Mathematically Clinching a Playoff Spot in the NHL and the Effect of Scoring Systems

Tyrel Russell and Peter van Beek

Cheriton School of Computer Science
University of Waterloo

Abstract. A problem of intense interest to many sports fans as a season progresses is whether their favorite team has mathematically clinched a playoff spot; i.e., whether there is no possible scenario under which their team will not qualify. In this paper, we consider the problem of determining when a National Hockey League (NHL) team has clinched a playoff spot. The problem is known to be NP-Complete and current approaches are either heuristic, and therefore not always announced as early as possible, or are exact but do not scale up. In contrast, we present an approach based on constraint programming which is fast and exact. The keys to our approach are the introduction of dominance constraints and special-purpose propagation algorithms. We experimentally evaluated our approach on the past two seasons of the NHL. Our method could show qualification before the results posted in the Globe and Mail, a widely read newspaper which uses a heuristic approach, and each instance was solved within seconds. Finally, we used our solver to examine the effect of scoring models on elimination dates. We found that the scoring model can affect the date of clinching on average by as much as two days and can result in different teams qualifying for the playoffs.

1 Introduction

As a season progresses, sports fans become intensely focused on the playoff race and the position of their team in the standings. Sports sections of major newspapers publish the results of the game and announce when teams have qualified for the playoffs and when they have been eliminated (e.g. the Globe and Mail [6]). However, the newspapers use a heuristic measure for determining when teams have qualified for or been eliminated from the playoffs and announcements are sometimes not made until several days after the team has clinched or been eliminated. Using optimization techniques like constraint programming, it is possible to determine exactly when teams have qualified for the playoffs.

In this paper, we present the first complete solution to the problem of determining when a National Hockey League (NHL) team has qualified for the playoffs. The problem of determining when a sports team has mathematically clinched a playoff spot has been well studied for several sports, including baseball [14, 12, 16, 1] and soccer [11]. Schwartz [14] first looked at this problem algorithmically for the historical baseball problem. He showed that this simple model can be solved with a polynomial algorithm. More recently, research has shown that a wide variety of qualification problems are NP-Complete, including the NHL problem [8]. Cheng and Steffy [4] looked at the problem of determining qualification for the NHL using an integer programming (IP) model but they found that they could not solve the model when secondary and tertiary tie breaking rules were applied. We model the problem as a constraint satisfaction problem and solve the problem using a constraint programming solver based on backtracking search and constraint propagation. By introducing dominance constraints and special-purpose propagation algorithms, our model can handle both secondary and tertiary tie-breaking constraints. Our approach is efficient and can solve all of the instances from the past two seasons within seconds.

The NHL is composed of thirty teams that are broken into two conferences, East and West, of fifteen teams. Each of these conferences is composed of three divisions with five teams in each division. Teams play eighty-two games spread unevenly between division, conference and inter-conference opponents. Each game in the NHL is formed from 3 periods, called regulation time, and, if the game is tied at the end of regulation time, a shorter overtime period and a shootout. The shootout continues until one team has won the game. A team makes the playoffs if they are a division leader or one of the top five teams that are not division leaders in their conference.

In the NHL, teams are placed in the standings by the number of points, for both divisional and overall standings. However, it is possible to have two teams with the same number of points. The NHL has introduced three different tie breaking measures. The first tie breaking measure is to compare the number of wins by each team. If the teams are still tied, the number of points earned against only those teams that are tied are compared. The third measure compares the number of wins earned against teams that are still tied.

The NHL qualification problem is defined with respect to a set of completed games, the number of games left between the teams, the scoring model used, the points accrued so far by each team and the number of wins. Given these factors, the NHL qualification problem is the problem of determining if a team k has qualified for the playoffs. We can find a solution to this problem, known as a wild-card qualification problem, by determining if there is a possible scenario for the remaining games that, according to the scoring model, awards sufficient points to teams other than k to ensure that k cannot be one of the division leaders or have enough points to earn one of the five wild card spots. If there is no scenario then k has qualified for the playoffs.

