Experimental Study of Preference-Based Assignment of University Students to Multiple Teaching Groups

Henrik Heitmann Wolfgang Brüggemann

23.12.2012

Based on the model formulation by Heitmann and Brüggemann (2012), an experimental evaluation is carried out for the student scheduling problem at the Business School of the University of Hamburg. The main purpose of this study is to demonstrate a better performance of an optimised student assignment without sub-cohorts compared to the status quo. Further studies about runtimes and sensitivity analyses are presented. The evaluation is carried out on the basis of real timetables and stochastically generated registration data. Subsequently, the experimental design, input data and the default parameter setting of the study are described. Afterwards, results are presented, findings are explained and discussed.

1 Experimental design

The overall framework in terms of a general problem description and a mixed-integer programming model can be found in Heitmann and Brüggemann (2012). There are 28 experiments, each consists of 20 independent optimisation runs. Each run is based on two components of input data. The first input data component — this data is the same in every run of an experiment — is deterministic which includes the sets of students, student time slot blocks, lectures, parallel teaching groups with group capacities and teaching times, discrete time slots and teaching days. This data is given by the study plan of the winter semester

2010/11 (University of Hamburg, Administration, 2010). The second input data component — this is different in every run — is stochastic. To this component belong the student enrolments, group preferences and fraction of students deviating from default weightings (100 points on group preferences). Overall, 560 optimisation runs are evaluated.

The experiments are carried out in GAMS 22.9.2 and solved by CPLEX 11.2.0 with two parallel threads in Windows 7 Enterprise on a 3.0 GHz Intel Core2 Duo machine with 4 GB RAM. For each run, the CPU runtime varies between 5 seconds and 11 minutes depending on the parameter settings and the stochastically generated data.

2 Input data

Subsequently, the part of deterministic input data, which does not change between the 20 independent runs of an experiment, is described. Thereafter, the generation of the non-deterministic part of input data is explained. The deterministic input data consist of the sets and information about teaching times and group capacities given by the study plan (University of Hamburg, Administration, 2010).

The sets are composed of

- 1,865 students, $I = \{1, 2, 3, \ldots, 1865\}$,
- 27 courses,
- 251 multiple teaching groups, thereby the number of multiple teaching groups varies between 2 and 32, e.g. for main lectures and tutorials,
- 12 discrete one-hour time slots (specified by their beginning), $H = \{8, 9, 10, \ldots, 19\}$ and
- five teaching days, $D = \{$Monday, Tuesday, Wednesday, Thursday, Friday$\}$.

1,865 students are made up from real students' enrolments of different study fields to classes offered by the Business School. The academic programmes and number of students are listed in Table 1. Also, the 27 courses and the 251 parallel teaching groups are given by the study plan. Table 2 shows the list of taught courses and

Table 1: Academic programmes and number of students.

Study programmes	Number of students
Business Management	840
Economics	190
Business Management & Computer Science	120
Economics & Culture of China	90
Business Management & Mathematics	120
Business Management for the Timber Industry	120
Business Management & Engineering	270
Teaching of Business Management & Economics	115
Total	1,865

Table 2: Offered courses, number of multiple teaching groups and mean number of student enrolments.

	Number of multiple teaching groups		Mean number of student enrolments
Course	Lecture	Tutorial	Lecture/Tutorial
Accounting	2	15	779
Private & Business Law	2	18	688
Basic Economics	2	28	668
Mathematics	2	32	582
Macroeconomics	2	27	545
Computer Science	2	13	486
Statistics	2	24	478
Investment	2	16	476
Computer Course	8	–	451
Balancing	2	18	398
Business Management	2	12	389
Industrial Econometrics	1	5	89
Applied Econometrics	1	6	72
Computer-aided Modelling	1	6	71
Total	251		11,891

the number of parallel offered groups of lectures and tutorials. Furthermore, all parallel groups are scheduled, such that the timetable is fixed and the group capacities result from the number of seats in the room, in which the groups are to be taught. Due to the fact that Business Management & Engineering is an interdisciplinary programme which is taught by different institutions, these students can only be assigned to parallel teaching groups here which are taught on Tuesdays. For these students i, the blocked-time-slot indicator is set to $b_{ihd} = 1$ for all other days of the week $d \in D \setminus \{\text{Tuesday}\}$ and all time slots h.

