Noname manuscript No.
(will be inserted by the editor)

International Timetabling Competition 2021: Sports Timetabling

Dries Goossens · Jeroen Beliën ·
Morteza Davari · David Van Bulck

Received: date / Accepted: date

Abstract This extended abstract discusses the organization of the most recent International Timetabling Competition (ITC 2021). This competition focused on sports timetabling, where the problem is to decide on a suitable date for each of the matches to be played in the tournament. This is a complex and challenging problem, even for tournaments with few contestants. As a consequence, state-of-the-art typically focuses on a particular season of a sports competition for which a tailored algorithm is developed which is then compared to a manual solution. The aim of this competition was therefore to promote and provide insights in the development of more generally applicable sports timetabling solvers. To this purpose, participants required to solve a rich and diverse set sports timetabling instances involving various constraints that are common in real life. We introduce the contours of the problem instances, as well as the data format. We give an overview of the competition rules and timeline, and conclude with an overview of the finalists.

Keywords Sport timetabling · Competition · File format · Rules

D. Goossens and D. Van Bulck
Ghent University, Faculty of Economics and Business Administration
Department of Business Informatics and Operations Management
E-mail: dries.goossens@ugent.be; david.vanbulck@ugent.be

J. Beliën
KU Leuven, Campus Brussels
Research Centre for Operations Management
E-mail: jeroen.belien@kuleuven.be

M. Davari
SKEMA Business School
E-mail: morteza.davari@skema.edu

1 Introduction

Creating timetables for sports competitions has been a topic of research since the 1970s (e.g., [1]). Ever since, academic papers about sports timetabling have increased considerably in numbers and sports timetabling has become a specialized field [10], which has been discusssed at most of the PATAT meetings. Sports timetabling is often complex and challenging, even for a small number of teams. While generating a schedule where each team plays against each other team once and no team is involved in simultaneous matches is easy [6], some rather basic sports timetabling problems are already NP-hard. For instance, Briskorn et al. [2] show that there is no constant-factor approximation (unless P=NP) for a sports timetabling problem where certain matches cannot be played on certain rounds. Furthermore, real-life sports timetabling problems are characterized by a wide diversity of constraints, and conflicting interests of many stakeholders. At the same time, in professional sports, the timetable has an impact on commercial interests and revenues of the clubs, broadcasters, sponsors, as well as an impact on society through resulting traffic and policing costs.

Since 2002, there have been frequent timetabling competitions, which have been benificial for the research community. The first international timetabling competition (ITC) was organized in 2002 and focused on (a simplified version of) the university course timetabling problem [11]. The next ITC competition (2007) aimed to further develop interest in the general area of educational timetabling and involved three problems: curriculum-based timetabling, examination timetabling, and post-enrollment timetabling [12,13]. With high-school timetabling, the ITC placed yet another educational timetabling problem in the spotlights in 2011 [16, 17]. The fourth ITC is again devoted to university course timetabling: it introduces the combination of student sectioning together with time and room assignment of events in courses [14, 15]. In between, PATAT has supported two international nurse rostering competitions in 2010 [9] and 2014 [4], as well as a cross-domain heuristic search challenge (CHeSC 2011), where the challenge was to design a high-level search strategy that controls a set of problem-specific low-level heuristics, which would be applicable to different problem domains [3].

Many of the sports timetabling contributions in the literature read as a case study, describing a single instance for which a tailored algorithm is developed (which is then compared to a manual solution). Moreover, the state-of-the-art does not offer a general solution method, or even much insight in which type of algorithm would work well for which type of problem [18]. One notable exception is the travelling tournament problem [7], which minimizes the total team travel in a timetable. For this problem, substantial algorithmic progress has been reported after Easton et al. [7] made a set of artificial benchmark instances publicly available, and for which best results can be submitted to a website maintained by professor Michael Trick (see `http://mat.tepper.cmu.edu/TOURN/`). Hence, an international timetabling competition could make a valuable contribution to the field of sports timetabling, and given the efforts done by Van Bulck et al. [18, 19] with respect to the development of an XML-based file format for problem instances and solutions, we believe the time was right for a timetabling competition on sports.

