



PUC

ISSN 0103-9741

Monografias em Ciência da Computação
nº 16/07

Referee Assignment in Sports Leagues

Alexandre Rocha Duarte
Edward Hermann Haeusler
Celso Carneiro Ribeiro
Sebastián Urrutia

Departamento de Informática

PONTIFÍCIA UNIVERSIDADE CATÓLICA DO RIO DE JANEIRO
RUA MARQUÊS DE SÃO VICENTE, 225 - CEP 22451-900
RIO DE JANEIRO - BRASIL

Referee Assignment in Sports Leagues

Alexandre R. Duarte and Edward H. Haeusler and Celso C. Ribeiro¹ and
Sebastián Urrutia²

¹ Departamento de Ciência da Computação - Universidade Federal Fluminense

² Departamento de Ciência da Computação - Universidade Federal de Minas Gerais

{aduarte,hermann,celso}@inf.puc-rio.br, surrutia@dcc.ufmg.br

Abstract. Optimization in sports is a field of increasing interest. Combinatorial optimization techniques have been applied e.g. to game scheduling and playoff elimination. A common problem usually found in sports management is the assignment of referees to games already scheduled. There are a number of rules and objectives that should be taken into account when referees are assigned to games. We address a simplified version of a referee assignment problem common to many amateur leagues of sports such as soccer, and basketball. The problem is formulated by integer programming and its decision version is proved to be NP-complete. To tackle real-life large instances of the referee assignment problem, we propose a three-phase heuristic approach based on a constructive procedure, a repair heuristic to make solutions feasible, and a local search heuristic to improve feasible solutions. Numerical results on realistic instances are presented and discussed.

Keywords: Referee Assignment, Optimization in Sports, Metaheuristics, Integer Programming

Resumo. Nos últimos anos, a otimização em esportes tem despertado o interesse de diversos grupos de pesquisa em todo o mundo. Técnicas da otimização combinatória foram aplicadas, por exemplo, no escalonamento de partidas esportivas. Um problema frequentemente enfrentado na gestão esportiva é a alocação de árbitros a partidas previamente agendadas. Diversas restrições e objetivos devem ser considerados ao se determinar quais juízes devem apitar quais jogos. Este trabalho considera um versão simplificada do problema de alocação de juízes comum a várias ligas amadoras de esportes como futebol e basquete. Foi desenvolvido um modelo de programação inteira para o problema e sua versão de decisão foi demonstrada ser NP-completa. Para resolver instâncias de tamanho real, uma heurística em três fases é proposta e se baseia num algoritmo construtivo, uma heurística reparadora para atuar sobre soluções inviáveis e uma busca local para aprimorar soluções viáveis. Resultados numéricos utilizando instâncias realistas são apresentados e discutidos.

Palavras-chave: Alocação de Juízes, Otimização em Esportes, Metaheurísticas, Programação Inteira

In charge of publications:

Rosane Teles Lins Castilho
Assessoria de Biblioteca, Documentação e Informação
PUC-Rio Departamento de Informática
Rua Marquês de São Vicente, 225 - Gávea
22453-900 Rio de Janeiro RJ Brasil
Tel. +55 21 3114-1516 Fax: +55 21 3114-1530
E-mail: bib-di@inf.puc-rio.br
Web site: <http://bib-di.inf.puc-rio.br/techreports/>

1 Introduction

Optimization in sports is a field of increasing interest. Some applications have been reviewed by Ribeiro and Urrutia [9]. Combinatorial optimization techniques have been applied e.g. to the traveling tournament problem [4, 11], to playoff elimination [10], and to the scheduling of a college basketball conference [8]. Easton et al. [5] reviewed scheduling problems in sports.

A common problem usually found in amateur sports management is the assignment of referees to games already scheduled. The number of referees to be assigned to each game may vary depending on the sport or the league: soccer games usually require three referees, while basketball games require two. There are a number of rules and objectives that should be taken into account when referees are assigned to games. Games in higher divisions may require higher-skilled referees. Since referees may officiate several games during the day, travel feasibility and travel times between the facilities where the games take place have to be considered. Additionally, and especially in some amateur children leagues, some of the referees are players or their relatives. In this case, a natural constraint is that a referee cannot officiate a game in which he/she or a relative is scheduled to play.

Real-life versions of this problem appear in regional amateur leagues in the United States. Amateur leagues of several sports, such as baseball, basketball and soccer, have hundreds of games every weekend in different divisions. In a single league in California there might be up to 500 soccer games in a weekend, to be refereed by hundreds of certified referees. In the MOSA (Monmouth & Ocean Counties Soccer Association) league, New Jersey, boys and girls of ages 8 to 18 make up six divisions per age and gender group with six teams per division, totalizing 396 games every Sunday.

Referee assignment problems in other contexts have been addressed in [6, 7, 12]. Dinitz and Stinson [3] considered a problem involving referee assignment to tournament schedules, connecting room squares and balanced tournament designs. We address a simplified version of a referee assignment problem common to many amateur leagues of sports such as soccer, basketball, and baseball, among others. In the next section, we state the problem considered in this work. Section 3 presents an integer programming formulation to this referee assignment problem. The decision version of the problem is proved to be NP-complete in Section 4. The proposed solution strategy is described in Section 5. In Section 6, computational results illustrating the application of the proposed approach to solve real-size randomly generated instances are shown. Concluding remarks and further extensions of this work are reported in the last Section.

