Constraint Definition in Student Scheduling with Adversarial Behaviour in Mind

A Case Study of a Multi-Objective Optimization Problem in the University of Porto

Miguel Sozinho Ramalho
E-mail: m.ramalho@fe.up.pt

Abstract In this paper, a multi-objective optimization solution is analysed for the student scheduling problem at the Faculty of Engineering of the University of Porto, in Portugal. A comparison was drawn with the previous algorithmic approach. Historical knowledge from various fields was gathered - from behavioural game theory to constraint definition - and then applied to that analysis. The outcome provided results that stem from non-partisan critical thinking, emphasizing on the fact that automated timetabling systems are open to game-play, if design flaws or discrepant stakeholders' awareness exist. These were identified in the case-study. The necessity for proper validation was highlighted and particular example improvements were suggested to the currently deployed system. Generalization of points learned was carried out to a sensible extent, so that the knowledge cloistered in here can encourage other researchers to consider the behavioural repercussions of their algorithms.

Keywords Student Scheduling · Educational Timetabling · Behavioural Game Theory · Multi-objective Optimization · Constraint Definition · University of Porto

1 Introduction

In Portugal, the Faculty of Engineering of the University of Porto (FEUP) is faced, semesterly and for each of its study programs, with the task of allocating students to individual classes according to previously constructed timetables. Hitherto, this process has been conducted using a priority queue of students ordered from highest to lowest cumulative Grade Point Average (GPA) - an average of a student's grades on a scale of 0 to 20, where 20 is the highest mark - each student specifies their schedule preferences, and then the system sequentially attempts to place a student into one of those preferences, with the limitation that each class has a maximum number of students. For the past two semesters (2017/2018), a new process has been deployed in a test environment, consisting of a single program at FEUP, the Master in Informatics and Computing Engineering (MIEIC). The new process targets the drawbacks of the aforementioned method, using a multi-objective optimization approach [17].

This innovative method has motivated the production of the present work, since the author was one of the students that had to experience first-hand the transition phase, and its adaptation challenges, as well as the group acceptance phase to the method, with the incurring implications into the behaviour and satisfaction at both the individual and general levels. The possibility of synthesizing existent research is also considered as a motive for this paper, as the generalization, extrapolation and innovation can only move forward once there is a ubiquitous, succinct and comprehensive knowledge of the state of the art.

In this article, a two-way analysis is conducted focusing on the comparison between the two models, with an emphasis on the newest approach, as it is very recent and lacks such an analysis. Considerations are made in regard to the constraints and weights selected for the multi-objective function, considering their improvement of previous statistics (Section 4) and also on the adversarial point of view that has not previously been considered in this and in most student scheduling problems. Thus, providing a more in-depth review of the problem. The author considers that both views must be considered, if a new system is to be generalized to the remaining courses at FEUP and eventually the entire University of Porto (UP), or even other institutions.

The manuscript first presents a contextualization of relevant concepts and an analysis of similar problems in Section 2. Then, a description of the implementation of the novel system is provided in Section 3. Afterwards, in Section 4, an aggregating analysis of the novel approach is performed, this section also contains concrete tuning operations to the multi-criterion search and real world requirements that are not yet met by the application of the process in the information system at FEUP [13]. Finally, a deliberation of the presented work is drawn and relevant future research topics are suggested, in Section 5.

2 Literature Review

2.1 Student Scheduling Problems

The problem in scope fits into the Student Scheduling Problem (SSP) category. SSPs are characterized as a timetabling process that sections students into particular class schedules according to a previously defined schedule structure - student sectioning is also another name for this set of problems [11]. Timetabling problems constraints fall on two main categories: hard and soft constraints [3].
Typical SSPs are consequentially constrained by factors such as existing blocks, limited number of vacancies or parallel classes [11]. These are labelled as *hard*

constraints, since violating one of them results in a non-existent solution for the problem, the problem becomes infeasible for the given characteristics [8]. Simultaneously, this set of problems can also be subject to *soft constraints*. Soft constraints can be violated. However, doing so results in a decrease in the quality of the solution found [8] - this concept can then be extended into a utility or cost function that describes a feasible solution and can therefore be used with a search algorithm to find a good or optimal solution. Consider the attempt of satisfying, as much as possible, the individual timetable preferences of all the students.

As with many timetabling problems, solutions are usually found *ad hoc* as they can take many different factors into account that vary amongst universities and learning institutions and, sometimes, amongst particular courses in the same institution [17]. Reported examples exist that describe real world problems where both constraint types are used, not just for higher education analysis [5] [8] [14] [16], but also for secondary schools [4] [19]. Despite the outreach of this and similar problems, none of the aforementioned solutions has made explicit efforts to accommodate for the competitive nature of humans, this paper is then an eye-opening attempt at such consideration. So, for the context of this paper, the author will consider only the soft constraints of the case study where adversarial behaviour is relevant.

2.2 Relevance of preferences

In essence, most soft constraints reflect personal or group preferences, desires even, and the better a solution can satisfy these requirements, the better it can serve the community it affects - student community, in this case.

