Reducing Contention in an Undergraduate Capstone Project Allocation System

Stephen Phillip Levitt(✉) and Kenneth John Nixon

University of the Witwatersrand, Johannesburg, South Africa
{stephen.levitt,ken.nixon}@wits.ac.za

Abstract. An online project bidding and allocation system used by the School of Electrical and Information Engineering, at the University of the Witwatersrand, is described and critically analysed. An important question that arose when designing the bidding system for projects was how to deal with the contention for highly popular projects. An approach, common in the literature, is to optimise the project allocation; however, the system implemented aims to incentivise students to consider less popular projects. Projects are allocated to groups using a quasi-random algorithm, as opposed to an optimal algorithm. Over the eleven year analysis period the percentage of first choices awarded has ranged from 57.5% to 80.8%, while the average number of unallocated groups per year is 8.2%. In addition to these outcomes project contention is shown to reduce during bidding window period in nine out of the eleven years of the study. This is due to novel aspects of this system including the transparency of the allocation algorithm and the fact that groups are able to see all competing bids for the projects they are bidding on and change their own bids in response. The allocation algorithm is, by necessity, sub-optimal in order to achieve the goals of transparency and fairness in that the stated winning probability, or a greater probability, for a bid always holds true when a project is allocated.

Keywords: capstone project · student project selection · allocation algorithm

1 Introduction

The Electrical/Information Engineering Laboratory is a final-year capstone course in the undergraduate degree programme of the School of Electrical and Information Engineering at the University of the Witwatersrand. The course is structured so that each academic member of staff, termed "supervisor" from here onwards, offers a small number of unique projects (between one and three). Students are required to work in pairs (as teamwork is one of the outcomes being assessed) and must undertake one particular project on offer.

Prior to the introduction of the system the allocation of projects to student groups was done in a manual fashion. This manual allocation process was

H. E. Van Rensburg et al. (Eds.): SACLA 2023, CCIS 1862, pp. 34–47, 2024.
https://doi.org/10.1007/978-3-031-48536-7_3

generally perceived as being unfair, especially in cases where there were highly contested projects, because there were no clear guidelines to students, and supervisors, as to how projects were being allocated. An additional issue was that this manual process was becoming unwieldy with increasing student numbers.

Given the importance of this course and the level of unhappiness that students were feeling at the time, the authors undertook to completely redesign and automate the project allocation process with the goals of still allowing student agency (in choosing projects) while promoting system transparency (in that all actions taken with regard to the allocation process are visible to all parties involved). A key consideration in the design was how to deal with the contention for highly popular projects. Rather than adopting the approach of optimising the project allocation, the solution adopted aims to incentivise students to consider less popular projects. This new system, and the way in which the above goals are achieved, is described in the following sections. Section 2 presents the relevant literature. Section 3 concentrates on the online bidding process, while Sect. 4 discusses how projects are allocated. In Sect. 5 analyses of various aspects of the system are provided. This is followed by the conclusion.

2 Existing Approaches

There are three key entities in student-project allocation problems which are *students*, *projects*, and *supervisors*. Approaches to allocating projects vary by allowing different kinds of constraints among the key entities. Some approaches solely allow students to rank projects in order of preference [10,15]. Others may additionally allow supervisors to specify their preferences over students, as in [1]. Supervisors may also specify their preferences over projects [12]. Finally, a fairly generalised approach is given wherein supervisors specify preferences with regard to which students they want, on the projects they want, and students specify their preferences with regard to the projects [7,13]. Many approaches also satisfy capacity constraints with regard to the number of projects that a supervisor may take on, and the number of students that may be allocated to a particular project.

Student-project allocation problems can be seen as specific instances of generalised assignment problems for which a body of mathematical and computational theory is applicable [4]. These problems are shown to be NP-hard in complexity, and more recently, certain formulations have been proven to be APX-hard [11,12] which means that approximate solutions within known bounds are attainable within polynomial-time.

A number of different techniques for solving these problems have been described. These include ad-hoc algorithms [15], integer programming [2,5], genetic algorithms [10,14], and two-sided matching algorithms including the classic Gale-Shapley stable matching algorithm [8] used by [9] and variants of this [1,7,12].

