Interfaces
Vol. 42, No. 3, May–June 2012, pp. 286–295
ISSN 0092-2102 (print) | ISSN 1526-551X (online)

http://dx.doi.org/10.1287/inte.1110.0561

A Decision Support System for Scheduling the Canadian Football League

Kent J. Kostuk
Federated Co-operatives Limited, Saskatoon, Saskatchewan S7K 3M9, Canada, kent.kostuk@usask.ca

Keith A. Willoughby
Department of Finance and Management Science, Edwards School of Business, University of Saskatchewan, Saskatoon, Saskatchewan S7N 5A7, Canada, willoughby@edwards.usask.ca

The eight-team Canadian Football League (CFL) features an 18-game regular season played between late June and early November. The regular-season schedule is manually created using an iterative process involving league management, teams, and a television sports broadcaster. We developed a mixed-integer program-driven decision support system that provided league officials with multiple schedule versions relatively quickly. Our approach informed the development of the regular-season schedule by offering league management a neighborhood of potentially solid solutions. We demonstrated the impact of specific schedule requests to decision makers. Such sensitivity analysis was impossible under the previous manual creation method.

Key words: decision support; optimization; OR practice; scheduling; sports.
History: This paper was refereed. Published online in *Articles in Advance* February 2, 2012.

Professional sports are big businesses. Successfully managing a sports league typically involves negotiating television contracts, training referees and other officials, achieving and maintaining a reputable brand image, overseeing merchandising opportunities, and scheduling games. Indeed, the use of analytical operations research (OR) models to assist in the scheduling of a professional sports league's regular-season matches is clearly relevant to academics and practitioners (Kendall et al. 2010). Our paper illustrates the development of a decision support system to inform the creation of a regular-season schedule for the Canadian Football League (CFL), a popular Canadian sports entity.

In Canada, the first football teams emerged in the 1860s. The Canadian Rugby Football Union was established in 1884, and led to the formation of the CFL in 1958. As North America's oldest professional football league, the CFL currently includes eight franchises. The league is grouped into two equally sized partitions known as divisions. This split is organized geographically as a "Western Division" and an "Eastern Division." The Western Division includes teams in British Columbia, Calgary, Edmonton, and Saskatchewan; the Eastern Division comprises Winnipeg, Hamilton, Toronto, and Montreal (see Figure 1).

Although the current structure of the league features eight teams, a brief foray into the United States resulted in the creation of a 13-team CFL in 1995, with teams in such locations as Baltimore, Memphis, and Shreveport. The league existed with nine teams from 2002–2005 (the current eight franchises plus a squad in Ottawa), after which the Ottawa team ceased operations. Rumours persist that the Ottawa franchise will resurface for the 2011 campaign (Brennan 2009). The league is also exploring an expansion interest in the Maritimes, as witnessed by the scheduling of a 2010 regular-season contest to take place in Moncton, New Brunswick.

The CFL operates an 18-game regular season played annually between late June and early November. Each year, 72 regular-season games take place (i.e., eight teams playing an 18-game schedule, then dividing this total in half because two teams participate in each contest). Except when the Canadian July 1st national holiday falls during the middle of the week, or in very rare cases in which stadium availability proves particularly problematic, regular-season matchups are