The non-deterministic component of input data consists of stochastic, independently generated data. This includes the three parts student enrolments to the courses, group preferences and different weightings. Firstly, with the known number of student enrolments to the different study fields, random enrolments to courses are generated. If a student is enroled to a course, the student must be assigned to the main lecture as well as to the corresponding tutorials. The mean numbers of enrolments to the courses over all 20 independent runs are listed in Table 2 and sum up to a total of 11,891 assignments.

Secondly, for each student enrolment a preference has to be generated. In Figure 1, the probabilities are shown to generate a preference depending on the start time of the course. In fact, because of the same contents of teaching in

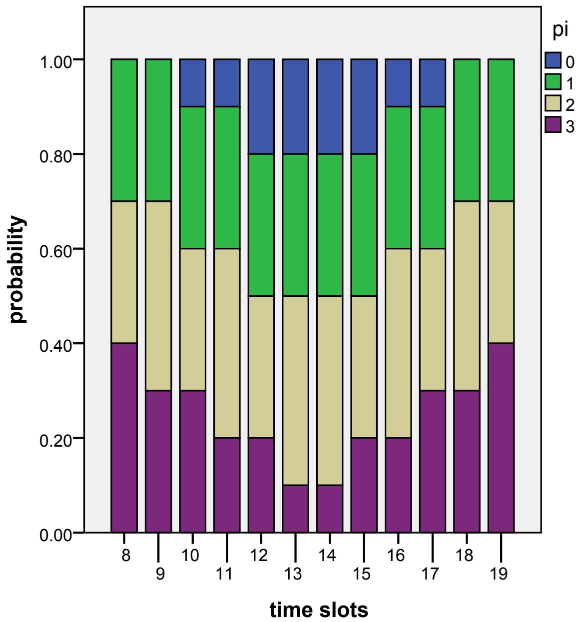

Figure 1: Probabilities for group preferences π_{ijk} to corresponding time slots.

each parallel group, the parallel groups mainly differ in start times (and possibly instructors), so that this distinction is a good predictor for different preferences which very likely conform with real student preferences. The probabilities for the preferences differ in time slots — the start time of each time slot is given on the abscissa. The blue, green, beige and violet parts of the bars are the probabilities for the selection of a 0, 1, 2 and 3–valued group preference, respectively. Thereby, the probability increases for 0–valued (highest) preference and decreases for 3–valued preference towards midday; assuming students would rather be assigned to courses which start in the middle of the day.

Thirdly, investigating the sensitivity of CPU-times, different weightings are generated. Thereby, the percentage of students without default weighting varies between 0% and 50%. For this purpose, a specified percentage of students in each run is randomly selected, each of these students gets a randomly generated set of different weightings as an individual preference. Therefore, a set of randomly rated goals is chosen with the same probability for each individual and for each selected goal, individual weights are again randomly generated.

3 Default parameter setting

A default parameter setting is defined for the following tests and sensitivity analyses. The experiments are then carried out ceteris paribus — changing one parameter at a time only. The default parameter setting is summarised in Table 3. In the default experiment, 30% of the students do not use default weightings for their preferences. The set of lunch break time slots is defined to $L_{id} = \{5, 6, 7\}$. The morning core times end at $e^M_{id} = 14$, and the afternoon core times begin at $s^A_{id} = 13$. The parameter for the preferred number of days at the university is set to $m_i = 4$. The individual objective value upper bound is set to $\overline{z}_i = 0.5$ which corresponds to 50% fulfilment of the individual preferences in the worst case. Violating the individual welfare upper bound is more penalised ($\alpha_i = 100$) than an unbalanced group utilisation ($\beta_j = 1$). Default parameter setting for the minimum room-capacity utilisation is set to $\kappa_{jk} = 0.75$. Furthermore, for the normalisation of each individual welfare violation, the parameters are set to the best and worst magnitude.

4 Results

First, the advantage of the students' assignments without sub-cohorts is shown. Then, a sensitivity analyses is presented. The purpose of this study is to analyse

Table 3: Default parameter setting.