2 Problem statement

We consider the general problem, in which each game has a number of refereeing positions to be assigned to referees. The games are previously scheduled and the facilities and time slots for each game are known beforehand. In our approach, referees are assigned to empty refereeing positions, not to games. This allows not only to handle referee assignment problems in sports requiring different number of referees, but also in tournaments where different games of the same sport may need different numbers of referees due to the game division or importance. Games with pre-assigned referees to some refereeing positions can also be handled by this approach. Each refereeing position to be filled by a referee is called

a refereeing slot.

Let $S = \{1, \dots, n\}$ be the set of refereeing slots. Each refereeing slot $j \in S$ has to be filled by a referee with a previously determined minimum skill level q_j . Let $R = \{1, \dots, m\}$ be the set of referees, represented by their indices. Each referee $i \in R$ has a certain skill level, denoted by p_i , defining the refereeing slots in which he/she can officiate. Referees may declare their unavailability to officiate at certain time slots. Furthermore, each referee $i \in R$ establishes M_i as the maximum number of games he/she is able to officiate and T_i as the target number of games he/she is willing to officiate. Travels are not allowed, i.e. referees that officiate more than one game in the same day must be assigned to games that take place at the same facility. Moreover, referees that are also players have a hard facility assignment constraint: they must officiate at the same facility where they play.

The *Referee Assignment Problem* (RAP) consists in assigning referees to all refereeing slots associated to games scheduled to a given time interval (typically, a day or a weekend), minimizing the sum over all referees of the absolute value of the difference between the target and the actual number of games assigned to each referee and satisfying a set of hard constraints listed below:

- (a) all refereeing slots must be filled for all games;
- (b) referees cannot be assigned to refereeing slots overlapping time slots where they are already scheduled to play or to officiate;
- (c) referees cannot officiate games at time slots where they declared to be unavailable;
- (d) referees must meet the minimum skill level established for each refereeing slot;
- (e) referees cannot officiate more than a given maximum number of games; and
- (f) referees cannot officiate in more than one facility.

3 Integer programming model

The problem described in the previous section can be formulated by integer programming. We denote by d_i the absolute value of the difference between the target and the actual numbers of games assigned to referee $i \in R$. The following variables are used in the formulation:

$$x_{ij} = \begin{cases} 1, & \text{if referee } i \in R \text{ is assigned to slot } j \in S \\ 0, & \text{otherwise.} \end{cases}$$

Furthermore, $C(j) \subseteq S$ denotes the set of refereeing slots conflicting with slot $j \in S$, i.e. refereeing slots that take place at different facilities than or overlapping with j . Also, $U(i) \subseteq S$ represents the set of refereeing slots to which referee $i \in R$ cannot be assigned due to a lower skill level or to his/her unavailability. The RAP integer programming model can be formulated as:

$$\text{minimize } \sum_{i=1}^m d_i \tag{1}$$

subject to:

$$d_i \geq T_i - \sum_{j=1}^n x_{ij} \quad \forall i = 1, \dots, m \quad (2)$$

$$d_i \geq \sum_{j=1}^n x_{ij} - T_i \quad \forall i = 1, \dots, m \quad (3)$$

$$\sum_{i=1}^m x_{ij} = 1 \quad \forall j = 1, \dots, n \quad (4)$$

$$\sum_{j=1}^n x_{ij} \leq M_i \quad \forall i = 1, \dots, m \quad (5)$$

$$x_{ij} + x_{i'j} \leq 1 \quad \forall i = 1, \dots, m, \quad \forall j = 1, \dots, n, \quad \forall j' \in C(j) \quad (6)$$

$$\sum_{j \in U(i)} x_{ij} = 0 \quad \forall i = 1, \dots, m \quad (7)$$

$$x_{ij} \in \{0, 1\} \quad \forall i = 1, \dots, m, \quad \forall j = 1, \dots, n. \quad (8)$$

The objective function (1) states that the sum over all referees of the slack between their target and actual numbers of scheduled games is minimized. Constraints (2) and (3) enforce that d_i is equal to the absolute value of the difference between the target and actual numbers of games assigned to referee $i \in R$. Constraints (4) ensure that every refereeing slot must be assigned to exactly one referee. Constraints (5) establish the upper bound to the number of refereeing slots that can be assigned to each referee. Constraints (6) ensure that refereeing slots with timetabling conflicts or taking place at different facilities cannot be assigned to the same referee. Constraints (7) prevent assignments that violate minimum skill level and unavailability restrictions (alternatively, all variables x_{ij} with $j \in U(i)$ may simply be removed from the model). Constraints (8) establish the integrality of the decision variables.

4 NP-completeness

We consider the following feasibility decision problem (DRAP):

Problem: REFEREE ASSIGNMENT

Input: Set S of refereeing slots, set R of referees, and the maximum number of games to be officiated by each referee.

Question: Is there an assignment of referees in R to refereeing slots in S satisfying constraints (a) to (f)?

Theorem 1. *DRAP is NP-complete.*

Proof. DRAP is clearly in NP, since the feasibility of any assignment can be checked in time polynomial in $|R|$ and $|S|$. To prove its NP-completeness, we use a transformation from the problem of Partition into Bounded Independent Sets on interval graphs.

