European Journal of Operational Research 238 (2014) 886–898

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Innovative Applications of O.R.

Decomposition and local search based methods for the traveling umpire problem

Tony Wauters *, Sam Van Malderen, Greet Vanden Berghe

CODeS, Department of Computer Science, KU Leuven, Gebroeders De Smetstraat 1, 9000 Gent, Belgium

ARTICLE INFO

Article history:
Received 5 November 2013
Accepted 29 April 2014
Available online 9 May 2014

Keywords:
OR in sports
Traveling umpire problem
Heuristics
Iterated local search
Decomposition

ABSTRACT

The Traveling Umpire Problem (TUP) is a challenging combinatorial optimization problem based on scheduling umpires for Major League Baseball. The TUP aims at assigning umpire crews to the games of a fixed tournament, minimizing the travel distance of the umpires. The present paper introduces two complementary heuristic solution approaches for the TUP. A new method called enhanced iterative deepening search with leaf node improvements (IDLI) generates schedules in several stages by subsequently considering parts of the problem. The second approach is a custom iterated local search algorithm (ILS) with a step counting hill climbing acceptance criterion. IDLI generates new best solutions for many small and medium sized benchmark instances. ILS produces significant improvements for the largest benchmark instances. In addition, the article introduces a new decomposition methodology for generating lower bounds, which improves all known lower bounds for the benchmark instances.

1. Introduction and problem description

Sports scheduling enjoys an ever increasing interest of the operations research community. The focus lies mainly on scheduling the games of a competition or tournament. Scheduling tournaments turns out to be a challenging task: a small number of participants and rounds results in a large number of possible combinations, while the fraction of desired solutions is small due to context specific constraints and objectives. To master this complexity, researchers apply a wide range of combinatorial optimization techniques.

Sports scheduling, be it to a lesser extent, also considers assigning officials to the games in a tournament for tennis (Farmer, Smith, & Miller, 2007), football (Alarcón, Durán, & Guajardo, 2014), cricket (Wright, 1991), etc. Duarte, Ribeiro, Urrutia, and Haeusler (2007b, 2007a) introduce the Referee Assignment Problem (RAP), which considers the assignment of a number of referees with different qualifications to the games in a fixed tournament. In an area other than sports, Lamghari and Ferland (2011) consider the assignment of judges for the John Molson International Case Competition. Kendall, Knust, Ribeiro, and Urrutia (2010) give a complete reference of the current state of sports scheduling in operations research.

The present paper focuses on the Traveling Umpire Problem (TUP), which is an academic version of the real world Major League Baseball umpire scheduling problem (MLB-USP). Trick and Yildiz (2007) introduce the TUP and describe it in more detail later (Trick, Yildiz, & Yunes, 2012), comparing the problem to the MLB-USP (Evans, 1988). MLB-USP defines the rules and regulations imposed by the baseball league and umpire union for assigning 17 umpire crews, each consisting of four umpires, to cover all 780 series of an MLB tournament. Each series contains two up to four consecutive games between the same two teams out of all 30 teams. Even though only taking into account the most important constraints, the academic problem retains the most important characteristics of the real world umpire scheduling problem.

The TUP is related to the Traveling Tournament Problem (TTP, Easton, Nemhauser, & Trick (2001)). The latter aims at finding a double round robin schedule for a season of Major League Baseball. Given $2n$ teams, the tournament consists of $4n-2$ rounds in which each team plays against exactly one other team in every round. The TUP considers assigning n umpire crews to the games in such a fixed TTP tournament. Its goal is obtaining a schedule which minimizes the travel distance of the umpire crews, while taking into account the following constraints:

C1. Every game in the tournament is officiated by exactly one umpire crew.
C2. An umpire crew officiates exactly one game per round.

* Corresponding author. Tel.: +32 92658704.
E-mail address: tony.wauters@cs.kuleuven.be (T. Wauters).

http://dx.doi.org/10.1016/j.ejor.2014.04.043
0377-2217/

C3. Every umpire crew should visit the home of every team at least once.

C4. An umpire crew must wait $q_1 - 1$ rounds before revisiting a team's home.

C5. An umpire crew must wait $q_2 - 1$ rounds before officiating the same team again.

With

$$q_1 = n - d_1 \quad (1)$$

$$q_2 = \left\lfloor \frac{n}{2} \right\rfloor - d_2 \quad (2)$$

whereby the values for parameters d_1 and d_2 range from 0 to n and 0 to $\lfloor \frac{n}{2} \rfloor$, respectively. Higher q values make the problem more constrained and drastically reduce the number of feasible solutions. Yildiz (2008) discusses the effect of different values of parameters q_1 and q_2 on the feasibility of the problem.

For simplicity, the above constraints and the remainder of this paper refer to umpire crews as a single umpire since a crew stays together throughout the whole season.