Although some implementations do not account for student preferences, and soft constraints all together [5], the fact is that soft constraints matter - research has even been made into the impact of poorly designed timetables on the students' life [2] [9]. As is natural, students manifest their ideal schedule into their preferences, thus these can: reflect their weekly availability, impact on their transportation costs, and on those of their families, limit their transportation means and the use of ride systems, to name a few. These are relevant factors for maintaining a *status-quo* and even socio-economic level. Differences in said factors can have a high impact on the student's performance [2] and should, the author believes, be taken into consideration when designing schedules. From a generalizing point of view, the subjects that are allocated to timetables that define their daily schedules depend highly on the ability of the timetable to reflect their needs, an example of such is the nurse scheduling problem [1].

2.3 Behavioural Game Theory

Student Scheduling Systems and other timetabling systems that take the stakeholders' desires into account must extend their consideration out of the ana-

lytical mindset when designing solutions and accepting the opinions of those affected by the generated schedules. As will be presented in section 4, a wrong or incomplete definition that outputs apparently expected results may hide fallacies with consequences in the real world implementation.
When an SSP considers the opinions of the participants - usually as soft constraints - an environment is created where all the participants interact indirectly through their preferences, this is even more noticeable in cases where the number of vacancies for a given time slot or complete schedule are limited. Participants become adversaries. Adversaries, in turn, display behaviours that will benefit them the most and, in many, harm others. The study of the interactions that flourish among such environments is encompassed in *game theory* [7]. The fact that adversaries tend to behave in a non-cooperative manner also helps identifying the expected damaging behaviours predicted above. Furthermore, there might also be scheduling problems were cooperation and coordination amongst the participants may result in a better outcome. However, as proved by [7], any cooperative game can be reduced to a non-cooperative bargaining game, which in turn reopens the paradigm of adversarial behaviour.
Another aspect to consider in game theory is that of one individual, or group of individuals smaller than the total number of participants, has privileged information on the mechanisms used in the assignment of schedules for a given SSP, such as knowledge of the search technique employed, the heuristics used and the priority given to soft constraints. This unbalanced environment is often referred to as information asymmetry and can also be described as incomplete information on one side [7]. Most research on it focuses on economic and geopolitical applications [12], nonetheless, timetabling oriented applications can also be found [10].
Lastly, a clear bridge needs to be created between the concept of adversarial settings and adversarial behaviour, especially because most of the behaviours portrayed in these types of environments fall in accordance with social decision making [15]. This idea includes concepts such as social exchange - protocols used in the evolution and stability of social perception in groups - ethic background, individual affective characteristics. Thus, typically less belligerent conditions exist in real world adversarial environments - meaning that given the chance to benefit in detriment of others, the choice to do so is not always taken.

3 Description of the Case Study

3.0 Histogram Interpretation

Figures 1, 2, 3 and 4 are histograms describing the allocation of students. In the x-axis we find the GPA of students on a discrete scale from 10 to 20. The y-axis describes the number of students for the given GPA on the x-axis. Furthermore, each stacked column describes, according to the label below the chart, how many students for a given GPA had their n^{th} option met. The labels follow a scale from blue to red, where blue is a perfect placement for the students (their preferred schedule) and red is the worst case (no schedule has been assigned), also purple represents students that have been assigned a different schedule, in any way, from their manifestation of preferences.

3.1 Previous Method

The novel approach to the SSP at FEUP has one main goal: to maximize the allocation of students to classes in a single distribution phase. The reason for this comes from the fact that up to now, students had to go through three sequential phases.

In the first phase, the algorithm described in Section 1 was employed to assign students to classes. Figure 1 depicts the allocation resulting from this first step, for the second semester of the school year of 16/17. In it, we see that, for that specific semester, the number of students that were left without a schedule was 104 out of 425, approximately 24.5%.

Secondly, the students left with no schedule would use the information system of the faculty to enrol in classes, one by one, until their schedule would be complete. The main point here is that, these choices would be on a first come first serve model. This was, evidently, a stressful and stringent moment for students. Furthermore, it showed complete disregard for the main visible goal in the previous phase: to benefit those who had better grades. Additionally, students who already had their schedules defined in the first phase could also partake in this phase in an attempt to improve their schedules. To demonstrate the exasperation felt by students, one of them actually developed and shared a computer script to automate the selection and submission of the forms in the information system interface, so as to beat the other students [6].

Lastly, the students that were left without a schedule after the second phase would need to go to the secretariat of the course and ask for available schedules, which sometimes required opening a few more vacancies in some classes, to avoid conflicts among parallel classes.