Given an undirected graph $G = (V, E)$ and integer numbers k and k' , this problem consists in deciding whether there exists a partition of V into k independent sets I_1, \dots, I_k , with $|I_i| \leq k'$ for $1 \leq i \leq k$. This problem is NP-complete even if G is an interval graph, see [1].

We build an instance of DRAP where the set $S = \{1, \dots, |V|\}$ has exactly $|V|$ refereeing slots, each of them associated with a different game. All games take place at the same date and facility. The minimum skill level associated to each refereeing slot $j \in S$ is set as $q_j = 1$. Let $R = \{1, \dots, k\}$ be the set of available referees and set $M_i = k'$, $p_i = 1$, and $U(i) = \emptyset$ for every $i \in R$. The linear time recognition algorithm of Corneil et al. [2] is used to build an interval representation of G . Each interval of the latter is mapped to one refereeing slot, whose starting and ending times coincide with the starting and ending points of the corresponding interval.

We now prove that given the interval graph G and the integer numbers $k, k' \in N$, there is a partition of V into independent sets I_1, \dots, I_k with $|I_i| \leq k'$ for $1 \leq i \leq k$ if, and only if, there is a feasible assignment of the referees in R to the set of refereeing slots S built as above, subject to constraints (a) to (f), with $q_j = 1$ for all $j \in S$, $M_i = k'$ and $p_i = 1$ for all $i \in R$.

First, suppose we have a partition of $G = (V, E)$ into independent sets I_1, \dots, I_k , with $|I_i| \leq k'$ for $1 \leq i \leq k$. The slots assigned to referee $i \in R$ are exactly those corresponding to the vertices in I_i . This association guarantees that constraints (a) and (b) are satisfied. Constraints (c) and (f) are trivially satisfied, since $U(i) = \emptyset$, for all $i \in R$, and all games take place at the same facility. Since $p_i = 1$, for all $i \in R$, and $q_j = 1$, for all $j \in S$, constraint (d) is also trivially satisfied. Finally, constraint (e) is satisfied since $|I_i| \leq k' = M_i$ for all $i \in R$.

We now consider a feasible solution to an instance of DRAP. We construct the interval graph $G = (V, E)$ by assigning each refereeing slot $j \in S$ to a vertex $j \in V$. There is an edge $(j, j') \in E$ for each pair j and j' of overlapping refereeing slots. The partition of V into the independent sets I_1, \dots, I_k is such that vertex j belongs to the independent set I_i if referee i is assigned to refereeing slot j . As the number of slots assigned to each referee is bounded by k' , this partition into bounded independent sets is feasible. \square

5 Solution approach

```

1 Algorithm
   RefereeAssignmentHeuristic(MaxIterations1, MaxIterations2)
2 Solution  $\leftarrow$  BuildGreedyRandomizedSolution();
3 if not isFeasible(Solution) then
4   | Solution  $\leftarrow$  RepairHeuristic(Solution, MaxIterations1);
5 end
6 if isFeasible(Solution) then
7   | Solution  $\leftarrow$  ImprovementHeuristic(Solution, MaxIterations2);
8   | return Solution;
9 else
10  | return no feasible solution was found;
11 end

```

Algorithm 1: Referee assignment heuristic.

We propose a three-phase heuristic approach to tackle real-life large instances of the referee assignment problem. The first phase consists in applying a greedy heuristic to find

an initial solution, possibly violating some constraints. The second phase is a repair heuristic, based on the Iterated Local Search (ILS) metaheuristic, applied whenever necessary to make the initial solution feasible. Finally, another ILS-based heuristic is used to improve the feasible solution. Algorithm 1 shows the general scheme of this approach.

5.1 Greedy randomized constructive heuristic

The first phase of our approach attempts to build a feasible solution. Its main principle consists in assigning first the referees that are also players to the facilities where they have a game. Next, while there are unassigned refereeing slots and unassigned referees, the heuristic greedily selects a facility with unassigned refereeing slots, obtains an unassigned referee and assign refereeing slots to this referee without violating any constraint. Finally, if any refereeing slot remains unassigned, the solution is completed with infeasible assignments.

The pseudo-code of this heuristic is presented in Algorithm 2. We denote by S^u the set of all unassigned refereeing slots, by R^{HF} the set of referees associated with a hard facility constraint, and by R^{NHF} the set of referees with no hard facility constraint, i.e. $R = R^{HF} \cup R^{NHF}$. These sets are initialized respectively in lines 2, 3, and 4. The loop in lines 5 to 15 is performed until all referees associated with hard facility constraints have been examined and assigned to as many refereeing slots as possible. A referee i with an associated hard facility constraint is randomly selected in line 6 and removed from R^{HF} in line 7. The facility f where referee i plays a game is identified in line 8. The loop in lines 9 to 14 investigates all empty refereeing slots corresponding to games that take place at this facility. If referee i can be assigned to refereeing slot j without violating any constraint (line 10), then the assignment is made in line 11 and refereeing slot j is removed from the set of unassigned refereeing slots in line 12.

