Interfaces
Vol. 37, No. 2, March–April 2007, pp. 187–196
ISSN 0092-2102 | EISSN 1526-551X | 07 | 3702 | 0187
informs®
DOI 10.1287/inte.1060.0259

Scheduling Umpire Crews for Professional Tennis Tournaments

Adam Farmer
Culverhouse College of Business, University of Alabama, Box 870226, Tuscaloosa, Alabama 35487-0226, afarmer@cba.ua.edu

Jeffrey S. Smith
College of Engineering, Auburn University, 207 Dunstan Hall, Auburn, Alabama 36849, jsmith@auburn.edu

Luke T. Miller
School of Business Administration, Fort Lewis College, 1000 Rim Drive, Durango, Colorado 81301, miller_l@fortlewis.edu

Professional tennis organizations, such as the United States Tennis Association (USTA), the Association of Tennis Professionals (ATP), the International Tennis Federation (ITF), and the Women's Tennis Association (WTA), host tennis tournaments throughout the world. At these tournaments, chief umpires assign and schedule line umpires for every match. For most tournaments, they perform this task manually, which can be cumbersome for large tournaments. For large tournaments, such as the US Open, they can use software developed to facilitate scheduling. Unfortunately, the software package currently available often creates suboptimal or infeasible schedules that must be manually adjusted. We developed a program based on optimization that automates the scheduling procedure. Our program consistently provides high-quality schedules in as little as 25 percent of the time taken with other methods.

Key words: recreation and sports; programming: integer.
History: This paper was refereed.

Any respectable tennis fan should be able to tell you who won the 2003 US Open. However, do you think the fan remembers the chair umpire of the women's or men's championship match? Probably not, and tournament officials prefer this anonymity. In any sport, the umpire or referee is responsible for calling a fair contest in agreement with the rules of the sporting event. The umpires' crucial role is to accurately and consistently oversee the match without drawing unnecessary attention to themselves or the governing process. Behind the scenes, an intricate system of hierarchies, experience, and qualifications dictates the proper assignment of umpires to tennis matches.

Chief umpires must follow guidelines that complicate the scheduling process in terms of time and effort needed to assign umpires properly. For example, during large tennis tournaments, up to 18 matches may be played simultaneously. On each court, up to 10 umpires may be calling one match—one chair umpire and nine line umpires. Furthermore, different criteria must be used in selecting chair umpires (who are in charge of each match) and line umpires for each match. The selection of chair umpires is based primarily on nationality, player histories, and experience. The selection of line umpires is based on skill level and experience at a particular position. In addition, line umpires must be allocated to courts to ensure that a minimum number of male and female umpires are on each court at all times. Furthermore, because line umpires cannot stay on court for an entire day while matches are played, the chief umpires must schedule rotation shifts to allow them lunch and rest breaks. Finally, tournaments run anywhere from one to three weeks, with the matches throughout that period increasing in visibility. Because of this visibility, requirements can become stricter as the tournament goes on, making scheduling even more difficult.

Typically, chief umpires assign line umpires to positions during a tournament. To facilitate scheduling, the United States Tennis Association (USTA) developed a software package to automate umpire assignments. The scheduling algorithm used in this program is a greedy heuristic that sequentially assigns the best available umpires to the highest priority lines. First, it evaluates umpires based on their historical performance and assigns them a skill-rating number ranging from 1 to 7 (1 represents the highest skill level). Next, it assigns each court a priority index and the required lines on each court a minimum and maximum skill-rating number. The sequential-assignment algorithm starts at the court with the highest priority and the first position on that court. If that position has a minimum rating of 1, the algorithm searches for the first available umpire with a rating of 1 and assigns him or her to that position. If no 1s are available, the algorithm searches for umpires with a rating of 2. After assigning an umpire to that position, the algorithm moves to the next position and repeats the process. After scheduling the entire court, it moves to the court with the next highest priority index and repeats the search procedure using that court's requirements. It continues until it has fully assigned all courts.