3.2 New Method

So as to avoid the chaotic nature of the process, the novel method was proposed. This section describes it in accordance to its definition in [17], highlighting relevant points for posterior discussion. For future reference, a student preference consists of a complete timetable description, and each student can have up to 10 of these preferences, ordered from most important (1^{st}) to least important. The new decision method is based on multi-objective optimization. The decision variables considered for the problem are binary variables described as follows: for each student $i \in I$, for each subject preference $j \in J_i$ and for each class $k \in K_j$, a decision variable X_{ijk} would be created:

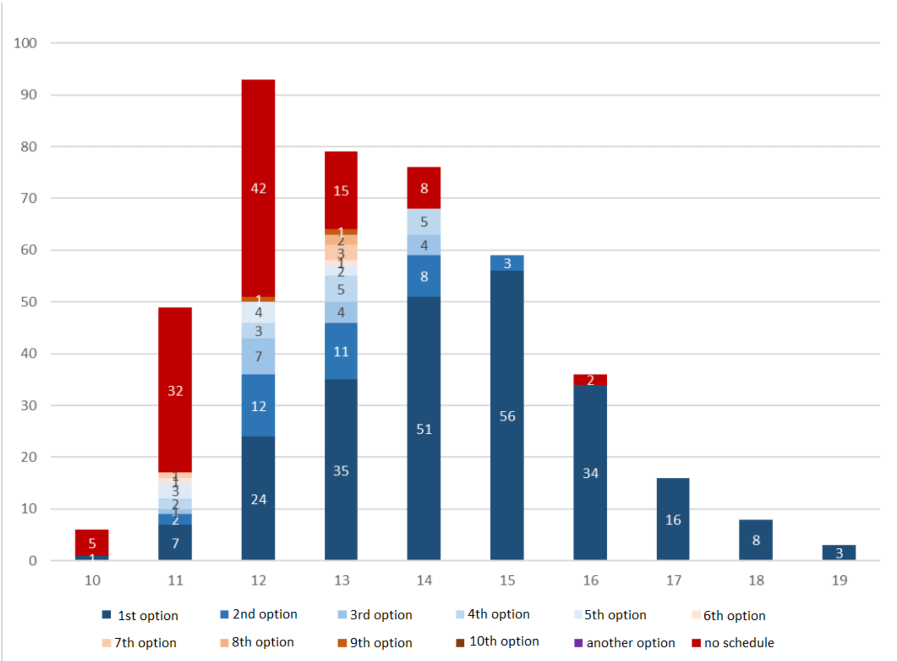

Fig. 1 Option assignment number for the different GPA for the old method for the MIEIC course - 2^{nd} semester 16/17 [17]

$$X_{ijk} = 1 \iff \textit{student } i \textit{ is assigned to class } k \textit{ of subject } j \tag{1}$$

The hard and soft constraints are defined on top of these decision variables. This paper's main focus is on the definition of soft constraints, so the hard constraints are only briefly mentioned so that the formulation of constraints becomes more intuitive for the reader:

1. Each student i is only assigned one class k for each of the subjects j.

$$\sum_{k \in K_j} X_{ijk} \leq 1, \qquad \forall i \in I, \forall j \in J \tag{2}$$

2. Prevent a student i from being assigned to overlapping classes in timeslot $t \in T$. Given K'_{jt} as the classes of $j \in J$ that occupy timeslot $t \in T$.

$$\sum_{j \in J_i} \sum_{k \in K'_{jt}} X_{ijk} \leq 1, \qquad \forall i \in I, \forall t \in T \tag{3}$$

3. Prevent each class from exceeding its predefined capacity. Given I_j as the set of students that have included $j \in J$ in one of their preferences, and q_{jk} as the capacity of class $k \in K$ belonging to the subject $j \in J$

$$\sum_{i \in I_j} X_{ijk} \leq q_{jk}, \qquad \forall j \in J, \forall k \in K_j \tag{4}$$

Given the unyielding satisfaction of the previous constraints, soft constraints need to be considered. Although an overview of the seven soft constraints used for the goal function is presented in Table 1, only the constraints the author deems worthy of discussion are explicitly defined in this paper, meaning the remaining ones cannot be influenced by the stakeholders' behaviour. All the formulas that are used in $maximize(x)$ or $minimize(x)$ refer to values that are normalized in a future step of the implementation so that $x \in [0, 1]$, these normalization is omitted here for simplification purposes. Furthermore, a block represents either a full morning or a full evening, in the context of the soft constraints.

The soft constraints relevant for this paper are briefly described below, and are referenced according to the ID in the first column of Table 1 (where the respective weights can be found):

- **ID 5**: Maximize student preferences.
  This constraint tries to ensure that students with a higher GPA have a better chance of getting the desired schedule, explicitly described through their preferences. This constraint is described in equation 5, given P_i as the preferences of student i; GPA_i as the GPA of student i; o_{ip} as the index of the preference p of the student i; and A_{ip} as an auxiliary boolean variable, such that $A_{ip} = 1$ if the preference p of student i is selected, $A_{ip} = 0$ otherwise.

$$maximize(\sum_{i \in I} \sum_{p \in P_i} (2^{GPA_i} \times (10 - (o_{ip} - 1)) \times A_{ip})) \tag{5}$$

- **ID 6**: Minimize occupied blocks for students without satisfiable preferences.
  This constraint aims at creating a schedule, for those students with no satisfiable preference, that is as grouped as possible, avoiding a scattered schedule, as is common to look for in most timetabling problems. Equation 6 describes this constraint, given O_{ib} as an auxiliary boolean variable, such that $O_{ib} = 1$ if student i, not having any satisfiable preference, is assigned to block b, and $O_{ib} = 0$ otherwise.

