Scheduling the Australian Football League Using the PEAST Algorithm

Jari Kyngäs[1], Kimmo Nurmi[1], Nico Kyngäs[1],George Lilley[2], Thea Salter[3]

[1] Satakunta University of Applied Sciences, Tiedepuisto 3, 28600 Pori, Finland

[2] Box Hill Senior Secondary College, 19 Dunloe Avenue, Mont Albert North 3129, Australia

[3] AFL HOUSE, 140 Harbour Esplanade, Docklands VIC 3008, Melbourne VIC 3001, Australia

Abstract**.** Generating a schedule for a professional sports league is an extremely demanding task. Good schedules have many benefits for the league, such as higher incomes, lower costs and more interesting and fairer seasons. This paper presents the 3-phase process needed to schedule the Australian Football League. The building of the schedule is very challenging and often requires computational intelligence to generate an acceptable schedule. There are a multitude of stakeholders with varying requests (and often requests vary significantly year on year). We used the PEAST (Population, Ejection, Annealing, Shuffling, Tabu) algorithm to schedule the 2013 season. The comparison showed that there are alternative solutions available that are comparable to the current scheduling outcome.

Keywords: Sports scheduling, Real-world scheduling, PEAST algorithm.

1 Introduction

Professional sports leagues have become big businesses. At the same time, the quality of the schedules has become increasingly important; the schedule has a direct impact on revenue for all involved parties. For instance, the number of spectators in the stadiums and the traveling costs for the teams are influenced by the schedule, and TV networks that pay for broadcasting rights want the most attractive games to be scheduled at commercially interesting times. Furthermore, a good schedule makes a season more interesting for the media and the fans, and fairer for the teams. Nurmi et al. (2010) report a growing number of cases in which academic researchers have been able to close a scheduling contract with a professional sports league owner. Excellent overviews of sports scheduling can be found in Easton et al. (2004) and Rasmussen, Trick (2008). An extensive bibliography can be found in Knust (2012) and an annotated bibliography in Kendall et al. (2010).

In a sports tournament, n teams play against each other over a period of time according to a given timetable. The teams belong to a *league*, which organizes *games* or *matches* between the teams. Each game consists of an ordered pair of teams, denoted *(i, j)* or *i-j*, where team *i* plays *at home* - that is, uses its own *venue* (stadium) for a game - and team *j* plays *away*. Games are scheduled in *rounds*, which are played on given *days*. A *schedule* consists of games assigned to rounds. A schedule is *compact* if it uses the minimum number of rounds required to schedule all the games; otherwise it is *relaxed*. If a team plays two home or two away games in two consecutive rounds, it is said to have a *break*. In general, for reasons of fairness, breaks are to be avoided. However, a team can prefer to have two or more consecutive away games if its stadium is located far from the opponent's venues, and the venues of these opponents are close to each other. A series of consecutive away games is called an *away tour*.

In a *round robin tournament* each team plays against every other team a fixed number of times. Most sports leagues play a double round robin tournament, where the teams meet once at home and once away. A *mirrored* double round robin tournament is a tournament where every team plays against every other team once in the first $n - 1$ rounds, followed by the same games with reversed venues in the last $n - 1$ rounds.

Sports scheduling involves three main problems. First, the easiest problem to solve is to find a schedule with the minimum number of breaks. De Werra (1981) has presented an efficient algorithm to compute a minimum break schedule for a single round robin tournament. If n is even, it is always possible to construct a schedule with $n - 2$ breaks.

Second, the problem of finding a schedule that *minimizes the travel distances* is called the Traveling Tournament Problem (TTP) (Easton et al., 2001). In TTP, the teams do not return home after each away game but instead travel from one away game to the next. However, excessively long away trips as well as home stands should be avoided.

Third, most professional sports leagues introduce many additional requirements in addition to minimizing breaks and travel distances. We call the problem of finding a schedule that *satisfies given constraints* (Nurmi et al., 2010) the Constrained Sports Scheduling Problem (CSSP). The goal is to find a feasible solution that is the most acceptable for the sports league owner - that is, a solution that has no hard constraint violations and that minimizes the weighted sum of the soft constraint violations. Scheduling the Australian Football League is an example of a CSSP.

Australian Rules football (officially Australian football) was invented in Melbourne, Australia. It has been played since 1858, when the first match between Melbourne Grammar School and Scotch College took place (Blainey, 2010). Originally it was invented to keep the cricketers fit during the winter time. The game has been played in some kind of league format since 1877. Some sources claim that the early history of Australian Rules football is more or less obscure, but the modern-day rules are well known (Sydney University, 2014)(OnlyMelbourne, 2014).

The game is most popular in Australia but is played practically all over the world. In all southern states of Australia it is the most popular of all sports. When measured by attendance, it is by far the most popular sport in Australia. The spectator average per match for the season 2013 was 33,500. The most popular matches in the regular season have more than 80,000 spectators.

Australian Rules football is a very physical game. What makes it different from other physical sports is the fact that the use of padding is not mandatory. Some players (ruckmen) wear shin and thigh pads, but in general, pads are rarely worn. This causes quite a lot of small injuries to thighs, hamstrings and calf muscles. The relatively high injury rates in the sport are taken into consideration by playing only once a week. There must be a minimum six-day break between the matches. This gives the players a chance to recover from minor injuries.