This system has multiple pitfalls. First, the program was designed for use only at the US Open. Most tournaments use different scheduling criteria. Next, the algorithm's greedy nature prevents any form of looking ahead. Similar in nature to a seating heuristic, the system schedules blindly without considering the global solution and is concerned only with the local assignment of an umpire to a location. In other words, it makes individual assignments with no regard for their effect on the schedule as a whole, often assigning too many highly skilled umpires to early courts. No more umpires may be available to meet the skill requirements for later courts, leaving no feasible solution. Furthermore, even if it finds a feasible schedule, the system does not account for gender requirements. The chief umpire must review every schedule for gender feasibility and adjust schedules manually when they do not meet gender requirements.

In 2002, Jeff Smith, an industrial and systems engineering professor at Auburn University, was at the US Open during a rain delay. During this tournament, extensive rain delays completely changed the scheduling requirements, increasing the number of courts needed, changing the umpire positions required on courts, and so forth. He observed the difficulties the chief umpires faced and realized that the problem of updating the schedule after a rain delay was identical to initially scheduling a tournament. He also realized the problem was related to other scheduling problems, such as scheduling airline crews. Smith then approached Adam Farmer (2004) and Luke Miller about developing a tool to provide quality schedules and generalizing the method to be applicable to all tournaments.

Various researchers have developed operations research tools for scheduling sports events. Yang et al. (2002) applied an evolutionary algorithm to minimize the cost associated with creating a baseball schedule while meeting all of the requirements for mandatory games. Urban and Russell (2003) addressed scheduling sports competitions at multiple venues. Schonberger et al. (2004) applied a memetic algorithm to create timetables for noncommercial sports leagues. These and other problems in the literature fall into two main categories: match scheduling for timing events and match locating or choosing physical venues for events. In our problem, we presume these two scheduling phases have been completed. We focus on assigning umpires of varying skill levels to courts that are part of a predetermined schedule.

Another problem that has attracted attention is scheduling airline crews. In general, scheduling airline crews consists of three main phases: (1) forecasting demand and creating flight paths and sequences to meet those demands and minimize costs; (2) assigning planes to flight paths while respecting restrictions on the types of planes allowed to fly into certain airports, the airlines assigned to particular flights, and the number of seats required; and (3) assigning pilots and attendant crews to flights, which is generally broken down into two problems: flight rotations and crew pairing, and crew rostering.

In crew rostering, one assigns pilots and attendants to the legs of flights. This phase of airline scheduling shares some characteristics with umpire scheduling. First, pilots must have specific qualifications to fly certain aircraft. Furthermore, pilot and attendant assignments must respect home-base restrictions and strict federal regulations governing rest periods. Similarly, tennis umpires must meet certain skill and experience qualifications to be assigned to particular lines on a court, and umpire assignments must adhere to strict regulations on gender and nationality issues.

Nicoletti (1975) developed a method for solving the month-long crew-rostering problem while equalizing the workloads of all crew members. Because of computational problems, Nicoletti found it difficult to create month-long schedules and instead modeled the problem as a series of daily minimum-cost network-flow problems. The flow weights included a penalty factor that linked consecutive days together. Nicoletti then solved individual network-flow problems using the Ford-Fulkerson out-of-kilter algorithm to obtain results comparable to manual solution methods.

Lucic and Teodorovic (1998, 1999) addressed a problem similar to Nicoletti's (1975). They modeled the crew-rostering problem as a series of daily problems with penalty factors linking consecutive days. Their method consisted of two phases: (1) constructing an initial solution using a sequential assignment method, and (2) improving that solution using a simulated annealing algorithm. These solutions were within two percent of ideal parameters, indicating that all crew members have equal workloads. No previous research considers all the characteristics of tennis-umpire scheduling, so we couldn't directly apply any previous solution method to our problem.

Scheduling Tennis Tournaments

The umpire-assignment system depends on classifying umpires and their skill levels. Chief umpires are at the highest level in the umpire chain and are responsible for scheduling. Chair umpires work on the court and are responsible for enforcing tournament guidelines and regulations, keeping score, and evaluating the line umpires. Line umpires also work on the court and are responsible for calling the shots hit in their locations. They are assigned to matches primarily based on skill and availability; we focused on assigning them.

Line umpires, which we call simply *umpires* from here on, are the key to ensuring a fair match. During a match, each umpire stands in an assigned position on the court. The umpires judge whether shots hit to their regions are inside or outside the boundary line. After each match, the chair umpires evaluate the line umpires based on the accuracy, confidence, and promptness of their calls. Chief umpires compile these evaluations and use them to assign numerical skill ratings to the umpires, ranging from 1 for the highest level of skill and experience to 7 for the lowest. Chief umpires determine which court location(s) umpires are qualified to judge based on these ratings.