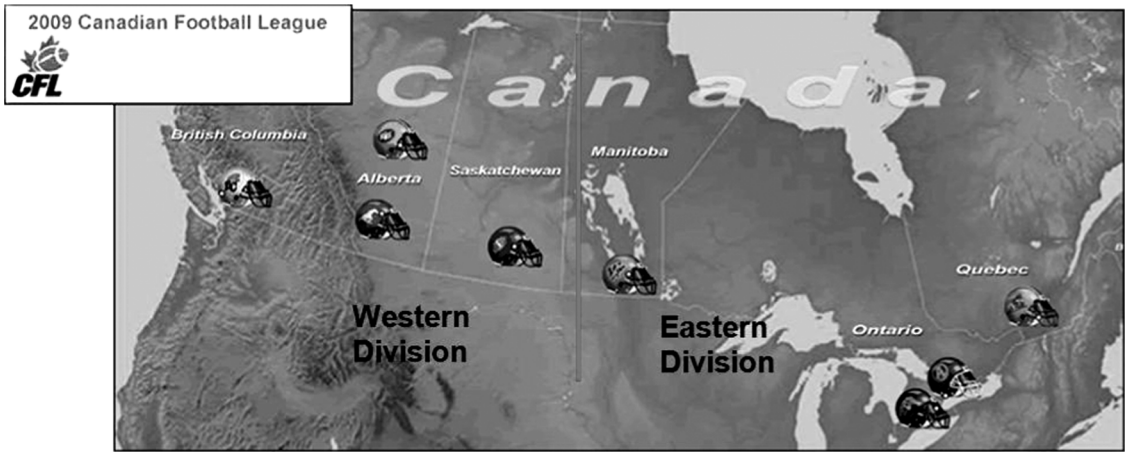

Figure 1: The CFL currently includes eight teams grouped into two equally sized partitions known as divisions. The helmets represent the logo and location of each franchise.

played exclusively between Thursday and Monday each week. Teams play one game per week, except for a period in August in which teams are given a "bye week" (i.e., a week with no matchups). The top six teams qualify for postseason competition. These play-offs subsequently occur over a three-week period, culminating in the awarding of the Grey Cup, North America's oldest professional football trophy. The Grey Cup broadcast is Canada's most-watched television event each year (The Sports Network 2009).

The regular-season schedule is manually created using an iterative process that involves league management, teams, and a television sports broadcaster. Specifically, a league manager uses a spreadsheet to slot each year's 72 required games into various days within specific weeks. This individual launches the process of creating a particular year's schedule about mid-September of the previous year. Typically, the manager is able to construct a manual version in one working day; his goal is to sift through his iterations and complete a reasonably suitable version by mid-December. The broadcaster and the respective teams subsequently offer feedback during the first two weeks of January; as a result, the manager might produce multiple changes to the schedule. A finalized version is generally released to the public at the beginning of February.

The problem of developing suitable schedules presents challenges for the CFL. In particular, league franchises submit specific days, known as "stadium blocks," in which they are unable to host games. This lack of availability could occur because of the reluctance of some franchises to host a game during times at which competing activities may stifle fan interest. Additionally, some teams play their home games in multiple-tenant facilities, thus suggesting that CFL management may have to revise schedules based on the requirements of other events, such as music concerts, professional baseball games, or trade shows.

Since 2008, a single Canadian broadcaster has televised each regular-season game. This serves to limit the number of games that may be scheduled on a specific day because showing two games simultaneously would be impossible. Moreover, time zone restrictions dictate that particular pairs of teams may be unable to simultaneously host games on a given day. For example, the league often schedules a doubleheader (i.e., two games) on Friday evenings; however, it would be infeasible for these three-hour contests to both take place in, for example, Calgary and Edmonton. Doing so would require that one game in this pairing begin exceptionally late, an unappealing scenario for teams and their fans.

League management must provide each team with an appropriate number of days off between games so as not to generate a competitive imbalance by forcing a team to play consecutive contests in quick succession. For example, scheduling Toronto to play on

Monday of week 1 and again on Thursday of week 2 would be unwise. Travel issues are also a consideration in setting up each year's schedule, although they are not a major driver. Unlike other professional sports, such as baseball or hockey, CFL teams do not undertake road trips in which multiple away games are combined within a single journey from a team's home location.

The league's 18-game regular-season schedule, split among eight teams, implies that teams will face some opponents more frequently than others. Franchises participate in 10 intradivisional games and 8 interdivisional games. Of the 10 games within its own division, teams will play one of their three opponents four times and face the remaining two squads three times each. For the eight games outside its division, a team will play each interdivisional opponent twice. Each team must play an equal number of home and away games. These structural requirements must be satisfied in each year's schedule.

Additional challenges result from the desire to schedule particular matchups around key holidays (e.g., Labour Day or Thanksgiving Day). The league also seeks to limit the number of times two teams face each other during consecutive weeks (i.e., back-to-backs); however, these instances are likely for an entity like the CFL, which has a relatively small number of teams. The league may also wish to slot particularly meaningful games (e.g., intradivisional contests) later in the season. This could ensure that the final weeks of the regular season involve matchups that are a determining factor in divisional play-off races, thereby intensifying fan interest.

Another hindrance to this process relates to human resource issues. Regular turnover within team management implies that different individuals may be interfacing with league officials on schedule creation on a year-to-year basis. This lack of continuity disrupts schedule development.