In arriving at a solution an objective function is defined and a search is conducted for an allocation which maximises this function. An optimal solution is often seen as one in which students are allocated their highest ranked projects as far as possible whilst minimising the number of unallocated students. In the case of two-sided matching problems, wherein, for example, both students and supervisors express preferences relating to one another, a stable matching is sought in which the number of unallocated parties is minimised. A stable matching is one in which no two parties who are not matched together would rather be matched to each other than their current assignees.

Existing literature on the student-project allocation problem has tended to focus narrowly on contributing to the theoretical understanding of the computational complexity of this type of assignment problem and it variants, and presenting algorithms for finding approximate or optimal solutions. The starting position, of much existing work, is seen to be a predetermined set of preferences and constraints, and the goal is to optimise the allocation given these initial conditions.

However, there are a few studies that specifically consider the broader social context in which the student-allocation problem is situated. In the work done by Greef et al. [9] an interesting approach is taken whereby students tender for projects, and express their project preferences by ranking their own tenders. They further participate in ranking the tenders of their peers. These rankings are combined and form the preference scores that are used by the stable matching algorithm. This promotes student engagement with the system and students are exposed to anonymised, competing tenders.

Briffa et al. [3] discuss an allocation system which makes use of a web application to provide real-time feedback to students concerning the popularity of topic choices. In this way, their system allows students to opt for less popular topics if they are at risk of not receiving their top preferences. They use a simulated annealing algorithm to seek a global optimum topic allocation.

The approach which is presented below is closest to those of [10,15], in that students express preferences over projects and there are no other constraints involved. However, a deliberate decision is taken to include the broader social context within which the project allocation is occurring in a somewhat similar manner to [3], and not to solely focus on an algorithmic solution to a set of expressed preferences.

3 Project Bidding

In the Electrical/Information Engineering Laboratory, students initially pair up by mutually selecting each other using the online system. Students are then given one week to obtain bidding rights followed by another week in which the actual bidding for projects takes place.

3.1 Obtaining Bidding Rights

Obtaining bidding rights is a important step which requires each group to meet with potential supervisors in order to discuss the projects that they are interested in. This ensures that groups have, at least, an initial engagement with supervisors regarding the requirements and expectations of projects that they are interested in. The supervisor, in turn, is obligated to grant the group bidding rights on these projects after this discussion. This will allow the group to place bids on these projects once the bidding window opens. Groups may also continue to obtain bidding rights on projects after the bidding window has opened.

Supervisors may not selectively grant bidding rights to groups that they would prefer. Unlike some other systems that have been documented in the literature, this system does not allow for supervisors to express preferences over students. This decision has been taken by the School because it supports the notions of fairness, in that supervisors may not prefer (cherry-pick) certain groups over others, and quality, in that any group in the final-year cohort, should be capable of doing on any project on offer that they may be interested in. One of the drawbacks of this decision is that it could potentially lead to supervisors being marginalised if students consistently avoid bidding on any of their projects. In practice, this has been found to be a minor issue with students only tending to avoid supervisors that they have never met before in a lecturing capacity. Within our School, this has applied to very few supervisors over the years and, in such cases, deliberate interventions can be put in place to introduce "unknown" supervisors to students.

All students and supervisors are able to see the number of groups that have obtained bidding rights on any particular project. This gives an idea as to the initial popularity of the project and is intended to encourage students to obtain bidding rights on additional projects if they see that all of the projects that they are interested in are popular.

3.2 Placing Bids

Each project has six bidding slots available. There are three first-choice slots, two second-choice slots, and one third-choice slot. This allows a maximum of six bids from six different groups for each project. Of course, projects may receive fewer bids or no bids at all.

Once the bidding window opens, each group must submit bids for three different projects *simultaneously*: a first-choice bid, a second-choice bid, and a third-choice bid. A group can only bid on projects for which they have obtained bidding rights. Furthermore, each project must have an open bidding slot for the particular choice of bid that they wish to make. The single third-choice slot is the limiting factor here, and groups often find that they need to secure bidding rights on additional projects because the third-choice slots for the projects on which they do have rights are already occupied.

The system guarantees that the probability of a first-choice bid winning a particular project is always greater than the probability of a second-choice bid, and that the probability of a second-choice bid winning a particular project is always greater than the probability of a third-choice bid. This is discussed in more detail in Sect. 4.