Next, the loop in lines 16 to 27 attempts to fill the remaining unassigned refereeing slots with referees without hard facility constraints. In line 17, \bar{p} is initialized with the highest skill level among all referees in R^{NHF} . The facility f with the strongest need for referees with skill level \bar{p} is greedily selected in line 18. A referee $i \in R^{NHF}$ with $p_i = \bar{p}$ is randomly selected in line 19, removed from R^{NHF} in line 20, and assigned to as many refereeing slots as possible. The loop in lines 21 to 26 investigates all empty refereeing slots corresponding to games that take place at facility f . If referee i can be assigned to refereeing slot j without violating any constraint (line 22), then the assignment is made in line 23 and refereeing slot j is removed from the set of unassigned refereeing slots in line 24.

If unassigned refereeing slots remain at this stage (line 28), then the loop in lines 29 to 34 makes infeasible assignments to complete the solution. The facility f where the game corresponding to the unassigned refereeing slot $j \in S^u$ takes place is determined in line 30. A referee $i \in R$ officiating at this facility is randomly selected in line 31. The assignment of referee i to refereeing slot j is made in line 32 and refereeing slot j is removed from the set of unassigned refereeing slots in line 33. The selection in line 31 ensures that each referee officiates all games at the same facility.


```

1 Algorithm BuildGreedyRandomizedSolution()
2  $S^u \leftarrow \{j = 1, \dots, n : \sum_{i=1}^m x_{ij} = 0\}$ ;
3  $R^{HF} \leftarrow \{i = 1, \dots, m : \text{referee } i \text{ plays at least one game}\}$ ;
4  $R^{NHF} \leftarrow R - R^{HF}$ ;
5 while  $R^{HF} \neq \emptyset$  do
6   Randomly select a referee  $i \in R^{HF}$ ;
7    $R^{HF} \leftarrow R^{HF} - \{i\}$ ;
8   Let  $f$  be the facility where referee  $i$  plays a game;
9   forall  $j \in S^u$  : refereeing slot  $j$  takes place at facility  $f$  do
10    if referee  $i$  can be assigned to refereeing slot  $j$  then
11       $x_{ij} \leftarrow 1$ ;
12       $S^u \leftarrow S^u - \{j\}$ ;
13    end
14  end
15 end
16 while  $S^u \neq \emptyset$  and  $R^{NHF} \neq \emptyset$  do
17    $\bar{p} \leftarrow \max_{i \in R^{NHF}} \{p_i\}$ ;
18   Let  $f$  be the facility with the strongest need for referees with skill level
   equal to  $\bar{p}$ ;
19   Randomly select a referee  $i \in R^{NHF}$  with  $p_i = \bar{p}$ ;
20    $R^{NHF} \leftarrow R^{NHF} - \{i\}$ ;
21   forall  $j \in S^u$  : refereeing slot  $j$  takes place at facility  $f$  do
22    if referee  $i$  can be assigned to refereeing slot  $j$  then
23       $x_{ij} \leftarrow 1$ ;
24       $S^u \leftarrow S^u - \{j\}$ ;
25    end
26  end
27 end
28 if  $S^u \neq \emptyset$  then
29   forall  $s_j \in S^u$  do
30    Let  $f$  be the facility where  $s_j$  takes place;
31    Randomly select a referee  $i \in R$  officiating at facility  $f$ ;
32     $x_{ij} \leftarrow 1$ ;
33     $S^u \leftarrow S^u - \{j\}$ ;
34   end
35 end
36 return Solution : referee  $i \in R$  is assigned to refereeing slot  $j \in S$ 
   if and only if  $x_{ij} = 1$ ;

```

Algorithm 2: Greedy randomized constructive heuristic.

A greedy criterion is applied in line 18 to select a facility with the strongest need for referees with a certain skill level. We recall that each refereeing slot is associated to one refereeing position in a game. For instance, a refereeing slot of a soccer game may be associated with the main referee or with one of the two linesmen. Therefore, a facility holding soccer games needs at least three referees: the main referee and two linesmen.

Suppose that one referee could be assigned to all games as the main referee and two other referees could be assigned as the first and second linesmen in all games. In this case, each of these three referees would be associated to the same refereeing position in all games. To be able to be assigned to all games in the same position, a referee must have a skill level greater than or equal to the minimum skill level of all refereeing slots associated to this position. In consequence, the skill level of this referee should be greater than or equal to the maximum minimum skill level over all refereeing slots.

The above reasoning is used to compute an estimate $c^1(f, \ell)$ of the minimum number of referees with skill level ℓ needed to officiate at facility f . We suppose that each refereeing slot is an abstraction of a real refereeing position. For instance, P_1 could be related to the main referee of a soccer game, P_2 to the first linesman, and P_3 to the second linesman. In addition, we assume that the refereeing positions are sorted by the non increasing order of their minimum skill levels: for the same game, the minimum skill level associated with a refereeing slot for position P_i cannot be smaller than the minimum skill level associated with a refereeing slot for position P_j if $j > i$. To illustrate this idea, we consider an instance with one single facility a where three games g_1 , g_2 , and g_3 take place, with three refereeing positions P_1 , P_2 , and P_3 each. Table 1 displays the minimum skill level requirements for each refereeing slot. A “-” denotes that the corresponding game does not require a referee in this position.

Facility		Refereeing positions		
a		P_1	P_2	P_3
	g_1	4	-	3
Games	g_2	5	4	2
	g_3	6	5	5

Table 1: Minimum skill levels.