2 Problem description and file format

The input of a sports timetabling problem consists of a set of rounds R, a set of teams T, and a set of games G. The set of games consists of ordered pairs (i, j) in which $i \in T$ is the home team providing the venue where the game is played, and $j \in T$ is the away team. Although many tournament formats are conceivable, in this competition we focus on so-called double round-robin tournaments (2RR), which are very common in practice [8]. In a double round-robin tournament, each team plays against each other team twice, typically once at home and once away. Although there is a line of research that focuses on the simultaneous scheduling of multiple leagues with dependencies [5], we focus on a single league. No team can play more than one game per round. In practice, rounds typically correspond to weekends, which may consist of several time slots (e.g., Saturday evening, or Sunday afternoon), each with their capacity. We focus on so-called time-constrained tournaments, i.e., tournaments that use the minimum number of rounds required to play all matches. In a 2RR with n teams, n even, the minimum number of rounds to play all games equals $2(n-1)$; if n is odd, the minimum number of rounds is $2n$.

A timetable maps each game in G to a round in R such that no team plays more than one game per round. An example of a timetable for a double round-robin tournament with 6 teams is given in Table 1.

Table 1 A compact double round-robin timetable for a league with 6 teams.

r_1	r_2	r_3	r_4	r_5	r_6	r_7	r_8	r_9	r_{10}
(1,2)	(2,5)	(2,4)	(2,3)	(6,2)	(2,1)	(5,2)	(4,2)	(3,2)	(2,6)
(3,4)	(4,1)	(1,6)	(5,1)	(4,5)	(4,3)	(1,4)	(6,1)	(1,5)	(5,4)
(5,6)	(6,3)	(5,3)	(6,4)	(1,3)	(6,5)	(3,6)	(3,5)	(4,6)	(3,1)

Sports timetables need to satisfy a usually large set of constraints, which is partitioned into hard constraints and soft constraints. Hard constraints represent fundamental properties of the timetable that can never be violated. Soft constraints, in contrast, rather represent preferences that should be satisfied whenever possible. The validation of each soft constraint c results in a vector D_c of n_c integral numbers, called the deviation vector $D_c = [d_1\ d_2\ \dots\ d_{n_c}]$. If a constraint is satisfied, all elements of its deviation vector are equal to zero. Contrarily, the deviation vector of a violated constraint contains one or more strictly positive elements. For hard constraints, any deviation renders the schedule infeasible. Each soft constraint features a cost function f_c and weight w_c. A violated soft constraint triggers a penalty $p_c = w_c f_c(D_c)$, equal to a weighted mapping of its deviation vector by its cost function. The objective we use for the competition instances sums over all violated soft constraint penalties.

The instances feature a variety of constraints from the classicification developed by Van Bulck et al. [18]. The authors distinguish capacity constraints, game constraints, break constraints, fairness/attractiveness constraints, and separation constraints. Capacity constraints force a team to play home or away and regulate the total number of games played by a team or group of teams. Game constraints enforce or forbid specific assignments of a game to rounds. Constraints to increase the fairness or attractiveness involve balancedness of, e.g., home advantage, travel distances, etc. Break constraints regulate the frequency and timing of breaks in a competition; we say that a team has a break if it has two consecutive home games, or two consecutive away games. Separation constraints regulate the number of rounds between consecutive games involving the same teams.

The problem instances are expressed using the standardized XML data format developed by Van Bulck et al. [18]. The main intention of this data format is to promote problem instance data sharing and reuse among different users and software applications, and this is exactly what the timetabling competition envisions. The XML data format is open, human readable (i.e., no binary format), software and platform independent, and flexible enough to store the problem instances.

Most of the sports timetabling constraints are easy to express in words but are hard to enforce within specific algorithms such as mathematical programming or metaheuristics. We believe this format minimizes the specification burden and maximizes the accessibility. The main advantage of XML over plain text-only file formats lies in the structured way of data storage. Indeed, an important motivation behind XML is to separate data representation from data content.

A detailed description of the the file format is available on the competition website (`http://itc2021.ugent.be`). The website also provides access to a validator, allowing participants to verify whether their solution satisfies all hard constraints and to determine its score on the objective function.

