

A simulated annealing approach to the curriculum-based course timetabling problem

Ruggero Bellio · Sara Ceschia · Luca Di Gaspero · Andrea Schaerf · Tommaso Urli

1 Introduction

Course timetabling is a typical problem that all universities around the world have to face every semester. There exist a number of variants of this problem, depending on the specific requirements of the institution involved [13].

Thanks also to the international timetabling competition ITC-2007 [11], two formulations have, to some extent, lifted up to the status of “standard”. These are the so-called Post-Enrolment Course Timetabling (PE-CTT) [8] and Curriculum-Based Course Timetabling (CB-CTT) [6]. The distinguishing difference between the two is the origin of the conflict matrix between courses, which is based on student enrolments and on predefined curricula, respectively. This is however only one of the differences, which actually include also many other features and cost components. Indeed, in the PE-CTT each event is self-standing, whereas in the CB-CTT a curriculum is composed by different courses that consist of multiple lectures and share the same students. In addition, the soft constraints are diverse: in the PE-CTT they are all related to events, penalising late, consecutive and isolated ones; in the CB-CTT they mainly involve curricula and courses, penalising isolated lectures in a curriculum, trying to minimise the number of days in which lectures of the same course are spread and the room transfers.

In this work we focus on CB-CTT, building upon our previous work on this problem [3]. Some of the ideas, however, are transferred from work on PE-CTT [5], on which we have obtained good results, including many new best-known solutions, applying a similar approach.

R. Bellio
DIES, University of Udine
via Tomadini, 30/A
E-mail: ruggero.bellio@uniud.it

S. Ceschia, L. Di Gaspero, A. Schaerf, and T. Urli
DIEGM, University of Udine
Via delle Scienze 206
E-mail: {sara.ceschia, l.digaspero, schaarf, tommaso.urli}@uniud.it

2 Search method

We propose a solution method based on Simulated Annealing (SA), along the lines developed for the PE-CTT [5]. The main modelling difference is that, in the CB-CTT formulation, the classrooms play an important role for the cost of a solution (mainly due to the `RoomStability` cost component which states that all lectures of the same course should be given in the same room), whereas in the PE-CTT all suitable rooms can be used interchangeably for the same lecture.

This observation has an impact on the definition of the neighbourhood, which is the crucial feature of all methods based on local search. In fact, while for PE-CTT a move could consider only the period of a lecture, leaving the room implicit to be selected a-posteriori, for CB-CTT the room is one of the attributes of the move.

Another difference with respect to previous work on the PE-CTT, is the use of a non-geometric cooling scheme for SA. In this version, the temperature decreases faster in the beginning of the search and slower towards the end, in order to speed-up the process at high temperatures. This allows to “save” iterations that can be put to better use in the final steps of the search, when intensification becomes more important. This technique, already described by Johnson *et al.* [7] as *cutoffs*, is achieved by triggering the cooling steps not only based on the number of iterations expired but also on the number of number of moves accepted.

We do not set a minimum temperature, but we base our stop criterion on the total number of iterations. This way, we have a running time which is fixed to a predefined value, so that we can have a fair comparison with previous work.

3 Experimental analysis and results

We have tuned the main parameters of our method on a set of instances, generated by Lopes and Smith-Miles [9], designed to mimic the features of real-world instances from the University of Udine.

Competing parameter assignments were compared against each other using an *F-Race(RSD)* approach [4] where the values for the starting parameter configurations were sampled from the Hammersley point sequence [15]. This choice has been driven by two properties of this sequence that make it favourable to parameter tuning. First, points from the Hammersley set exhibit *low discrepancy*, i.e. they are well spread across the parameter space despite being random-like. Second, the sequence is *scalable* both with respect to the number of points to be sampled and to the number of dimensions of the sampling space. This allows to choose the number of racing configurations in advance, avoiding an full-factorial approach, while still tuning several parameters at once. Note that this process is generally applicable to any search method involving multiple parameters, and has been automated through *json2run* [14], a tool for experiment design and analysis developed and open-sourced by our research group.

All the experiments were executed on an Ubuntu Linux 12.04 machine with 16 Intel® Xeon® CPU E5-2660 (2.20GHz) cores. The total number of allowed iterations is set to $3.6 \cdot 10^8$ which corresponds to the running time prescribed by the ICT-2007 (about 6 minutes on our machine).