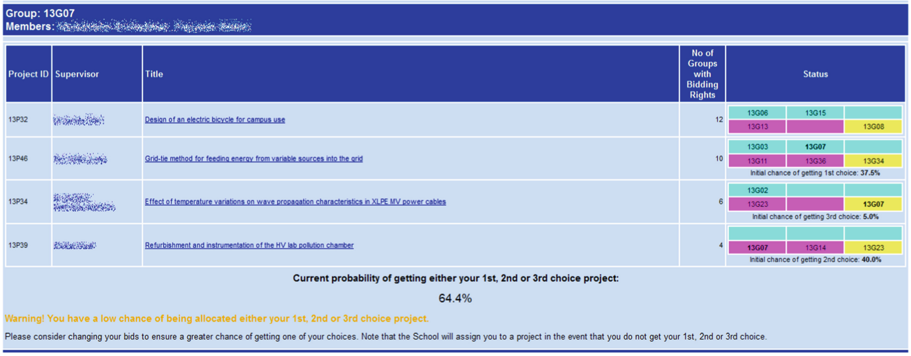

Fig. 1. Student view of bidding status (anonymised). This particular group has bidding rights on four projects and have placed bids for the bottom three projects in the list.

Groups may repeatedly change their bids throughout the bidding period. The system is entirely transparent in that students can see exactly which other groups are bidding for the projects that they are interested in. They can also see their probability of being awarded a project that they are making a bid on, assuming that it is processed independently of all the other projects. Finally, their overall probability of being allocated any of the projects that they have bid on by the allocation algorithm is calculated using the per-project probabilities. If the overall probability is low for the group then they are issued with a warning that they may not, in fact, be allocated any project by the system. This is illustrated in Fig. 1.

After a week of bidding, the bidding window is closed and the projects are then allocated using an automated allocation algorithm (coded in Octave [6]) which is described in Sect. 4. The source code of the algorithm is made available for students to download, along with instructions on how to run the algorithm against the current bidding state. In addition to this, *throughout the bidding period* students are able to run the algorithm directly from the course website and view the output (which varies due to its inherent randomness). The algorithm output is detailed, showing not only the final outcome of the allocation, but also clearly documenting each step that is taken in reaching those results.

If any group fails to make their bids within the bidding period they are allocated a project manually, from the remaining, unallocated projects.

4 Project Allocation

In order to allocate a given project, one of the bids for the project is randomly selected by the allocation algorithm and the group that placed that bid wins the project. Not all of the bids stand an equal chance of being selected and this is described in more detail below.

Table 1. Processing Order for Project Categories

Sequence	Project category (based on bids present)	Bidding pool selection probability
1	First-choice bids only	100 : - : -
2	First- and second-choice bids only	79 : 21 : -
3	First-, second- and third-choice bids	75 : 20 : 5
4	Second-choice bids only	- : 100 : -
5	Second- and third-choice bids only	- : 80 : 20
6	Third-choice bids only	- : - : 100

4.1 Weighted Bidding Pool Probabilities

All the first-choice bids for a particular project form the first-choice pool. All the second-choice bids form the second-choice pool and the single third-choice bid forms the third-choice pool. If all three bidding pools have bids present, then the probabilities are weighted such that the first-choice pool has a 75% probability of being selected, the second-choice pool, a 20% probability, and the third-choice pool a 5% probability. In other words, the likelihood that a bid in a particular pool will be selected can be expressed using the following ratios:

$$\begin{array}{ccccc} \text{First-Choices} & : & \text{Second-Choices} & : & \text{Third-Choice} \\ 15 & : & 4 & : & 1 \end{array}$$

A first-choice bid is fifteen times more likely to be selected by the algorithm than a competing third-choice bid.

In cases where a bidding pool is not occupied, the empty pool is omitted from selection. The ratios given above are preserved for the remaining bidding pools by redistributing the empty pool's probability of selection proportionately. In cases where there are two unoccupied pools, the remaining pool will have a 100% probability of being selected. The probability of each bidding pool being selected in these different scenarios is illustrated in Table 1.

4.2 Project Processing Order

A critical aspect of the allocation algorithm is the order in which the projects are allocated to groups. In order to honour the ranking of choices it is essential to process projects, for example, which have first-choice bids before projects that only have third-choice bids. The projects are therefore grouped into categories based on the type of bids present. The category are processed in the order shown in Table 1.