3 Competition rules

We are much indebted to the various organizers of the previous international timetabling competitions. Their experience has crystallized into the rules that were used for the ITC 2019 competition [15], and to which we will largely adhere for this competition. In particular, we enforce no bound on the computation time. In fact, the objective function value of the solution is the only criterion that matters. While computation time is obviously not unimportant, a fair comparison in terms of computation time is quite challenging, and it could easily lead to disputes that we as organizers prefer to avoid. Moreover, from a practical point of view, sports timetabling problems are often not so

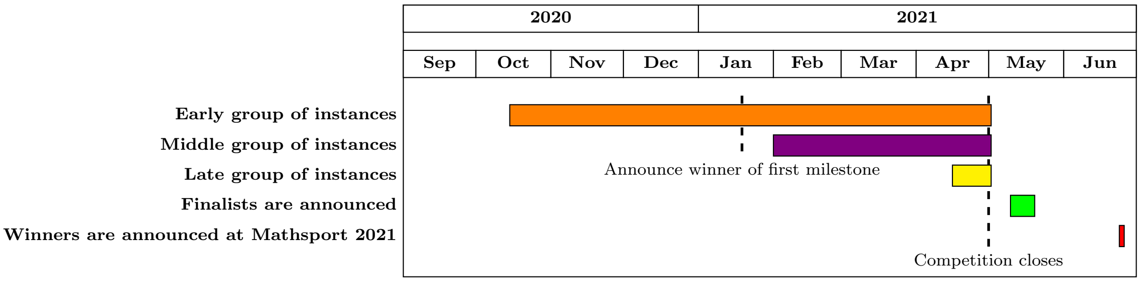

Fig. 1 Timeline for the International Timetabling Competition 2021

time-critical, as there are often several days or even weeks available to obtain a good solution.

We also allow to make use of any commercial solver. In this way, we would like to lower the threshold to participate, and reach out to the largest possible research community. Obviously, to keep it interesting, the instances for the competition will not solve to optimality with any straightforward formulations on e.g., state-of-the-art IP solvers.

Although we allow parameter tuning, we require that the same version of the algorithm is used for all instances. In other words, the algorithm should not "know" which instance it is solving. While the algorithm may analyze the problem instance and set parameters accordingly, it should apply this same procedure for all instances. The programmer should not set different parameters for different instances, however, if the program is doing this automatically, then this is acceptable. We will be asking for the source code of the finalists, in order to check whether the participants comply with this rule.

We believe these rules are efficient (in the sense that they do not require the organizer to run the participant's code) and fair/simple (in the sense that the only thing that matters is the obtained objective value; it avoids all discussion about measuring, e.g., computation time, the impact of random seeds, etc.).

4 Competition timeline and results

An overview of the competition timeline is given in Figure 1. In total, we released three groups of 15 artificially generated problem instances each: early, middle, and late instances. While all instances contributed to the final ranking of participants, instances that were released later in the competition had a higher weight. For instance, the overall best found solutions was respectively awarded 10, 15, and 25 instances for an early, middle, and late problem instance. The early group of instances were already available from our website at the time the competition was officially announced (mid October 2020), while the middle group of instances were only released in February 2021. The late instances followed half April 2021, which gave the participants two weeks to come up with solutions.

Team name	Research institute	Participants
TU/e	Eindhoven University of Technology	F. Spieksma, H. Christopher, R. Lambers, and J. van Doornmalen
Saturn	HSE University	S. Daniil and R. Ivan
MODAL	Zuse Institute Berlin	T. Koch, T. Berthold, and Y. Shinano
GOAL	Federal University of Ouro Preto	G. H. G. Fonseca and T. A. M. Toffolo
UoS	University of Southampton	T. Martínez-Sykora, C. Potts, C. Lamas-Fernández
Udine	University of Udine	R. M. Rosati, M. Petris, L. Di Gaspero, and A. Schaerf

Table 2 Overview of the 6 finalists (randomly ordered)

Around half January 2021, we organized a first milestone event where participants had the possibility to submit their best solutions found at that time. Although optional, participation in the first milestone was strongly encouraged as it provided participants with the feedback on where their algorithms ranked among their peers as well as a chance to win a small prize (free registration for Mathsport 2022). The first milestone was won by team UoS, followed by team Udine and TU/e (see Table 2).

At the time of the final submission deadline, 13 research teams from over 10 different countries successfully submitted solutions. As a comparison, the cross-domain heuristic search challenge attracted 17 teams, the two international nurse rostering competitions each attracted 15 teams, and the third and fourth international timetabling competition each attracted 5 teams that submitted one or more solutions by the final submission deadline.

