The invigilator assignment problem

Hannah Verplancke[1,2 [0000−0001−9104−3169]], David Van Bulck[1,2 [0000−0002−1222−4541]], Veronique Limère[1,2 [0000−0001−6858−6309]], and Nico André Schmid[1,3 [0000−0002−2432−5913]]

[1] Ghent University, Ghent, Belgium
[2] Core lab CVAMO, FlandersMake@UGent, Belgium
hannah.verplancke@ugent.be
[3] IESEG School of Management, Univ. Lille, CNRS, UMR 9221 - LEM - Lille Economie Management, F-59000 Lille, France

Keywords: Invigilator assignment, Examination scheduling.

1 Introduction

The invigilator assignment problem (IAP) aims at providing an invigilator schedule, appointing sufficient invigilators to all exams of a predetermined examination schedule. These invigilators perform the required logistic and administrative tasks when students take exams, such as distributing copies, preventing students from cheating, or answering questions. Even though the need for research on the IAP had already been raised in the nineties (see, e.g., [2]), only recently the problem gained more importance in literature (e.g., [5] and [6]).
As stated above, the IAP uses a predetermined examination schedule, the outcome of the widely studied examination timetabling (ETT) problem, as an input. A first variant of the ETT, known as uncapacitated ETT, assigns a set of examinations to a fixed number of periods while avoiding students having to take two exams at the same time [3]. A second variant, the capacitated ETT, also considers the assignment of exams to examination rooms with individual sizes, thereby limiting the number of exams that can be taking place concurrently [7]. The assignment of exams to timeslots and rooms is managed on the university level, while the assignment of invigilators is situated on the faculty level. Moreover, the ETT problem is typically solved long before the actual exams take place, so students can optimally plan, whereas the IAP can be solved until a few days before the exams, to capture the invigilator availabilities as accurate as possible. Therefore, a sequential solution approach has been chosen to address those problems. For an extensive overview of educational timetabling problems, including the ETT problem, we refer to [4].
Previous research on the IAP mostly focused on providing formulations and heuristic solutions for a variety of case-specific problems. A general formulation seems to be missing, as well as insights on the complexity of the problem at hand. In this research, we want to summarize common elements identified through the different previous papers and provide insights on the complexity of the proposed base model and some interesting variants.

2 Problem description

We identify the following base problem, which forms the foundation of almost any IAP studied in literature. In the IAP, we are presented with a set of examination periods P, invigilators I, and exams E. Each exam $e \in E$ is preassigned to a period $p \in P$ and requires d_e invigilators. Assigning invigilator i to exam e results in cost $c_{i,e}$. This cost represents the preference of the invigilator, where a more preferred time slot induces a lower cost and a less preferred time slot induces a higher cost. An understaffing cost σ is incurred for each invigilator allocated less than needed. Notice that this cost will be much higher than any of the preference costs $c_{i,e}$. The IAP is to assign invigilators to exams in a way that minimizes total costs related to invigilator assignments and potential understaffing. Additionally, the solution must satisfy two hard constraints: C1 and C2.

C1 Each invigilator is limited to supervising at most one exam in any given period, ensuring there are no scheduling conflicts.
C2 The total number of exams assigned to invigilator i should be between $\underline{w}_i$ and $\overline{w}_i$.

These constraints may result in incomplete assignments for some exams, as captured in the model by the understaffing costs. In practice, these remaining invigilators are staffed manually after personal correspondence with invigilators who have less than $\overline{w}_i$ assignments and verifying their availability. If desired, this decision could also be included in the model by removing the understaffing cost and including a preference cost equal to the understaffing cost for those time slots outside the initially provided invigilators' preferences. In practice, this also requires manual verification of the invigilators' availability. A final alternative could be to treat C2 as a soft constraint, but according to our experience this could result in very unbalanced solutions.

Theorem 1. *The IAP can be solved in polynomial time.*

Proof. We show that IAP can be modelled as a special case of the Minimum Cost Network Flow Problem (MCNFP), which can be solved in polynomial time [1].
MCNFP
Input. A directed graph $G(V, A)$ with a net supply b_i for each vertex $i \in V$, and a capacity $u_{i,j}$ and cost $c_{i,j}$ for each arc $(i, j) \in A$.
Output. A minimum cost flow respecting the net supply at each vertex and the capacity at each arc.

To construct the graph G, we start by creating a node for each invigilator $i \in I$ with a net supply of $\overline{w}_i$ and for each invigilator also $|P|$ invigilator-period nodes with zero net supply. Each invigilator node is connected to its corresponding period nodes with arcs of zero cost and unit capacity.

Next, for each exam e scheduled in period p, an exam node is created with net supply $-d_e$, linked to all invigilator-period nodes associated with period p with unit capacity arcs costing $c_{i,e}$.