Projects are allocated in sequence starting with those belonging to the first category listed in Table 1. If there is more than one project in any category, then one of the projects in the category is selected randomly and the project allocated. Once a project is allocated it is removed from the project list. All of the remaining bids of the "winning group" are deleted from the project list, the probabilities for each project are recalculated, and the list is re-categorised. Project allocation continues in this fashion until there are no more projects remaining in the list.

If a group makes all of its bids for contested projects there is a chance that they may not be allocated any of their choices. In other words, competing bids may be selected for each of the projects on which they have bid. If this occurs the group concerned is manually allocated a project (in consultation with the course co-ordinator) from those that remain after the allocation has been finalised. In spite of being warned about this consequence, some groups decide to take the risk for the chance of doing a project that they are really interested in and they enter the final allocation with a low overall chance of being allocated a project by the system.

There may be multiple first and second-choice bids for a project. The probability of any *one of the bids* winning the pool is equal to the bidding pool probability divided by the number of bids in the pool. So, for example, assume that a project has two first-choice bids, two second-choice bids and one third-choice bid present. The probability of either of the first choices winning is 37.5%, the probability of either of the second choices winning is 10%, and the probability of the third choice winning is 5%.

4.3 Allocating a Project

Selecting a winning bid from the bidding pools for any given project is readily implemented. An array representing the bids for the project is created. Seventy-five elements of the array are used to represent the first choice bidding pool, twenty the second-choice bidding pool, and five the third-choice.

In the scenario where there is one first-choice, one second-choice and one-third choice bid for the project, all seventy-five elements representing the first choice will contain the identifier of the group which made the project their first-choice, twenty will contain the identifier for the group making the project their second-choice, and five will contain the identifier of group making it their third-choice. In order to allocate the project, one of the array elements is randomly selected.

In cases where there are multiple groups within a single bidding pool (first- and/or second-choice bids) the array elements are divided evenly among the bidding groups. For example, if there happen to be three first-choice bids then each of the bidding groups will be allocated twenty-five elements of the array. Likewise, if there are two second-choice bids then each the groups will be allocated ten elements.

The array is sized to only represent bidding pools which have bids present. For example, if there are no third-choice bids, the array will consist of only ninety-five elements with seventy-five being reserved for first-choice bids and twenty being reserved for second-choice bids. As before, after the array is populated with the bidding groups an element is randomly selected to determine which group is given the project.

4.4 Finalising the Allocation

Owing to the deliberate random nature of the allocation algorithm, it may happen that a single run of the algorithm produces an unfavourable result. An unfavourable result would be where a relatively large number of groups are not allocated a project or the number of first choice allocations is relatively low. To deal with this possibility, the allocation algorithm is run three times in the presence of the student class representatives, the course co-ordinator, and a senior administrative officer. The class representatives then select one of the three runs as the final allocation. In doing this they need to decide which of the runs represents the best trade-off between honouring student preferences and minimising the number of unallocated groups. At no time are they privy to which groups have received which projects as this could bias their decision. Once the final allocation has been decided upon it is published to the class, along with the entire output of the selected run.

5 System Performance and Analysis

In analysing the system that has been presented, the statistics are given first which provide insight into various aspects of the allocation algorithm's performance over an eleven year period. This is followed by a detailed discussion of the two different ways in which this system reduces the contention for projects. Lastly, the optimality of the allocation algorithm is considered.

5.1 System Statistics

Table 2 presents the statistics for the system that have been gathered since its inception. The system's components were not changed over this time period and therefore these yearly statistics are comparable with each other. The first column designates the year. The second column indicates the number of groups that were involved in the bidding process. This is followed by the number of projects that were available for students to bid on, and the number of excess projects (total projects on offer minus groups bidding) expressed as a percentage. The next two columns show the outcome of the allocation algorithm, in particular the number of winning first-choice bids (column 5) and the number of unallocated groups (column 6) expressed as percentages.

