The Student-Project Allocation Problem

Frances Cooper and David Manlove, University of Glasgow, Scotland

We have a **matching problem** when we wish to assign one set of agents to another set of agents. Examples include assigning **children to schools**, **students to projects** and **kidney transplant patients to kidney donors**. Here we consider the **Student-Project Allocation problem** and describe how to cope with NP-hard variants.

The Student-Project Allocation Problem

The Student-Project Allocation problem with lecturer preferences over Students with Ties (SPA-ST)

- A set of students, projects and lecturers.
- Each project is offered by a unique lecturer.
- Students have preferences over projects, whilst lecturers have preferences over students.
- Projects and lecturers have upper quotas.
- Ties are allowed in student and lecturer preference lists.

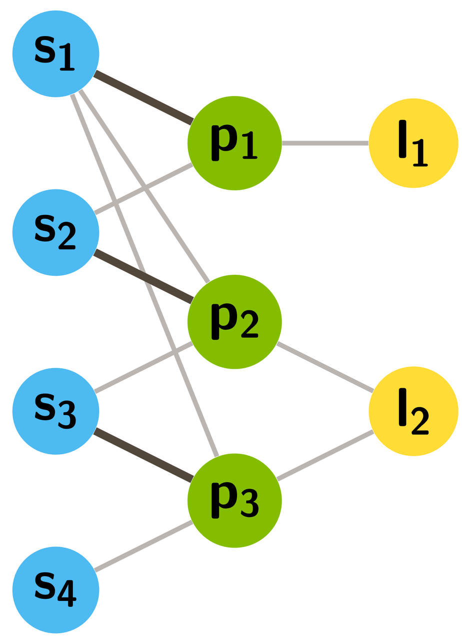

An instance can either be viewed as a graph (above) or using preference lists (matching is underlined):

s_1: ($\underline{p_1}$ p_2) p_3	p_1: UQ: 2	l_1: ($\underline{s_1}$ s_2)	UQ: 2
s_2: $\underline{p_2}$ p_1	p_2: UQ: 1	l_2: s_4 ($\underline{s_3}$ s_1) $\underline{s_2}$	UQ: 2
s_3: p_2 $\underline{p_3}$	p_3: UQ: 2		
s_4: p_3			

An extension to this problem allows lecturers to have **targets** which indicate the preferred number of allocations.

Optimisations

- A **stable matching** in SPA-ST is an assignment of students to projects such that upper quotas are respected and there is no student-project pair (s_i,p_j) where s_i and l_k, the lecturer offering p_j, have an incentive to deviate from the assignments (if any) and form a pairing.

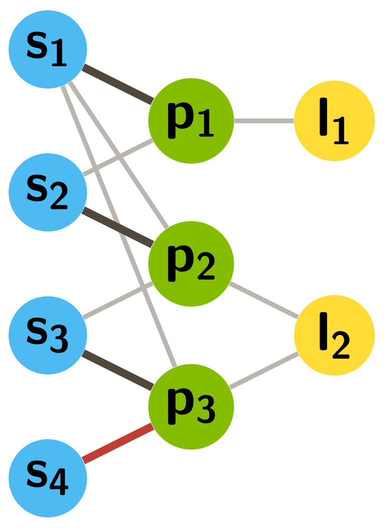

This matching **is not** stable. Student s_4 would prefer to be assigned to project p_3 and lecturer l_2 would also prefer this change.

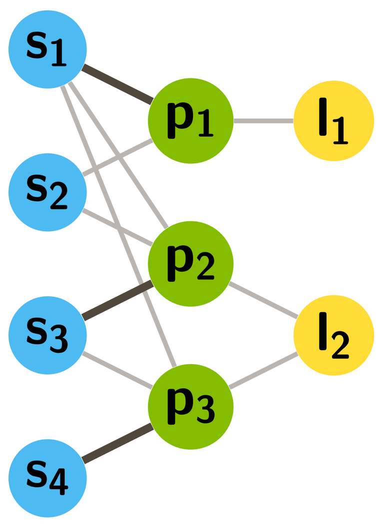

This matching **is** stable. There are no student-lecturer pairs who have an incentive to deviate from their assigned partners.

- **Maximum size** - maximises the number of students assigned.
- **Greedy** - maximises the number of students assigned to their first choice project and, subject to this, their second choice, and so on.
- **Generous** - minimises the number of students assigned to their last choice and, subject to this, their second to last choice, and so on.
- Extension: **load balancing** - variety of comparisons between the number of lecturer allocations and targets.

Important Results

- Every instance of SPA-ST admits a stable matching [1].
- A stable matching can be found in linear time [1].
- Stable matchings can have different sizes and finding a maximum sized stable matching in an instance of SPA-ST is NP-hard [2].