The calculation of the points for a team is dependent on the scoring model that is used. The NHL has used several different scoring models over the years. A scoring model is defined as a set of tuples defining the possible outcomes, and subsequent reward, of the games. Each of the NHL scoring models can be viewed in Table 1. Referring to Table 1, the current or Shootout model allows for two pairs of outcomes. Either one team wins in regulation time and earns two points while the other receives none or one team wins in overtime and earns two points while the other receives one point.

Table 1. The scoring models that have been used by the NHL. The possible outcomes of an NHL game between two teams (i, j) are A) j wins in regulation time, B) j wins in overtime, C) tied game, D) i wins in overtime and E) i wins in regulation time.

Scoring Model	A	B	C	D	E
Historic Era	{ $(0, 2)$		$(1, 1)$		$(2, 0)$ }
Overtime	{ $(0, 2)$	$(0, 2)$	$(1, 1)$	$(2, 0)$	$(2, 0)$ }
Extra Point	{ $(0, 2)$	$(1, 2)$	$(1, 1)$	$(2, 1)$	$(2, 0)$ }
Shootout	{ $(0, 2)$	$(1, 2)$		$(2, 1)$	$(2, 0)$ }
Proposed	{ $(0, 3)$	$(1, 2)$		$(2, 1)$	$(3, 0)$ }

We experimentally evaluated our constraint programming approach on instances from the 2005-06 and 2006-07 seasons. For these seasons, we can show qualification of teams up to five days earlier than the Globe and Mail [6]. As well, we studied all of the different scoring models that have been used by the NHL as well as a model recently proposed at the 2006 meeting of general managers of the NHL. The NHL has tried various scoring models to increase the competitiveness of the league though there exists no data on how these changes affected the clinching dates. We show that adjusting the scoring model can change the average qualification date by more than two days. Interestingly, the Toronto Maple Leafs would have qualified for the playoffs under any other model than that used by the NHL in 2006-07 season. Another example where the scoring matters is the Edmonton Oilers in 2005-06 who were Western conference champions but would not have made the playoffs if the proposed scoring model had been used.

The remainder of the paper breaks down the various parts of the research. Section 2 introduces the constraint programming model for solving the NHL elimination problem. Also in this section, the specific dominance constraints and custom propagation algorithms are introduced. Section 3 presents the experimental results from our study. Section 4 discusses some concluding remarks.

2 The Model

In this section, we describe the constraint programming model that we used to solve the NHL qualification problem. We also introduce the techniques that we used to make the model solvable for realistic problem sizes.

Constraint programming is a technique for solving combinatorial problems. A problem is defined as a set of variables with a given domain of values and constraints on the values that those variables can take. The goal of constraint programming is to find an assignment of values to those variables such that each of the constraints is satisfied. Typically, these problems are solved with backtracking search or other search techniques that explore the possible set of solutions. A more detailed explanation of constraint programming can be found in either Rossi, van Beek and Walsh [13] or Marriott and Stuckey [10].

A constraint programming approach can often be dramatically improved by adding dominance constraints and special purpose constraint propagation algorithms. Dominance constraints are constraints that remove feasible solutions that can be shown to be equal to or worse than another feasible solution [15]. One application of these constraints is in online scheduling for photo-copiers [5]. Constraint propagation is the process of inference by which domains of a variable is updated due to the constraints on that variable and the values of other variables referenced by the constraints [3]. Laborie [9] presents one example where a special purpose propagation algorithm is used in scheduling.

2.1 Basic Notation

To present our constraint programming approach, we first define some notation. We define a set of teams T with $|T| = n$. For any team $i \in T$, we define the number of points to be p_i and the number of wins to be w_i. We define the initial points—that is, the points earned before a given moment in the season—as p_i^0 and initial wins as w_i^0. We define the number of wins by a team i against a team j as a w_{ij}. The number of games remaining for a team i is denoted g_i and between two teams i and j is denoted g_{ij}.