It is clear from the table that this system is able to grant the majority of groups their first choices and this is a direct result of using weighted bidding pool probabilities. Looking at the results in more detail, it is evident that Year 3 is an outlier in terms of the number of groups bidding. The forty groups bidding in that year is far higher than any of the other years. Additionally, the number of excess projects in that year (7%) was the smallest. So, essentially, year 3 represented a stress test for the system in that one would expect a high amount of contention for projects. The allocation results show that the system was still able to grant a majority of first choices at the expense of a relatively high number

Table 2. Statistics for the System

Year	Groups bidding	Available projects	Excess projects (%)	First choices granted (%)	Groups unallocated (%)
1	23	39	41.0	60.9	4.4
2	28	37	24.3	62.1	10.7
3	40	43	7.0	57.5	20.0
4	26	29	10.3	80.8	11.5
5	25	29	13.8	72.0	8.0
6	31	46	32.6	71.0	9.7
7	25	40	37.5	73.1	4.0
8	27	40	32.5	63.0	0.0
9	24	39	38.5	75.0	0.0
10	27	31	12.9	63.0	7.4
11	27	34	20.6	59.3	14.8

of unallocated groups (20%). In other years, with fewer groups, the percentage of unallocated groups is far less.

This suggests that the system is unable to deal effectively with larger group numbers if there are insufficient excess projects. From the table, it appears that the number of excess projects needs to be at least greater than 10% in order to keep the number of unallocated groups small. Where there is a larger number of extra projects (years 6 through 9) it is possible to grant a high number of first choices with few unallocated groups, in some cases zero unallocated groups. Of course, the effort in this regard needs to be balanced because it entails more work on the part of the supervisors and coordinators who have to propose and review additional projects.

The allocation system appears to have been well received by students from the informal feedback that has been elicited over the decade that the system has been employed. There have been less than a handful of complaints and these have originated from groups that have tried to game the system. Instances of attempting to game the system have occurred during the bidding process when a group has observed the outcome of a number of algorithm runs and seen that in these runs they do not receive their first-choice project, so in order to try and rectify this they change their first-choice bid to a second-choice bid and (mistakenly) hope this will improve their chances.

It is arguable, that the success of the system lies not in its ability to award the majority of groups their first-choice projects, which it does, but rather in that it is transparent and accustoms students to the idea that they may not receive their first-choice project.

5.2 Project Contention

The level of projection contention is a crucial factor for any technique which attempts to solve the student-project allocation problem. The more highly contested projects are, the more disappointing the results of the allocation will be for the majority of students irrespective of which allocation technique is adopted and whether or not the allocation is optimal, according to the objective function, or sub-optimal (approximate). This simply attests to the fact that many groups/students in highly-contested scenarios will not be able to receive any of their top-ranked projects.

Another issue that is of concern, is that groups may in fact not receive any projects at all after the allocation process has run, having lost to competing groups on all of the projects that they had ranked. Different approaches to handling this situation are documented in the literature. In [15], for example, unassigned groups re-rank the remaining, unassigned projects. In [10] students may be a assigned to a "non-choice" which is a project that they did not actually rank. A solution to this is to increase the number of projects that students must initially rank in order to give the allocation algorithm a greater possibility of finding an allocation in which all students receive one of their preferences.

This system deals with project contention in two different ways. The first is structurally, and the second is through the social dynamics at play.

Reducing Contention Through Bidding Slot Structure. The bidding slots have been structured in the manner described above (three first-choice slots, two second-choice slots and one-third choice slot per project) limiting the contention for any one project to a maximum of six different groups.

The motivation for such a structure is twofold. On the one hand, popular projects exist and it is desirable to have a system which takes cognisance of this by allowing multiple bids, in particular first-choice bids. If the system was not structured in this manner, and there was, for example, only a single first choice slot for projects, then whichever group was the quickest in making their bids at the opening of the bidding session would most likely receive the project. It is not the intention that projects are awarded on a "first-come, first-served basis". Given the total number of groups in the School, having three first choices gives groups a reasonable chance of making a first-choice bid on the projects that they are interested in.