Facility a in this example needs at least one referee with skill level greater than or equal to 6 to be the main referee in all games, one referee with skill level greater than or equal to 5 to be the first linesman in all games, and another referee with skill level greater than or equal to 5 to be the second linesman in all games. Therefore, $c^1(a, 6) = 1$, $c^1(a, 5) = 2$, and $c^1(a, 1) = c^1(a, 2) = c^1(a, 3) = c^1(a, 4) = 0$.

For any two facilities a and b , we say that a has a stronger need than b for referees with skill level ℓ if $c^1(a, \ell) > c^1(b, \ell)$. The number $c^2(f, \ell)$ of refereeing slots with minimum skill level less than or equal to ℓ to be assigned at each facility f is used to break ties regarding the first criterion, i.e. whenever $c^1(a, \ell) = c^1(b, \ell)$. Table 2 displays the values of $c^2(a, \ell)$, $\ell = 1, \dots, 6$, for the same example.

Once \bar{p} is computed in line 17 of Algorithm 2, the facility with the strongest need for referees with skill level equal to \bar{p} is determined in line 18 as that which lexicographic maximizes the following criteria:

- Criterion 1: $c^1(f, \bar{p})$ (an estimate of the minimum number of referees with skill level \bar{p} needed to officiate at facility f).
- Criterion 2: $c^2(f, \bar{p}) - c^2(f, \bar{p} - 1)$ (number of refereeing slots with minimum skill level equal to \bar{p} to be assigned at facility f , if $\bar{p} > 1$).

Skill level	Refereeing slots to be assigned at facility a
≤ 1	$c^2(a, 1) = 0$
≤ 2	$c^2(a, 2) = 1 + 0 = 1$
≤ 3	$c^2(a, 3) = 1 + 1 = 2$
≤ 4	$c^2(a, 4) = 2 + 2 = 4$
≤ 5	$c^2(a, 5) = 3 + 4 = 7$
≤ 6	$c^2(a, 6) = 1 + 7 = 8$

Table 2: Computations of the metric c^2 .

- Criterion 3: $c^2(f, \bar{p})$ (number of refereeing slots to be assigned with minimum skill level less than or equal to \bar{p} at facility f).
- Criterion 4: $c^2(f, HSL)$ (number of refereeing slots with minimum skill level less than or equal to HSL at facility f , where HSL denotes the highest required skill level in the problem domain).

5.2 ILS-based scheme

Both the repair and the improvement heuristics use similar ILS schemes. They start by applying a first improving local search to the initial solution. Since the local search involves moves that change referee assignments for only one facility at a time, it should be applied to every facility.

Then, for a given number of iterations, a perturbation involving one pair of facilities is applied to the current solution. Each perturbation is followed by two applications of the local search procedure, once to each of the facilities of the pair involved in the perturbation. The solution obtained by local search is accepted if it satisfies a given acceptance criterion. This scheme is illustrated by the pseudo-code of Algorithm 3.

```

1 Algorithm ILS_Scheme(Solution, MaxIterations)
2 foreach facility  $f$  do
3   | Solution  $\leftarrow$  LocalSearch( $f$ , Solution);
4 end
5 for  $i = 1, \dots, \text{MaxIterations}$  do
6   | NewSolution  $\leftarrow$  Perturbation(Solution);
7   | Let  $f_1$  and  $f_2$  be the facilities involved in the perturbation;
8   | NewSolution  $\leftarrow$  LocalSearch( $f_1$ , NewSolution);
9   | NewSolution  $\leftarrow$  LocalSearch( $f_2$ , NewSolution);
10  | Solution  $\leftarrow$  AcceptanceCriterion(Solution, NewSolution);
11 end
12 return Solution;

```

Algorithm 3: ILS scheme.

We describe next the local search procedure and its associated neighborhoods, followed by the repair and improvement heuristics.

5.3 Local search and neighborhoods

Solutions built by the constructive heuristic are not necessarily optimal or even feasible. A local search algorithm successively replaces the current solution by a better one in a neighborhood of the first, terminating at a local optimum. In a *first improving* strategy, the current solution is replaced by the first neighbor whose cost function value improves that of the current solution. In this work, we consider two neighborhoods:

- swap moves: referees assigned to two refereeing slots are swapped (such moves do not change the number of games assigned to each referee) and
- replace moves: the referee assigned to a refereeing slot is replaced by another referee (such moves increase by one the number of games assigned to one referee and decrease by one the number of games assigned to the other).

As referees cannot be assigned to games at different facilities (hard constraint), only moves involving referees that officiate at the same facility (or do not officiate at all) are allowed (otherwise, and unless two referees were assigned to exactly one game each, a move involving referees that officiate at different facilities would imply on at least one constraint violation). Such restricted neighborhoods considering only moves involving the same facility allow the acceleration of the local search.

The local search procedure performed within the ILS scheme is divided into two phases, both of them using a first improving strategy. In the first phase, only improving moves are accepted. The second phase also accepts moves leading to solutions at least as good as the current one, using a tabu list to prevent cycles. Each phase is separated in two parts: first, only swap moves are considered; next, only replace moves.