We define C_i to be the set of teams in the conference that includes i and $C_i = C_j$ if and only if i is in the same conference as j. We define D_m to be the set of teams in division m, where m is one of the six divisions.

2.2 The Basic Model

To specify a constraint programming model, one must specify the variables and their associated domains along with the constraints on the values those variables may take. The variables used in our constraint model, their domains and a short description are shown in Table 2. More complete descriptions of the variables are presented as necessary through the remainder of this section. The following model is described with respect to the current or Shootout scoring model. Small modifications are made to make this constraint model work for the other scoring models. For the remainder of this section, we will refer to the team being tested for qualification as k.

The most basic variable in our model is x_{ij} with domain equal to $[0 \ldots g_{ij}]$ to represent the number of wins by a team i over a team j. Cheng and Steffy [4] observe that if elimination is possible then it always occurs when the worst possible outcome happens to k. Under the current scoring model and with the notion that an elimination solution is found in the worst case for k, we know that a team will always win games against k earning at least one point and, when playing other teams, the worst case situation is for both teams to earn at least one point [4]. Using this observation, we note that it is sufficient under this model to keep track of the number of wins to calculate the number of points earned.

Given that there is a symmetric variable for wins by j over i with the same domain, we define the following constraint on the x variables,

$$x_{ij} + x_{ji} \leq g_{ij}. \tag{1}$$

Table 2. Variables within the constraint model.

Variable	Domain	Description
x_{ij}	$[0 \dots g_{ij}]$	The number of wins by i over j
b_i	$[0, 1]$	1 if i is better than k else 0
d_m	D_m	The division leader for division m
div_m	$[0, 1]$	1 if the leader of division m is not better than k else 0

Note that constraint 1 is an inequality. The reason for this is to simplify the addition of dominance constraints.

We must represent the quantities used in the elimination problem in terms of our decision variables. We define the number of wins by a team i against j as,

$$w_{ij} = w_{ij}^0 + x_{ij}, \tag{2}$$

and the total number wins for a team i as,

$$w_i = \sum_{j \in T} w_{ij}. \tag{3}$$

We can also represent the number of points earned by a team i in a similar manner. One exception is team k, which must not win any games. Therefore, team k does not earn any more points. The number of points earned by a team i against j is denoted as,

$$p_{ij} = \begin{cases} p_{ij}^0 + x_{ij} + g_{ij} & i \neq k \\ p_{ij}^0 & i = k \end{cases}, \tag{4}$$

and the number of points earned by a team i as,

$$p_i = \begin{cases} p_i^0 + \sum_{j \in C_i}(x_{ij} + g_{ij}) + 2\sum_{j \notin C_i} g_{ij} & i \neq k \\ p_i^0 & i = k \end{cases}. \tag{5}$$

Now that we have defined the wins and points in terms of x_{ij}, we can look at the constraints on the variable in terms of the qualification problem. The first constraint that can be derived is that k should win no games in the worst case solution. This can be denoted as,

$$\forall i \;\; w_{ki} = 0. \tag{6}$$

The objective function that we are using in this case is the wild card objective, so in order to determine if a team can be eliminated we must keep track of the number of teams better than k. We define a binary variable b_i to be true if and only if team i is better than k. A team is better than k if the team has more points than k or better tie breakers than k. Before we show this constraint, we have to define two sets that are used in the tie-breaking procedure. The secondary tie-breaking set is the set of teams that have identical numbers of points and wins and is denoted as,

$$TB_i^2 = \{j \mid j \in C_i \wedge p_i = p_j \wedge w_i = w_j\}.$$

The tertiary tie-breaking set is the set of teams that are in the secondary tie-breaking set and have identical number of points earned between teams in the secondary set. The tertiary tie-breaking set is defined as,

$$TB_i^3 = \{j \mid j \in TB_i^2 \wedge \sum_{y \in TB_i^2} p_{iy} = \sum_{y \in TB_i^2} p_{jy}\}.$$