During a particular tournament session, all courts do not necessarily have the same umpire requirements; they can require three to nine umpires per match. On the tennis court, there are five types of lines—*serve*, *base*, *near long*, *far long*, and *center* (Figure 1). The terms *near* and *far* indicate the distance from the chair umpire, who sits above the net on one side of the court.

The US Open uses a 7-point rating system for line umpires (Table 1). The skill levels required run from highest to lowest for the following positions: serve line, baseline, far line, centerline, and near line. Overlaps in skill levels between court locations provide flexibility for assigning umpires, allowing slightly under- or overqualified umpires to call matches.

The most common court configurations consist of a chair umpire plus three, five, or seven line umpires (Figure 1). Chair-and-three configurations are used

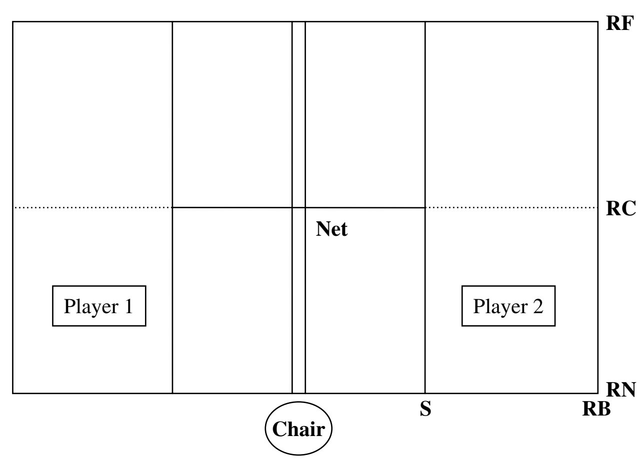

Figure 1: In this overview of a tennis court, we show the locations of the line umpires: S = the serve-line umpire, RB = the right baseline umpire, RN = the right near long-line umpire, RC = the right centerline umpire, RF = the right far-side long-line umpire. The left side of the court has corresponding line umpires, except for the serve-line umpire who moves to the left side when the right side player has the serve.

Line umpire location	Rating (1 to 7)
Serve line	1, 2
Baseline	3, 4
Far line	4, 5
Centerline	5, 6
Near line	6, 7

Table 1: Assuming umpire-rating numbers running from 1 (best) to 7 (as used at the US Open), we assign each line umpire location an umpire with a desirable rating.

only during early rounds of small tournaments. In this configuration, one umpire calls the serve, and the other two call the near-side lines, the far-side lines, and the center serve lines. In this configuration, the side-line umpires are forced to call through the net, that is, one umpire calls both right and left near-side lines, and the other calls right and left far-side lines. The chair umpire is responsible for the baseline calls. No other configuration requires the chair umpire to call a line. In the chair-and-five configuration, the added baseline umpires take some responsibility away from the chair umpire. The chair-and-seven configuration is probably the most common configuration at larger events. The two additional umpires are assigned to the near-side and far-side lines so that no umpire has to call through the net. Only the centerlines lack dedicated umpires in this configuration. In any configuration with no dedicated centerline umpires, the long-line umpires move to the center during serves and then return to their long-line positions once the ball is in play. When rain delays matches, more matches are squeezed into shorter time slots. Because the chief umpire has no additional umpires, he or she must spread the existing umpires among a larger number of matches than previously intended. They often use chair-and-five configurations instead of the chair-and-seven initially scheduled.

Another court requirement concern is gender. Because coaching during a match is a rule violation, any player asking to leave the court during a match (for example, for a restroom break) must be accompanied by a tournament official. Therefore, at least two male umpires must be on court at all times during a men's match, and at least two female umpires must be on court during a women's match. The gender requirement complicates assignments.