$$minimize(\sum_{i \in I} \sum_{b \in B} (O_{ib})) \tag{6}$$

- **ID 7**: Minimize blocks assigned that are not in the student's preferences for students without satisfiable preferences.
  This constraint represents an attempt at not assigning classes to a student that has implicitly omitted some blocks of classes from his or her preferences. This is expressed through equation 7, given B_i as the blocks that belong to the preferences of student i, and O_{ib} as in equation 6.

$$minimize(\sum_{i \in I} \sum_{b \in B \setminus B_i} (O_{ib})) \tag{7}$$

ID	Goal	Criteria	weight	GPA
1	Maximize assignments	Maximize assigned pairs student-subject	30%	yes
2		Maximize students with complete schedule	10%	yes
3		Maximize occupied timeslots	10%	no
4	Balance class occupation	Minimize underemployment of class vacancies	20%	no
5	Meet student preferences	Maximize student preferences	10%	yes
6	Improve schedule quality	Minimize occupied blocks for students without satisfiable preferences	10%	no
7		Minimize blocks assigned that are not in the student's preferences for students without satisfiable preferences	10%	no

Table 1 Soft constraints overview: goals, description, weight and consideration of GPA

The validation of the designed method was done by feeding the data gathered for years where the previous method was in place through this new algorithm. The results presented in Figure 2 show these assignment values for the validation of this algorithm. There is an unmentioned difference between the approaches - the new approach assigns a schedule even if none of the student's preferences were satisfiable, whereas the old did not, as such the *another option* label on Figure 2 should be viewed with this knowledge.

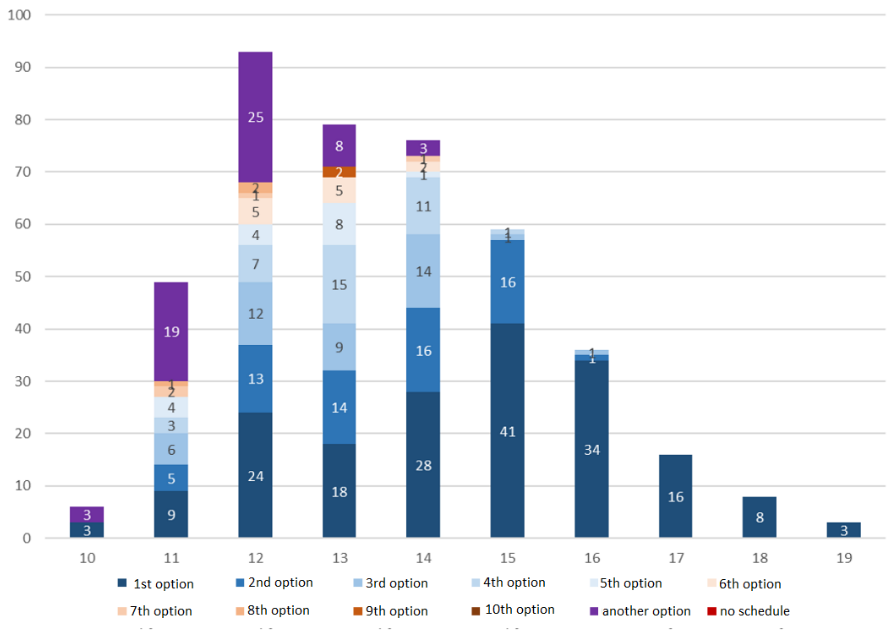

Fig. 2 Option assignment number for the different GPA for the new method for the MIEIC course - 2^{nd} semester 16/17 [17]

4 Analysis of the Case Study

Given the above background, we are now ready to analyse the student scheduling problem at FEUP. In this section, a sequential identification of drawbacks is made on the new model for the solution of the SSP at FEUP. Simultaneously, some of the implementation's benefits are also identified and explicitly explained, for the cases that are seen as lacking discussion in [17]. Furthermore, suggestions of improvement are presented so as to empower this novel approach and to improve the results of the future developments it may undertake.

4.1 Behavioural Game Theory

One of the least analytical topics presented here is that of adversarial behaviour in the preference selection process. However, the author believes it is of paramount importance and that it is precisely because of its insubstantial nature that no noticeable mention of it is made in the definition of the new method [17].
Consider the following simplistic abstraction exercise:

> Alice and Bob are students at FEUP and both want the same schedule. The problem is such that only one of them can get it. Alice, having a greater GPA, has a higher chance of having her preferences met, she then specifies more than one possible schedule ordered by decreasing utility. Bob, on the other hand, has a lower GPA (not so low as to make this example void). However, Bob, unlike Alice, is aware that the assignment program's first priority is to minimize the number of people left with no schedule (see Table 1), if he only presents a single option for his preferences, this will play in his favour, and so he does. Eventually, the schedules are distributed to each individual and Bob gets the schedule he wanted, whereas Alice got one of her other options. Since the schedules are private, Alice is unable to ascertain whether her first option was not met due to someone with higher priority having chosen it, even less that our malevolent Bob (also unaware that he stole Alice's preferred schedule) got the best of the system.

The author knows, from direct observation of the environment, that there are both *Bobs* and *Alices*, in the sense that only a few students are aware of the definition of the new method, or even of its implications in the selection process. This means that we are dealing with an information asymmetric environment. One fact to support this argument is that if we compare the results of the validation methods employed for the new algorithm (see Figure 2) with the same plots, but this time with data posterior to the implementation of the system (see Figure 3 for the first semester of 17/18 and Figure 4 for the second semester of the same year), some patterns start to arise.