Once we have the tie-breaking sets defined, we can define the constraint between b_i and x_{ij} as,

$$\begin{aligned}
\forall i \quad & p_i > p_k \\
& \vee \ (p_i = p_k \wedge w_i > w_k) \\
& \vee \ (p_i = p_k \wedge w_i = w_k \wedge \sum_{j \in TB_k^2} p_{ij} > \sum_{j \in TB_k^2} p_{kj}) \\
& \vee \ (p_i = p_k \wedge w_i = w_k \wedge \sum_{j \in TB_k^2} p_{ij} = \sum_{j \in TB_k^2} p_{kj} \wedge \\
& \quad \sum_{j \in TB_k^3} w_{ij} > \sum_{j \in TB_k^3} w_{kj}) \\
& \Leftrightarrow b_i = 1.
\end{aligned} \tag{7}$$

A division leader is a team that has the most points in its division and beats any other division member with sufficient points on tie-breakers. We define a variable d_m with domain D_m to be the division leader for division m. We define the maximum number of points in a division m to be $maxp_m$. For conciseness, the secondary and tertiary tie-breakers are omitted for constraints 8, 10, 13 and 14. We define the constraint between x_{ij} and d_m as,

$$\forall m \ \ [p_i = maxp_m \wedge \nexists j (j \neq i \wedge p_j = maxp_m \wedge w_j > w_i) \Leftrightarrow d_m = i]. \tag{8}$$

We need to ensure that the team that must be eliminated does not become a division leader as this guarantees a playoff spot even if there are many teams with more points. Therefore, we add the following constraint,

$$\forall m \ \ d_m \neq k. \tag{9}$$

Once we have division leaders and none of them are k, we need to determine if we need to count that team in elimination, by setting $div_m = 1$, as a division leader always makes the playoffs. The division constraints ensure that a div_m variable is true only if the division leader is not better than k. There are two situations where div_m is true. The first is when the division leader has less points than k and the constraint is denoted as,

$$\forall m \ \ maxp_m < p_k \Rightarrow div_m = 1. \tag{10}$$

The second situation is when the division leader has the same number of points as k but is placed lower on tie-breakers. The constraint ensuring this situation is defined as,

$$\forall m \;\; maxp_m = p_k \wedge w_{d_m} < w_k \Rightarrow div_m = 1. \tag{11}$$

The last constraint on the model ensures that the team is eliminated under the wild card objective. There must be eight teams in front of k either by placing better than k or by being a division leader. The following constraint captures this relationship,

$$\sum_{i \in C_k} b_i + \sum_{m=1}^{\ell} div_m \geq 8, \tag{12}$$

where ℓ is the number of divisions.

The basic model is a complete and correct model for the NHL qualification problem. If there does not exist a solution to the CSP then k has qualified for the playoffs. However, the model as it stands is inefficient and could not solve realistic instances within a time bound of fifteen hours. Once we improved the model using dominance constraints and special purpose constraint propagation algorithms, we could solve all of the instances within seconds.

2.3 Dominance Constraints

It can be observed that if a team has earned more points than k then it is unnecessary for that team to win any more games because under the worst case assumption k will not win any more games and thus earn no more points. This means that once a team has passed k either by points or tie-breakers, that team could lose the remainder of their games to other opponents and still be better than k. This leads to the following dominance constraint,

$$\begin{aligned} &\forall i \; p_i > p_k \vee (p_i = p_k \wedge w_i > w_k) \Rightarrow \\ &\qquad \forall_j [(\neg bound(x_{ij}) \Rightarrow x_{ij} = 0) \wedge (x_{ji} = g_{ij} - x_{ij})]. \end{aligned} \tag{13}$$

By a similar token, if a team can no longer earn more points than k then there is no need for that team to win any more games against opponents that could earn more points than k. Therefore, we can introduce another constraint that deals with this situation. But first, we introduce notation for maximum number of points and wins possible for a team i. We denote the maximum number of possible points and possible wins as mp_i and mw_i, respectively. The dominance constraint is given by,

$$\begin{aligned} &\forall i \; mp_i < p_k \vee (mp_i = p_k \wedge mw_i < w_k) \Rightarrow \\ &\qquad \forall_j [(\neg bound(x_{ij}) \Rightarrow x_{ij} = 0) \wedge (x_{ji} = g_{ij} - x_{ij})]. \end{aligned} \tag{14}$$

