Computers & Operations Research 43 (2014) 36–49

Contents lists available at ScienceDirect

Computers & Operations Research

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

A Two-Stage Decomposition of High School Timetabling applied to cases in Denmark

Matias Sørensen [a,b,*], Florian H.W. Dahms [c]

[a] Section of Operations Research, Department of Management Engineering, Technical University of Denmark, DK-2800 Kgs. Lyngby, Denmark
[b] MaCom A/S, Vesterbrogade 48, 1., DK-1620 Copenhagen V, Denmark
[c] Chair of Operations Research, RWTH Aachen University, Kackertstrasse 7, 52072 Aachen, Germany

ARTICLE INFO

Available online 5 September 2013

Keywords:
High school timetabling
Integer programming
Decomposition
Bipartite matching

ABSTRACT

Integer Programming (IP) has been used to model educational timetabling problems since the very early days of Operations Research. It is well recognized that these IP models in general are hard to solve, and this area of research is dominated by heuristic solution approaches. In this paper a *Two-Stage Decomposition* of an IP model for a practical case of high school timetabling is shown. This particular timetabling problem consists of assigning lectures to both a timeslot and a classroom, which is modeled using a very large amount of binary variables. The decomposition splits this model into two separate problems (Stage I and Stage II) with far less variables. These two separate problems are solved in sequence, such that the solution for the Stage I model is given as input to the Stage II model, implying that irreversible decisions are made in Stage I. However, the objective of the Stage II model is partly incorporated in the Stage I model by exploiting that Stage II can be seen as a *minimum weight maximum matching* problem in a bipartite graph. This theoretically strengthens the decomposition in terms of global optimality. The approach relies on Hall's theorem for the existence of matchings in bipartite graphs, which in its basic form yields an exponential amount of constraints in the Stage I model. However, it is shown that only a small subset of these constraints is needed, making the decomposition tractable in practice for IP solvers. To evaluate the decomposition, 100 real-life problem instances from the database of the high school ERP system Lectio are used. Computational results show that the decomposition performs significantly better than solving the original IP, in terms of both found solutions and bounds.

1. Introduction

Integer Programming (IP) has been used to model educational timetabling problems since the very early days of Operations Research (see e.g. [16,19]). It is well recognized that these IP models in general are hard to solve (most forms of educational timetabling are in fact $\mathcal{NP}$−hard [4]), and this area of research is dominated by heuristic solution approaches.

In this paper a large IP model for a real-world case of high school timetabling is considered, which has previously been shown to be a challenge for state-of-the-art MIP solvers. We consider a basic version of this IP, which includes the essential constraints of most timetabling problems. An innovative decomposition of this model is shown, which proves to be more efficient to solve.

When facing a hard IP model, decomposition is a commonly used tool to help speed up the solution procedure. Perhaps the most successful decomposition method in recent years is Column Generation (CG). However, not many papers on CG and timetabling models are found in the literature, and it seems that only relatively small instances have been attempted. Papoutsis et al. [23] use CG to solve a Greek case of high school timetabling, with the largest instance containing 9 class section, 21 teachers and 306 teaching hours. Santos et al. [28] handle larger instances, but only generate lower bounds. Qualizza and Serafini [27] describe a CG procedure for a university timetabling problem with 63 courses and 25 timeslots. The real-world instances considered in this paper are of much larger size.

A crucial part of a CG procedure is the identification of a block-diagonal structure in the problem, otherwise the CG procedure is most likely not efficient. For the high school timetabling problem described in this paper, it has not been possible to identify such a structure. Therefore this paper shows a different type of decomposition, a *Two-Stage Decomposition* (TSD). Such an approach was first used for timetabling applications in Lach and Lübbecke [17,18] with great success for the *curriculum-based university course* timetabling problem. The goal of this paper will be to modify the aforementioned approach to be applicable for the high school timetabling problem – giving special attention to the high school system in Denmark.

* Corresponding author at: Section of Operations Research, Department of Management Engineering, Technical University of Denmark, DK-2800 Kgs. Lyngby, Denmark. Tel.: +45 45254800.
E-mail addresses: msso@dtu.dk, sorensen.matias@gmail.com (M. Sørensen).

0305-0548/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.cor.2013.08.025

The considered timetabling problem essentially consists of assigning lectures to rooms and timeslots, which is commonly modeled using a very large amount of binary variables. There are three key points to the TSD:

- By substitution, the total amount of variables is significantly reduced, while linearity is maintained.
- Instead of solving the entire model at once, it can be solved in a two-stage fashion. i.e. both the set of variables and constraints are divided into two distinct sets, corresponding to two smaller IPs (denoted *Stage I* and *Stage II*, respectively).
- It will be evident that, except for two soft-constraints, this decomposition maintains optimality of the original model.