Table 2 shows metrics gathered from both aforementioned conditions: data from before and after the deployment of the new system. The gap is clear between the two conditions that the algorithm was applied two, for instance, the percentage of students with a GPA of 16 that were not placed in either their first or second option (preference) was 2.8% in the system validation and 29% after the system deployment. This tendency is actually persistent across all other metrics for these two GPAs, others were not depicted for simplicity and because some do not denote any discrepancies, as happens with the higher GPAs - this actually makes the underlying validation problems be camouflaged. Some factors may influence these results, but the author believes there is indeed a high correlation between information spread and the results obtained, and that the impact of information asymmetry and adversarial behaviour is is very near to undeniable.

	Not first option		**Not first nor second option**	
	GPA of 15	GPA of 16	GPA of 15	GPA of 16
Data used for validation	$18/59 = 30.5\%$	$2/36 = 5.5\%$	$2/59 = 3.4\%$	$1/36 = 2.8\%$
Data from 1st semester of 17/18	$46/80 = 57.5\%$	$7/42 = 16.7\%$	$35/80 = 43.8\%$	$3/42 = 7.1\%$
Data from 2nd semester of 17/18	$36/69 = 52.2\%$	$10/31 = 32.3\%$	$17/69 = 24.6\%$	$9/31 = 29\%$

Table 2 Comparison of some metrics on the new method with data used in the validation and data gathered after the deployment of the system

An easier error to fix is also worth mentioning. During the school year of 17/18 students were presented this new system and they were instructed that the minimum number of preferences the system accepted were three - this was an attempt to reduce the impact of the weights and soft constraints chosen. However, recently, a flaw in the system was found - students could, and to the author's best knowledge, still can, choose three equal preferences. This flaw should be addressed immediately due to havoc it can bring and also to the fact that only some *Bobs* knew about this flaw, resulting in even more information asymmetry and unfairness in the outcome.

To conclude on the topic of behavioural game theory, there is a need to make this process as informed throughout as possible, not just at FEUP but also across other institutions implementing similar methods.

4.2 Priority Across Preferences

From equation 5 we gather that the weight of each of the preferences of the students decreases in a linear fashion. Furthermore, it is also possible to deduce that between consecutive integer GPAs the weight of corresponding (by order) preferences doubles as the GPA increases by one. Both of these notions are visible in Figure 5 that plots the weights of each preference order (from

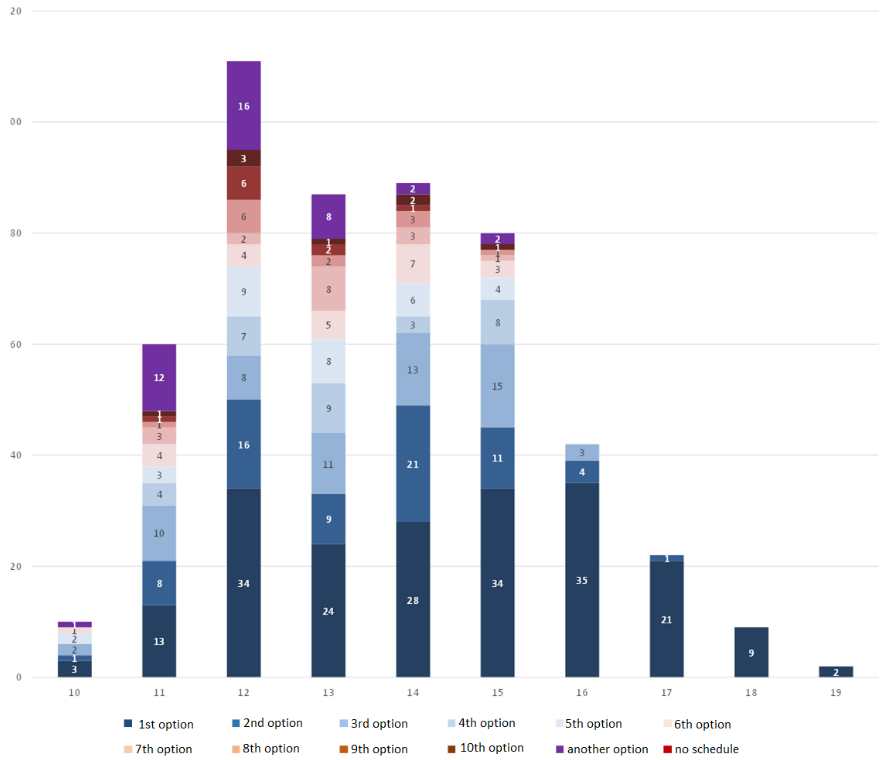

Fig. 3 Option assignment number for the different GPA for the new method for the MIEIC course on the first semester of 2017/2018 - after the deployment of the new method

