigation of how sensitive is the difference between two sets of features, e.g. can two schedules with similar feature values yield completely contradictory outcomes, i.e. one poor and one good schedule? Or will they more or less follow the same path? This essentially answers the question of whether is is in fact feasible to discriminate between *good* and *bad* schedules using the currently selected features as a measure. If results are contradictory, it is an indicator the features selected are not robust enough to capture the essence of the data structure. Additionally, there is also the question of how can one define 'similar' schedules, what measures should be used? This paper describes some preliminary experiments with the aim of investigating the feasibility of finding distinguishing features corresponding to *good* and *bad* schedules in JSSP.

Instead of searching through a large set of algorithms (creating an algorithm portfolio) and determining which algorithm is the most suitable for a given subset of the instance space, as is generally the focus in the current literature [2, 4, 5], our focus is rather on a single algorithm and understanding *how* it works on the instance space – in the hopes of being able to extrapolate where it excels in order to aid its failing aspects.

The outline of the paper is as follows, in section 2 priority dispatch rules for the JSSP problem are discussed, what features are of interest and how data is generated. Followed by a discussion of how to map the relationship between problem structure and heuristic efficiency in section ??. A preliminary experimental study is presented in section 3. The paper concludes with a summary of main findings and points to future work.

2 Job-shop scheduling

The job-shop scheduling task considered here is where n jobs are scheduled on a set of m machines, subject to the constraint that each job must follow a predefined machine order and that a machine can handle at most one job at a time. The objective is to schedule the jobs so as to minimize the maximum

φ	Feature description
ϕ_1	processing time for job on machine
ϕ_2	start-time
ϕ_3	end-time
ϕ_4	when machine is next free
ϕ_5	current makespan
ϕ_6	work remaining
ϕ_7	most work remaining
ϕ_8	slack time for this particular machine
ϕ_9	slack time for all machines
ϕ_{10}	slack time weighted w.r.t. number of operations already assigned
ϕ_{11}	time job had to wait
ϕ_{12}	size of slot created by assignment
ϕ_{13}	total processing time for job
ϕ_{14}	total processing time for all jobs
ϕ_{15}	mean processing time for all jobs
ϕ_{16}	range of processing times over all jobs

Table 1. Feature space \mathcal{F} for JSSP. Features 1–13 can vary throughout the scheduling process w.r.t. tasks that can be dispatched next, however features 14–16 are static.

completion times, also known as the makespan. For a mathematical formulation of JSSP the reader is recommended [9].

2.1 Single-priority dispatching heuristic

Dispatching rules are of a construction heuristics, where one starts with an empty schedule and adds on one job at a time. When a machine is free the dispatching rule inspects the waiting jobs and selects the job with the highest priority. A survey of more than 100 of such priority rules was given in 1977 by [6]. In this paper however, only most work remaining (MWRM) dispatching rule will be investigated.

In order to apply a dispatching rule a number of features of the schedule being built must be computed. The features of particular interest were obtained from inspecting the aforementioned single priority-based dispatching rules. The temporal scheduling features applied in this paper are given in Table 1. These are not the only possible set of features, they are however built on the work published in [5,9] and deemed successful in capturing the essence of a JSSP data structure.

2.2 Data generation

Problem instances were generated stochastically by fixing the number of jobs and machines and sampling a discrete processing time from the uniform distribution U(1, 200). The machine order is a random permutation of $\{1, ..., m\}$. A total of 1,500 instances were generated for a six job and six machine job-shop problem.

In the experimental study the performance of the MWRM, μ_{MWRM} , and compared with its optimal makespan, μ_{opt} . Since the optimal makespan varies between problem instances the following performance measure is used:

$$\rho = \frac{\mu_{\rm MWRM}}{\mu_{\rm opt}}.$$
 (1)



Table 2. Features for *easy* and *hard* problems are drawn from the same data distribution (denoted by \cdot).