Some CFL franchises, especially those in the Prairie provinces, are prone to suffering extremely cold weather during the latter stages of the regular season. Given that these teams host games in outdoor facilities, league officials recognize that fan interest may be threatened by severe weather issues. This presents a further challenge to the scheduling of regular-season games.

Providing an OR-based approach to schedule development would be beneficial to CFL stakeholders. The current practice of manually developing each year's schedule—along with the concomitant lengthy times for incorporating various modifications—suggests that league officials may be hampered in terms of the overall number of iterations they could conceivably consider. This limitation may thwart the attempt to identify the best schedule. In addition, management may be unclear as to what constitutes a good schedule, because several competing criteria exist. Under the current schedule-development approach, identifying opportunities for improving the final product may be difficult.

The league's plans for future expansion provide another compelling reason why an analytical approach could be useful. Ratcheting up to an odd number of teams (e.g., from eight to nine) would precipitate an unbalanced weekly schedule in which not all teams would be assigned to play during a particular week. Developing regular-season schedules in this particular scenario may be beyond the current (manual) capabilities of the league's official scheduler.

An analytical approach could allow a manager to quickly generate multiple schedule versions, thus permitting league officials the access to these multiple versions more rapidly than they do using the current process. An OR model could also demonstrate the interaction between several scheduling criteria and illustrate the restrictions that may be driving solution infeasibility.

The remainder of our paper is structured as follows. The *Literature Review* section provides a review of pertinent sports-scheduling literature. We then develop the OR model in the *Method* section and follow this with an illustration of our experiences using this approach to inform the creation of the 2010 CFL regular-season schedule. We offer some concluding remarks in the *Conclusions* section.

Literature Review

Researchers have applied OR approaches in the scheduling of games for several professional sports. The ongoing interest in this field is perhaps driven by the sheer size of the economic outlays devoted to sports management or the growing power of

analytical methods to investigate increasingly complex sports-scheduling applications. The worldwide appeal associated with sports may also help to explain the popularity of this research. Whatever the reason, OR applications in sports scheduling have enjoyed a rich history over the past 40 years. Kendall et al. (2010) offer a thorough review of this research; we document a few applications in a variety of sports below.

To the best of our knowledge, no prior published work on the development of regular-season schedules for Canadian or American football is available. We note here that although American football bears a sketchy similarity in rules and style of play to its Canadian counterpart, the scheduling of games in the CFL represents a markedly more complex task than that of scheduling matchups for the National Football League (NFL), the organization charged with managing U.S. football. In the 32-team NFL, up to 15 of its weekly 16-game slate could be scheduled for Sundays (one game each week is traditionally scheduled for a Monday night telecast). In most cases, NFL teams will have exactly six days off between consecutive contests. In the CFL, games could be spread out among five days (Thursday through Monday), thus suggesting that greater care would need to be exercised to ensure that teams have comparable days off between games. A multiple number of American television entities means that NFL teams in the same time zone could easily cohost games, unlike the situation that the single Canadian football telecaster experiences.

The scheduling of European and South American football (called soccer in North America) is a fertile field for the application of OR methods. Della Croce and Oliveri (2006) developed a double round-robin tournament schedule for the Italian Major Football League. Their approach was able to balance specific team requests and cable television requirements. Duran et al. (2007) used an integer linear programming model to schedule Chile's professional soccer league. Chilean soccer officials implemented their final model because it created schedules that were more attractive to fans. The Danish football league features a triple round-robin tournament; Rasmussen (2008) illustrated an integer programming (IP) model to schedule games in this particular context. Bartsch et al. (2006) crafted an OR model for professional leagues in Austria and Germany. Kendall (2008) analyzed the intriguing problem of creating travel efficiencies over the Christmas and New Year periods for English football clubs. His approach resulted in schedules that obeyed all league restrictions but cut travel distances by 25 percent.

In a seminal piece, Nemhauser and Trick (1998) merged IP with enumerative techniques to schedule the double round-robin season for the Atlantic Coast Conference (ACC), a popular U.S. college basketball entity. The ACC implemented their approach in the 1997–1998 season. Devoting their attention to professional basketball, Bean and Birge (1980) used heuristics based on a travelling salesman problem to reduce airline travel costs.

Russell and Leung (1994) used heuristics to create the Texas Baseball League schedule. Using integer linear programming, Fleurent and Ferland (1993) explored the sport of ice hockey. In particular, they analyzed the scheduling impact associated with league expansion and a larger number of regular-season contests.