1^{st} to 10^{th}) to each of the valid GPAs (10 to 20).
This constraint combines two of the most important factors for the students: the impact of their GPA on the outcome of the system and the relative weight between each of their individual preferences. The author considers this constraint to be rather stringent and believes that the formula in equation 8 would be more flexible for both factors in consideration. The new equation aims at avoiding the linearly decreasing weight of each preference for the same student, it is converted into a logarithmic slope which makes the sorting of the preferences a more meaningful and realistic action, by taking into consideration the fact that most preferences, but the first, are simply safety fallbacks options that students present, instead of linearly decreasing in importance for the student. Figure 6 depicts equation 8.

$$maximize(\sum_{i \in I} \sum_{p \in P_i} \times (\alpha^{GPA_i}(1 - \log_\beta o_{ip}) \times A_{ip})), \qquad \beta \geq |P| + 1 \tag{8}$$

Where $|P|$ is the maximum number of preferences allowed, 10 for the current case; α and β represent two parameters that are fixed for each iteration of

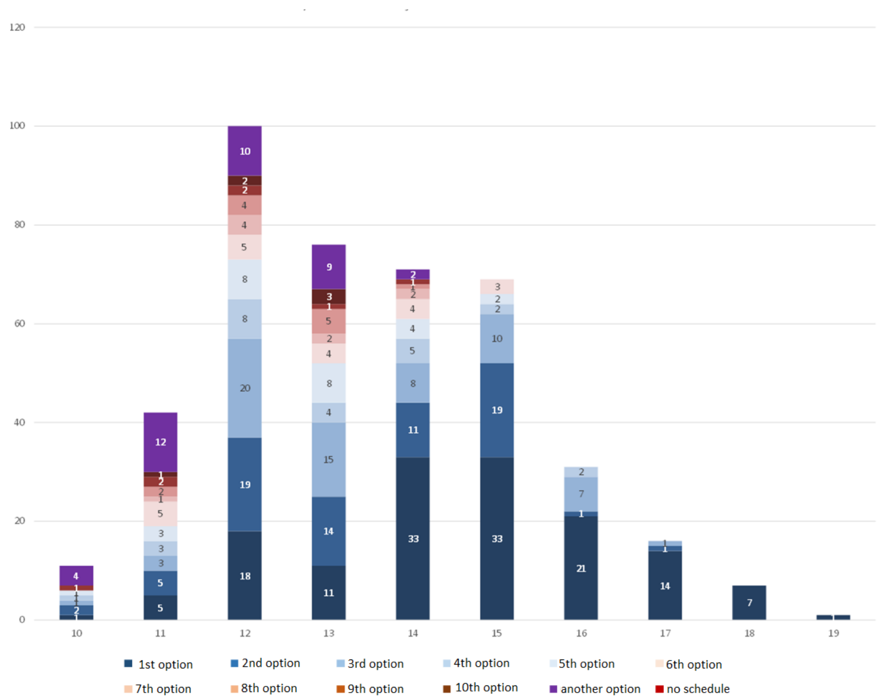

Fig. 4 Option assignment number for the different GPA for the new method for the MIEIC course on the second semester of 2017/2018 - after the deployment of the new method

the program but should be tuned accordingly. α describes the rate at which students with higher GPAs are favoured. β describes the manner in which individual preferences' relative weights decrease. As β increases the discrepancy between consecutive GPAs increases as well, some β benchmarks are worth mentioning. Consider *Student A* with a GPA of x and *Student B* with a GPA of $x+1$. For the case of $\alpha = 2$: if $\beta = 11$ (minimum value) then the first preference of *Student A* has as much weight as the fourth preference of the *Student B* and less than *Student B*'s third preference. If we were to consider $\beta = 20$ (see Figure 6), then the equivalence would be located between the fifth and fourth preferences of *Student B*. Also, $\beta = 100$ would make the weight of the first preference of *Student A* weight exactly as much as the tenth preference of *Student B*.
In addition, the author believes that much could also be gained by using a two-dimensional representation of the student preferences, where preferences in the same column would have the same level of importance for the student and preferences in the same row would have decreasing - ideally logarithmically - importance, and therefore weights, for the students. equation 8 would not even need to change, as it already encompasses the would-be row index and accommodates the column index seamlessly.