Rotation scheduling is another factor to be considered in dealing with multiple concurrently scheduled courts. An umpire crew consists of umpires who travel together from court to court during a session and meet the requirements of those courts. Rotation scheduling consists of grouping courts and assigning crews. In general, chief umpires group with similar requirements for crew configuration, skill, and gender. The scheduler then assigns crews to groups or teams. The number of crews in a team is generally greater than the number of courts in the group to which it is assigned. Because crews outnumber the courts, some crews will be on break while other crews work. The common combinations of crew teams and courts are two crews for one court, three crews for two courts, five crews for three courts, and seven crews for four courts (Figure 2).

A rotation in which two crews rotate on and off one court each hour is common for high-profile matches in the final rounds of a tournament. In a rotation schedule in which five crews alternate time on three courts, all five crews must meet the requirements of all three courts. If one court has all women's matches on that day and either of the other courts hosts a men's match, then all five crews must include at least two men and two women to ensure that all crews in the team are interchangeable. Also, even though both court groupings are for the same day, crews may have to meet the chair-and-seven requirements, whereas crews on the second team must meet only chair-and-five requirements. Furthermore, if one court is hosting a high-profile match, its umpire-skill requirements would probably be higher than those of the other three courts.

The number of umpires available during a session determines the number of crews that can be created. Because rotations of court and crew teams generally require a ratio of at least 1.5 crews per court, the number of crews limits the potential groupings that can be used. It would be impossible to set up a two-for-one rotation for four courts with fewer than eight crews. With seven crews, one could schedule two three-for-two rotations (requiring six crews), a single seven-for-four rotation, or a three-for-two rotation combined with a four-for-two rotation.

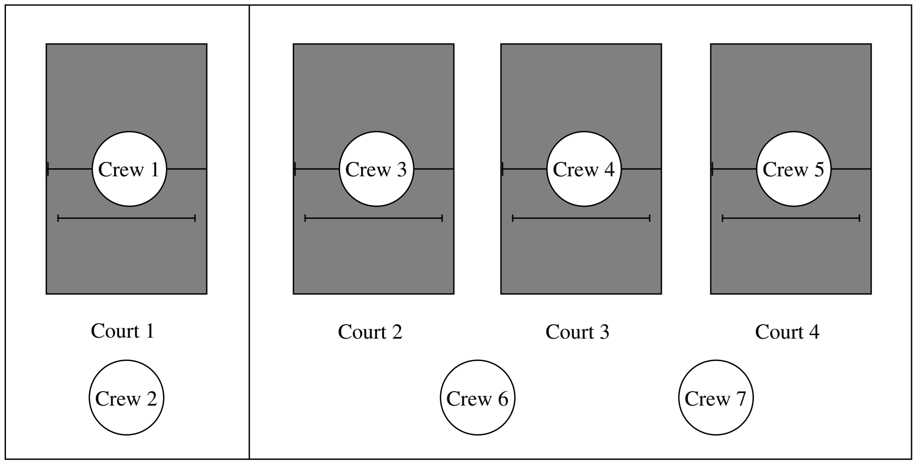

Figure 2: In this crew and court rotation schedule, two crews are assigned to one court and five crews are assigned to three courts. Crew 2 rests while crew 1 calls a match, similarly to crews 6 and 7 resting while crews 3, 4, and 5 call matches. Eventually, crews 2, 6, and 7 will rotate in and call the matches, with crew 1 and two of the three crews, 3, 4, and 5, resting.

A schedule must meet certain stringent quality requirements. All schedules must assign umpires to the proper number of crews based on the selected rotation schedules. All the umpires within a crew must be qualified for their positions. A good schedule should maximize the number of umpires qualified for their positions.

Skill requirements for individual lines are twofold. First, every line on a court has a maximum allowable rating number for assigned umpires (the lower the rating number, the more skilled the umpire). For example, umpires with skill-rating numbers above 2 may not be allowed assignment to a serve line, but umpires with rating numbers up to 4 may be allowed assignment to a baseline. In addition to having a maximum allowable rating number, each line also has a minimum desirable rating number. This limit is not a strict requirement; however, it is a key factor in differentiating between two feasible schedules.