Table 3. Significant correlation (denoted by \cdot) for *easy* (left) and *hard* (right) problems and resulting ratio from optimality, ρ defined by (1).

3 Experimental study

In order to differentiate between problems, a threshold of a $\rho < 1.1$ and $\rho > 1.3$ was used to classify *easy* and *hard* problems. Of the 1500 instances created, 271 and 161 problems were classified *easy* and *hard*, respectively.

Table 2 reports where data distributions are the same (denoted by \cdot). From the table one can see that distribution for ϕ_1 , ϕ_6 , ϕ_{12} and ϕ_{16}) are (more or less) the same throughout the scheduling process. However there is a clear time of divergence for distribution of slacks; step 6 for ϕ_8 and step 12 for ϕ_9 and ϕ_{10} .

In order to find defining characteristics for *easy* and *hard* problems, a (linear) correlation was computed between features (on a step-by-step basis) to the resulting ratio from optimality. Significant features are reported in Table 3 for *easy* and *hard* problems, (denoted by \cdot). As one can see from the tables, the significant features for the different difficulties are varying. Some are commonly significant features across the tables (denoted by \bullet).

4 Discussion and Conclusion

From the experimental study it is apparent that features have different correlation with the resulting schedule depending in what stage it is in the scheduling process, implying that their influence varies throughout the scheduling process. And features constant throughout the scheduling process are not correlated with the end-result. There are some common features for both difficulties considered which define JSSP on a whole. However the significant features are quite different across the two difficulties, implying there is a clear difference in their data structure. The amount of significant features were considerably more for easy problems, indicating their key elements had been found. However, the features distinguishing hard problems were scarce. Most likely due to their more complex data structure their key features are of a more composite nature.

The feature attributes need to be based on statistical or theoretical grounds. Thus scrutiny in understanding the nature of problem instances is of paramount importance in feature engineering for learning. Which yields feedback into what features are important to devout more attention to, i.e. features that result in a failing algorithm. In general, this sort of investigation can undoubtedly be used in better algorithm design which is more equipped to deal with varying problem instances and tailor to individual problem instance's needs, i.e. a footprint-oriented algorithm.

Although this methodology was only implemented on a simple single-priority dispatching rule heuristic, the methodology is easily adaptable for more complex algorithms. The main objective of this work is to illustrate the interaction of a specific algorithm on a given problem structure and its properties.

References

- 1. Wolpert, D.H., Macready, W.G.: No free lunch theorems for optimization. IEEE Transactions on Evolutionary Computation 1(1) (1997) 67–82
- Corne, D., Reynolds, A.: Optimisation and generalisation: footprints in instance space. Parallel Problem Solving from Nature, PPSN XI (2011) 22–31
- 3. Pfahringer, B., Bensusan, H.: Meta-learning by landmarking various learning algorithms. on Machine Learning (2000)
- Smith-miles, K., Lopes, L.: Generalising Algorithm Performance in Instance Space: A timetabling case study. In Coello Coello, C.A., ed.: Learning and Intelligent Optimization, 5th International Conference, LION 5. Lecture Notes in Computer Science (2010)
- Smith-Miles, K., James, R., Giffin, J.: A knowledge discovery approach to understanding relationships between scheduling problem structure and heuristic performance. Learning and Intelligent (2009) 89–103
- Panwalkar, S., Iskander, W.: A Survey of Scheduling Rules. Operations Research 25(1) (1977) 45–61
- 7. Gu, Z., Rothberg, E., Bixby, R.: Gurobi Optimizer Version 4.5.2. software (2011)
- Rice, J.R.: The algorithm selection problem. Advances in Computers 15 (1976) 65–118
- Ingimundardottir, H., Runarsson, T.P.: Supervised Learning Linear Priority Dispatch Rules for Job-Shop Scheduling. In Coello Coello, C.A., ed.: Learning and Intelligent Optimization, 5th International Conference, LION 5. Lecture Notes in Computer Science (2010)