Conflicts in Examination Timetabling under Uncertainty

Bernd Bassimir · **Rolf Wanka**

Abstract In the literature the examination timetabling problem (ETTP) is mostly described as a post enrollment problem (PE-ETTP). As such it is known at optimization time how many students will take an exam and consequently how big a room is needed for the exam and which exams should not be held at the same time because of overlapping student lists. In contrast some universities start their scheduling process before students register. As such the model is subject to uncertainty in respect to the number of students per exam and the conflicts between exams. In this work we focus on the uncertainty of conflicts between exams and introduce two soft criteria for handling this uncertainty. We show results for two real world instances taken from the School of Engineering at the Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU).

Keywords Examination Timetabling, Curriculum-Based Timetabling, Robust Optimization

1 Introduction

In every academic term, universities are faced with a number of different aspects of academic timetabling. Academic timetabling can be divided into two distinct, but similar problems, namely the *Course Timetabling Problem* (CTTP) prior to each term and the *Examination Timetabling Problem* (ETTP) at the end of the term.

In the literature the ETTP is often treated as a timetabling problem, where all data is provided as input to a scheduling algorithm and no uncertainty is present. This approach is sometimes called *post-enrollment* ETTP (PE-ETTP) or is at least treated as PE-ETTP, e. g., see [5,6,8]. First the students register for their different exams and after registration is finished the optimization takes place. In this approach the complete input is known, as we have the exact number of students registered for an exam and exams are in conflict, i. e., should not be scheduled at the same

Department of Computer Science, University of Erlangen-Nueremberg, Germany
E-mail: bernd.bassimir@fau.de · rolf.wanka@cs.fau.de

time, if there are students taking both exams. In practice however some universities generate their schedules for the exams before this registrations takes place to provide a time and date for the exams when students register, e. g., see [2,3,4]. Most often the values for the number of students per exam are taken from the last year as well as the conflicts that were present in the previous year.

In this work we focus on the implications of the uncertainty present in the conflict data. We discuss pitfalls when using an estimation approach for handling this uncertainty and propose new soft criteria that avoid these issues and show promising results on two real world instances of the School of Engineering at the FAU.

2 Model

Most of the following model is similar to the model presented in [7] except that we allow multiple rooms per exam, which already makes the subproblem of assigning rooms to exams in a fixed timeslot NP-complete, which can be shown by a reduction from 3-partition.

Definition 1 *An instance of the Examination Timetabling Problem is represented as follows.*

- $\mathscr{E}$*: A set of exams*
- $\mathscr{R}$*: A set of rooms*
- $\kappa : \mathscr{R} \to \mathbb{N}$*: the capacity available for each room and* $\kappa : \mathscr{P}(\mathscr{R}) \to \mathbb{N}$ *as the canonical extension such that for* $X \subseteq \mathscr{R} : \kappa(X) := \sum_{r\in X} \kappa(r)$
- $\nu \in \mathbb{N}^{\mathscr{E}}$*: A vector specifying for each exam how many students are attending*
- *Conflict matrix* $C \in \mathbb{N}^{\mathscr{E}\times\mathscr{E}}$*: specifying the number of overlapping students of two exams*
- *A set of hard constraints and a set of soft constraints with associated weights*
- *A number of available timeslots*

Note that the vector ν *and the conflict matrix* C *are subject to uncertainty in our case.*

To be feasible, a timetable must meet the following hard constraints:

(H1) each exam is assigned to exactly one timeslot
(H2) each exam is assigned to one or more rooms
(H3) two exams that are in conflict are not scheduled at the same time
(H4) no room is used at the same time by two different exams
(H5) the sum of capacities of the assigned rooms is larger than the number of students taking the exam

In our reduced model we use the soft constraints *two-in-a-row*, *two-in-a-day* and *period-spread* as defined in [7] and the robustness soft constraint introduced in [2]

(S1) *two-in-a-row*: If two exams are in conflict according to C they should not be assigned to two adjacent timeslots on the same day.
(S2) *two-in-a-day*: If two exams are in conflict according to C they should not be assigned to two timeslots on the same day. Note that we exclude the directly adjacent timeslot to avoid double counting.

(S3) *period-spread*: If two exams are in conflict according to C they should not be assigned to timeslots less than λ apart.