Finding Optimal Results with Integer Programming

Integer Programming (IP) is a computational technique that can deal with NP-hard optimisation problems. Finding a maximum stable matching in an instance of SPA-ST is NP-hard and so an **IP model** was developed.

- New integer inequalities and objective functions were created for stability constraints and load balancing optimisations.
- A desktop application and web application were developed to run allocations. Optimisations can be run in any order.

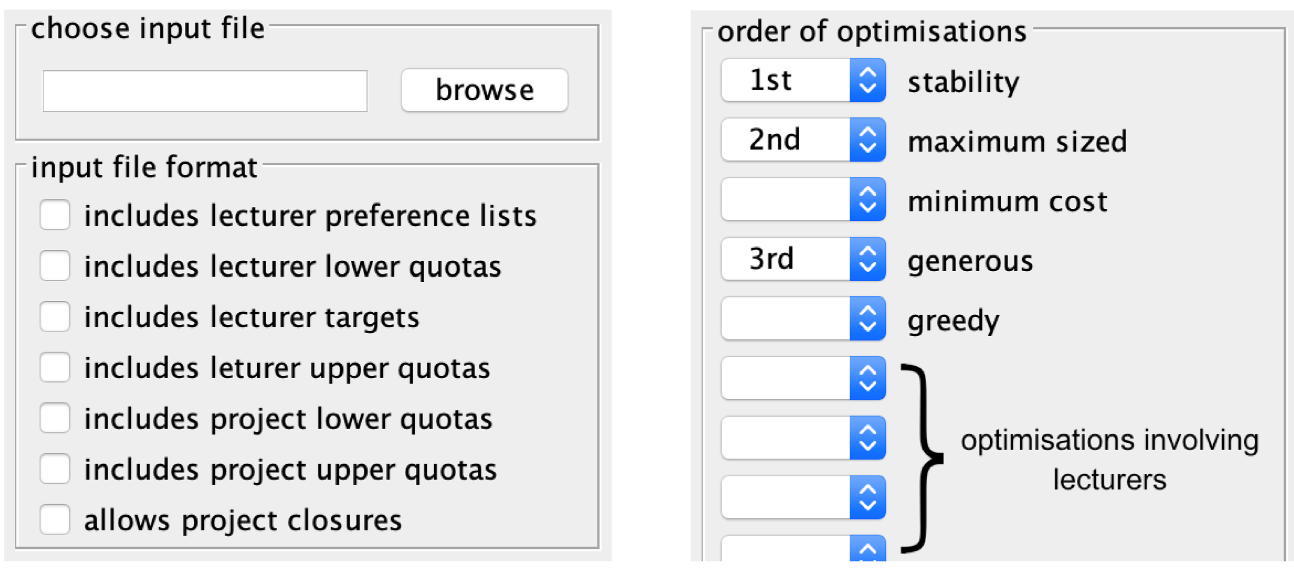

Real-World Results

The IP model has been used in **several universities and organisations around the world** including student-project allocations in Ireland, the UK, China and Singapore. Each scenario had varying requirements but in several instances the IP model replaced a manual allocation process which was both time-consuming and unlikely to result in an optimal outcome.

Finding Approximate Results with Linear Time Algorithms

Integer programming solvers use algorithms that run in exponential time in the worst case. Therefore, depending on priorities, it might be preferred that an **approximate** solution is found **quickly**.

Important Results

- For a simplified version of SPA-ST in which projects are supervised by a unique lecturer with identical capacity, it is possible to find a matching at least 2/3 the size of the maximum stable matching in linear time [3].

- We developed a linear time **approximation algorithm** that generalises the above result, and finds a stable matching at least 2/3 the size of the maximum stable matching for SPA-ST [4].
- This work was presented at SEA 2018: the 17th International Symposium on Experimental Algorithms in L'Aquila, Italy.

References

[1] D.J. Abraham, R.W. Irving, and D.F. Manlove. Two algorithms for the student-project allocation problem. Journal of Discrete Algorithms, 5:73-90, 2007.

[2] D.F. Manlove, R.W. Irving, K. Iwama, S. Miyazaki and Y. Morita. Hard variants of stable matchings. Theoretical Computer Science, 276(1-2):261-279, 2002.

[3] Z Király. Linear time local approximation algorithm for maximum stable marriage. Algorithms, 6:471-484, 2013.

[4] F. Cooper, D.F. Manlove. A 3/2-Approximation Algorithm for the Student-Project Allocation Problem. In Proceedings of SEA 2018, vol 103 of LIPIcs, pages 1-13, 2018.

Supervised by David Manlove and Kitty Meeks. Supported by EPSRC as DTA funding.