Fairness-aware opponent assignment in sports timetabling with fixed home dates

Ali Hassanzadeh[1] [0000-0001-8475-8676], Dries Goossens[2] [0000-0003-0224-3412], and John Turner[3] [0000-0002-5835-1247]

[1] Centre for AI and Decision Sciences, Alliance Manchester Business School, University of Manchester, UK. ali.h@manchester.ac.uk
[2] Ghent University, Department of Business Administration and Operations Management, Belgium. dries.goossens@ugent.be
[3] Paul Merage School of Business, University of California, USA. john.turner@uci.edu

1 Introduction

Sports timetables are often judged by whether they satisfy logistical and commercial requirements, but they are also judged by whether they are fair. A common complaint from teams, fans, and the media is that some clubs face unusually difficult stretches of opponents within a short period, while others experience a smoother sequence of fixtures. This issue is especially important in leagues where short runs of difficult matches can strongly affect performance, momentum, and league position.

In this extended abstract, we study a fairness-aware timetabling problem in which the dates of home games are treated as fixed, but the assignment of opponents to those dates is to be optimized as the season progresses. This setting is motivated by the fact that, in many real leagues, home dates need to be fixed long beforehand because of venue availability, city operations, broadcast planning, and other logistical constraints. We therefore focus on a more practical question: if the home calendar is fixed, how much can fairness still be improved by dynamically choosing the opponent assignment?

2 Related Literature

Optimizing sports schedules with the requirement that a team should not play against extremely weak or extremely strong teams in consecutive rounds has been linked with fairness and studied from a theoretical perspective (see [2], [3], and [7]). In an incomplete round robin tournament (where teams face different sets of opponents), a team's strength of schedule is calculated based on the set of opponents each team faces, and their performance (see e.g. [5] for a discussion of metrics). However, even in (complete) round robin tournaments, where each team faces each opponent the same number of times, the order in which the opponents are faced has been shown to be significant (see e.g. [4]). Dynamic scheduling is rather uncommon in sports, but has been discussed for e-sports ([6]) and the NFL ([1]).

3 Our Contributions

Our framework measures fairness through the idea of tough stretches. Instead of looking only at average strength of schedule across a whole season, we examine rolling windows of time or games and ask whether some teams are exposed to clusters of especially strong opponents. In a deterministic formulation, each team's strength is assumed to be known at the time of scheduling (typically estimated from historical data, preseason projections, or current performance indicators). Given the team strength measures, as a first step, we optimize the schedule to reduce the worst stretch of strong opponents across the league. This gives a natural min-max fairness objective and provides a useful benchmark for understanding how much imbalance is present in an existing schedule.

In reality, team strengths are stochastic and manifest themselves as the results unfold. Rather than minimizing an *expected unfairness* measure, we propose that the league fixes the home dates for the entire season (ensuring logistical stability for venues, cities, and broadcasters), but announces the opponent assignments in smaller batches (e.g., 5–10 games at a time), instead of releasing the entire timetable at once. At each batch release, the most recent estimates of team strength are treated deterministically, the fairness model is re-solved for the remaining games, and the new batch of matchups is published. This effectively forms a stochastic program with recourse, allowing the schedule to adapt as uncertainty unfolds, but without disrupting existing logistical commitments.

We apply this idea to two different settings. The first is double round-robin scheduling, motivated by association football leagues such as the English Premier League, where each team plays every other team once at home and once away. The second is asymmetrical round-robin scheduling, motivated by professional basketball, where the number and pattern of meetings between pairs of teams are not fully symmetric. In both settings, the set of all matches to be played is fixed; what changes is the timing of which opponent is assigned to each fixed home slot. Studying these two tournament formats helps us understand in which type of competition dynamic fairness-aware scheduling has the greatest value.

Our early computational results suggest that meaningful fairness improvements are possible even under strong practical restrictions. At the same time, the work highlights an important trade-off between flexibility and computation. Allowing more freedom in reassignment generally improves the quality of the best timetable found, but it also increases model difficulty. This makes the problem interesting both from a practical timetabling point of view and from an optimization perspective.

At PATAT, we aim to present the modeling framework, discuss alternative fairness definitions, and share initial computational evidence from basketball and football leagues. We hope this opens a broader discussion on fairness-aware timetabling in sports, especially in settings where the full schedule cannot be rebuilt freely but targeted improvements may still be realistic and valuable.

4 Mathematical Model

We present a simple deterministic model for a double round-robin league. Let T be the set of teams. For each team $i \in T$, let H_i be the set of fixed home slots of team i, and let r_{is} denote the round in which home slot $s \in H_i$ is played. Let q_j be the strength value of team $j \in T$, obtained for example from historical data or current performance indicators. We define a binary decision variable x_{isj} which is one if team i hosts team j in home slot s, and zero otherwise. The basic assignment constraints are:

$$\sum_{j \in T \setminus \{i\}} x_{isj} = 1 \qquad \forall i \in T,\ \forall s \in H_i, \tag{1}$$

$$\sum_{s \in H_i} x_{isj} = 1 \qquad \forall i, j \in T,\ i \neq j. \tag{2}$$

Constraint (1) ensures that each fixed home slot receives exactly one opponent. Constraint (2) ensures that each team i hosts each other team j exactly once over the season. To measure fairness, we consider rolling windows of L rounds. For each team i and each starting round k, let W_{ik} be the set of home slots of team i whose rounds belong to $\{k, k+1, \dots, k+L-1\}$. The toughness of this window is

$$\sum_{s \in W_{ik}} \sum_{j \in T \setminus \{i\}} q_j x_{isj}.$$

We introduce a variable z representing the worst window toughness across the league, and minimize it:

$$\min \quad z \tag{3}$$

$$\text{s.t.} \quad \sum_{s \in W_{ik}} \sum_{j \in T \setminus \{i\}} q_j x_{isj} \le z \qquad \forall i \in T,\ \forall k, \tag{4}$$

$$x_{isj} \in \{0, 1\} \qquad \forall i \in T,\ \forall s \in H_i,\ \forall j \in T \setminus \{i\}. \tag{5}$$