Willis and Terrill (1994) used simulated annealing to schedule Australian cricket matches. The schedule produced by their approach was subsequently fine-tuned using manual adjustments and implemented for the 1992–1993 season.

Formulating a model to develop the CFL schedule is somewhat different from the modeling efforts devoted to determining schedules for other sports. Whereas applications for such sports as hockey and European or South American football attempt to minimize total distance travelled, our approach does not consider these factors because CFL teams return to their home cities after a road game. Moreover, baseball and hockey teams have the flexibility to play multiple games spread out on any day of the week. In the CFL, each team plays at most one weekly game on any day between Thursday and Monday.

The above instances document examples that have been published in the academic literature. We note that commercial software providers have also developed scheduling applications that are being successfully implemented in both professional and amateur sports organizations. Examples of such companies include the Bortz Media and Sports Group, Optimal Planning Solutions, and the Sports Scheduling Group.

Apparently, the CFL had been approached in previous years by different organizations interested in applying computer-based procedures to developing the league's regular-season schedules. The league's reluctance to enter into agreements with such parties may have been driven by the costly nature (in the CFL's estimation) of these services. Our involvement with the league occurred in an unexpected fashion. Both authors are sports fans—one, a season ticket holder with a CFL team, had written a letter to league officials requesting clarification of the schedule creation process and indicating that an analytical approach could prove advantageous. When we indicated our willingness to develop an OR model simply "for the love of the game," we sparked the interest of league management.

Method

As is the case for any project, the tool selected is driven by the end goal. Our objective was to develop a regular-season schedule for the CFL. Unlike most scheduling problems, there was no explicit objective function. As we mentioned above, minimizing travel was not applicable because most teams return to their home cities after each game. We initially considered the following objective functions: maximizing the number of Friday games (the CFL and the television broadcaster have developed a compelling brand identity around its weekly "Friday Night Football") or minimizing the number of Sunday games (because the television network has a contractual obligation to broadcast one Sunday NFL game each week, and would therefore not want a CFL game to potentially interfere with its NFL commitment). We also contemplated maximizing the number of intradivisional games played in the last four weeks of the season because this could spark fan interest and strengthen play-off races. Based on the league scheduler's perception and our experience as football followers, we arbitrarily selected four weeks as a suitable time frame, judging that this represented a period at which media pundits, teams, and fans began to contemplate the heated races for postseason competition.

After much discussion, it became apparent that although the decision maker could not explicitly define an objective function, he could differentiate between a good schedule and a bad one. Both parties (i.e., the decision maker and the researchers) decided that a simple approach would be best. As more models were generated, the decision maker would become comfortable developing a schedule using mathematical methods, and we would become increasingly familiar with the attributes of good and/or acceptable schedules. We could then better incorporate that knowledge into the model. Ultimately, we wanted an OR tool that could capitalize on the decision maker's ability and experience while quickly generating new schedules so that alternatives could be compared and contrasted. Indeed, the understanding of both parties was enhanced as this work evolved; the researchers gained an appreciation of schedule intricacies and the decision maker gained a better understanding of mathematical modeling.

As we developed and evaluated additional alternatives, we were able to highlight promising or weak aspects of each solution. The decision maker had no previous experience with OR analysis, and we initially observed a dose of benign skepticism about the value of analytical modeling. This perspective ultimately transformed to an appreciation of how these models could contribute to effective decision making. For example, our modeling approach demonstrated the interplay between different performance criteria. League officials were able to observe that, for example, maximizing the number of intradivisional games scheduled over the final few weeks of the season could lead to a larger number of back-to-back contests. This strengthened their acceptance of the contributions of our analytical approach to developing schedules.

In another case, one team requested a substantial change in its available stadium dates at a relatively late juncture in the schedule creation process. Under the league's previous manual approach, this setback would have created frustration and a tremendous amount of rework as the schedules were overhauled. Using the OR approach, we simply incorporated the revised stadium dates and reran the model.

We had also provided the league with several schedule offerings; however, because both parties scrutinized the versions, we observed that some teams were allocated a large number of home or away games during each half of the season. Initially, a requirement to balance the home and away contests in the respective halves of the season was not included in our discussions with league officials. However, it became

apparent that the more promising schedules were the ones that provided such a game-allocation balance. Overall, our experiences reinforce the value of overall modeling flexibility combined with solid user interaction.