The outline of the TSD is to first solve Stage I, which provides a solution where lectures are assigned to timeslots. This partial solution is given as input to Stage II, which will assign rooms to the lectures, obtaining a solution for the original problem. The drawback of this decomposition is that the timeslots assigned to lectures in Stage I are considered as fixed by the Stage II model, which might prevent an optimal allocation of rooms to lectures. However, by exploiting the structure of the Stage II model, the Stage I model can be constrained in such a way that some penalties for assigning rooms to lectures are handled implicitly. Note that if all penalties for room assigning could be handled implicitly, the approach would be exact. However, two soft-constraints are not fully incorporated, so only a lower bound on the room penalties is known by the Stage I model. In fact, one of these soft-constraints are not handled at all by the described approach. Despite this, it seems likely that incorporating this lower bound in the Stage I model will provide better results overall (assuming that computing the lower bound does not have very bad influence on the computational efforts of the used IP solver). So instead of the Stage I model being completely unaware of the penalties for room allocation, it seems better to at least incorporate some of them. Furthermore, the decomposition of the problem into two smaller problems presents a big advantage in terms of reduction in the number of variables. Therefore the overall benefits of the TSD out-weight the downsides, and computational results will show that it is indeed way more effective than solving the original IP.

The contributions of this paper are the following: (1) It is shown that the approach from Lach and Lübbecke [17] can also be applied to a high school timetabling problem originating from a practical setting, and by extensive computational results it is argued that the TSD is more effective than solving the original IP. Notice that a similar decomposition is briefly mentioned in Sørensen and Stidsen [31] for the same high school timetabling problem, but this paper enhances the approach such that the theoretical maximum gap from optimality is narrowed. The presented approach turns out to be the most efficient exact algorithm for the problem so far. (2) Generally, it is shown how this type of decomposition can be applied to models with set-packing structure, by modifying the underlying equations originating from Hall's Theorem for matchings in bipartite graphs. (3) It is shown how the room-priorities of lectures can be handled, by adding a lower bound on the corresponding penalties to the Stage I model. This facilitates the quality of the solutions found, as shown by the computational results.

We expect that the basic structure required for applying the TSD can be found in other timetabling problems as well, and therefore the decomposition can potentially be used more broadly than the case of high school timetabling shown in this paper. This seems likely because the essential constraints used in the decomposition are among the most common ones found in timetabling problems.

The paper is structured as follows. First related papers are described in Section 2. The basic IP model is introduced in Section 3, including the essential constraints. Section 4 shows the TSD of this model, and derives the lower bound on room allocation penalties for the Stage I model. Section 5 extends the model so it encapsulates a practical version of the high school timetabling problem, defined by the online high school administration system Lectio. Section 6 shows computational results, comparing the decomposition to previous approaches for 100 problem instances taken from the Lectio database. Section 7 concludes on our findings.

2. Related work

Integer Programming has been used to model various educational timetabling problems. However, heuristics are still the most popular method for these problems, see surveys [29,24]. In terms of IP, de Werra [12] describes what is called 'a simple model' for the class-teacher problem, and existence of solutions is proven under certain circumstances using graph theoretical models. The problems considered are feasibility problems, and soft constraints are not added to the models. Birbas et al. [5] describes a 'fully defined' IP model for Greek secondary schools, which is evaluated on five different schools with success. Avella et al. [1] formulates an IP model which is used to solve small instances of various origin. The IP is solved within a VLSN algorithm, with good results.

For the related university course timetabling problem, Daskalaki et al. [11] presents a model which schedules courses to timeslots and classrooms, using many so called *operational rules*. Three different problem instances of significant size are all solved to optimality using CPLEX. MirHassani [21] describes the problem for an Iranian university, and reports good results by applying the XA solver. In Dimopoulou and Miliotis [13] an IP model is used to solve the timetabling problem for The Athens University of Economics and Business.

Decomposition of IP models for educational timetabling is not a very well researched topic. Burke et al. [7] state that: *In the timetabling community, the "times first, rooms second" decomposition is a standard procedure*. However, it seems that this procedure has not been applied much in context of IP models. Burke and Newall [8] apply the procedure in context of an Evolutionary Algorithm for Examination Timetabling. In terms of multistage-decompositions, the importance of Lach and Lübbecke [17,18] has already been discussed. Carter [9] presents an interesting decomposition algorithm for course timetabling with elective courses. Stating the problem in terms of a vertex coloring problem facilitates the decomposition of the graph by cliques, such that the subproblem defined by each clique is solved separately.

In Burke et al. [7], experiments are conducted on disabling different combinations of soft-constraint penalties of the *Udine Course Timetabling Problem*, including one where all room penalties on room allocation are disabled. Thereby a similar decomposition to that of Lach and Lübbecke [18] is obtained.

Daskalaki and Birbas [10] presents an approach for university timetabling, where courses are first assigned to days (skipping some requirements for compactness), and in the following stage the timetable for each day is treated locally (enforcing the compactness). Convincing computational results are shown. In Birbas et al. [6], a high school timetabling problem is solved by first allocating 'work shifts' to teachers, and then solving the actual timetabling problem. This is related to the type of decomposition performed in this paper. Badri [2] uses a related approach for university course timetabling, where *faculties* are first assigned to courses, and then faculties are assigned to timeslots. However the problems solved are tiny.