(S4) $\min\left\{\frac{\kappa(RP(e))}{\nu(e)} \,\middle|\, e \in \mathscr{E}, RP(e) \text{ is room pattern of } e \text{ in current timetable}\right\}$

Each violation of a soft criterion induces a penalty corresponding to the number of students that are in conflict given by conflict matrix C. For a feasible timetable, i. e., a timetable that satisfies hard constraints (H1) - (H5), we can formulate the objective value to be minimized as the weighted sum of the penalties induced by the soft constraints (S1) - (S3) and the weighted negative value of soft constraint (S4), as (S1) - (S3) have to be minimized and (S4) maximized.

3 New contribution

Hard constraint (H3) ensures for a feasible timetable, that for each student the chosen exams do not overlap. However in a pre-enrollment setting these choices are not known at scheduling time. It therefore becomes necessary to account for this uncertainty in the scheduling process. We call a conflict *active* iff after registration there is a student that takes both exams, otherwise the conflict is called *inactive*.

For a given major available lectures and therefore exams can usually be classified into two distinct categories. The first group consists of exams a student has to take, often called *mandatory*. Furthermore there might exist a portfolio of choices from which a student has to select a certain amount of lectures and exams respectively, called *elective*.

The strict robustness approach for this setting is to have no distinction between these two categories and enforce conflict freeness, i. e., (H3), for all possible combinations of exams of a major. However this approach leads to large numbers of conflicts per major and therefore no feasible timetable might exist for the model. This is the case for the School of Engineering at FAU.

It becomes necessary to limit the conflicts that are taken into account in the scheduling process. The resulting question is which conflicts to consider and which to ignore. As for *mandatory* exams all students of the major have to take these exams in a specified term and therefore all conflicts involving these exams have to be considered such that all students can take the exam in this term. Consequently there is no uncertainty if such a conflict is active or inactive.

For conflicts between *elective* exams there is uncertainty whether the conflict is active or inactive. The first possible solution for this uncertainty is to use data from previous terms to estimate if a conflict is active and to take all conflicts that are estimated to be active into account in the scheduling process. This was the approach we used in previous works, when talking about the uncertainty in the number of students per exam.

However there are a few inherent problems with this approach. The first problem when using data for estimations, especially when taken from more than one previous term is that almost all conflicts might be active and the model therefore becomes infeasible again. The other problem, which is far more problematic is that the resulting timetables might get biased against certain possible choices. If a conflict is considered inactive the exams might be scheduled in the same timeslot, therefore no student can attend both exams. In the following year the estimation will again treat this conflict as inactive as no student has taken both exams in the previous term. This problem might not even be visible to the responsible scheduler, however still reducing the acceptance of the calculated timetables by the student body.

Instead of estimating which *elective* conflicts are active and enforcing them via hard constraint (H3) we consider all elective conflicts to be inactive in regard to feasibility. For these conflicts we introduce a new soft criterion to minimize the number of students that cannot take all exams they want.

Given a feasible timetable TT as a set of timeslots we want to minimize the following soft constraint.

$$\sum_{T \in TT} \sum_{\substack{e_1, e_2 \in T \\ e_1 \neq e_2}} \bar{\mathrm{E}}[\#\text{students attending } e_1 \text{ and } e_2] \tag{S5}$$

$\bar{\mathrm{E}}$ is an estimator for how many students induce a given conflict between two exams. Therefore (S5) measures the estimated number of students with a conflict in timetable TT. Using this soft criterion instead of enforcing all estimated conflicts between elective exams through hard constraint (H3) there always exits a feasible timetable.

However this soft criterion will not prevent the issue of bias in regard to the estimations. To also address this issue we formulate a second soft criterion.

$$\sum_{T \in TT} \sum_{\substack{e_1, e_2 \in T \\ e_1 \neq e_2}} \max\{\bar{\mathrm{E}}[\#\text{students attending } e_1 \text{ and } e_2], \mathbb{1}_{\{e_1, e_2\} \text{ is elective conflict}}\} \tag{S6}$$

Instead of using only the sum of estimated students for elective exam conflicts we consider all elective exam conflicts with a value of at least 1. Therefore even if in the last terms no student did choose both exams our optimization will try to schedule the exams conflict free and thus enabling students to choose both.

4 Preliminary Results

To test the performance of the introduced soft criteria and to evaluate the impact on the overall objective value, we used two real world data instances taken from the School of Engineering at FAU. In our experiments we used a simulated annealing algorithm, with a kempe-exchange neighborhood. A more detailed description can be found in [2].

