High School Course Scheduling: Student Preferences and Fairness Constraints (Student Abstract)

Mitsuka Kiyohara

University of Toronto, Canada
mitsuka.kiyohara@mail.utoronto.ca

Abstract

Increasing student populations and diverse course offerings have led to perceived inequities in U.S. high school course scheduling. Traditional integer programming (IP) methods for the High School Scheduling Problem (HSSP) fail to address these fairness concerns. This research introduces the Fair High School Scheduling Problem (FHSSP), an extension of the HSSP that incorporates student preferences and fairness principles from market design. We develop an IP model to generate course schedules that are both feasible and equitable. Tested on real course request data from a California high school, our model successfully produces schedules that ensure fairness without compromising feasibility. These results demonstrate the potential of our approach to enhance fairness in high school scheduling and its applicability to various real-world scheduling challenges. Additionally, this study highlights the feasibility of integrating human preferences and emotions into mathematical models, promoting more inclusive and balanced allocation systems.

Introduction

High schools in the United States are increasingly challenged by scheduling complexities due to a growing student population and a broader array of course offerings. Over the past decade, the number of public and private high schools has risen by 912, accommodating an additional 1 million students (National Center for Education Statistics 2022). This growth, within a decentralized education system where each state sets its own standards, makes creating efficient and equitable schedules more difficult. Diverse course offerings and resource constraints further complicate the development of generalized scheduling solutions.

Currently, high school schedules are typically generated using constraint programming (CP) techniques, which prioritize efficiency but often overlook student preferences and emotions. This results in course allocations that students perceive as unfair. Some schools attempt to manually adjust schedules after using CP solvers to address fairness, but this method is inefficient and prone to errors. While IP approaches have been explored (Kristiansen, Sørensen, and Stidsen 2015; Ribić, Turčinhožić, and Muratović-Ribić 2015), they generally do not incorporate fairness constraints necessary for equitable student outcomes.

In contrast, market design research has extensively addressed fairness in various matching markets, such as daycare placement and university course allocation (Sun et al. 2022; Behrenk, Güçlükol Ergin, and Toy 2023; Son and Ngan 2021). However, these solutions have not been directly applied to high school scheduling. To address this gap, our research introduces the FHSSP, an extension of the HSSP. By adapting principles from market design to the unique context of US high schools, we develop an IP model to solve the FHSSP, overcoming the limitations of existing CP models and promoting fair scheduling practices.

Proposed Method

Problem Formulation. The HSSP involves assigning students to courses in a manner that satisfies various school and logistical constraints. Formulated as a Constraint Satisfaction Problem (CSP), the HSSP aims to create feasible weekly schedules by appropriately assigning students to lectures while considering the availability of instructors, rooms, and time slots.

The primary inputs to the HSSP include $|C|$ courses, $|I|$ instructors, $|S|$ students, $|T|$ time slots, and $|R|$ rooms. The decision variables consist of lectures $L \subseteq C \times I$, which pair courses with instructors; assignments $A \subseteq L \times S$ representing student enrollments in lectures; and units $U \subseteq T \times R \times L$, which schedule lectures in specific rooms and time slots.

To ensure a schedule is feasible, it must adhere to key constraints: Time slot conflicts prevent double-booking of instructors, students, or rooms; eligibility ensures that only qualified instructors teach courses, students enroll in courses they are eligible for, and rooms are appropriately equipped; numerical constraints maintain minimum and maximum limits on the number of units an instructor teaches per week, the number of courses a student can enroll in, and room capacities; and lecture frequency guarantees that each lecture is scheduled the required number of times per week. By integrating these elements, the HSSP seeks to produce efficient and equitable course schedules that accommodate the diverse needs of high schools.

While traditional approaches to the HSSP using IP and CP solvers focus on efficiency, they often omit fairness, leading to potential inequities in course allocations among students.

To address this limitation, we introduce the FHSSP, an extension of the HSSP that incorporates fairness constraints. The FHSSP defines two novel concepts:

- Degree of interest $d_{s,c}$: Quantifies a student's preference for a particular course on a scale from zero to one.
- Priority order $p_{s,c}$: Determines a student's priority for enrolling in a course based on their academic performance and degree of interest. For courses with prerequisites, priority is calculated using the student's average grades in those prerequisites; otherwise, it relies solely on the degree of interest.