On the other hand, in order to have viable system in which most groups are allocated projects, it is necessary to force groups to spread their bids. The "bidding spread" can be defined as the number of projects which have at least one bid. The spread when considering only first choice bids, in the worst-case, maximally-contested, scenario (smallest spread) will be equal to the number of groups divided by three. However, for the third-choice bids the spread will always equal number the groups, which is the widest spread possible. The two second-choice slots enforce an intermediate spread. Having such a structure leads to the following bounds on the average number of bids per project:

$$\text{Average bids per project} = \frac{\text{total bids}}{\text{number of projects bid upon}} = \frac{\text{groups bidding} \times 3}{\text{number of projects bid upon}}$$

Table 3. Bidding System Social Dynamics

Year	Bidding state changes (all bidding)	Change in project contention (all bidding)
1	11	0.000
2	28	−0.071
3	142	−0.045
4	192	−0.037
5	49	−0.060
6	56	0.000
7	15	−0.056
8	14	−0.040
9	46	−0.020
10	19	−0.058
11	78	−0.073

The upper bound of the average is thus equal to three because the number of projects that are bid upon must, at a minimum, be equal to the number of groups bidding as each group is required to make a third-choice bid and there is only one such slot per project. So although the maximum number of bids for any one project is six bids, the average bids per project will, at most, be three which is low.

Reducing Contention Through Social Interaction. Given the transparent nature of the system, in that groups can see competing bids, it is interesting to examine the social interaction which is at play. Two metrics that provide some insight into the bidding behaviour are state and contention changes as detailed in Table 3. Groups often desire and compete for the same projects; however, they run the real risk of not being awarded any projects if they only place bids on popular projects. It is therefore reasonable to assume that, on the whole, the class will attempt to self-optimise by moving away from popular projects to some extent.

In order to examine this hypothesis, it is necessary to analyse the bidding activity that takes place. The time period of interest for this analysis begins at the point at which all groups are on board and have made bids and ends at the close of bidding. This particular time span is chosen for two reasons. Firstly, as groups join the bidding over time they introduce "turbulence" into the system by supplying three completely new bids and altering the bidding landscape. Therefore, by only considering the period from which all groups have bids present this effect is negated and the students are able to optimise against a relatively stable system state. The second reason is that students at the tail-end of the bidding period are forced to "show their hand" before the bidding closes so there should be less attempts to play the system and a concerted effort to try and improve their chances of receiving a project.

The number of bidding state changes which occurred during this time span, over the years since the system's inception can be seen in Table 3. The bidding state changes whenever a group updates its bids. The simultaneous submission of all three of the group's bids (a first, second and third choice) is counted as a single state change. It is evident that there is a fair amount of activity that takes place as groups manoeuvre in response to each other's bids but there does not appear to be a correlation with the total number of groups that are bidding.

Lastly, in order to support this analysis a *project contention measure* is defined as being the total number of bids that have been placed divided by the total number of occupied bidding pools (bidding pools which have at least one slot occupied). This measure attempts to capture how widely students are spreading their bids. If the bids cluster around popular projects then fewer bidding pools will be occupied as popular projects will have many bids per pool. If the bids are widely spread among different projects than more pools will be occupied with the pools containing less bids on average.

This measure is calculated at the start and end of the time period of interest and difference is given in the last column of Table 2. In nine out of the eleven years that the system has been operation, the project contention has dropped by the close of bidding. For remaining two years (years 1 and 6), it stayed the same. These statistics offer strong evidence that the class is self-optimising and that the social dimension to the system has a small but tangible effect in causing the bids to spread.

It is reasonable to expect that as the project contention is decreased there will be an increase in the number of first choices granted, in other words, a negative correlation. From the limited number of data points there is a weak positive correlation ($r = 0.23$) which seems counter-intuitive. However, the number of excess projects should also have an effect on the number of first choices granted. Ultimately many more data points are needed to have confidence in the correlation between the variables.

5.3 Optimal Versus Non-optimal Allocation

Given the relatively small search space it is possible to use a brute-force search, or one of the other optimisation techniques mentioned in Sect. 2, to find an optimal, or near-optimal, allocation of projects, for a given bidding state, by maximising the percentage of first-choice allocations (or any other cost function related to the system outputs).

However, this presents a conundrum in that certain groups will consistently, throughout the bidding period, be granted second or third choice projects in order for an optimal allocation to be achieved. In a sense, these groups have to be "sacrificed" for the greater good of the entire cohort. Given the transparent nature of this system with the ability to run the allocation algorithm at any point in time on the current bidding state, the inability for some groups to ever achieve their first choices would become clear and this would effect the perceived fairness of the system. Therefore, a sub-optimal algorithm is used and because of its random nature the allocation of projects is not predetermined based on the bidding state.