Table 1 and 2 show the arithmetic mean rounded to the nearest integer of 24 runs using the simulated annealing algorithm for the summer term 2018 and the winter term 2017 at the School of Engineering at FAU. The solutions are based on our introduced model using estimations for the number of students attending an exam as described in [1] and the arithmetic mean over the previous years for the elective conflicts and the minimum of the attending students for the mandatory conflicts. Each solution is then evaluated with the actual values of the corresponding term. The values for the soft criteria (S1) - (S3) are shown, with the final value being the number of students that have elective exam conflicts.

Version	(S1)	(S2)	(S3)	#Conflicting
All Conflicts $\neq 0$	70	171	1720	73
All Conflicts $\neq 0$ + (S6)	132	146	1919	20
Only Mandatory + (S5)	78	163	1792	74
Only Mandatory + (S6)	142	141	1937	20

Table 1: Results on the Summer Term 2018 instance of the School of Engineering at FAU.

Version	(S1)	(S2)	(S3)	#Conflicting
All Conflicts $\neq 0$	90	170	1832	109
All Conflicts $\neq 0$ + (S6)	168	169	2112	33
Only Mandatory + (S5)	106	182	1854	101
Only Mandatory + (S6)	173	178	2146	38

Table 2: Results on the Winter Term 2017 instance of the School of Engineering at FAU.

Preliminary results show only a slight decrease in the objective values when using the soft criterion (S5) compared to the solution when considering the estimated elective conflicts as active and enforcing them with hard constraint (H3). For the soft criterion (S6) preliminary results show that we can reduce the number of students that could not choose their intended exams by a factor of around 3. For this soft criterion the objective value shows a larger increase. Partly this increase is a result of the more robust solution, as the soft constraints measure the number of students with conflicting exams in different timeslots, therefore if two exams are in conflict in the same timeslot they do not distribute to the objective value. As such we can argue that the price of robustness is tolarable.

5 Conclusion

In this work we address the issue of uncertainty in the ETTP model in regard to the conflicts, when using a pre-enrollment approach. We discuss pitfalls when using estimations for the conflicts between exams and introduce two soft criteria (S5) and (S6) to address this uncertainty and the resulting issues. We present a case study for two real world instances at FAU and give preliminary results that show only a moderately large increase in the objective value, while providing a large decrease in the number of conflicts.

References

1. Bernd Bassimir and Rolf Wanka. Probabilistic Curriculum-based Examination Timetabling. In *Proc 12th International Conference on the Practice and Theory of Automated Timetabling (PATAT)*, pages 273–285, 2018.

2. Bernd Bassimir and Rolf Wanka. Robustness Approaches for the Examination Timetabling Problem under Data Uncertainty. In *Proc. 9th Multidisciplinary International Conference on Scheduling: Theory and Applications (MISTA)*, pages 381–395, 2019.
3. Alejandro Cataldo, Juan-Carlos Ferrer, Jaime Miranda, Pablo A. Rey, and Antoine Sauré. An integer programming approach to curriculum-based examination timetabling. *Annals of Operations Research*, 258(2):369–393, Nov 2017.
4. Peter Demeester, Burak Bilgin, Patrick De Causmaecker, and Greet Vanden Berghe. A hyperheuristic approach to examination timetabling problems: benchmarks and a new problem from practice. *Journal of Scheduling*, 15(1):83–103, Feb 2012.
5. Michael Eley. Ant algorithms for the exam timetabling problem. In *Proceedings of the 6th International Conference on Practice and Theory of Automated Timetabling VI*, PATAT'06, pages 364–382, Berlin, Heidelberg, 2007. Springer-Verlag.
6. Christos Gogos, Panayiotis Alefragis, and Efthymios Housos. An improved multi-staged algorithmic process for the solution of the examination timetabling problem. *Annals of Operations Research*, 194(1):203–221, Apr 2012.
7. Barry McCollum, Paul McMullan, Andrew J. Parkes, Edmund K. Burke, and Rong Qu. A new model for automated examination timetabling. *Annals of Operations Research*, 194(1):291–315, Apr 2012.
8. Tomáš Müller. Itc2007 solver description: A hybrid approach. 172:429–446, 01 2008.