2.4 Calculating Tie-Breakers and Division Leaders

To calculate the sets needed for both the tie-breakers and division leaders requires the values for x_{ij} to be set for all teams. However, it is often possible to determine the result of the calculation earlier and avoid calculating the sets. This means that special purpose propagation algorithms must be used to ensure constraints are propagated as early as possible.

Observe in constraint 7, secondary and tertiary tie-breakers need the sets of teams tied after primary and secondary tie-breaking, respectively. However, it is not always impossible to calculate these quantities during search and thus these tie-breakers can not be fully resolved during search. When either the points or wins changes for a team i, instead of waiting for all x variables to be completely set to propagate values, the solver uses only the first two disjunctive clauses to determine if b_i values should be set. In many cases, this negates the need to calculate the sets explicitly. Any decisions on teams still tied are delayed until all of the x variables have been assigned values, if necessary. If there are still teams tied, the secondary and, if needed, the tertiary tie-breakers are applied.

Observing constraints 9, 10 and 11, it can be seen that it is only necessary to know which team is a division leader if the elimination team is possibly a division leader or it is possible for a division leader to have an equivalent or lesser number of points. That means that in all other situations we can avoid calculating which team is actually a division leader. This also means that anytime a team in a given division acquires more points than the elimination team, the division leader will not be a factor. This relationship can be described as the following constraint,

$$\exists i \;\; i \in D_m \wedge i \in C_k \wedge (p_i > p_k \vee (p_i = p_k \wedge w_i > w_k)) \Rightarrow div_m = 0. \quad (15)$$

While constraint 10 uses the maximum number of points earned by any team in a given division, it is also possible to look at a more relaxed version of this constraint. Instead of determining the maximum number of points once all of the games have been assigned, it is also possible to determine the maximum possible number of points earned by any team in that division at any time. We denote this quantity mp_m. Using this new quantity, we can reformulate a relaxed version of constraint 10 that can be used at any point during search and is defined as,

$$\forall m \;\; mp_m < p_k \Rightarrow div_m = 1. \quad (16)$$

2.5 Different Scoring Models

The previous section described the constraint model for the NHL's current scoring model. Only minor changes must be made to adapt the other models that have been used by the NHL. Two new variables must be introduced to deal with ties for the historic and overtime models and with overtime wins for the proposed model. Also, the calculation of the wins and points must be modified to deal with the new variables when appropriate.

3 Experimental Results

3.1 Overview

We implemented our model using the constraint programming solver built by ILOG [7]. To test the model and to ensure that it can solve practical instances, we tested the solver on all possible elimination instances from the 2005-06 and 2006-07 seasons. We compared our calculated results against those shown in the Globe and Mail [6]. Finally, we modified the constraint model to handle all of the scoring models that have been used by the NHL and compared the results.

3.2 Problem Instances

To generate the qualification problem instances, we took the results from the 2005-06 and 2006-07 seasons and broke down the results into separate days. To create instances from the results for each day of the season, we calculated the number of points and wins earned up to that point in the season. Given the win-loss records, the points accrued so far and the remaining games, we created instances for each day of the season.

3.3 Comparison against the Globe and Mail

We ran each of the problem instances for the 2005-06 and 2006-07 seasons and determined the date of qualification for each of the teams. We could determine for every team whether they had qualified for the playoffs at every point in the season in just under two and half minutes. For most instances, the time it took was a fraction of a second. We found that, when comparing the results to the results posted in the Globe and Mail, we were able to show qualification earlier for nine teams during the 2005-06 season and for four teams during the 2006-07 season while never, of course, being later than the Globe and Mail. The results of this experiment are shown in Table 3. In Table 3, entries with multiple dates are due to Sunday editions or unreported results in which it was unclear on which day the result would have been reported.