Ultimately, we determined that in combination with some basic structural constraints (i.e., 72-game regular season in which each CFL squad plays nine home and nine away contests), simply maximizing the number of games played was sufficient to find an initial solution. This initial solution would act as a base case to determine the sensitivity of the schedule to different objective functions.

With a clearly defined goal (i.e., to create a schedule) and a relatively straightforward objective (i.e., a feasible schedule), we selected a basic 0-1 IP formulation. The natural way to describe the problem was that team i would visit team j on day k. Mathematically, the decision variable $x(i, j, k)$ was set to 1 when this was true. Given this decision variable, the simple objective function was

$$\text{Max} \sum_i \sum_j \sum_k x(i, j, k).$$

As an aside, for the case in which we explore maximizing the number of Friday games, the objective function can be easily modified by summing over all dates (k) that are Fridays.

With the objective function defined, modeling constraints could be introduced. As the constraints were listed out, it became apparent that they could be broadly classified as: structural, stadium blocks, preassignments, and pattern assignments.

Structural Constraints

These constraints represent the basic logic of the schedule:

- The schedule comprises 19 weeks (from late June to early November);
- Games can be played Thursday through Sunday (except for the Mondays following two long weekends);
- A team cannot play against itself;
- All teams play nine games at home and nine games on the road;
- At least one game must be scheduled every Friday;
- Four games are played each week, except for weeks eight and nine ("bye weeks") in which only two games are slated;
- Teams play once per week;
- A maximum of two games can take place each day.

Expressing the structural constraints was relatively straightforward using our 0-1 decision-variable formulation. For example, the weekly Friday game commitment was accomplished by summing the decision variables over all team pairings (i, j) and forcing this total to be at least one for any date (k) that was a Friday. We guaranteed that two (or fewer) games could occur each day by summing the decision variables over all team pairings (i, j) and ensuring that the total was no more than two for any date (k).

Stadium Blocks

Each CFL franchise submits stadium blocks to league headquarters. These blocks represent a list of dates during the season in which any of three mutually exclusive outcomes could occur: the team prefers to play a home game (i.e., preferred dates), the stadium is not available (therefore, the team cannot be scheduled to play at home), and the stadium is available but the team would prefer not to play. For the 2010 campaign, CFL teams indicated stadium unavailability for about half of the playing dates throughout the regular season. The remaining dates were roughly equally split between preferred dates and those on which the stadium would be open, but the team preferred not to host a matchup.

Typically, stadium blocks present challenges to the league scheduler. Some franchises are prone to providing more constrained availabilities, thereby restricting scheduling flexibility. For each home team (j) and its set of dates (k) in which it could not host a game, we simply forced the sum over any other competitor (i) to be equal to 0.

Preassignments

In some cases, the CFL has predetermined which teams will play against each other on particular dates. For example, the league features several rivalries wherein specific pairings are an annual ritual at set times (e.g., Saskatchewan plays Winnipeg around the Labour Day weekend). These games are promoted heavily, help teams to maximize attendance, and enhance league ratings.

Pattern Assignments

These assignments involve pairings that follow a pattern but are not constrained to occur on specific dates or between specific teams. Examples of such assignments include:

- East Division teams visit West Division teams once during the season (and vice versa);
- teams that play on a Sunday do not play on the following Thursday or Friday (to ensure a sufficient number of days of rest before the next scheduled game);
- teams playing on a Monday do not play the following Thursday, Friday, or Saturday;
- teams do not play on the road more than two weeks in a row, or two weeks at home (exceptions could occur because of restrictive stadium availability);
- teams should not finish the season with a home-home or away-away pattern;
- when an East Division team plays a West Division opponent, the first game should occur prior to Labour Day (the Monday of this long weekend is used as a reference point to divide the season into two parts); the second game should take place after Labour Day;
- because the league comprises eight teams and the teams play an 18-game schedule, teams do not play each other an equal number of times; consequently there are asymmetric pairings (i.e., teams play some opponents twice, some three times, and some four times);
- because the league schedules many games on Thursday and Friday evenings, and would like to maximize the number of televised contests, it cannot allow certain pairs of teams to host a game on the same night because they are in the same (or adjacent) time zones.

As an example of conveying a pattern-assignment constraint in our model, consider the requirement that teams cannot play more than two consecutive weeks on the road. For each road team (i), we modeled this restriction by summing the decision variable over all home opponents (j) and dates (k) corresponding to a three-week period, and forcing this total to be no more than two.

