Information Sciences 373 (2016) 476–498

Contents lists available at ScienceDirect

Information Sciences

journal homepage: www.elsevier.com/locate/ins

Combining Monte-Carlo and hyper-heuristic methods for the multi-mode resource-constrained multi-project scheduling problem

Shahriar Asta[a], Daniel Karapetyan[a,b,*], Ahmed Kheiri[a,c], Ender Özcan[a], Andrew J. Parkes[a]

[a] University of Nottingham, School of Computer Science, Jubilee Campus, Wollaton Road, Nottingham, NG8 1BB, UK
[b] University of Essex, Institute for Analytics and Data Science, Wivenhoe Park, Colchester, CO4 3SQ, UK
[c] Cardiff University, School of Mathematics, Senghennydd Road, Cardiff, CF24 4AG, UK

A R T I C L E I N F O

Article history:
Received 29 October 2015
Revised 20 August 2016
Accepted 5 September 2016
Available online 7 September 2016

Keywords:
Metaheuristics
Hybrid heuristics
Hyper-heuristics
Monte Carlo tree search
Permutation based local search
Multi-project scheduling

A B S T R A C T

Multi-mode resource and precedence-constrained project scheduling is a well-known challenging real-world optimisation problem. An important variant of the problem requires scheduling of activities for multiple projects considering availability of local and global resources while respecting a range of constraints. A critical aspect of the benchmarks addressed in this paper is that the primary objective is to minimise the sum of the project completion times, with the usual makespan minimisation as a secondary objective. We observe that this leads to an expected different overall structure of good solutions and discuss the effects this has on the algorithm design. This paper presents a carefully-designed hybrid of Monte-Carlo tree search, novel neighbourhood moves, memetic algorithms, and hyper-heuristic methods. The implementation is also engineered to increase the speed with which iterations are performed, and to exploit the computing power of multicore machines. Empirical evaluation shows that the resulting information-sharing multi-component algorithm significantly outperforms other solvers on a set of "hidden" instances, i.e. instances not available at the algorithm design phase.

1. Introduction

Project scheduling has been of long-standing interest to academics as well as practitioners. Solving such a problem requires scheduling of interrelated activities (jobs), potentially each using or sharing scarce resources, subject to a set of constraints, and with one or several of a variety of objective functions. There are various project scheduling problems and many relevant surveys in the literature, e.g. see [4,20–22,38,47,64]. The best-known problem class is the Resource Constrained Project Scheduling Problem (RCPSP) in which activities have fixed usages of the resources, there are fixed precedence constraints between them, and often the objective is simple minimisation of the makespan (completion time of last activity). These problems have been proven to be NP-hard [2], and a well-known benchmark suite, PSPLIB, is provided in [28].

* Corresponding author.
E-mail addresses: sba@cs.nott.ac.uk (S. Asta), daniel.karapetyan@gmail.com (D. Karapetyan), kheiria@cardiff.ac.uk (A. Kheiri), exo@cs.nott.ac.uk (E. Özcan), ajp@cs.nott.ac.uk (A.J. Parkes).

http://dx.doi.org/10.1016/j.ins.2016.09.010
0020-0255/

A generalisation of the RCPSP is to also consider 'Multi-mode RCPSP' (MRCPSP) in which activities can be undertaken in one of a set of modes, with each mode potentially using different sets of resources. Furthermore, there are many options besides makespan for the objective function(s); a typical one is that a weighted sum of completion times is minimised. As common in optimisation problems, exact methods perform best on smaller instances and on larger instances heuristics and metaheuristics become necessary. Recent works on the MRCPSP range from exact approaches, such as, MILP [32], and branch-and-bound [58], to metaheuristics, such as, differential evolution [11], estimation of distribution algorithms [60], evolutionary algorithms [14,17,54], swarm intelligence methods [30], and others [8,59].