Two completely feasible schedules may meet all gender requirements, have the correct number of umpires assigned to each crew, and umpires with acceptable skill ratings for all positions. The first could include several umpires with ratings of 1 and 2, and the second schedule includes several umpires with ratings between 3 and 5. Which schedule is better depends on what minimum ratings are desirable for each line. A chief umpire seeking a very strong field of line umpires would choose the first schedule. The maximum allowable rating numbers would not change, but the minimum rating numbers would be set very low. On the other hand, a chief umpire wishing to allow inexperienced umpires to gain valuable experience might set the minimum desired rating number on a serve line at 1 and raise the minimum rating numbers on the far and near lines to 3 and 5, respectively. If necessary, umpires with ratings of 1 and 2 could be assigned to the far and near lines, but those with ratings of 3 and 5 would have higher priority. The chief umpire can control the skill distribution among courts by specifying different sets of minimum ratings. The chief umpire's overall measure of quality changes from court to court and day to day as a tournament progresses, partly because of media exposure.

Problem Formulation and Solution Methodology

We developed our integer-programming formulation (appendix) for the USTA. Our overall objective for the crew-scheduling problem was to minimize the weighted deviation of assigned skill ratings from target skill ratings for each position. For lines with positive assignment weight factors, we assigned umpires with lower skill-rating numbers, and for lines with negative assignment weight factors, we assigned umpires with high skill-rating numbers. In minimizing this objective, we must also adhere to several constraints. First, each available umpire can be assigned only once during a scheduling session. Second, each crew must have the correct number of umpires, and each umpire must have a skill-rating number below the maximum allowable for that position. Finally, we must maintain gender feasibility within each crew.

We coded the integer program (IP) in AMPL and solved it using CPLEX optimization software. We tested the IP extensively at the 2003 US Open. By the end of that tournament, USTA officials confirmed that our solution method provided highly acceptable schedules with a minimal amount of effort. The USTA was interested in having the system installed on several PCs in different locations, but it did not have the funds to maintain multiple CPLEX licenses. The USTA officials asked us if we could develop an application that did not rely on using CPLEX's optimization engine.

To eliminate the need for CPLEX, we developed a two-phase heuristic. In the first phase, we use a semisequential method to construct an initial set of umpire assignments. We assign umpires in order of decreasing skill-rating numbers. We use this worst-first assignment method to minimize the number of skill-related infeasibilities. In creating a feasible schedule, maintaining gender feasibility is difficult. The construction heuristic first tries to meet all gender requirements. Because the pool of available umpires generally contains fewer women than men, the procedure first tries to meet all female-gender requirements. Next, it tries to meet all male-gender requirements. After it meets all gender requirements, it fills any remaining gaps in the schedule using male or female umpires.

Once we have created an initial schedule, we use a simulated-annealing algorithm to improve on it (Figure 3). We developed several data sets to test the performance of the simulated-annealing heuristic. It consistently produces schedules whose objective values are within two percent of those we found using CPLEX. Whereas two percent may be a considerable gap for some problems, it isn't for the umpire-scheduling problem; reassigning only one or two umpires can create a two percent difference.

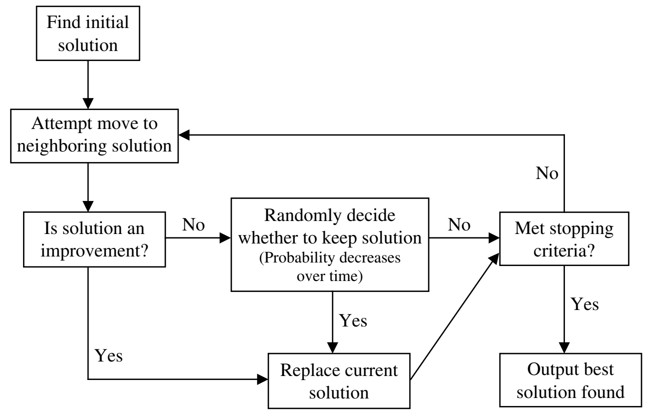

Figure 3: We use this process flow for our simulated-annealing algorithm. Results are within two percent of those we found using CPLEX.

Scheduler and Results

The umpire-crew-scheduling tool we developed is straightforward and simple to use. We structured the data inputs in steps that build on one another. Step 1 involves entering the raw data on an umpire's rating, gender, availability, and the number of courts to schedule. We maintain the umpire ratings between tournaments and update them only when an umpire's rating changes. In Step 2, the user selects the session to schedule. In Steps 3 and 4, the user specifies the number of teams to assign to the courts and the teams to schedule. In Step 5, the user identifies the team configurations (for example, chair-and-nine), the maximum and minimum constraints for the umpire ratings per court location, and the gender requirements. After setting up the required sessions in Steps 2 through 5, the user can manually assign a particular umpire to a particular court and position in Step 6. After finishing the setup, the user runs the scheduler, which produces an itemized list of the crews, the crew members' positions, and the court-rotation schedules for the crews.