Although classifying the constraints was unnecessary from an analytical perspective, it was appreciably invaluable from an organizational viewpoint. As we spent time with the decision maker, we were able to better communicate each other's requirements and develop a more sophisticated model. The classification system provided a basis for expanding and reevaluating the model. Initially, we introduced structural constraints. If the model could not be solved by only considering the structural constraints, this would indicate that the league's basic requirements were infeasible. Fortunately, this was never the case.

Once we established that the structural constraints provided for a feasible schedule, the next step was to introduce the stadium blocks. This permitted us to factor in the requirements of the teams and the facilities they shared with other tenants. Any infeasibilities identified at this time resulted in a review of team preferences and of stadium availability.

We next introduced preassignment constraints. Recall that these assignments are league-mandated. If their introduction induced infeasibilities, then trade-offs associated with these preassignments and the resulting impossibilities would need to be addressed. In practice, this was not an issue. The most likely reason the preassignments proved feasible is that most of these game pairings are long-held rivalries based on rich CFL tradition. Had they been problematic, they would not have survived the annual process associated with the scheduling efforts of previous years.

Admittedly, the pattern-assignment constraints were the ones most likely to generate conflicts with previously introduced restrictions. Given stadium availabilities, it was not always possible for a team playing on a Sunday to be given a rest period until the following Saturday. The requirement that specific team pairs play one of their two interdivisional games in each half of the season could be annulled by virtue of the league's desires to schedule a large number of intradivisional contests during the latter parts of the season. At this point, the decision maker must draw upon past experience and determine how best to trade-off these various inconsistencies. For example, he could use an alternative set of rules or constraints that would generate a schedule without compromising the intended pattern of games.

Informing the 2010 CFL Schedule

We began collaborating with the CFL on the 2010 schedule development in October 2009. This repre-

sented the time by which the league obtained stadium blocks from all eight franchises. We set up our decision support system as a 0-1 integer optimization model (consisting of 5,320 binary variables) using the mathematical programming language (MPL) modeling system with the 2008 version of FICO's XPressMP on a Lenovo ThinkPad T42–Pentium M735. In this MPL-XPressMP environment, we were able to solve the model in roughly two minutes, including the time required for data transfer activities. The actual optimization period was in the range of 30 seconds or less. These rapid solution times—a far cry from the one-day effort required under the league manager's manual approach—enabled us to feed multiple versions in quick succession to CFL headquarters.

We did not need to tweak any optimization parameters to solve the model. We relied on the XPressMP's presolve routine to identify conflicting combinations of constraints. When we addressed these constraints, we were able to solve the model to optimality. Reformulation of the model was unnecessary because we approached the solution in a hierarchical fashion. We structured the model in a hierarchy representing constraints sequenced from the fundamental "must have" to the "nice-to-have" restrictions. As we introduced infeasibilities by incorporating additional constraints, we could identify trade-offs and present them to the decision maker, who could then choose which constraints to include or exclude in future iterations.

When solving our models, we found it particularly helpful to initially optimize on a specific performance measure, for example, the number of intradivisional matchups during the season's final four weeks. After determining the optimal value for this criterion, we would subsequently set this value as a constraint and then reoptimize on another performance measure (e.g., the minimum number of Thursday games). This enabled us to identify how we could obtain the best performance on a specific measure while ensuring reasonably good outcomes on the first criterion.

Ultimately, our work with the CFL's official scheduler followed an iterative process. After we provided him with a revised version, we would deliberate with him about the benefits and drawbacks of the latest offering. Frequently, the manager would identify issues he wanted us to explore in the next set of versions. This drove our modeling efforts.

During an intense period between December 10, 2009 and January 20, 2010, we provided the official scheduler with 22 schedule versions (this total does not include several versions we produced during our earlier deliberations in October and November). The respective versions resulted from the modeling of different objective functions, or from incorporating additional constraints the league provided to us in our real-time iterative work. Examples of such constraints could be new requirements that the league wanted to include in a version (e.g., restricting Team A from hosting two home games in consecutive weeks) or revised stadium availabilities. During our deliberations, the importance attached to each team experiencing a rough balance of home and away games during either half of the season became readily apparent. League officials were critical of any version that provided a team with, for example, six home games during its first nine matchups.