Parameter for ...	Parameter	Value
Fraction of students deviating from default preferences		30%
Set of lunch break periods	L_{id}	$\{5, 6, 7\}$
Morning core time end	e_{id}	14
Afternoon core time start	f_{id}	13
Maximum presence days	m_i	4
Individual welfare upper bound	$\overline{z}_i$	0.5
Penalty for individual welfare upper bound	α_i	100
Penalty for unbalanced parallel group utilisation	β_j	1
Group utilisation factor	κ_{jk}	0.75
Best magnitude of group preferences	$\underline{\Pi}_i$	0
Worst magnitude of group preferences	$\overline{\Pi}_i$	$3 \cdot \lvert J_i \rvert$
Best magnitude for lunch break violations	$\underline{l}_i$	0
Worst magnitude for lunch break violations	$\overline{l}_i$	5
Best magnitude of maximum presence days violation	$\underline{r}_i$	0
Worst magnitude of maximum presence days violation	$\overline{r}_i$	1
Best magnitude of assignment outside core times	$\underline{w}_i^M$, $\underline{w}_i^A$	0
Worst magnitude of assignment outside core times	$\overline{w}_i^M$, $\overline{w}_i^A$	75

and to identify the joint effects of different settings of the parameters for individual welfare upper bound ($\overline{z}_i$) and group utilisation factor (κ_{jk}). Then, the results are summarised and the most important joint effects are presented and discussed. Each experiment with a different parameter setting consists of 20 independent runs. Mean values and interpolations are calculated over the results of these runs.

The mean number of π-valued assignments with and without sub-cohorts are depicted in Figure 2. The two different experiments — the assignment of the students with and without sub-cohorts — are indicated by the two different groups of bars on the abscissa. The performance of the experiments — measured by the mean number of π-valued assignments — is shown on the ordinate. As expected, the assignments without sub-cohorts lead to better results than the assignments with sub-cohorts. In contrast to the assignments with sub-cohorts, the mean number of $\pi = 0$ assignments (blue bars) is increased and the mean numbers of $\pi = 1, 2$ and 3 assignments (green, beige and violet bars) are decreased without sub-cohorts. This is reflected in the mean objective function values with sub-cohorts of 728.9 and without sub-cohorts of 443.5. Thus, the student preferences are better met by the assignment without dividing the students and parallel of-

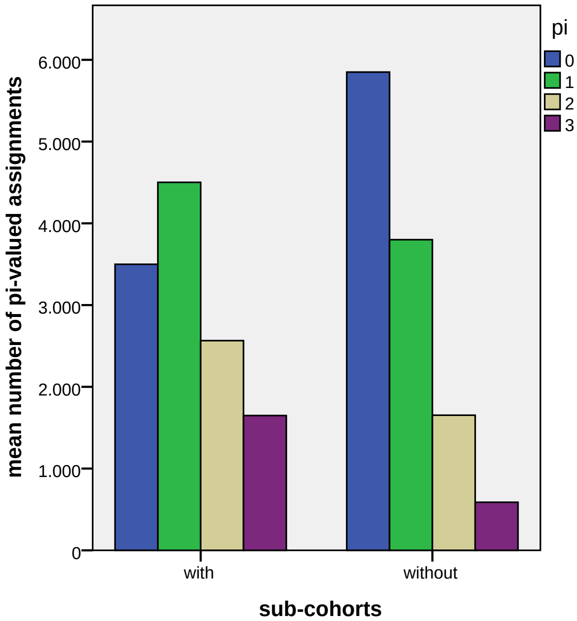

Figure 2: Mean number of π-valued assignments with and without sub-cohorts.

fered groups into sub-cohorts. Another key argument for the assignment without sub-cohorts is the number of infeasible enrolments. Without sub-cohorts, all enrolments in each of the 20 runs are feasibly possible, such that all students can be assigned without time clashes. In contrast, on average 67 student enrolments are infeasible with the division into sub-cohorts. However, these infeasibilities would, for example, result in additional expenditure for the administration, because such infeasible enrolments have to be resolved. These results suggest the assignment of the students to the parallel offered groups without sub-cohorts, where student preferences can be accommodated and fewer infeasible enrolments need to be resolved.

The main question of the feasibility study is, whether the student-scheduling problem can be solved in operable time for the Business School in Hamburg. Experiments are carried out to find out how CPU-times depend on the percentage of students without default weighting. The CPU-times of these experiments are shown in Figure 3. The experiments differ in the percentage of students without default weighting (100 points on group preferences) under the assumption, that in the worst case 50 % of the students assign positive weights to other goals than their group preferences. The different percentages, which are examined, are shown on the abscissa. The CPU-times of these experiments are given on the ordinate. The bars mark the range of solution times, in which the 20 independent

