The Fifth International Timetabling Competition (ITC 2021): Sports Timetabling

D. Van Bulck*, D. Goossens**, J. Beliën***, and M. Davari****

*Ghent University, david.vanbulck@ugent.be

**Ghent University, dries.goossens@ugent.be

***KU Leuven, jeroen.belien@kuleuven.be

****SKEMA Business School, morteza.davari@skema.edu

Abstract

This paper 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 of 45 sports timetabling instances involving up to 9 different constraints that are common in real life. We discuss the format of these instances, how the instances were released during the competition, and conclude with an overview of the finalists.

1 Introduction

Creating timetables for sports competitions has been a topic of research since the 1970s (e.g., Ball and Webster (1977)). Ever since, academic papers about sports timetabling have increased considerably in numbers and sports timetabling has become a specialized field (Kendall et al., 2010). Sports timetabling is often complex and challenging, even for a small number of teams. While generating a timetable where each team plays against each other team once and no team is involved in simultaneous matches is easy (e.g. de Werra (1980)), some rather basic sports timetabling problems are already NP-hard. For instance, Briskorn et al. (2010) show that there is no constant-factor approximation (unless $\mathsf{P} = \mathsf{NP}$) for a sports timetabling problem where certain matches cannot be played on a set of predefined 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 beneficial for the research community. The first international timetabling competition was organized in 2002 and focused on (a simplified version of) the university course timetabling problem (see Paechter et al. (2003)). 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-enrolment timetabling (see McCollum et al. (2007, 2010)). With high-school timetabling, the ITC placed yet another educational timetabling problem in the spotlights in 2011 (see Post et al. (2011, 2016)). 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 (see Müller et al. (2018, 2019)). In between, there have been two international nurse rostering competitions in 2010 (see Haspeslagh et al. (2014)) and 2014 (see Ceschia et al. (2019)), 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 (see Burke et al. (2011)).

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 typically 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 (see Van Bulck et al. (2020)). One notable exception is the travelling tournament problem (Easton et al., 2001), an artificial and somewhat simplified sports timetabling problem where the objective is to minimize the total team travel in a timetable. For this problem, substantial algorithmic progress has been reported after Easton et al. (2001) 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/`). A long standing obstacle to benchmark algorithms for sports timetabling problems that are real-world-like was the absence of a file format to express the wide amount and variety of constraints that are typically present in real-life problem instances. Given the recent efforts by Van Bulck et al. (2020) to overcome this obstacle, we believed the time was right to organize an international timetabling competition on sports.

The remainder of this paper is as follows. Section 2 provides a general description of the type of problems offered in the competition, and Section 3 outlines the competition rules. We conclude in Section 4 with a short discussion of the competition timeline, the prizes that could be earned, and the announcement of the finalists.

2 Problem description and file format

In essence, sports timetabling is deciding on a suitable round for each of the matches to be played in the tournament. In practice, rounds typically correspond to weekends, and consist of several time slots (e.g., Saturday evening, or Sunday afternoon), however, each team plays at most once per round. The competition focuses on the construction of round-robin timetables, meaning that each team plays against every other team a fixed number of times. Although many other tournament formats are conceivable (e.g. the knock-out tournament), round-robin tournaments are probably the most researched format (see Knust (2021)) and are very common in practice (see e.g. Goossens and Spieksma (2011)). Most sports competitions organize a double round-robin tournament (2RR) where teams meet twice but single, triple, and even quadruple round-robin tournaments also occur. Existing literature distinguishes two types of round-robin tournaments: time-constrained timetables and time-relaxed timetables. A timetable is time-constrained (also called compact) if it uses the minimal number of rounds needed, and is time-relaxed otherwise. In this competition, we only consider time-constrained double round-robin tournaments with an even number of teams. Under this setting, the total number of rounds is exactly equal to the total number of games per team, and hence each

Table 1: A time-constrained double round-robin timetable for a single league with 6 teams. Each game is represented by an ordered pair in which the first element is the home team, and the second element is the away team.