Our versions, although feasible according to the constraints included in a particular model, occasionally delivered "unintended consequences" that had to be fully addressed before the league could submit the version to its teams and the broadcaster. As a case in point, we produced one version in which Team A played Team B during weeks two and four of the season. Although this was not a back-to-back matchup (both teams played another opponent in week three), the league perceived that this game sequence was unnecessarily repetitive. It could also stifle fan interest, especially for Team B's followers who would be exposed to the same visitor in two consecutive home games. Issues such as these underscored the value of continually communicating with league officials during the schedule creation process.

Table 1 compares eight of our versions to the actual 2010 schedule based on six respective performance measures. We chose these various versions to demonstrate the range of options we were able to offer the decision maker.

The six criteria include measures by which league management evaluates schedule adequacy. The CFL preferred to reduce the number of "back-to-back" contests as much as possible, although it recognized that a few such matchups were bound to occur because of preassignments involving specific teams around holiday periods. Optimizing the number of intradivisional

Performance measure	Final 2010 schedule	Versions produced using OR model							
		A	B	C	D	E	F	G	H
Number of "back-to-back" games	9	9	8	12	16	8	7	9	9
Intradivisional games during final four weeks	12	16	16	14	14	10	12	14	10
Thursday games	6	7	8	6	8	7	7	6	4
Friday games	23	24	25	25	23	26	29	28	29
Sunday games	13	10	8	10	13	10	14	19	10
Number of games played on preferred dates	54	62	60	59	56	54	51	55	54

Table 1: This table represents a sampling of the schedule versions we provided to league management. Using six different criteria, we compare the performance of each version to the actual 2010 CFL schedule. Trade-offs become evident. For example, the league could reduce the number of back-to-back matchups (see version F); however, this would necessitate fewer preferred dates and more Friday games.

games during the season's final four weeks (up to a maximum of 16) was a laudable goal. The league favored a relatively low number of Thursday games, while ensuring that the number of contests on Fridays and Sundays were at reasonable levels. Finally, the league and its teams wanted to schedule as many of the 72 games as possible for preferred dates.

The luxury of assessing a fair number of acceptable versions was appealing to the league scheduler. Trade-offs became evident with our different versions. Achieving benefits along one dimension came at the expense of another measure.

To illustrate, consider the following examples. In each case, we will compare our version to the actual 2010 schedule. Should the league want to increase the number of games scheduled on preferred dates (see version A in Table 1), this would have the fortunate outcome of improving the intradivisional game allotment; however, it would require one more Thursday game, an additional game on Friday, and a reduction in the number of Sunday contests.

The league could lessen the number of back-to-back matchups (see version F in Table 1), but this would necessitate fewer preferred dates and more Friday games. In version H, the league could reduce the number of Thursday contests, but this would entail fewer intradivisional matchups, more Friday games, and a drop in the number of Sunday contests.

A further example of our approach's advantages involved its ability to inform league management of proper courses of action to mitigate several stumbling blocks that emerged as the league began composing the schedule. Because of stadium availability issues, one franchise (Montreal) was forced to begin the 2010 season with three consecutive road games. League management duly recognized that Montreal's situation was inevitable. However, it was concerned that such a scenario would oblige another team to start its season with three consecutive home games. This would entail a team playing one-third of its entire home game complement by the middle of July, hardly a result that would engender ongoing fan interest throughout the remainder of the season. By virtue of our model's results, we demonstrated that a team would be forced to start its season with a three-game home stand, given the current (stringent) stadium availabilities throughout the league. Only by obtaining relaxed stadium availabilities could the number of home games be sufficiently distributed during the first three weeks of the season so that no team would start with three consecutive home contests. The league used this finding to request various franchises to submit more flexible stadium availabilities for the start of the season.

One team asked to be scheduled for a road game during a particular week. We were able to illustrate the trade-offs that would have to be made—assuming current schedule restrictions—if the league acquiesced to this team's request. In this instance, it would have yielded the inadvertent consequence of compelling another team to play home games in three straight weeks. The official scheduler used this knowledge to negotiate more favorable game assignments (from a league perspective) during this particular week of the season.

In another case, one franchise was scheduled to play six road games over an eight-week period (this involved a repeating pattern of two consecutive road

games followed by a home contest). Through our modeling efforts, we were able to show league officials that such a situation could be avoided by negotiating with another team to "free up" its restrictive home-date desires by being willing to host one of its games on a Thursday.