5.4 Repair heuristic

The repair heuristic follows the ILS scheme in Algorithm 3, based on local search and perturbations. It attempts to make feasible the initial solution obtained by the greedy randomized constructive heuristic. Constraint violations in the initial solution may concern time conflicts, referee unavailabilities, inadequate skill levels, or maximum numbers of games. The repair heuristic minimizes the number of constraint violations of an infeasible initial solution. A modified solution is feasible if and only if it has no constraint violations.

Solutions built by the constructive heuristic have the property that all referees officiate in at most one facility. Therefore, the local search considers only moves involving referees that officiate at the same facility (or do not officiate at all) and attempts to find a feasible solution by minimizing the number of constraint violations. Ties with respect to the number of constraint violations are broken in favor of the solution with the smaller objective function value (i.e., the absolute value of the difference between the target and the actual number of games assigned to each referee involved in a move).

The perturbation procedure within the repair heuristic changes the facility where one of the referees officiates, according to the following steps:

1. select a facility f with infeasible referee assignments;
2. select the highest minimum skill level ℓ^* over all refereeing slots in facility f assigned to referees with at least one violation;

3. determine a referee r that officiates at another facility f' (or does not officiate at all) whose skill level is greater than or equal to ℓ^* ;
4. randomly select referees other than r that officiate at facility f' and assign them to the refereeing slots currently assigned to r ;
5. assign referee r to any refereeing slot at facility f currently assigned to a referee with at least one violation.

The solution `NewSolution` obtained after a perturbation followed by local search is accepted by procedure *AcceptanceCriterion* if and only if it has fewer constraint violations or the same number of violations and a smaller objective function value than the current solution. The maximum number of iterations was set at one thousand.

5.5 Improvement heuristic

Once a feasible solution is known, the improvement heuristic is performed as an attempt to reduce the current value of the objective function, i.e. to minimize the sum over all referees of the absolute value of the difference between the target and the actual number of games assigned to each of them. The improvement heuristic is also based on the ILS scheme presented in Section 5.2.

The local search used in the improvement heuristic differs slightly from that used in the repair heuristic: swap moves are not performed (because they cannot improve the objective function) and only moves and perturbations that preserve feasibility are considered.

The perturbations applied to the current solution within the improvement heuristic select two referees that officiate at different facilities and swap all their assignments, according to the following steps:

1. Choose a possible perturbation: select two referees officiating at different facilities. If the swap of all their assignments do not preserve feasibility, then go to the final step. Otherwise, temporarily perform the swap of all assignments of the two selected referees.
2. Look ahead: for each of the two selected referees whose target number of games is greater than the new (after the swap) number of games he will officiate, check if there are other refereeing slots in which he could officiate at the new facility. This look ahead procedure only checks refereeing slots that are currently assigned to referees officiating more games than their targets and only until the referee under investigation does not officiate more games than his/her target. Whenever possible, temporarily replace the previously assigned referee by the new referee involved in the perturbation.
3. Accept the perturbation: if the perturbation applied to these two referees (swap of all their assignments) followed by all possible replace moves in the destination facility for each referee (look ahead) decreases the objective function value, then accept the perturbation by making final all temporary changes. Otherwise, go to step 4.

4. Return: go back to the first step to select a new pair of referees, unless all pairs of referees have already been considered. In this case, perform the swap of all assignments of the pair of referees that preserves feasibility and increases the least the objective function.

Solution `NewSolution` obtained after a perturbation followed by local search is always accepted, because the heuristic chooses either an improving perturbation or the one that deteriorates the least the current solution. The maximum number of iterations was set at 200.

6 Computational results

Only very small instances with up to 40 games and 60 referees could be exactly solved by a commercial solver such as CPLEX 9.0, applied directly to the integer programming model presented in Section 3. In this section, we report computational results obtained on realistic, real-size randomly generated instances.

6.1 Test problems

Test instances have been randomly generated following patterns similar to those observed in real-life soccer instances. They have up to 500 games and up to 1000 referees, with different number of referees, different numbers of facilities, and different patterns of the target number of games each referee is willing to officiate. Each game has three refereeing positions.

Two different patterns were used to generate the target number of games T_i each referee $i \in R$ is willing to officiate. According to pattern P_0 , T_i is an integer randomly selected in the interval $[0, M_i]$. According to pattern P_1 , T_i is proportional to $1/p_i$, i.e. the higher the referee skill level is, the lower his/her target number of games is. Each game is associated to a division selected randomly. The higher is a game division, the higher are the minimum skill levels of its refereeing slots.

Table 3 presents the parameter values used to generate the test problems. Five different instances were generated for each of the 36 parameter combinations, in a total of 180 test problems.

Games	Referees	Facilities	Patterns
300	450, 525, 600	40, 50	P_0, P_1
400	600, 700, 800	55, 65	P_0, P_1
500	750, 875, 1000	65, 85	P_0, P_1

Table 3: Instances dimensions combinations.

6.2 Numerical results

The experimental results reported in this section were obtained on a 2.0 GHz Pentium IV processor with 512 Mbytes of RAM memory running Windows 2000TM. All codes were implemented in C.

Tables 4 to 9 summarize the results obtained by our heuristic approach for some classes of parameter combinations. Initial solutions are computed by the greedy randomized heuristic. Computation times (in seconds) and objective function values are average results over ten runs for each instance. For each phase of the heuristic (construction, repair, improvement), we present its computation time (in seconds) and the objective function value of the solutions found. For the construction and repair phases, we also report the number of runs where a feasible solution was found.