Our scheduling tool improves on the USTA's sequential assignment system. (1) It considers the skill rating of the umpires and their locations on the court. The optimization process minimizes the difference between the skill rating of the umpire and the minimum desired rating for the umpire for a particular court location. (2) The scheduler ensures that the schedule meets gender requirements, if possible. If too

Tournament size (dollars)	Number of tournaments (per year)	Current manual system (hours per year)	Umpire-scheduling tool (hours per year)	Savings (%)
Large (>2.5 million)	6	378	35	91
Medium (75,000–2.5 million)	19	532	76	86
Small (<75,000)	229	1,603	515	68
Total	254	2,513	626	75

Table 2: Based on discussions with chief umpires of the four tennis governing bodies (USTA, Association of Tennis Professionals, International Tennis Federation, and Women's Tennis Association), we developed this table summarizing the impact of the umpire-scheduling tool. By using the umpire-scheduling tool instead of the preexisting manual system, the USTA can make time savings of 68 to 91 percent.

few male or female umpires are available to meet the gender requirements, the scheduling tool will provide the next-best solution. (3) The scheduler allows the chief umpire to manually assign umpires to particular positions within a crew while optimizing the remaining assignments. (4) The scheduler provides the chief umpire flexibility in umpire assignments, allowing the umpire to assign new umpires to difficult lines to give them on-court experience.

We have tested the umpire-scheduling tool at many professional tournaments, including the 2003 US Open and 2004 US Open in New York, the 2003 Nasdaq 100 in Miami, the 2003 RCA Championship in Indianapolis, and the 2004 Western/Southern Bank Financial Group Masters in Cincinnati. The user-friendly tool shaves hours off the umpire-scheduling process (Table 2).

Based on discussions with USTA officials, we categorized tennis tournaments as small, medium, and large according to the size of the prize pools and tournaments' duration. In general, we classified tournaments with prize pools less than $75,000 as small, between $75,000 and $2.5 million as medium, and greater than $2.5 million as large. Although more than 1,100 tournaments are scheduled worldwide, our scheduler has been used only for US tournaments. For the 2004 season, there were six large, 19 medium, and 229 small US tournaments. With the manual system, chief umpires spent an average of three hours scheduling for each day of a large tournament, two hours for each day of a medium tournament, and one hour for each day of a small tournament. They spend about 2,513 hours scheduling all US tournaments.

Because of a shared umpire database, setting up our program takes less than 30 minutes. Once we input umpire data into the database, we can readily access and use it to schedule tournaments. After setting up the program, users can schedule individual tournaments in about 15 minutes. We base our estimates of 30 minutes for setup and 15 minutes for scheduling on our test cases. The umpire-scheduling tool reduces the total time spent on scheduling tournaments by 68 to 91 percent. By using our system, umpires can divert about 75 percent of the time they anticipate spending on manual scheduling to managing tournaments.

Concluding Remarks

The three pillars of a successful professional tennis tournament are superior athletes, supportive fans, and qualified umpires. The umpire scheduler ensures that qualified umpires are assigned to facilitate a fairly and consistently called match to keep the athletes and the fans happy. We designed our scheduler to meet the umpire-scheduling requirements, and it is flexible enough to allow chief umpires to actively manage assignments. Once trained, the chief umpires can greatly reduce the time they allot to scheduling umpires and focus more on managing the matches.

USTA officials have inquired about extensions to the current version of our umpire scheduler. In the current problem formulation, we assume a single-session scheduling process because we assume that all sessions are independent of each other. Solving a multiple session problem would be the equivalent of solving a series of single-session assignment problems. However, we could incorporate some umpire-workload issues into the model as constraints, such as limiting the number of consecutive sessions an umpire could work or spend on one set of courts. In such cases, single-session models would be useful only if we could generate weights relating previous session assignments to ideal scenarios (Nicoletti 1975; Lucic and Teodorovic 1998, 1999). Another possible approach would be to schedule an entire tournament at once while incorporating workload constraints. This approach would increase the problem size and decrease the flexibility in scheduling because it would not treat umpires with equivalent ratings equally.