Although the complexity of the TUP is still open at the time of writing, the problem appears to be hard to solve. Minimizing the travel distance puts pressure on constraint C4 and C5 whereas enforcing constraints C4 and C5 increases the travel distance. Moreover, the assignment of one umpire influences the schedule of other umpires due to constraints C1 and C2. In addition, the problem description does not make a distinction between the different umpires. However, symmetrical solutions can be avoided by fixing the umpire assignments within a certain round. Trick et al. (2012) mention that the TUP can be seen as a special case of the vehicle routing problem with time windows.

The present paper presents two complementary heuristic approaches to the Traveling Umpire Problem. Enhanced iterative deepening search with leaf node improvements (IDLI) generates all partial schedules for a window of W rounds. The algorithm then greedily picks the best partial schedules to complete in subsequent stages. The second approach is a custom iterated local search algorithm (ILS) with a step counting hill climbing acceptance criterion. A steepest descent algorithm ensures all solutions are local optima before the ILS invokes the acceptance criterion. Finally, the article introduces a new decomposition methodology for generating tight lower bounds.

The structure of this paper is as follows: Section 2 presents an overview of the existing approaches to the TUP. Section 3 introduces the solution strategies and lower bound methodology. Section 4 reports and discusses computational results for benchmark instances. The last section summarizes conclusions and presents pointers for future research.

2. Related research

Trick and Yildiz (2007) introduce TUP and formulate it as an Integer Program (IP) and a Constraint Program (CP). The same paper presents a greedy matching heuristic (GMH) with Bender's based modifications (GBNS). GBNS constructs a solution one round at a time by matching the umpires to games within the considered round. If the matching heuristic does not obtain a feasible matching at a certain round, GBNS examines the cause of the infeasibility and generates Benders' cuts. A very large neighborhood search algorithm then uses the Benders' cuts to resolve the infeasibility. The paper also tests the performance of the IP and CP formulations on benchmark instances and compares the results to those obtained by GBNS. The IP and CP formulation obtain solutions for relaxations of the benchmark instances, with $q_1 < n$ and/or $q_2 < \lfloor \frac{n}{2} \rfloor$. GBNS obtains solutions of higher quality in a shorter timespan, even for the most constrained versions of the benchmark instances.

Trick et al. (2012) present a simulated annealing algorithm with k-umpire neighborhood to both MLB-USP and TUP. The initial solution is constructed using GMH. The MLB-USP version of the algorithm obtains solutions of much higher quality than those constructed manually in previous years. The TUP version is capable of generating feasible solutions for relaxations of the larger benchmark instances.

Trick and Yildiz (2011) re-evaluate the performance of the algorithms presented by Trick and Yildiz (2007) on new benchmark instances. GBNS obtains the best results for a majority of the benchmark instances. New solver settings improve the performance of the CP and IP formulation.

Trick and Yildiz (2012) propose a genetic algorithm (GA) with a locally optimized crossover operator. Given two schedules and a round as crossover point, the crossover operator matches the rounds appearing before the crossover point in the first schedule to those appearing after the crossover point in the second schedule. The GA improves several of the best results for the TUP benchmark instances compared to the results obtained by Trick et al. (2012).

de Oliveira, de Souza, and Yunes (2014) strengthen the original IP formulation of Trick and Yildiz (2007) by removing one of the variables and some redundant constraints and by adding new valid inequalities. This noteworthy formulation improves all known lower bounds and is the first one capable of obtaining lower bounds for the larger problem instances. A relax-and-fix heuristic then uses this formulation to obtain solutions for the TUP. The relax-and-fix heuristic improves all best known solutions for the benchmark instances.

3. Approaches

The following sections present two new approaches to the TUP. Before going into detail, the first section clarifies the choice behind the final approaches. The last section discusses a methodology for generating tight lower bounds for the TUP.

3.1. Exploratory experiments

Initial experiments have been conducted using both branch and bound and local search. The branch and bound algorithm assigns umpires round by round to the games of the tournament. It reassigns previous umpires when no feasible assignment for an umpire has been found. The local search algorithm improves a given initial solution by randomly exchanging umpire assignments within a given round. It explores infeasible solutions by adding violations of the hard constraints as a penalty term to the objective function.

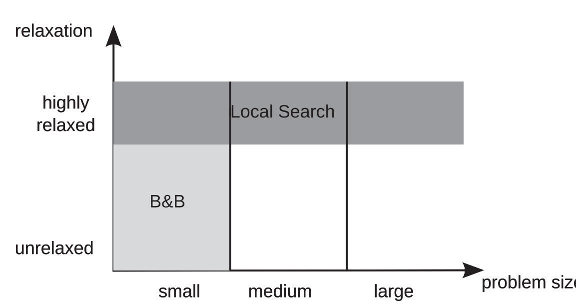

Fig. 1. Summary of the performance of both the initial local search and the branch and bound algorithm on benchmark instances as a function of the problem size and the relaxation level of constraints C4 and C5.