Finally, we add two dummy nodes. A first dummy node, x, has net supply equal to $s = \sum_{e \in E} d_e$ and is connected to each of the exam nodes using an arc with capacity d_e and cost σ, indicating potential understaffing for exams. Secondly, to balance the network, we include a dummy exam node y with a net supply of $-t = \sum_{i \in I} \overline{w}_i$ and link it with invigilator nodes using a zero-cost arc with capacity $\overline{w}_i - \underline{w}_i$. Finally, we add a zero-cost arc from the dummy-invigilator node to the dummy exam node with capacity v. □

In order to solve the associated MCNFP instance, we use a standard linear programming formulation, which is known to be totally unimodular (i.e., it is sufficient to consider the LP-relaxation to obtain integral solutions) [1]. This way, we obtain optimal solutions within seconds, even for instances involving hundreds of exams. Although, this base problem recurs in most of the studied literature, different objectives and constraints are needed to tackle context-specific problems, making it hard to combine all of them in a single model or definition.
In this talk, we show how the addition of certain constraints or objectives turns the problem NP-hard, e.g., including fairness costs to limit the difference in satisfaction of different invigilators. We contribute to the literature by discussing the complexity of the base model and some interesting variants.

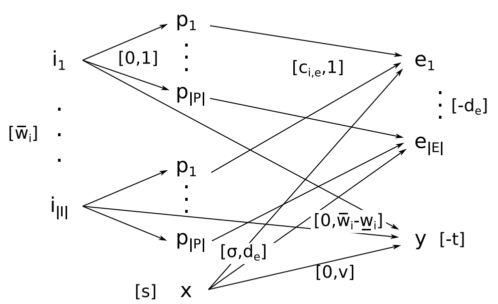

Fig. 1: Network G used to transform IAP into an instance of MCNFP. Numbers between brackets denote the net supply at each nonde, and the cost and capacity at arcs.

3 Case study

A variant of the model above was used to schedule invigilators atthe faculty of Economics and Business Administration at Ghent University. The tests were conducted on an instance with 195 exams and a total demand of 1140 invigilators, spread over 67 time periods and 321 invigilators. Prior to the implementation of our model, invigilators had to fill in a shared file to immediately register for specific exams and rooms in a First Come, First Serve (FCFS) manner, oftentimes resulting in unfulfilled demand. To assess the benefit of an optimization-based solution, we compare it to various simulation runs using the FCFS principle that varied the arrival sequence of invigilators for a predefined set of available time slots per invigilator. Invigilators would be available for 5, 10, or 20 time slots and register for 3-4 exams as long as there were exams during such time slots that still required invigilators. This approach is compared to the optimization-based approach where the invigilators' preferences are used to find an optimal solution. Table 1 compares the optimal solution to the simulation runs by reporting the percentage of demand that is covered by available invigilators (Coverage %) and the average preference values per assignment (Avg. pref.) with lower preference values indicating more preferable assignments. One can observe the high variability in performance depending on the random sequence created for the FCFS approach, indicating a low robustness of this approach. The optimal approach significantly reduces the number of missing invigilators compared to any result obtained by the FCFS principle. In some cases, this increases average preference values, but this effect is nullified when the number of preferences per invigilator is sufficiently large. All experiments were solved optimally on a laptop with 8GB RAM and an Intel m5-6Y57 processor with 1.1 GHz in less than 3 seconds.

During the pandemic period we implemented an adapted version of this base model, minimizing invigilator contacts instead of preference values. A contact would occur for each pair of supervisors assigned to the same examination room at the same time. To solve this computationally intractable problem, we developed a fix-and-optimize heuristic which will be presented during this talk along with results from the real-life case study at the faculty of Economics and Business Administration of Ghent University, Belgium.

Table 1: FCFS results over 10,000 simulation runs compared to optimal solution.

Approach	Pref./Inv.	Coverage %	Avg. pref.
Optimal	5	75.70	2.59
FCFS	5	[73.07-75.18]	[1.8-1.95]
Optimal	10	96.74	3.85
FCFS	10	[86.49-90.70]	[3.02-3.34]
Optimal	20	100	4.22
FCFS	20	[92.89-97.71]	[4.15-4.86]

References

1. Ahuja, R., Magnanti, T., Orlin, J.: Network flows: theory, algorithms, and applications. Prentice-Hall (1993)
2. Burke, E., Elliman, D., Ford, P., Weare, R.: Examination timetabling in British universities: A survey. In: Burke, E., Ross, P. (eds.) Practice and Theory of Automated Timetabling. pp. 76–90. Springer Berlin Heidelberg, Berlin, Heidelberg (1996)
3. Carter, M.W.: OR Practice – a survey of practical applications of examination timetabling algorithms. Oper. Res. **34**, 193–202 (1986)
4. Ceschia, S., Di Gaspero, L., Schaerf, A.: Educational timetabling: Problems, benchmarks, and state-of-the-art results. Eur. J. Oper. Res. **308**, 1–18 (2022)
5. Çimen, M., Belbağ, S., Soysal, M., Sel, Ç.: Invigilators assignment in practical examination timetabling problems. International Journal of Industrial Engineering **29**, 351–371 (2022)
6. Kahar, M.N.M., Kendall, G.: Universiti Malaysia Pahang examination timetabling problem: scheduling invigilators. Journal of the Operational Research Society **65**(2), 214–226 (2014)

7. McCollum, B., McMullan, P., Burke, E.K., Parkes, E.J., Qu, R.: The second international timetabling competition: Examination timetabling track. Tech. rep., Queen's University, Belfast (UK) (2007)