The objective (3)–(5) minimizes the maximum strength accumulated by any team over any rolling window of length L. This provides a simple min-max notion of fairness based on avoiding clusters of strong opponents. This formulation is intentionally kept simple. It is meant to illustrate the main idea of fairness-aware opponent assignment when home dates are fixed, rather than to capture every practical feature of a real sports timetable. In particular, the model uses fixed team strength values, focuses only on opponent assignment, and measures fairness through a basic rolling-window criterion.

A few additional points are worth noting. First, in a fuller version of the problem, fairness should ideally be evaluated over the full sequence of matches for each team, including both home and away games, even though the decision variables act only on the assignment of opponents to fixed home slots. Second, in a deterministic setting, the strength parameters are treated as fixed inputs at the time of optimization; dynamic updates of team strength belong more naturally to a rolling-horizon or stochastic extension. Third, one may also want to limit how far a match can move away from its original date or opponent pattern, in order to preserve practical similarity to the published schedule.

We do not include these extra details here because of the short format of the conference extended abstract. They are not ignored, but deferred. A more complete treatment, including full home-and-away sequence effects, richer fairness definitions, and dynamic or stochastic extensions, will be discussed in the presentation and developed further in the full paper and the broader research project.

5 Preliminary Results

To illustrate the potential of the model, we applied it to three English Premier League (EPL) seasons (2021–2022, 2022–2023, and 2023–2024) using a 14-day rolling window, which can be viewed as roughly corresponding to two or three rounds of matches. Across all nine tested instances (three seasons and three strength definitions), the optimized schedule, as shown in Figure 1, improved the worst 14-day stretch relative to the original schedule. The reduction ranges from 27.9% to 42.4%, with an average improvement of 34.6%. This suggests that even under the strong restriction that home dates remain fixed, there is still meaningful room to reduce concentrated runs of strong opponents.

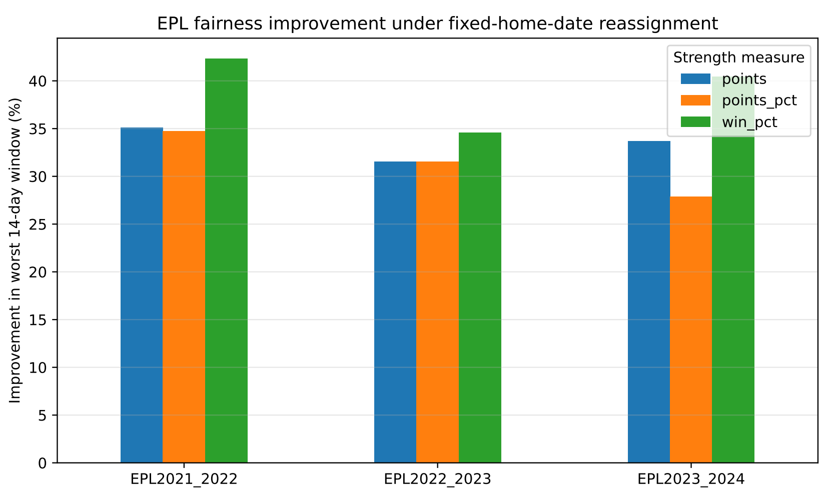

Fig. 1. Percentage improvement in the worst 14-day window for three EPL seasons under three team-strength definitions. In every tested case, the optimized assignment improves upon the original schedule.

References

1. Bouzarth, E.L., Cromer, A.W., Fravel, W.J., Grannan, B.C., Hutson, K.R.: Dynamically scheduling NFL games to reduce strength of schedule variability. Journal of Sports Analytics **6**(4), 281–293 (2020)
2. Briskorn, D.: Combinatorial properties of strength groups in round robin tournaments. European Journal of Operational Research **192**(3), 744–754 (2009)
3. Briskorn, D., Knust, S.: Constructing fair sports league schedules with regard to strength groups. Discrete Applied Mathematics **158**(2), 123–135 (2010)
4. Csató, L., Molontay, R., Pintér, J.: Tournament schedules and incentives in a double round-robin tournament with four teams. International Transactions in Operational Research **31**(3), 1486–1514 (2024)
5. DeHollander, A., Karwan, M.: Improving strength of schedule metrics in sports scheduling. Journal of Quantitative Analysis in Sports **22**(1), 53–78 (2026)
6. Dong, Z.L., Ribeiro, C.C., Xu, F., Zamora, A., Ma, Y., Jing, K.: Dynamic scheduling of e-sports tournaments. Transportation Research Part E: Logistics and Transportation Review **169**, 102988 (2023)
7. Zeng, L., Mizuno, S.: Constructing fair single round robin tournaments regarding strength groups with a minimum number of breaks. Operations Research Letters **41**, 506–510 (2013)