Furthermore, we assume that chief umpires have already determined rotation schedules. Although the same person creates the rotation schedules and the umpire assignments, he or she does not consider the two problems simultaneously. It might be possible to develop rotation schedules and assign umpires concurrently to maximize overall quality. An additional objective could be to minimize the number of umpires used or maximize the use of the available umpires while meeting a workload requirement.

Appendix

In our mathematical model of the umpire-scheduling problem, we use the following notation:

I = Total number of umpires available.
J = Maximum number of umpire positions required on a court.
K = Maximum number of crews on a team.
L = Total number of crew rotation teams.
T = Total number of sessions being scheduled.
h_{it} = 1 if umpire i is available during session t, 0 otherwise.
r_i = Skill rating of umpire i, the lower the better, generally between 1 and 8 or 1 and 5.
g_i = 1 if umpire is male, 0 if female.
n_{klt} = Number of umpires needed at position k for team l during session t.
y_{lt} = Number of crews on rotation team l during session t.
a_{jlt} = Minimum desired rating for umpire at position j on team l during session t.
b_{jlt} = Maximum allowable rating for umpire at position j on team l during session t.
f_{lt} = Number of female umpires required for crews on team l during session t.
m_{lt} = Number of male umpires required for crews on team l during session t.
x_{ijklt} = Decision variable: 1 if umpire i is assigned to team l, crew k, position j, and session t.
w_{jlt} = Weighting factor used to schedule less-skilled umpires.
ϕ_{klt} = Number of female umpire shortages.
μ_{klt} = Number of male umpire shortages.

The formulation consists of three types of decision variables. The primary decision variable is x_{ijklt}. This binary variable represents each umpire assignment. For each session t, it takes on a value of 1 if umpire i is assigned to team i and crew k and is working on line j. The variables or (μ_{klt} and ϕ_{klt}) represent the number of male and female umpire shortages on team l and crew k during session t.

Objective Function

The objective includes components for which the user can use a subjective weighting scheme to set priorities. The user can thus assign umpires with lower skill ratings (higher rating numbers) to some lines for training purposes instead of umpires with better skill ratings. Finally, all feasible schedules may have gender shortages. We use an arbitrarily high penalty factor of 100 in the objective function to heavily penalize shortages. The final objective function is

$$\text{Min}Z = \sum_{i}^{I}\sum_{j}^{J}\sum_{k}^{K}\sum_{l}^{L}\sum_{t}^{T} x_{ijklt} w_{jlt} \max(r_i - a_{jlt}, 0) + 100 \sum_{k}^{K}\sum_{l}^{L}\sum_{t}^{T} (\phi_{klt} + \mu_{klt}). \quad (1)$$

The most important part of the function is the maximum( ) component, which assigns umpires to positions so that skill ratings match target ratings exactly (that is, we would prefer $r_i - a_{jlt} = 0$). The maximum function eliminates the desirability of negative $r_i - a_{jlt}$ values while enabling the model to find a minimizing solution.

The objective function can be broken down into two parts. The first part calculates the total weighted deviation of assigned umpires' skill ratings from the target skill ratings for each position. Because the objective is to minimize, a positive weight factor forces an assigned umpire's skill rating to be as close to the target rating as possible. We use a maximum( ) function, which provides no advantage for assigning an umpire with a skill-rating number lower than the target.

On the other hand, a negative weight factor forces an assigned umpire's skill rating to be as far from the target as possible. By using these weight factors, the scheduler can control the assignment procedure.

The second part of the objective function concerns the shortage variables. Although we would like zero gender shortages, we sometimes have too few male or female umpires of the right skill levels to meet all gender requirements. In such cases, we still want to find a feasible schedule. So, rather than building a hard constraint into the formulation, we used the Big-M method and added the constraint to the objective function. By placing a large penalty factor on shortages, the minimization objective forces the model to minimize shortages. With a large penalty factor, the model will find a feasible solution if one exists. We don't sacrifice gender feasibility to reduce the weighted skill deviation.