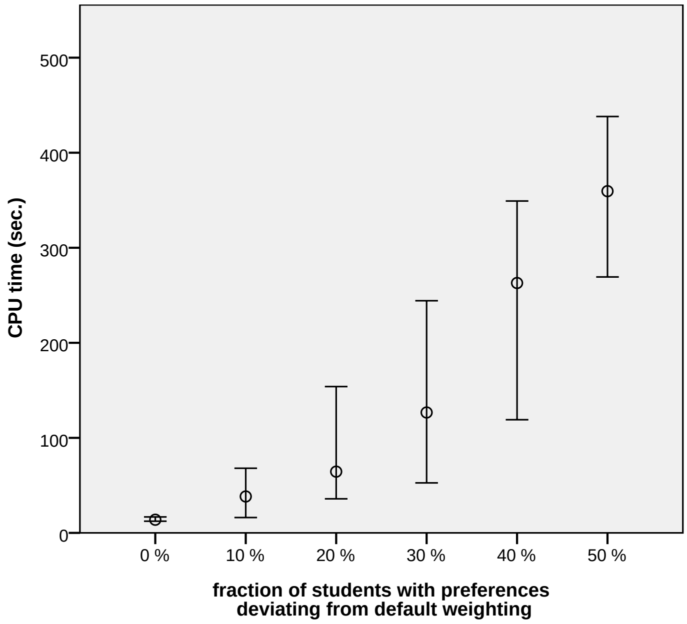

Figure 3: CPU times (sec.) of experiments depending on the percentage of students with preference weightings different from the default weighting scheme.

runs are solved while the circles on the bars identify the mean CPU-time of each experiment. Inspection of Figure 3 strongly suggests a super linear relation between the percentage of students with different weightings and CPU-times. Firstly, mean CPU-times increase with the number of students without default weighting. Secondly, the variance of CPU-times appears to increase as well. Thirdly, the maximal CPU-time in the worst case of this analysis (50 % of the students without default weighting) is 438 seconds. It follows that in a real application, the student-scheduling problem at the School of Business in Hamburg can be optimally solved in operable time.

Joint effects of different individual objective value upper bounds $\overline{z}_i$ are illustrated in Figure 4. It should be noted that these individual objectives are normalised to be between 0 and 1. The figure shows on the left side the sum of all individual objective values (green circles) and on the right side the number of violations (blue crosses) resulting from the individual upper bound in 20 experiments. The corresponding averages from the 20 experiments are also shown and

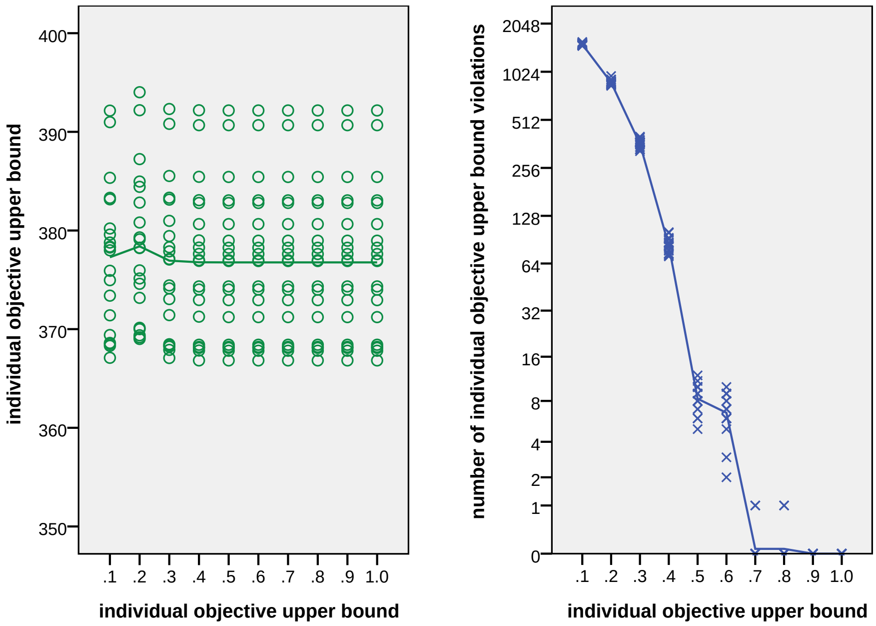

Figure 4: Joint effects of different individual objective value upper bounds $\overline{z}_i$.