Overall, we were able to depict the features one would have to sacrifice to achieve particular benefits. It was then up to league officials to determine if the benefits of a specific version justified the concomitant costs. Eventually, the league considered our various versions and subsequently fine-tuned them manually to create the final 2010 schedule. We anticipate continuing our relationship with the league in its development of the 2011 regular-season schedule.

Concluding Remarks

We designed a rigorous OR model to inform the CFL on its 2010 schedule development process. Exploring an admittedly messy problem, we helped a notable Canadian sports entity gain an enhanced perspective on the interaction between different restrictions. We introduced various objective functions so the decision maker could obtain versions that optimized respective criteria.

This work proved especially valuable to league managers because it provided them with multiple acceptable versions that they could consider in crafting the final product. Obviously, regardless of our assistance, the league would have generated and released an actual schedule. However, our analysis reduced the labour hours required by the scheduler to create and disseminate the final schedule. In previous seasons using his manual schedule creation method, he readily admitted that he would only end up with two or three suitable versions that could be circulated to the teams and television broadcaster for further scrutiny. With our approach, he enjoyed the luxury of several acceptable versions that league stakeholders could evaluate. Our experiences with the CFL demonstrate the successful role that OR modeling can play in informing actual practice.

References

Bartsch, T., A. Drexl, S. Kröger. 2006. Scheduling the professional soccer leagues of Austria and Germany. *Comput. Oper. Res.* **33**(7) 1907–1937.

Bean, J. C., J. R. Birge. 1980. Reducing travelling costs and player fatigue in the National Basketball Association. *Interfaces* **10**(3) 98–102.

Brennan, D. 2009. Rough time over riders. Accessed April 26, 2010, http://slam.canoe.ca/Slam/Football/CFL/Ottawa/2009/04/25/9244391-sun.html.

Della Croce, F., D. Oliveri. 2006. Scheduling the Italian Football League: An ILP-based approach. *Comput. Oper. Res.* **33**(7) 1963–1974.

Duran, G., M. Guajardo, J. Miranda, D. Sauré, S. Souyris, A. Weintraub, R. Wolf. 2007. Scheduling the Chilean Soccer League by integer programming. *Interfaces* **37**(6) 539–552.

Fleurent, C., J. A. Ferland. 1993. Allocating games for the NHL using integer programming. *Oper. Res.* **41**(4) 649–654.

Kendall, G. 2008. Scheduling English football fixtures over holiday periods. *J. Oper. Res. Soc.* **59**(6) 743–755.

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

Nemhauser, G. L., M. A. Trick. 1998. Scheduling a major college basketball conference. *Oper. Res.* **46**(1) 1–8.

Rasmussen, R. V. 2008. Scheduling a triple round robin tournament for the best Danish soccer league. *Eur. J. Oper. Res.* **185**(2) 795–810.

Russell, R. A., J. M. Leung. 1994. Devising a cost effective schedule for a baseball league. *Oper. Res.* **42**(4) 614–625.

The Sports Network. 2009. Grey Cup delivers a record audience of 6.1 million viewers. Accessed April 26, 2010, http://www.tsn.ca/cfl/story/?id=300593.

Willis, R. J., B. J. Terrill. 1994. Scheduling the Australian state cricket season using simulated annealing. *J. Oper. Res. Soc.* **45**(3) 276–280.

Michael Copeland, Chief Operating Officer, Canadian Football League, Toronto, Ontario, Canada, writes: "This letter is to verify the authenticity of the work that Dr. Kostuk and Dr. Willoughby describe in their *Interfaces* manuscript entitled 'A Decision Support System for Scheduling the Canadian Football League.'

"These two individuals developed a decision support system that informed the creation of our 2010 CFL regular season schedule. During our schedule deliberation process, they provided us with multiple versions in considerably short turnaround times. As we observed in previous years, creating feasible schedules by hand or modifying versions to incorporate additional requirements was a day-long exercise. Because of the automated process developed by Dr. Kostuk and Dr. Willoughby, multiple schedule variations could be turned around in minutes. Further, their approach demonstrated the impact of enforcing particular schedule restrictions. Such knowledge and understanding would not have been possible under our previous schedule creation method.

"Although we cannot attach a specific dollar value to this work, their model did improve the efficiency with which our schedule was created. Further, we gained additional perspectives on schedule development that truly benefited our overall process."