We show qualification earlier than the Globe and Mail by as much five and four days for the 2005-06 and 2006-07 seasons, respectively. Note that the largest discrepancy between our results and their results is found for the first team to qualify in both seasons. It is possible that this is due to the paper not realizing that teams had clinched playoff spots. Even disregarding these results, we still found discrepancies as large as four days and two days for the 2005-06 and 2006-07 seasons, respectively.

3.4 Comparisons of Scoring Models

Once the solver has been designed for one scoring model, it is relatively straightforward to translate the model to other scoring systems. The instance generator was modified to calculate the points correctly for each of the different scoring

Table 3. The results of qualification compared against the published Globe and Mail results. Only results that differ are shown.

2005-06			2006-07		
Team	Optimal	Globe and Mail	Team	Optimal	Globe and Mail
Ottawa	Mar 22	Mar 27	Buffalo	Mar 18	Mar 22
Montreal	Apr 14	Apr 18	Atlanta	Apr 2	Apr 4
Buffalo	Apr 4	Apr 5	Detroit	Mar 24	Mar 25 or 26
New Jersey	Apr 12	Apr 14	Nashville	Mar 23	Mar 24
Calgary	Apr 8	Apr 9 or 10			
Colorado	Apr 13	Apr 15			
San Jose	Apr 13	Apr 14			
Dallas	Mar 31	Apr 1, 2 or 3			
Nashville	Apr 9	Apr 10			

Table 4. Average days remaining by scoring system for the 05-06 and 06-07 seasons

	Average Days Remaining	
Scoring Model	2005-06	2006-07
Historic	12.75	11.75
Overtime	11.56	12.81
Extra Point	11.75	10.63
Current	11.94	10.75
Proposed	13.13	10.13

models and we ran the instances for all of the scoring models used by the NHL as well as the recently proposed model. The results of this experiment can be found in Table 4. We found that the average number of days remaining in a given season can vary by more than two days in a given season depending on the scoring model used.

The most interesting effect of changing the scoring model was to observe which teams qualified for the playoffs. In the case of the 2006-07 season, we can note that the Toronto Maple Leafs would have qualified for the playoffs under any other scoring model then the one used that season all other things being equal (see Table 5). Another interesting note is that Edmonton, who were the Western Conference Champions in 2005-06, would not have made the playoffs if the proposed scoring model had been imposed and their spot would have gone to divisional rivals Vancouver (see Table 6). It should be noted that these results assume that the games would end in the same way regardless of the scoring system. However, research in economics has shown that teams will modify the way they play depending on the scoring model [2]. Even given this fact, it is interesting to note that some team strategies would have been good enough to secure a playoff spot under any scoring model while other teams strategies only ensure qualification under the current scoring model, for example, Tampa Bay in 2006-07.

Table 5. Date of clinching under different scoring models in the Eastern conference of the NHL in 2006-07

Team	Historic	Overtime	Extra Point	Current	Proposed
East					
Toronto	Apr 8	Apr 8	Apr 8	—	Apr 9
Ottawa	Mar 19	Mar 18	Mar 21	Mar 25	Mar 21
Montreal	—	Apr 8	—	—	Apr 7
Buffalo	Mar 24	Mar 21	Mar 23	Mar 18	Mar 22
Boston	—	—	—	—	—
NY Islanders	Apr 8	—	—	Apr 9	—
NY Rangers	—	—	Apr 9	Apr 6	Apr 7
...	...	...	...	...	...
Tampa Bay	—	—	—	Apr 6	—
Florida	Apr 6	—	—	—	—
Atlanta	Apr 7	Apr 1	Apr 2	Apr 2	Apr 7
Carolina	—	Apr 2	Apr 6	—	—

Table 6. Date of clinching under different scoring models in the Western conference of the NHL in 2005-06

Team	Historic	Overtime	Extra Point	Current	Proposed
Vancouver	—	—	—	—	Apr 14
Edmonton	Apr 14	Apr 14	Apr 18	Apr 14	—
Calgary	Apr 6	Apr 6	Apr 6	Apr 8	Apr 6
...	...	...	...	...	...