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)	(4,2)	(5,2)	(2,1)	(3,2)	(2,6)
(3,4)	(4,1)	(1,6)	(5,1)	(4,5)	(6,1)	(1,4)	(4,3)	(1,5)	(5,4)
(5,6)	(6,3)	(5,3)	(6,4)	(1,3)	(3,5)	(3,6)	(6,5)	(4,6)	(3,1)

team plays exactly one game per round. Although there is a line of research that focuses on the simultaneous scheduling of multiple leagues with dependencies (Davari et al., 2020), we focus on the construction of a 2RR for a single league. For an example of a time-constrained 2RR timetable, we refer to Table 1.

The constraints that appear in real-life problem instances are extremely diverse: apart from some basic constraints, each competition has its own requirements. In this competition, we assume that there are two types of constraints: hard constraints that represent fundamental properties of the timetable that can never be violated, and soft constraints that represent preferences that should be satisfied whenever possible. While many possible optimization objectives appear in the literature (e.g. the minimization of travel), this competition considers problem instances only where the objective is to minimize the penalties from violated soft constraints. This assumption makes the problem formulation more attractive for a wider timetabling community, while retaining the empirical complexity of the problems. In total, nine types of constraints were considered that can be categorized into the following five constraint classes as introduced by Van Bulck et al. (2020). 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. Finally, 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. (2020). 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 envisioned. 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 minimized the specification burden and maximized the accessibility. The main advantage of XML over plain text-only file formats lies in the structured way of data storage which separates data representation from data content.

3 Competition rules

Prior to the competition, all rules and a number of sample instances were made available at the competition website (`itc2021.ugent.be`). The website provides more details on the rules of the competition, the problem instances and their XML format, the awards for the winners, and intermediate results. The website

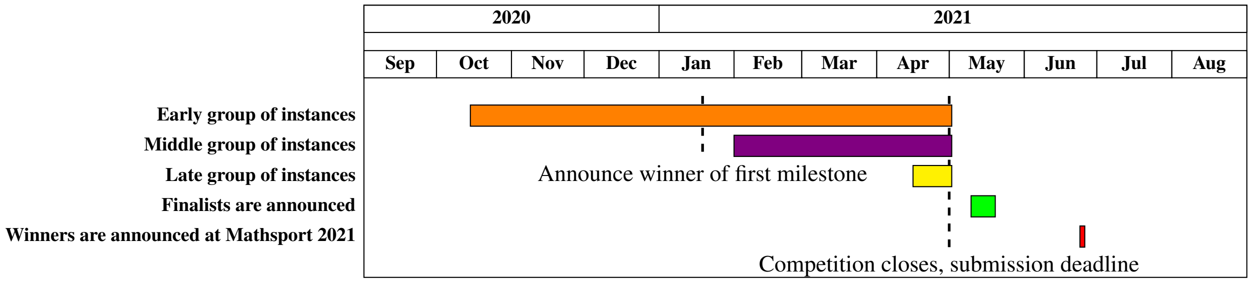

Figure 1: Timeline for the International Timetabling Competition 2021

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.

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 (Müller et al., 2019), and to which we largely adhered for this competition. In particular, we enforced no bound on the computation time. In fact, the objective function value of the best submitted solution was the only criterion that mattered. 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 time-critical, as there are often several days or even weeks available to obtain a good solution.

We also allowed to make use of any commercial solver. In this way, we tried 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 were designed such that a straightforward implementation on e.g., state-of-the-art integer or constraint programming solvers, could not solve the problem instances to optimality. In fact, for most problem instances, a straightforward integer programming formulation could not even generate a feasible solutions within a reasonable amount of time.

Although we allowed parameter tuning, we required that the same version of the algorithm was 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 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,

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 (ordered randomly)

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.

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. Note that the order of this list was generated at random and hence is unlikely to represent the final ordering. The prize fund is 1,750EUR 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. The final ordering of the finalists together with an overview of the best found solutions will be announced at the Mathsport International 2021 conference.