given in interpolation lines. It may be concluded that, with a decreasing upper bound $\overline{z}_i$ from 1.0 to 0.4, the sum of all individual objective values is almost constant. The non-monotonous course of the green interpolation line between the range from 0.3 to 0.1 can be explained by the focus on only the sum of all individual objective values and not the objective values which additionally include the individual upper bound and group utilisation compensations. Moreover, the increase is only very small in relation to the whole sum of individual objective values with upper bounds in the range from 0.3 to 0.1. However, the number of upper bound violations increases rapidly — nearly up to the whole number of students — if the individual objective upper bound is decreased to 0.1. But for the default parameter setting of the individual upper bound $\overline{z}_i = 0.5$, the number of violations is less than 1% of all possible. As a consequence, the individual preferences are entirely fulfilled for all other students in this case.

Next, the group utilisation factor κ_{jk} is analysed. Figure 5 displays the effects of different factor values which are printed on the abscissas. On the ordinate in the left figure, the mean (of 20 experiments) sum of all individual objective values $z_i + u_i$ is given, and on the ordinate in the right figure the variance of parallel group utilisations is shown. The means are again shown with their interpolation lines. As expected, the mean sum of the individual objective values (green circles)

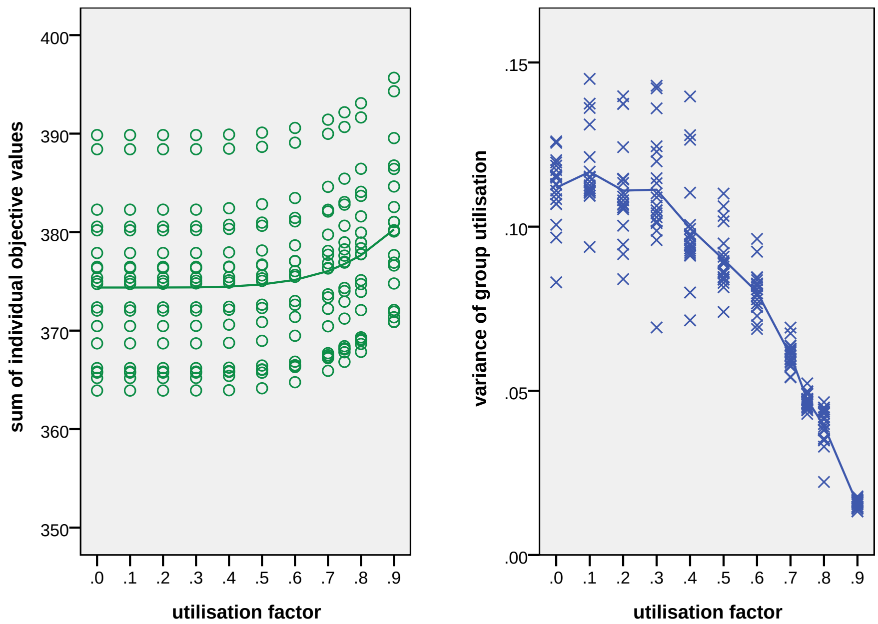

Figure 5: Joint effects of different utilisation factors (κ_{jk}).

exhibits a nearly constant behavior up to utilisation factor $\kappa_{jk} = 0.4$; only then the curve shows a moderate, clearly nonlinear increase. In contrast, the mean variance of group utilisations increases almost monotonously with a decreasing utilisation factor. This is particularly evident in the sloping curve (blue crosses).

5 Summary and conclusion

An experimental study is carried out at the Business School of the University of Hamburg for the student scheduling problem which is modelled as a mixed-integer programme by and can be found in Heitmann and Brüggemann (2012). The design, input data and default parameter setting for this experimental study are described. The problem is easily solvable in operable time for typical instances generated randomly on the basis of a real-life problem. Furthermore, the sub-cohorts (i. e., the division of students and multiple teaching groups) used up to now become obsolete. Finally, the mutual influence between students' individual welfare and balanced group utilisations is analysed.

References

Heitmann H, Brüggemann W (2012) Preference-Based Assignment of University Students to Multiple Teaching Groups. Submitted to OR Spectrum

University of Hamburg, Administration (2010) Winter Semester 2010/11 1st and 3rd Semester Bachelor Tutorial Groups. URL `http://www.uni-hamburg.de/fachbereiche-einrichtungen/pa2/Wintersemester_1011_/UEBERSICHT_Bachelor-Uebungsgruppen_WiSe1011.pdf`, last visit: 2012/04/02 4:00 p.m.