6 Conclusion

The student-project allocation approach that is presented here allows students to specify preferences for projects by making first-, second- and third-choice bids. No constraints of any other kind are incorporated. The deliberate random nature of the allocation algorithm produces approximate solutions for a given bidding state rather than a single optimal or near-optimal solution.

This system has a strong social dimension and is highly transparent. During the bidding process, groups are able to see competing bids on the projects that they are interested in and change their own bids in response. They are presented with their probability of being awarded a project by the system. Additionally, they are able to run the allocation algorithm from the course website and view the results, which will vary on each run due to algorithm's randomness. This results in a system that is fair, or egalitarian, in that no groups are precluded from being awarded their first choice project, which would potentially be the case for some groups if an optimal solution was generated.

Particular care has been taken to design the allocation system to reduce project contention but still accommodate popular projects. The system is structured in such a way that there are six bidding slots per project. However, the single third-choice bidding slot per project enforces groups to spread their bids widely. The social dynamics of the system are shown to further reduce project contention as groups are generally risk averse and would rather move some of their bids to less popular projects than risk not being allocated a project at all. A reduction in project contention, in turn, makes it is easier for the allocation algorithm to perform effectively. Ultimately, such a system affords students a strong sense of agency in that not only are students able to decide on which projects to bid for, they are also able to decide for themselves on the level of risk that they wish to undertake with respect to being awarded a project by the system.

References

1. Abraham, D.J., Irving, R.W., Manlove, D.F.: Two algorithms for the student-project allocation problem. J. Discrete Algorithms **5**(1), 73–90 (2007)
2. Anwar, A.A., Bahaj, A.S.: Student project allocation using integer programming. IEEE Trans. Educ. **46**(3), 359–367 (2003)
3. Briffa, J.A., Lygo-Baker, S.: Enhancing student project selection and allocation in higher education programmes. In: 2018 28th EAEEIE Annual Conference (EAEEIE), pp. 1–6. IEEE (2018)
4. Burkard, R., Dell'Amico, M., Martello, S.: Assignment Problems, Revised Reprint. Society for Industrial and Applied Mathematics (2009)

5. Chiarandini, M., Fagerberg, R., Gualandi, S.: Handling preferences in student-project allocation. Ann. Oper. Res. **275**(1), 39–78 (2019)
6. Eaton, J.W.: GNU Octave. https://octave.org/. Accessed 11 Apr 2023
7. El-Atta, A.H.A., Moussa, M.I.: Student project allocation with preference lists over (student, project) pairs. In: 2009 Second International Conference on Computer and Electrical Engineering, vol. 1, pp. 375–379. IEEE (2009)
8. Gale, D., Shapley, L.S.: College admissions and the stability of marriage. Am. Math. Mon. **69**(1), 9–15 (1962)
9. Greeff, J.J., Heymann, R., Nel, A., Carroll, J.: Aligning student and educator capstone project preferences algorithmically. In: 2018 IEEE Global Engineering Education Conference (EDUCON), pp. 521–529. IEEE (2018)
10. Harper, P.R., de Senna, V., Vieira, I.T., Shahani, A.K.: A genetic algorithm for the project assignment problem. Comput. Oper. Res. **32**, 1255–1265 (2005)
11. Iwama, K., Miyazaki, S., Yanagisawa, H.: Improved approximation bounds for the student-project allocation problem with preferences over projects. J. Discrete Algorithms **13**, 59–66 (2012)
12. Manlove, D.F., O'Malley, G.: Student-project allocation with preferences over projects. J. Discrete Algorithms **6**(4), 553–560 (2008)
13. Moussa, M.I., El-Atta, A.H.A.: A visual implementation of student project allocation. Int. J. Comput. Theory Eng. **3**(2), 178–184 (2011)
14. Salami, H.O., Mamman, E.Y.: A genetic algorithm for allocating project supervisors to students. Int. J. Intell. Syst. Appl. **8**(10), 51 (2016)
15. Teo, C.Y., Ho, D.J.: A systematic approach to the implementation of final year projects in an electrical engineering undergraduate course. IEEE Trans. Educ. **41**(1), 25–30 (1998)