The constructive heuristic ran in less than 0.1 second for all instances and found feasible solutions for most of them. The repair heuristic succeeded to find a feasible solution in most cases in less than one minute, whenever the constructive heuristic failed. We stress the importance of quick procedures for finding initial solutions for hard combinatorial problems in sports, as already noticed by Ribeiro and Urrutia [11].

The improvement phase improved the objective function value of feasible initial solutions by up to 63%. Instances with more facilities or fewer referees were harder in terms of computation times and building feasible solutions.

Instance	Construction			Repair			Improvement	
	secs.	value	feas.	secs.	value	feas.	secs.	value
I ₁	0.02	1286.20	10	—	—	—	32.34	619.60
I ₂	0.02	1360.00	5	0.47	1338.00	10	31.81	623.40
I ₃	0.02	1269.00	2	0.60	1247.00	10	33.87	621.60
I ₄	0.03	—	—	1.14	1303.20	10	30.28	627.20
I ₅	0.03	1302.67	3	1.40	1259.14	10	33.73	654.00

Table 4: 500 games, 750 referees, 65 facilities, and pattern P_0 .

Instance	Construction			Repair			Improvement	
	secs.	value	feas.	secs.	value	feas.	secs.	value
I ₁	0.02	1752.75	8	0.66	1709.00	10	31.59	1022.60
I ₂	0.02	1669.57	7	0.02	1675.67	10	30.34	888.60
I ₃	0.02	—	—	5.91	1569.80	10	29.55	942.00
I ₄	0.03	1777.00	1	1.53	1725.00	10	31.81	1033.80
I ₅	0.02	1704.80	5	0.49	1704.80	10	28.17	952.00

Table 5: 500 games, 750 referees, 65 facilities, and pattern P_1 .

Instance	Construction			Repair			Improvement	
	secs.	value	feas.	secs.	value	feas.	secs.	value
I ₁	0.03	—	—	11.27	1111.60	10	22.74	612.60
I ₂	0.03	—	—	6.69	1231.60	10	24.18	715.20
I ₃	0.03	—	—	11.33	1182.40	10	22.29	672.60
I ₄	0.03	—	—	4.61	1229.00	10	23.45	692.80
I ₅	0.03	—	—	3.39	1234.60	10	19.50	646.00

Table 6: 500 games, 750 referees, 85 facilities, and pattern P_0 .

Instance	Construction			Repair			Improvement	
	secs.	value	feas.	secs.	value	feas.	secs.	value
I ₁	0.03	—	—	2.75	1670.60	10	25.85	1043.80
I ₂	0.02	—	—	19.29	1649.50	8	26.15	1147.00
I ₃	0.03	—	—	14.77	1586.60	10	24.65	1107.60
I ₄	0.03	—	—	1.22	1602.80	10	25.59	1007.40
I ₅	0.03	—	—	2.69	1611.20	10	24.60	1002.80

Table 7: 500 games, 750 referees, 85 facilities, and pattern P_1 .

Instance	Construction			Repair			Improvement	
	secs.	value	feas.	secs.	value	feas.	secs.	value
I ₁	0.03	1582.80	10	—	—	—	45.07	574.40
I ₂	0.02	1627.40	10	—	—	—	42.92	609.20
I ₃	0.03	1535.40	10	—	—	—	44.62	558.20
I ₄	0.03	1655.60	10	—	—	—	43.28	576.60
I ₅	0.02	1626.00	10	—	—	—	43.31	619.20

Table 8: 500 games, 875 referees, 65 facilities, and pattern P_0 .

Instance	Construction			Repair			Improvement	
	secs.	value	feas.	secs.	value	feas.	secs.	value
I ₁	0.03	2195.20	10	—	—	—	39.64	1091.80
I ₂	0.03	2040.20	10	—	—	—	42.33	955.40
I ₃	0.02	2153.40	10	—	—	—	41.34	1032.20
I ₄	0.03	2173.40	10	—	—	—	42.54	1069.00
I ₅	0.03	2137.60	10	—	—	—	39.73	1035.00

Table 9: 500 games, 875 referees, 65 facilities, and pattern P_1 .

In another experiment, we compared the results obtained by the heuristic with those found by CPLEX 9.0 when applied to formulation (1)-(8) for some small instances that could be exactly solved within reasonable computation time. The heuristic always received the same computation time that CPLEX took to find the optimal solution. Some numerical results are summarized in Table 10. For each instance, we first give its identification and the pattern used for its generation. The two next columns display the optimal solution value and the computation time in seconds taken by CPLEX (and, consequently, given to the heuristic). Next, the table shows the average and the best solution values found by the heuristic over ten runs. The last column gives the time taken to find the best solution in the corresponding run. These results show that the heuristic was able to find the optimal solution for three out of the five test instances considered in this table. Furthermore, the times taken by the heuristic are significantly smaller than those observed with CPLEX, even for the small instances that the latter was able to solve to optimality.