Out of all 13 participating teams, the 6 finalists given in Table 2 were selected. The prize fund is 1,750 EUR to be split between the first, second, and third place competitors. Moreover, a discount on registration for the upcoming PATAT conference is awarded to the top three overall. We thank our sponsors OR in Sports and PATAT for their generous contribution to the rewards we could distribute over the winners. At the Mathsport International 2021 conference, team UoS (University of Southampton) was announced as the winner of the ITC 2021.

References

1. Ball, B. C., Webster, D. B. (1977). Optimal scheduling for even-numbered team athletic conferences. AIIE Transactions 9(2):161–169.
2. Briskorn, D., Drexl, A., Spieksma F.C.R. (2010). Round robin tournament and three index assignments. 4OR - A Quarterly Journal of Operations Research, 8:365–374.
3. Burke, E.K., Gendreau, M., Hyde, M., Kendall, G., McCollum, B., Ochoa, G., Parkes, A.J., Petrovic, S. (2011). The Cross-Domain Heuristic Search Challenge – An International Research Competition. In: C.A.C. Coello (editor) Learning and Intelligent Optimization, pp. 631–634. Springer.
4. Ceschia, S., Dang, N., De Causmaecker, P., Haspeslagh, S., Schaerf, A. (2019). The second international nurse rostering competition. Annals of Operations Research, 274:171-–186.
5. Davari, M., Goossens, D., Beliën, Lambers, R., Spieksma, F.C.R. (2020). The multi-league sports scheduling problem, or how to schedule thousands of matches. Operations Research Letters, forthcoming.

6. de Werra, D. (1980). Geography, games and graphs. Discrete Applied Mathematics 2:327–337.
7. Easton, K., Nemhauser, G., Trick, M., (2001). The traveling tournament problem description and benchmarks. In: Walsh, T. (editor), Principles and Practice of Constraint Programming — CP 2001. Springer, Berlin, Heidelberg, pp. 580-–584.
8. Goossens, D., Spieksma, F.C.R. (2012). Soccer schedules in Europe: an overview. Journal of Scheduling 15(5): 641–651.
9. Haspeslagh, S., De Causmaecker, P., Schaerf, A., Stølevik, M. (2010). The first international nurse rostering competition 2010. Annals of Operations Research 218(1):221–236.
10. Kendall, G., Knust, S., Ribeiro, C.C., Urrutia, S. (2010). Scheduling in sports: An annotated bibliography. Computers & Operations Research 37(1):1–19.
11. Paechter, B., Gambardella, L.M., Rossi-Doria, O. (2002). International timetabling competition (ITC2002). `http://sferics.idsia.ch/Files/ttcomp2002`
12. McCollum, B., Schaerf, A., Paechter, B., McMullan, P., Lewis, R., Parkes, A.J., Di Gaspero, L., Qu, R., Burke, E.K. (2007). The second international timetabling competition (ITC2007). `http://www.cs.qub.ac.uk/itc2007/index.htm`
13. McCollum, B., Schaerf, A., Paechter, B., McMullan, P., Lewis, R., Parkes, A.J., Di Gaspero, L., Qu, R., Burke, E.K. (2010). Setting the research agenda in automated timetabling: The second international timetabling competition. INFORMS Journal on Computing 22(1):120–130.
14. Müller, T., Rudová, H. Müllerová, Z. (2018). University course timetabling and International Timetabling Competition 2019. In: E. K. Burke, L. Di Gaspero, B. McCollum, N. Musliu, and E. Özcan (editors), Proceedings of the 12th International Conference on the Practice and Theory of Automated Timetabling, Vienna, pp. 5–31. PATAT.
15. Müller, T., Rudová, H. Müllerová, Z. (2019). International Timetabling Competition (ITC2019). `http://www.itc2019.org/home`
16. Post, G., Di Gaspero L., Kingston, J.H., McCollum, B., Schaerf, A. (2011). International timetabling competition (ITC2011). `http://www.utwente.nl/en/eemcs/dmmp/hstt/itc2011/`
17. Post, G., Di Gaspero L., Kingston, J.H., McCollum, B., Schaerf, A. (2016). The third international timetabling competition. Annals of Operations Research, 239:69—75.
18. Van Bulck D., Goossens D., Schönberger J., Guajardo M. (2020a), RobinX: A three-field classification and unified data format for round-robin sports timetabling. European Journal of Operational Research, 280:568–580.
19. Van Bulck, D., Goossens, D., Schönberger, J., Guajardo, M. (2020b). An instance data repository for the round-robin sports timetabling problem. Management and Labour Studies, 45:184–200.