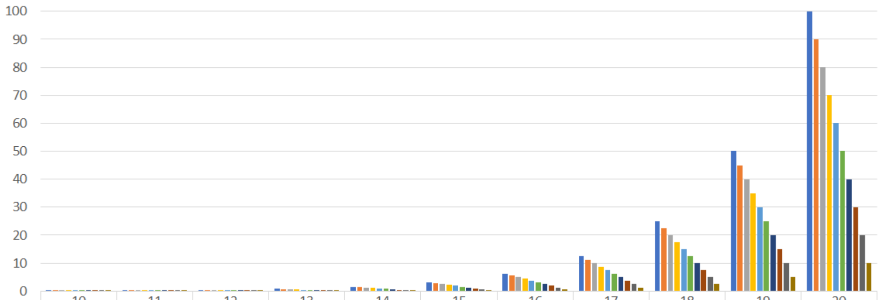

	10	11	12	13	14	15	16	17	18	19	20
1	0.098	0.195	0.391	0.781	1.563	3.125	6.25	12.5	25	50	100
2	0.088	0.176	0.352	0.703	1.406	2.813	5.625	11.25	22.5	45	90
3	0.078	0.156	0.313	0.625	1.250	2.500	5	10	20	40	80
4	0.068	0.137	0.273	0.547	1.094	2.188	4.375	8.75	17.5	35	70
5	0.059	0.117	0.234	0.469	0.938	1.875	3.75	7.5	15	30	60
6	0.049	0.098	0.195	0.391	0.781	1.563	3.125	6.25	12.5	25	50
7	0.039	0.078	0.156	0.313	0.625	1.250	2.5	5	10	20	40
8	0.029	0.059	0.117	0.234	0.469	0.938	1.875	3.75	7.5	15	30
9	0.020	0.039	0.078	0.156	0.313	0.625	1.25	2.5	5	10	20
10	0.010	0.020	0.039	0.078	0.156	0.313	0.625	1.25	2.5	5	10

Fig. 5 Plotting of the weights of each preference order against discrete GPAs, according to equation 5

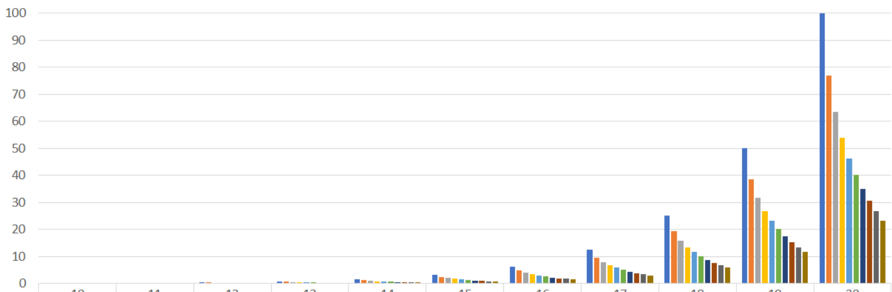

	10	11	12	13	14	15	16	17	18	19	20
1	0.098	0.195	0.391	0.781	1.563	3.125	6.250	12.500	25	50	100
2	0.075	0.150	0.300	0.600	1.201	2.402	4.804	9.608	19.216	38.431	76.862
3	0.062	0.124	0.247	0.495	0.989	1.979	3.958	7.916	15.832	31.664	63.327
4	0.052	0.105	0.210	0.420	0.839	1.679	3.358	6.716	13.431	26.862	53.724
5	0.045	0.090	0.181	0.362	0.723	1.446	2.892	5.784	11.569	23.138	46.276
6	0.039	0.078	0.157	0.314	0.628	1.256	2.512	5.024	10.047	20.095	40.190
7	0.034	0.068	0.137	0.274	0.548	1.095	2.190	4.380	8.761	17.522	35.044
8	0.030	0.060	0.119	0.239	0.478	0.956	1.912	3.823	7.647	15.293	30.587
9	0.026	0.052	0.104	0.208	0.416	0.833	1.666	3.332	6.664	13.327	26.655
10	0.023	0.045	0.090	0.181	0.362	0.723	1.446	2.892	5.784	11.569	23.138

Fig. 6 Plotting of the weights of each preference order against discrete GPAs, according to equation 8 ($\alpha = 2$ and $\beta = 20$)

4.3 Handling Students without a Preference-Based Allocation

Soft constrains 6 and 7 (see Table 1) dictate how the schedule for students without satisfiable preferences is created. Sot constraint 6 is a schedule quality improvement constraint, and this represents such an important notion - gapped schedules are nightmarish - that every student should have an equal right at having a gapless schedule, at this stage. However, the author believes the same cannot be said for soft constraint 7. This is no longer a question of schedule quality, but student preferences propagation. As such, it would only make sense that, even though these students had no preference that could yield a feasible solution for the problem, there would be some balance in which students with better grades could have a better chance of being assigned a schedule that reflected the implicit preferences they have expressed. The author argues that if meeting student's preferences according to their GPA is one of the goals of the novel system, this premise should be carried out to its full extent and, therefore, soft constraint 7 should take the student's GPA into account.

4.4 GPA for Life

Although some students aim at good grades naturally, something can also be done to further push motivation and to encourage short-term dedication. Both systems described in this paper consider a student's GPA, that is, an average of **all** the grades a student has been given since their school path began in the program. This may work as a relief for those students with high performance from the beginning, as their advantage can almost be taken for granted. The opposite is also true - students that perform poorly on the first semesters or years are then cursed with having little to no chance of raising in the priority ladder. This can be a motivation barrier as one of the reasons that good grades are useful is actually not accessible to them.
Nevertheless, this can be fixed. Imagine, for instance, that the value of the GPA used for the calculations is not that of the overall course, but rather of the last one or two semesters. By knowing this, two things happen: students with better grades in the beginning stop taking that advantage for granted and students with worse grades are actually fuelled to improve, to do better at school, as they do not loose the race in the first minutes, but can actually win if they just start running faster.
It is worth mentioning that considering, for instance, only a student's best semester GPA would fall into the same kind of bias as the complete GPA.