In the last computational experiment, we replaced by a quadratic penalization the linear objective function penalizing the absolute value of the difference between the target and the actual numbers of games assigned to each referee. Figure 1 displays results for instance I₃ with 500 games, 750 referees, 85 facilities, and pattern P_1 . This figure shows that

Instance		CPLEX		Heuristic		
identification	pattern	optimum	secs.	average	best	secs.
I ₂	P_0	43	164.00	47.00	44	18.99
I ₃	P_0	18	200.00	20.80	18	3.05
I ₅	P_0	44	137.00	45.20	44	11.45
I ₂	P_1	65	128.00	67.20	65	15.32
I ₅	P_1	72	47.00	82.40	75	8.83

Table 10: 33 games, 57 referees, 5 facilities, patterns P_0 and P_1 .

more balanced solutions can be obtained, in which the occurrences of larger differences are replaced by those of smaller differences concentrated at only one unit. The computation times of the constructive, repair, and improvement heuristics were not affected by the change of the objective function.

Figure 1: Linear vs. quadratic objective functions.

Table 11 details the differences between the target and the actual numbers of games assigned to each referee in the solutions obtained with the linear and quadratic cost functions. We observe that 76 extremely privileged referees (i.e., those officiating exactly their target number of games) in the solution obtained with the linear cost function lose their privileges in the solution obtained with the quadratic cost function. Also, 23 referees that were far from their targets are now very close to them (i.e., their differences are now equal to one). The new solution obtained with the quadratic cost function is certainly more fair than that associated with the linear costs.

Difference	Linear	Quadratic
0	255	179
1	182	281
2	156	149
3	67	66
4	50	43
5	23	18
6	13	10
7	3	3

Table 11: Differences in the linear and quadratic objective functions.

7 Concluding remarks

We introduced in this paper the referee assignment problem, a new optimization problem in sports. The problem was formulated as an integer model and the NP-completeness of its decision version was proved.

A three-phase heuristic was proposed and implemented. Computational results on realistic instances showed the effectiveness of the greedy randomized constructive heuristic

combined with the repair heuristic to build feasible solutions. The improvement procedure used in the third phase was able to substantially improve solution quality. We also illustrated the importance of a quick construction procedure to build initial solutions.

We also compared the solutions obtained by the heuristic with those produced by CPLEX for some small instances that could be solved to optimality in reasonable computation times. The heuristic not only was able to find the optimal solutions for several instances, but also the computation times to find the best solution were significantly smaller than those observed with CPLEX.

Finally, we investigated and compared the behavior of an alternative quadratic objective function, which was able to find more fair solutions than the formulation with a linear cost function.

We are currently working on some extensions addressing further constraints of real-life applications, such as the existence of hard and soft links between some referees. In these situations, some referees may want to work with the same referees as partners in every game they officiate. This is the case when they are more confident to officiate together, but also when they want to travel in car pools or to officiate with relatives. Decision makers may also want referee assignments matching preferences regarding the facilities, divisions, and time slots where the referees officiate.

Another extension occurs when referees are able to officiate games in different facilities. In this case, travel times between facilities should also be considered for feasibility matters. They can also be incorporated to the objective function, so as that the minimization of the total traveling time turns out to be another objective. The minimization of the waiting times between consecutive games assigned to the same referee is also relevant.

The referee assignment problem has clearly the flavor of a multi-criteria optimization application. We are also investigating the use of multi-criteria methods coupled with a decision support system for its solution in practice.

References

- [1] H.L. Bodlaender and K. Jansen. Restrictions of graph partition problems - Part I. *Theoretical Computer Science*, 148:93–109, 1995.
- [2] D.G. Corneil, S. Olariu, and L. Stewart. The ultimate interval graph recognition algorithm? In *Proceedings of the Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 175–180. SIAM, 1998.
- [3] J.H. Dinitz and D.R. Stinson. On assigning referees to tournament schedules. *Bulletin of the Institute of Combinatorics and its Applications*, 44:22–28, 2005.
- [4] K. Easton, G. L. Nemhauser, and M. Trick. The traveling tournament problem: Description and benchmarks. In *Proceedings of the 7th International Conference on Principles and Practice of Constraint Programming*, pages 580–584, London, 2001. Springer-Verlag.
- [5] K. Easton, G.L. Nemhauser, and M. Trick. Sports scheduling. In J.T. Leung, editor, *Handbook of Scheduling: Algorithms, Models and Performance Analysis*, pages 52.1–52.19. CRC Press, 2004.

- [6] J.R. Evans. A microcomputer-based decision support system for scheduling umpires in the American baseball league. *Interfaces*, 18:42–51, 1988.
- [7] J.R. Evans, J.E. Hebert, and R.F. Deckro. Play ball - The scheduling of sports officials. *Perspectives in Computing*, 4:18–29, 1984.
- [8] G.L. Nemhauser and M.A. Trick. Scheduling a major college basketball conference. *Operations Research*, 46:1–8, 1997.
- [9] C.C. Ribeiro and S. Urrutia. OR on the ball: Applications in sports scheduling and management. *OR/MS Today*, 31:50–54, 2004.
- [10] C.C. Ribeiro and S. Urrutia. An application of integer programming to playoff elimination in football championships. *International Transactions in Operational Research*, 12:375–386, 2005.
- [11] C.C. Ribeiro and S. Urrutia. Heuristics for the mirrored traveling tournament problem. *European Journal of Operational Research*, to appear.
- [12] M.B. Wright. Scheduling English cricket umpires. *Journal of the Operational Research Society*, 42:447–452, 1991.