In our formulation, the weighting factor, w_{jlt}, can take on positive or negative values, and it becomes a multiplier in the objective function. Negative values flip the objective function. That is, negative values of w make it attractive to increase the distance between r and a. Thus, they tend to lead to umpires with high skill-rating numbers (and lower skills) being assigned before those with low skill-rating numbers, facilitating their training.

Constraints

$$\sum_{j}^{J}\sum_{k}^{K}\sum_{l}^{L} x_{ijklt} \le h_{it}, \quad i=1,\ldots,I,\ t=1,\ldots,T. \quad (2)$$

We established a constraint to ensure that available umpires are assigned to at most one position during each session. If an umpire is not available during a session, the right-hand side of the constraint becomes zero, preventing the umpire's assignment during that period. If the umpire is available, the right-hand side takes on a value of one, making the umpire available for assignment to up to one position on one crew. The following constraint ensures the assignment of umpires in the exact amount required, n_{jlt}:

$$\sum_{i}^{I} x_{ijklt} = n_{jlt}, \quad j=1,\ldots,J,\ k=1,\ldots,y_{lt},\ l=1,\ldots,L,\ t=1,\ldots,T. \quad (3)$$

In combination, these two constraints ensure that umpires' assignments meet the line requirements of the rotation schedule without exceeding their availability. The next constraint ensures that every umpire assigned has a skill-rating number at or below the maximum number allowed for that position:

$$x_{ijklt} r_i \le b_{jlt}, \quad i=1,\ldots,I,\ j=1,\ldots,J,\ k=1,\ldots,y_{lt},\ l=1,\ldots,L,\ t=1,\ldots,T. \quad (4)$$

The following constraints ensure that each crew has at least the minimum number of female umpires required and at least the minimum number of male umpires required. They allow for a solution even if there are too few male or female umpires to meet the requirements. We denote male and female shortages with ϕ and μ, respectively:

$$\phi_{lkt} = f_{lt} - \sum_{i}^{I}\sum_{j}^{J}(1-g_i)x_{ijklt}, \quad (5a)$$

$$\mu_{lkt} = m_{lt} - \sum_{i}^{I}\sum_{j}^{J} g_i x_{ijklt}, \quad (5b)$$

where $l=1,\ldots,L$, $k=1,\ldots,y_{lt}$, $t=1,\ldots,T$.

References

Farmer, A. 2004. Automation of the umpire crew scheduling process for tennis tournaments. Master's thesis, Industrial and Systems Engineering Department, Auburn University, Auburn, AL.

Lucic, P., D. Teodorovic. 1998. Fuzzy set theory approach to the aircrew rostering problem. *Fuzzy Sets Systems* **95**(3) 261–271.

Lucic, P., D. Teodorovic. 1999. Simulated annealing for the multi-objective aircrew rostering problem. *Transportation Res.* **33**(1) 19–45.

Nicoletti, B. 1975. Automatic crew rostering. *Transportation Sci.* **9**(1) 33–42.

Schonberger, J., D. C. Mattfeld, H. Kopfer. 2004. Memetic algorithm timetabling for non-commercial sports leagues. *Eur. J. Oper. Res.* **153**(1) 102–116.

Urban, T. L., R. A. Russell. 2003. Scheduling sports competitions on multiple venues. *Eur. J. Oper. Res.* **148**(2) 302–311.

Yang, J. T., H. D. Huang, J. T. Horng. 2002. Devising a cost-effective baseball schedule using evolutionary algorithms. *Proc. 2002 Congress on Evolutionary Comput.* **2**(2) 1660–1665.

Rich Kaufman, director of officials, United States Tennis Association, 70 West Red Oak Lane, White Plains, New York 10604, writes: "I am writing to verify that the USTA worked closely with Dr. Smith and his students at Auburn University in the development of the Umpire Scheduling Application described in their research paper. We are pleased with the resulting software and were able to use the application at the 2004 US Open, following testing that took place at the tournaments specified in the research paper. The system performed as described and we expect to use it at future tournaments. We are also considering expanding the application to include additional features related to umpire scheduling and the general administrative functions of the chief umpire. For purposes of clarity, nothing contained in this letter shall be construed as the USTA's sanctioning, endorsement or official approval of the Umpire Scheduling Application for any purpose."