New Insights into Local Search for Curriculum-Based Course Timetabling

Sara Ceschia, Francesca Da Ros, Luca Di Gaspero, and Andrea Schaerf✉

Intelligent Optimization Laboratory, University of Udine, Udine, Italy
andrea.schaerf@uniud.it

Abstract. Curriculum-Based Course Timetabling (CB-CTT) is a long-standing benchmark problem in educational timetabling, yet obtaining robust performance across diverse datasets under strict time limits remains challenging. In this work, we revisit the role of local search for CB-CTT and present a Simulated Annealing (SA) method that uses the LectureKick neighborhood, a generalized version of the classic LectureMove neighborhood. The method combines moves preserving feasibility, sampling biases for soft constraints, efficient delta evaluation, and a cut-off mechanism for the cooling schedule. The algorithm is tuned on a large artificial training set and evaluated on the International Timetabling Competition 2007 instances and additional benchmark datasets. Experimental evidence shows that, under the original ITC 2007 time limits, the proposed approach clearly outperforms previously published local search methods and remains competitive with the state-of-the-art methods reported in the literature. These results indicate that a careful redesign of neighborhoods and move-selection biases can still significantly improve local search performance on a mature benchmark such as CB-CTT.

Keywords: Course Timetabling · Simulated Annealing · Neighborhood Relations

1 Introduction

Curriculum-Based Course Timetabling (CB-CTT) is one of the most widely studied university timetabling formulations, largely due to its adoption in Track 3 of the Second International Timetabling Competition (ITC 2007) and the continued use of its benchmark datasets [6]. In CB-CTT, lectures must be assigned to periods and rooms over a weekly horizon while satisfying hard constraints on room occupancy, conflicts, and teacher availability, and minimizing violations of soft constraints related to room capacity, room stability, minimum working days, and isolated lectures [4,6].

Over the years, CB-CTT has been addressed using local search, mathematical programming, and hybrid approaches [2]. Under the strict time limits of ITC 2007, local search has historically played a central role, from the competition-winning hybrid solver by Müller [7] to later tabu search, simulated annealing, and hybrid large neighborhood approaches. Nonetheless, the large number of neighborhoods proposed in the literature has made it difficult to identify which components are truly responsible for performance gains.

The aim of this work is to revisit local search for CB-CTT through a systematic consolidation of neighborhood ideas previously proposed in the literature. Our goal is to show that a careful redesign of move generation, neighborhood biases, and parameter tuning can substantially improve the effectiveness of a classic Simulated Annealing (SA) framework. SA was chosen based on both the strong performance reported in the literature and the extensive expertise developed by our research group through the solution of other timetabling problems. The resulting method is evaluated under the original ITC 2007 time limits and compared against strong local search baselines and state-of-the-art reference methods. All instances and best solutions are available on our website OptHub for inspection and future comparison.

2 Solution Method

The initial solution is constructed in three phases. First, a randomized construction procedure generates a complete assignment that satisfies all hard constraints except, possibly, curriculum and teacher conflicts. A short hill climbing phase based on lecture relocations and lecture swaps is then used to reach feasibility. A second hill climbing stage, still using the same neighborhoods, quickly improves the feasible solution before the main search begins.

The core of the method is an SA procedure operating on a neighborhood called LectureKick. The move extends Müller's LectureMove [7] in the following way: a lecture is moved to a target slot, and if that slot is occupied, the conflicting lecture is "kicked" to another admissible slot instead of being forced back to the original one (as in LectureMove). This generalization increases flexibility while preserving feasibility. Move generation is incremental and supported by auxiliary data structures storing admissible periods, conflicts per course-period pair, and room-related information, so that infeasible partial moves can be discarded early and delta costs can be evaluated efficiently.

To bias the search toward high-quality solutions, the neighborhood is not sampled uniformly. Four mutually exclusive biases target the soft constraints: one favors moves that improve minimum working days, one focuses on critical courses that are likely to affect isolated lectures, and two select lectures involved in room capacity and room stability violations. Two additional independent biases are used to avoid introducing new room capacity violations and to favor swap-like moves when beneficial. A preliminary screening phase showed that not all biases are equally relevant; in the final configuration, the most effective components have been retained and tuned. The final parameter configuration was obtained with `irace` on 996 artificial instances selected from the generator of De Coster et al. [3].

The SA schedule follows the classical Metropolis acceptance rule, combined with a geometric cooling schedule and a cut-off mechanism that accelerates the early stages of the search without reducing the total computational budget.

3 Experimental Results

First, we considered the `comp` dataset introduced at ITC 2007, following the competition protocol (single core and time limit). On the 19 non-trivial instances of this dataset, the proposed SA improves the best previously reported average result among local search methods on 18 instances and ties on the remaining one. The method achieves an average rank of 1.11, clearly outperforming the best local search solvers from the literature.When compared with the state-of-the-art methods [1,5] under the competition time limit, our SA obtains the best overall average rank (1.79). In this broader comparison, however, the differences are not statistically significant, which suggests that the three methods are very close on the standard benchmark.

The experimental analysis was extended to additional benchmark datasets proposed in the literature, under the same time limit as in the competition setting. On these datasets, our method generally improves the average results previously reported in the literature, with only marginal exceptions on a very small number of instances (Figure 1). In terms of best-known solutions, the approach improves the reported upper bounds for `test3` and `test4`, and matches the best-known values on most of the retained Udine instances.

More generally, the results suggest that mature timetabling benchmarks can still benefit from a careful re-examination of classical metaheuristics when supported by modern algorithm configuration tools and efficient implementations.