This paper presents our winning approach submitted to MISTA 2013 challenge[1] on a further extension called 'multi-mode resource-constrained multi-project scheduling' (MRCMPSP) and the results on the associated benchmark/competition instances. The full description of this problem domain can be found on the competition website and in [61]; however, for completeness we also summarise it in Section 2. The broad aim is to schedule a set of different and partially interacting projects, with each project consisting of a set of activities. There are no precedence constraints between the activities of different projects however they can compete for resources. Also, the objective function is extended to be a mix of a kind of weighted completion time and makespan. The MRCMPSP is hence interesting in that it has a mix of structures and requirements that are a step towards modelling the complexity of real-world scheduling problems. The high real-world relevance of the multi-project version of scheduling is well-known, e.g. a survey [34] found that "84% of the companies which responded to the survey indicated that they worked with multiple projects". However, the majority of scheduling work is on the single project version, though there is some existing work on the multi-project case, e.g. see [18,31,34,35].

Our approach searches the space of sequences of activities, from which schedules are constructed and then the quality of each schedule is evaluated using the objective function. The search process on the set of sequences operates in two phases in a "construct and improve" fashion. In the first phase, a heuristic constructor creates initial sequences of activities. A novel proposal in this paper is to investigate the overall global structure of the solutions and use this to motivate constructing the initial sequences using a Monte-Carlo Tree Search (MCTS) method, e.g. see [3]. This construction phase is followed by an improvement phase which makes use of a large and diverse set of heuristic neighbourhood moves. The search process during the improvement phase is carefully controlled by a combination of methods arising from a standard metaheuristic, namely memetic algorithm, and also an extension of existing hyper-heuristic components [27,45].

There is an interesting potential for dual views of the overall problem. It is defined as a multi-project problem, but it can be also viewed as a single project (multi-mode) RCPSP, in which the precedence graph has a particular structure, consisting of disjoint clusters. There is a sense in which we work with both views together. Some neighbourhood moves treat the problem in a single-project fashion and work on the constituent activities; other neighbourhood operators explicitly consider the multi-project nature of the problem, and focus on moves of projects. Both views, and kinds of operators, are used and work together to improve the overall project-level structure as well as the detailed activity level structure. A discussion and a computational study on both approaches can be found in [34].

The primary contributions of this paper are:

- Observation and investigation of how the primary objective function being essentially a "sum of project completion times" leads to good solutions having inherently different structure to those with makespan as the primary objective. In particular, minimisation of project completion times subject to limited global resource results in partial ordering of projects; this does somewhat reduce the effective size of the search space, but also may lead to good solutions being more widely separated. Understanding of this significantly affected our algorithm design, including an MCTS construction method aiming to create solutions having such structure.
- Novel neighbourhood moves, including those that are designed specifically for smoother navigation through the search space of the multi-project extension of MRCPSP – reflecting our observation of the effect that the main objective function has on the solution structures.
- An adaptive hybrid hyper-heuristic system to effectively control the usage of the rich set of neighbourhood moves.
- Evidence of the effectiveness based on successful results on a range of benchmark problems. This includes winning a competition, in which some problems were hidden at the algorithm design/tuning phase. We also tested our algorithm on single-project instances from PSPLIB. Although our algorithm was not designed to work on single-project instances, it demonstrated good performance in these tests, and was competitive with the state-of-the-art methods tailored to the single-project case. Furthermore, it improved 3 best solutions on these PSPLIB instances during these experiments.

These contributions are directed towards a system that is both robust and flexible; with the potential to be effective at handling a wide variety of problem requirements and instances. Arguably, one of the lessons of this paper is that greater complexity and richness of such scheduling problems needs to be matched with a greater complexity and richness of the associated algorithms; especially when not all instances are known in advance, and so algorithms should not over-specialise to a particular data set.

Regarding the structure of the paper; in Section 2 we describe the problem to be solved. In Section 3 we discuss how we have carefully chosen the appropriate data structures and implemented algorithms operating with those data structures efficiently in order to construct the schedule from a given sequence as fast as possible. (To build an effective system, one has

[1] http://gent.cs.kuleuven.be/mista2013challenge/.