Inst.	Müller		Lü & Hao		Abdullah <i>et al.</i>		Asin Achá & Nieuwenhuis	SA		Best known
	ITC 07	[12]	[10]	[1]	[2]	z		avg	best	
comp01	5.0	5	5.0	5	5.0	5	5	5.00	5	5
comp02	61.3	43	60.6	34	53.90	39	24	53.0	40	24
comp03	94.8	72	86.6	70	84.20	76	111	79.0	70	66
comp04	42.8	35	47.9	38	51.90	35	35	38.3	35	35
comp05	343.5	298	328.5	298	339.5	315	1343	365.20	326	290
comp06	56.8	41	69.9	47	64.40	50	27	50.4	41	27
comp07	33.9	14	28.2	19	20.20	12	6	23.8	17	6
comp08	46.5	39	51.4	43	47.90	37	37	43.6	40	37
comp09	113.1	103	113.2	99	113.90	104	171	105.0	98	96
comp10	21.3	9	38.0	16	24.10	10	4	20.5	11	4
comp11	0.0	0	0.0	0	0.0	0	0	0.00	0	0
comp12	351.6	331	365.0	320	355.90	337	977	340.5	325	300
comp13	73.9	66	76.2	65	72.40	61	59	71.3	64	59
comp14	61.8	53	62.9	52	63.30	53	51	57.9	54	51
comp15	94.8	–	87.8	69	88.00	73	111	78.8	70	66
comp16	41.2	–	53.7	38	51.70	32	18	34.8	27	18
comp17	86.6	–	100.5	80	86.20	72	56	75.7	67	56
comp18	91.7	–	82.6	67	85.80	77	83	80.8	69	62
comp19	68.8	–	75.0	59	78.10	60	57	67.0	61	57
comp20	34.3	–	58.2	35	42.90	22	4	38.8	33	4
comp21	108.0	–	125.3	105	121.50	95	86	100.1	89	75
	87.2	–	91.2	74.2	88.13	74.52	155.48	82.3	73.4	63.7

Table 1 Comparison of results

The best parameter configuration has been tested on the `comp` instance set, which has been used extensively in the literature. Table 1 shows the results in comparison with the best available ones. The table shows that our SA solver obtains very good results, despite being a rather simple method, but with an extensive and principled tuning phase.

We have highlighted in bold the lowest average and best costs. The results labeled with z (Asin Achá and Nieuwenhuis [2]) have not been compared with the others as they are obtained by allotting a much longer runtime (from 10'000 to 100'000 seconds) than the one allowed in the ITC competition rules.

4 Conclusions and Future Work

In this paper we have presented an ongoing study on the effectiveness of a simple but properly tuned SA variant on a popular scheduling problem, namely the Curriculum-Based Course Timetabling. Our preliminary results are encouraging, comparing favorably with the current bests and suggesting that parameter tuning should be carried out in a more statistically principled way, in light of the major role it plays in solvers' performance. This is, of course, not only limited to SA as all meta-heuristics are, to various extents, sensitive to parameter values. For this reason we plan to carry out similar investigations starting with approaches similar to SA, such as Great Deluge, and moving on to different meta-heuristics, such as Tabu Search.

References

1. Salwani Abdullah, Hamza Turabieh, Barry McCollum, and Paul McMullan. A hybrid metaheuristic approach to the university course timetabling problem. *Journal of Heuristics*, 18(1):1–23, 2012.

2. Roberto Asín Achá and Robert Nieuwenhuis. Curriculum-based course timetabling with SAT and MaxSAT. *Annals of Operations Research*, February 2012.
3. Ruggero Bellio, Luca Di Gaspero, and Andrea Schaerf. Design and statistical analysis of a hybrid local search algorithm for course timetabling. *Journal of Scheduling*, 15(1):49–61, 2012.
4. Mauro Birattari, Z. Yuan, P. Balaprakash, and Thomas Stützle. *F-Race and iterated F-race: An overview*. Springer, Berlin, 2010.
5. Sara Ceschia, Luca Di Gaspero, and Andrea Schaerf. Design, engineering, and experimental analysis of a simulated annealing approach to the post-enrolment course timetabling problem. *Computers & Operations Research*, 39:1615–1624, 2012.
6. Luca Di Gaspero, Barry McCollum, and Andrea Schaerf. The second international timetabling competition (ITC-2007): Curriculum-based course timetabling (track 3). Technical report, Queen’s University, Belfast (UK), August 2007.
7. D. S. Johnson, C. R. Aragon, L. A. McGeoch, and C. Schevon. Optimization by simulated annealing: an experimental evaluation; part I, graph partitioning. *Operations Research*, 37(6):865–892, 1989.
8. Rhyd Lewis, Ben Paechter, and Barry McCollum. Post enrolment based course timetabling: A description of the problem model used for track two of the second international timetabling competition. Technical report, Cardiff University, Wales, UK, 2007.
9. Leo Lopes and Kate Smith-Miles. Pitfalls in instance generation for Udine timetabling. In *Proceedings of the LION₄ Conference*, 2010.
10. Zhipeng Lü and Jin-Kao Hao. Adaptive tabu search for course timetabling. *European Journal of Operational Research*, 200(1):235 – 244, 2009.
11. Barry McCollum, Andrea Schaerf, Ben Paechter, Paul McMullan, Rhyd Lewis, Andrew J. Parkes, Luca Di Gaspero, Rong Qu, and Edmund K. Burke. Setting the research agenda in automated timetabling: The second international timetabling competition. *INFORMS Journal on Computing*, 22(1):120–130, 2010.
12. Tomáš Müller. ITC2007 solver description: a hybrid approach. *Annals of Operations Research*, 172(1):429–446, 2009.
13. Andrea Schaerf. A survey of automated timetabling. *Artificial Intelligence Review*, 13(2):87–127, 1999.
14. Tommaso Urli. json2run: a tool for experiment design & analysis. *arXiv preprint arXiv:1305.1112*, 2013.
15. Tien-Tsin Wong, Wai-Shing Luk, and Pheng-Ann Heng. Sampling with Hammersley and Halton Points. *Journal of Graphics Tools*, 2(2):9–24, January 1997.