References

1. Abdullah, S., Turabieh, H.: On the use of multi neighbourhood structures within a Tabu-based memetic approach to university timetabling problems. Information Sciences **191**, 146–168 (2012)
2. Ceschia, S., Di Gaspero, L., Schaerf, A.: Educational timetabling: Problems, benchmarks, and state-of-the-art results. European Journal of Operational Research **308**(1), 1–18 (2023)
3. De Coster, A., Musliu, N., Schaerf, A., Schoisswohl, J., Smith-Miles, K.: Algorithm selection and instance space analysis for curriculum-based course timetabling. Journal of Scheduling **25**, 35–58 (2022)
4. Di Gaspero, L., Schaerf, A.: Multi-neighbourhood local search with application to course timetabling. In: Proceedings of the 4th International Conference on the Practice and Theory of Automated Timetabling (PATAT-2002), selected papers. Lecture Notes in Computer Science, vol. 2740, pp. 262–275. Springer (2003)
5. Kiefer, A., Hartl, R.F., Schnell, A.: Adaptive large neighborhood search for the curriculum-based course timetabling problem. Annals of Operations Research **252**(2), 255–282 (2017)
6. McCollum, B., Schaerf, A., Paechter, B., McMullan, P., Lewis, R., Parkes, A.J., Di Gaspero, L., Qu, R., Burke, E.K.: Setting the research agenda in automated timetabling: The second international timetabling competition. INFORMS Journal on Computing **22**(1), 120–130 (2010), https://www.eeecs.qub.ac.uk/itc2007/
7. Müller, T.: ITC2007 solver description: a hybrid approach. Annals of Operations Research **172**(1), 429–446 (2009)

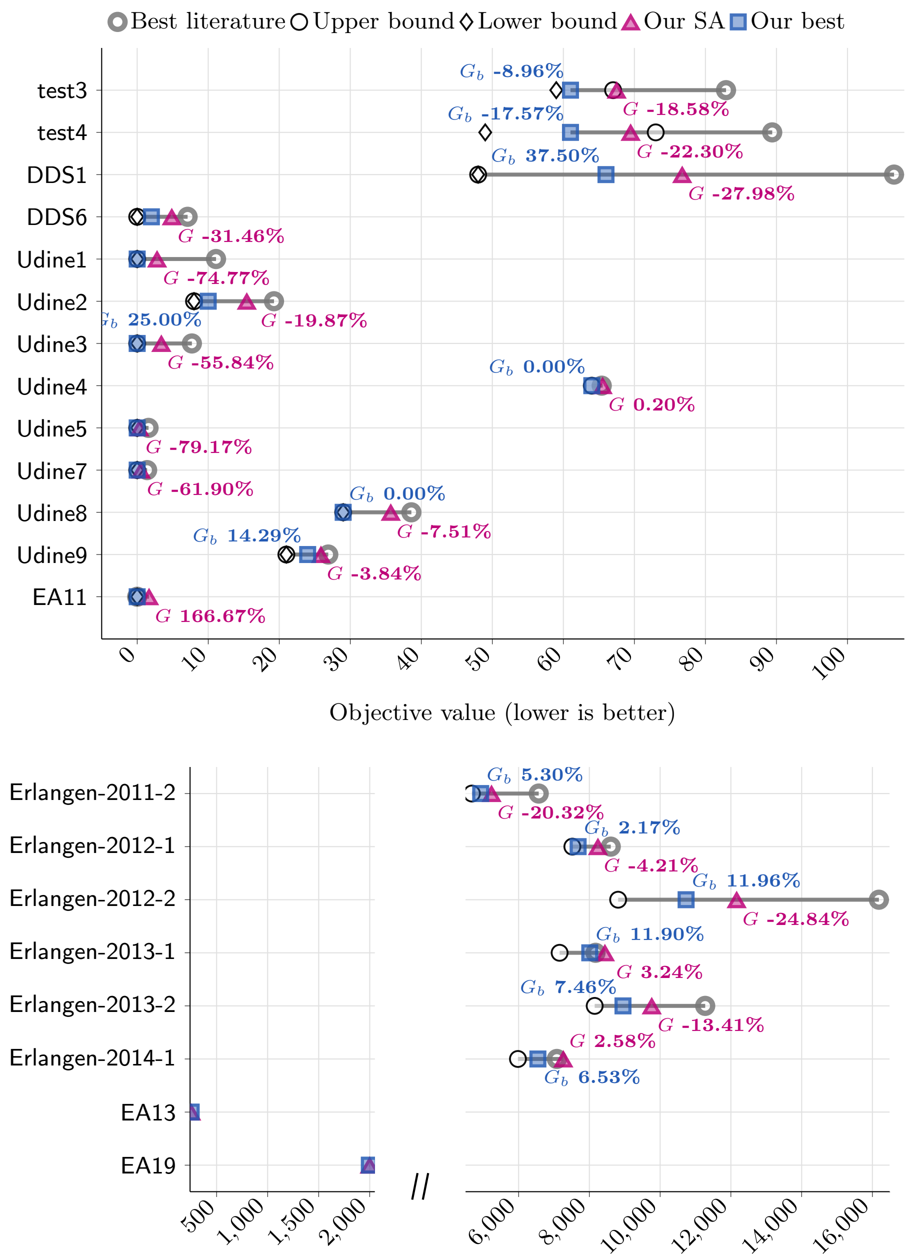

Fig. 1: Comparison plots for the test, DDS, Udine, EasyAcademy (EA), and Erlangen datasets. Symbols denote literature/known best values, lower bounds, and the average and best values of our SA; G and G_b report the corresponding gaps when available.