Given the large number of teams that participated in the competition and the fact that feasible solutions were found for all problem instances, we conclude that the ITC 2021 competition was a huge success.

References

B. C. Ball and D. B. Webster. Optimal scheduling for even-numbered team athletic conferences. *Aiie T.*, 9: 161–169, 1977.

D. Briskorn, A. Drexl, and F. C. R. Spieksma. Round robin tournaments and three index assignments. *4OR-Q J Oper Res*, 8:365–374, 2010.

Edmund K. Burke, Michel Gendreau, Matthew Hyde, Graham Kendall, Barry McCollum, Gabriela Ochoa, Andrew J. Parkes, and Sanja Petrovic. The cross-domain heuristic search challenge – an international research competition. In Carlos A. C. Coello, editor, *Learning and Intelligent Optimization*, pages 631–634, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.

S. Ceschia, N. Dang, P. De Causmaecker, S. Haspeslagh, and A. Schaerf. The second international nurse rostering competition. *Ann. Oper. Res.*, 274:171–186, 2019.

M. Davari, D. Goossens, J. Beliën, R. Lambers, and F. Spieksma. The multi-league sports scheduling problem, or how to schedule thousands of matches. *Oper. Res. Lett.*, 48:180 – 187, 2020.

D. de Werra. Geography, games and graphs. *Discrete Appl. Math.*, 2:327–337, 1980.

K. Easton, G. Nemhauser, and M. Trick. The traveling tournament problem description and benchmarks. In T. Walsh, editor, *Principles and Practice of Constraint Programming — CP 2001*, pages 580–584, Berlin, Heidelberg, 2001. Springer.

D. R. Goossens and F. C.R. Spieksma. Soccer schedules in Europe: an overview. *J. Sched.*, 15:641–651, 2011.

S. Haspeslagh, P. De Causmaecker, A. Schaerf, and M. Stølevik. The first international nurse rostering competition 2010. *Ann. Oper. Res.*, 218:221–236, 2014.

G. Kendall, S. Knust, C. C. Ribeiro, and S. Urrutia. Scheduling in sports: An annotated bibliography. *Comput. Oper. Res.*, 37:1–19, 2010.

S. Knust. Classification of literature on sports scheduling, 2021. URL `http://www2.inf.uos.de/knust/sportssched/sportlit_class/`.

B. McCollum, A. Schaerf, B. Paechter, P. McMullan, R. Lewis, A.J. Parkes, L. Di Gaspero, R. Qu, and E.K. Burke. The second international timetabling competition (ITC2007), 2007. URL `http://www.cs.qub.ac.uk/itc2007/index.htm`.

B. McCollum, A. Schaerf, B. Paechter, P. McMullan, R. Lewis, A. J. Parkes, L. Di Gaspero, R. Qu, and E. K. Burke. Setting the research agenda in automated timetabling: The second international timetabling competition. *Informs J. Comput.*, 22:120–130, 2010.

T. Müller, H. Rudová, and S. Müllerová. 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, 2018. PATAT.

T. Müller, H. Rudová, and S. Müllerová. International timetabling competition (ITC2019), 2019. URL `http://www.itc2019.org/home`.

B. Paechter, L.M. Gambardella, and O. Rossi-Doria. International timetabling competition (ITC2002), 2003. URL `http://sferics.idsia.ch/Files/ttcomp2002/`.

G. Post, L. Di Gaspero, J.H. Kingston, B. McCollum, and A. Schaerf. International timetabling competition (ITC2011), 2011. URL `http://www.utwente.nl/en/eemcs/dmmp/hstt/itc2011/`.

G. Post, L. Di Gaspero, J. H. Kingston, B. McCollum, and A. Schaerf. The third international timetabling competition. *Ann. Oper. Res.*, 239:69–75, 2016.

D. Van Bulck, D. Goossens, J. Schönberger, and M. Guajardo. RobinX: A three-field classification and unified data format for round-robin sports timetabling. *Eur. J. Oper. Res.*, 280:568 – 580, 2020.