5 Conclusions and Future Work

Overall, the presented work takes a glance into what a single university is doing and is able to find a few important improvements and come forward with a few important suggestions. Yet, it does beg the question: What other flaws are out there to find across timetabling systems in universities, worldwide? The student scheduling systems might not be making the best of their resources to find fair and just solutions. The universe of problems can stem from things so little as tuning weights to something as large as a community behaviour in an adversarial environment. Therefore, researchers are challenged to question their work and that of others, but more importantly they are **dared** to look

past what might seem like the most obvious solution.
As future work, three ideas are shared. Firstly, the design of post-scheduling processes that can be applied to almost any SSP that considers a trade system amongst students after the results are published. Imagine, that each student specifies pairs of classes they would be willing to exchange, with this information a graph visualization can be conceived. Eulerian circuits, which describe graphs where a path that transverses each arc only once and returns to the original node exist [18], could then be applied for improving the quality of the solution (system administrators might need to slightly increase the number of vacancies in some classes to improve the eulerization of the graph). Secondly, to consider not just blocks of classes but also preferred teachers, either explicitly - each student could specify preferences - or implicitly - deducible in the same way that students' block preferences are, as described in Section 4.3. Thirdly, the search for innovative ways of tuning the weights used for soft constraints and also for specific parameters for each soft constraint, an example would be α and β as described in equation 8 and the approaches could go from brute-force simulations to hyper-parameter tuning in machine learning classifiers.
Some questions remain unanswered. Is the fact of fixing the number of preferences a student can make feasible in a real world scenario, where students may only find one or few plausible schedules for them? Are the weights used for the soft constraints tuned? What other considerations could be made to diminish information asymmetry in the case of FEUP? One crucial setback for the production of this work was the lack of access to real data, as it could not be provided in a manner expedient enough for the pressing speed of this report, this could shine some light on different aspects and also help uncover other hidden *minutiae*.
As the reader must have realized by now, this paper pushes against blind approaches through a mixed (some might say multi-objective) approach in order to emphasize that SSP systems should make sound and realistic assumptions, and escape the desire to break the proverbial *plateau* all too quickly, for validation must reflect the scope of the problem as well as its peculiarities. Metaphorically, if a guitar is missing strings, they should be added, but the process is not finished until they have been tuned. In accordance, scheduling problems should not let silence be replaced by tuneless sounds, rather by mellifluous chords that please the audience and meet their tastes, as much as possible - such is proper optimization.

References

1. Berrada, I., Ferland, J.A., Michelon, P.: A multi-objective approach to nurse scheduling with both hard and soft constraints. Socio-Economic Planning Sciences **30**(3), 183–193 (1996)
2. Bouzarth, E.L., Forrester, R., Hutson, K.R., Reddoch, L.: Assigning students to schools to minimize both transportation costs and socioeconomic variation between schools. Socio-Economic Planning Sciences pp. 1–8 (2017)

3. Burke, E.K., Jackson, K.S., Kingston, J.H., Weare, R.F.: Automated Timetabling: The State of the Art. The Computer Journal **40**(9), 565–571 (1997)
4. Cheng, E., Kruk, S., Lipman, M.: Flow Formulations for the Student Scheduling Problem. Practice and Theory of Automated Timetabling IV pp. 299–309 (2003)
5. Desroches, G.L., Sylvain: The problem of assigning students to course sections in a large engineering school. Computers & Operations Research **13**(4), 387 – 394 (1986)
6. Fraga, P.: bot-horarios (2017). URL https://github.com/pedrofraga/bot-horarios
7. Harsanyi, J.C., Selten, R.: A General Theory of Equilibrium Selection in Games **1** (1988)
8. Heitmann, H., Brüggemann, W.: Preference-based assignment of university students to multiple teaching groups. OR Spectrum **36**(3), 607–629 (2014)
9. Henebry, K.: The impact of class schedule on student performance in a Financial Management Course. Journal of Education for Business **73**(2), 114 (1997)
10. Li, X., Gao, L., Li, W.: Application of game theory based hybrid algorithm for multi-objective integrated process planning and scheduling. Expert Systems with Applications **39**(1), 288–297 (2012)
11. Müller, T., Murray, K.: Comprehensive approach to student sectioning. Annals of Operations Research **181**(1), 249–269 (2010)
12. Owen, S., Yawson, A.: Information asymmetry and international strategic alliances. Journal of Banking and Finance **37**(10), 3890–3903 (2013)
13. Ribeiro, L.M., David, G.: SiFEUP (1996)
14. Rudová, H., Murray, K.: University Course Timetabling with Soft Constraints. Practice and Theory of Automated Timetabling IV pp. 310–328 (2003)
15. Sanfey, A.G.: Social Decision-Making - Insights from Game Theory and Neuroscience. Science Magazine **338**(June), 1541–1545 (2011)
16. Schimmelpfeng, K., Helber, S.: Application of a real-world university-course timetabling model solved by integer programming. OR Spectrum **29**(4), 783–803 (2007)
17. Silva, G.T.N.d.: Atribuição automática de estudantes universitários a turmas baseada em otimização multicritério (2017)
18. Thimbleby, H.: The directed Chinese postman problem. Software - Practice and Experience **33**(11), 1081–1096 (2003)
19. Willemen, H.M.M.t.E.J.: Some Complexity Aspects of Secondary School Timetabling Problems. In: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) (2001)