4 Conclusion and Future Work

This paper looks at the problem of determining qualification in the NHL. We present the first complete and efficient solution to the NHL qualification problem. The key to scaling up our constraint programming approach was a combination of dominance constraints and special purpose algorithms for dealing with tie-breakers and division leaders. By introducing the dominance constraints to the model, we significantly reduced the amount of search necessary to find a solution or to confirm that no solution was possible. As well, we avoided costly work associated with tie-breakers and division leaders by strategically delaying the calculation of secondary and tertiary tie-breakers and by avoiding the calculation of the actual division leaders as much as possible.

We implemented the proposed model using a constraint programming package developed by ILOG [7]. Using the 2005-06 and 2006-07 seasons, we developed a set of problem instances to verify that our solver could solve realistic problems. We found that solving an individual instance only took a fraction of a second and that all instances of the 2005-06 and 2006-07 seasons could be solved in a several minutes while respecting all tie-breaking rules.

Using the data from the problem instances, we found that we could announce results as much as five and four days earlier than the Globe and Mail [6] for the 2005-06 and 2006-07 seasons, respectively. We also extended our work to look at the effect of scoring model on playoff qualification for the first time. We found that qualification under scoring models could vary by by more than two days. We also found that changing the scoring model caused different teams to qualify for the playoffs.

In the future, we will extend this work to look at not just the decision problem of deciding qualification but also the optimization problem of determining the number of games that must be won to earn a playoff spot. This idea was first proposed by Ribeiro and Urrutia [11] for the Brazilian Soccer championship, which uses a different scoring model and playoff objective.

References

1. I. Adler, A. L. Erera, D. S. Hochbaum, and E. V. Olinick. Baseball, optimization and the world wide web. *Interfaces*, 32:12–22, 2002.
2. A. N. Banerjee, J. F. M. Swinnen, and A. Weersink. Skating on thin ice: rule changes and team strategies in the NHL. *Can. J. of Economics*, 40:493–514, 2007.
3. C. Bessiere. Constraint propagation. In *The Handbook of Constraint Programming*, chapter 3. Elsevier, 2006.
4. E. Cheng and D. Steffy. Clinching and elimination of playoff berth in the NHL. *International Journal of Operations Research*. To appear.
5. L. Getoor, G. Ottosson, M. Fromherz, and B. Carlson. Effective redundant constraints for online scheduling. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence (AAAI)*, pages 302–307, 1997.
6. Globe and Mail. Hockey scoreboard. Metro ed., 2006-2007.
7. ILOG S. A. ILOG Solver 4.2 user's manual, 1998.
8. W. Kern and D. Paulusma. The computational complexity of the elimination problem in generalized sports competitions. *Discrete Optimization*, 1:205–214, 2004.
9. P. Laborie. Algorithms for propagating resource constraints in AI planning and scheduling: Existing approaches and new results. *Artificial Intelligence*, 143:151–188, 2003.
10. K. Marriott and P. Stuckey. *Programming with Constraints: An Introduction*. MIT Press, 1998.
11. C. C. Ribeiro and S. Urrutia. An application of integer programming to playoff elimination in football championships. *International Transactions in Operational Research*, 12:375–386, 2005.
12. L. W. Robinson. Baseball playoff eliminations: an application of linear programming. *Operations Research Letters*, 10:67–74, 1991.
13. F. Rossi, P. van Beek, and T. Walsh, editors. *Handbook of Constraint Programming*. Elsevier, 2006.
14. B. Schwartz. Possible winners in partially completed tournaments. *SIAM Review*, 8:302–308, 1966.
15. B. M. Smith. Modelling. In *Handbook of Constraint Programming*, chapter 11. Elsevier, 2006.
16. K. D. Wayne. A new property and a faster algorithm for baseball elimination. *SIAM Journal on Discrete Mathematics*, 14:223–229, 2001.