To mitigate student envy, FHSSP introduces a fairness constraint ensuring that a student with higher priority or greater interest in a course does not feel envious towards another student assigned to that course. Formally, student s envies student s' if all the following conditions are met:

- $c \in C_{s'} \setminus C_s$: Course c is assigned to s' but not to s.
- $p_{s,c} > p_{s',c}$: s has a higher priority for c than s'.
- $d_{s,c} > d_{s',c}$: s has a higher degree of interest in c than s'.
- $|C_s| < |C_{s'}|$: s is enrolled in fewer courses than s'.

Methodology. We use binary encoding for all decision variables and constraints to develop an IP solver for the HSSP and FHSSP. The solver utilizes two distinct objective functions: one to maximize overall course assignments for the HSSP and another to minimize instances of student envy for the FHSSP.

Results and Contributions

We evaluated the performance of our solver by running experiments on a real-life data set provided by a private high school located in the San Francisco Bay Area, California, which contained course requests from 295 students for 121 courses. All experiments were conducted on a laptop with an Apple M2-Max processor and 32 GB of memory. Our IP model was implemented using Google OR-Tools' CP-SAT solver [1] and was written in Python.

To manage the large problem size, we divided the data into six roughly equal student subsets, using artificial data for instructors and rooms based on key assumptions missing from the real-life data set such as dedicated rooms per course, fixed teaching capacities for instructors, six daily periods, multiple weekly course meetings, and a maximum of six courses per student. Priority orders and student preferences were randomized in the data.

For each student subset, we first verified the algorithm's feasibility in solving the HSSP, then progressively added fairness constraints to evaluate the model's ability to generate optimal and equitable schedules to solve the FHSSP.

Our experiments showed that, across all problem settings, our model consistently produced a feasible solution for the HSSP and an optimal solution for the FHSSP. Additionally, there was no instance of student envy for all problem settings. A summary of the experimental result can be found in Table 1.

[1]https://developers.google.com/optimization/cp/cp_solver

Group	Students	Courses Per Student	Envy	Time(s)
1	48	5	0	25.2s
2	48	5	0	26.2s
3	49	5	0	26.9s
4	50	5	0	20.2s
5	50	5	0	18.4s
6	50	5	0	19.0s

Table 1: Results of experiments for each problem group.

Conclusion

By unifying methodologies from market design, optimization, and planning, we developed an IP model that can solve both the HSSP and FHSSP. This research demonstrates the potential of combining multiple disciplines to address issues of disparity and transparency in various scheduling problems. Importantly, it highlights the necessity of incorporating human preferences and emotions into mathematical models to achieve equitable outcomes in real-world applications. We hope that this research is one small step towards this goal.

Acknowledgments

This work was supported by the JST Global Science Campus Experts in Science Information program (NII, IPSJ, JCIOI). I would like to extend my deepest gratitude to Dr. Masakazu Ishihata for his invaluable guidance and mentorship, as well as the Menlo School for providing the course request data set and helpful insights into the scheduling problem.

References

Behrenk, A. B.; Güçlükol Ergin, S.; and Toy, A. O. 2023. Course Scheduling Problem and Real-Life Implementation. In *Towards Industry 5.0*, 749–758. Cham: Springer International Publishing. ISBN 9783031244568 9783031244575.

Kristiansen, S.; Sørensen, M.; and Stidsen, T. R. 2015. Integer programming for the generalized high school timetabling problem. *Journal of Scheduling*, 18(4): 377–392.

National Center for Education Statistics. 2022. Enrollment in elementary, secondary, and degree-granting postsecondary institutions, by level and control of institution: Selected years, 1869-70 through fall 2031.

Ribić, S.; Turčinhožić, R.; and Muratović-Ribić, A. 2015. Modelling constraints in school timetabling using integer linear programming. In *2015 XXV International Conference on Information, Communication and Automation Technologies (ICAT)*, 1–6.

Son, T. A.; and Ngan, N. T. 2021. An Integer Programming Model for Course Timetabling Problem at HUST. In *Intelligent Systems and Networks*, volume 243, 345–354. Singapore: Springer Singapore. ISBN 9789811620935 9789811620942.

Sun, Z.; Takenami, Y.; Moriwaki, D.; Tomita, Y.; and Yokoo, M. 2022. Daycare Matching in Japan: Transfers and Siblings. arXiv:2212.02751.