Section 2 presents the 3-phase process needed to schedule the Australian Football League. The section introduces the requirements, requests and other constraints the format implies. In Section 3 we describe the PEAST algorithm that has been used to schedule professional sports leagues. Section 4 reports the computational results for the 2013 season.

2 The Schedule

The Australian Football League (hereafter, AFL) has 18 teams: Adelaide Crows and Port Adelaide (Southern Australia), Brisbane Lions and Gold Coast Suns (Queensland), Fremantle and West Coast Eagles (Western Australia), Greater Western Sydney Giants and Sydney Swans (Sydney, New South Wales) and Carlton, Collingwood, Essendon, Geelong Cats, Hawthorn, Melbourne, North Melbourne, Richmond, St Kilda and Western Bulldogs (Victoria region).

AFL is trying to expand the game throughout the country and even to New Zealand. Therefore, some of the matches are played in cities and stadiums that do not have a permanent home team. Such cities are Darwin (Northern Australia), Hobart and Launceston (Tasmania), Cairns (Northern Queensland) and Wellington (New Zealand).

Australia is a big country, which causes extensive travel loads for the teams, especially for the teams from Queensland and Western Australia. The Victoria teams travel the least, about 12,000-20,000 kilometers each season, while the teams from Western Australia have to travel about 70,000 kilometers per season. Unfortunately, not much can be done to reduce the total travelling distance since no away tours (two or more consecutive away games) can be arranged due to a regular match being scheduled in each city per week.

The schedule for AFL in 2013 had a complicated structure. It consisted of each team playing against every other team once - i.e., a single round robin. In addition, each team had to play 5 additional matches. This adds up to 22 matches, 11 home and 11 away matches, for each team. The combination of a single round robin and the additional matches makes the schedule different from most of the other professional sports league schedules. The schedule consists of 23 rounds, and each team has one bye during rounds 11-13. There are 20 rounds of 9 matches and 3 rounds of 6 matches.

Building the schedule is a process comprising three phases. First, the host teams of the single round robin tournament and the five additional matches are decided. The second phase includes building the actual schedule. The schedule is built based on rounds – not the actual match days. In most cases the matches of a round are played on Friday, Saturday and Sunday. In the last phase the exact weekdays and venues of the matches are decided. These decisions are mostly constrained by various broadcasting and venue requirements.

2.1 Phase 1: Deciding the Host Teams

In the first phase the host teams of the single round robin tournament and the five additional matches are decided. The league authorities defined 12 selection rules, which we converted to constraints. Ten of the rules were strict and two had some flexibility in them. In the sports scheduling literature, a strict rule is defined as a hard constraint and a flexible rule is defined as a soft constraint.

The basic selection of the matches in the single round robin is subject to the following constraints: (H denotes a hard constraint and S denotes a soft constraint)

- H1. All teams have to play a minimum of 5 matches in Victoria.
- H2. Victorian teams should travel outside Victoria a minimum of 5 times.
- H3. Each team must have at least one home match against Collingwood or Essendon.
- H4. Each team has to travel to Western Australia at least once. If a team travels to Western Australia twice, there must be at least six rounds between the visits (this is a constraint of phase two).
- H5. All clubs have to play at least one match at the MCG stadium.

The selection for the remaining 5 matches for each team is as follows:

- H6. “Blockbuster” matches must be included. These are the matches between the big six teams from Victoria – Carlton, Collingwood, Essendon, Geelong Cats, Hawthorn and Richmond. These are fixed by the league authorities (there are some exceptions, for instance Collingwood and Richmond did not play against each other twice in 2013).
- H7. Matches between local rivals (Adelaide Crows and Port Adelaide, Brisbane Lions and Gold Coast Suns, Fremantle and West Coast Eagles, Greater Western Sydney Giants and Sydney Swans) should be respected.
- H8. The top four teams from 2012 can have only one meeting with the bottom four teams from 2012, with the exception of Sydney rivals.
- S1. The top eight teams should play the top eight teams twice, a minimum of three times.
- S2. The bottom ten teams should play the bottom ten teams twice, a minimum of three times.
- H9. The bottom two teams from 2012 should not meet the top eight teams from 2012 twice (Sydney rivals are an exception).
- H10. No team that travelled in round 23 in 2012 is to travel in round 23 this year.

According to the league authorities, the travel load is the second most important thing to consider. Therefore, we added two restrictions (as hard constraints):

- H11. All teams should travel 2-3 times to either Western Australia or Queensland. This constraint somewhat equalizes the travel load between the teams (of course, only those teams not from Western Australia or Queensland). Without this restriction, some teams might visit Western Australia once and Queensland not at all, while some other teams might make a maximum number of visits to these areas (two visits to both areas).
- H12. The local rivals’ travel loads should be as equal in length as possible. We recognized that in previous years’ schedules there were big differences in the local rivals’ travel loads.

2.2 Phase 2: Building the Schedule

The second phase consists of actually building the schedule. A basic round consists of 9 matches, of which 1 is played on Friday, 5 are played on Saturday and 3 are played on Sunday. In rounds 11-13, each team has one bye (i.e. there are 6 matches in these rounds) and the distribution of matches is 1 match on Friday, 3 matches on Saturday and 2 matches on Sunday. Moreover, in rounds 1, 7 and 10, one Sunday match is actually played on Monday.

Anzac Day (25 April) is also special because no matter which weekday it happens to be, at least one match is played. The match played is between Essendon and Collingwood. As mentioned earlier, there should be a six-day break between the matches, but Anzac Day is an exception. All the teams playing on Anzac Day must be prepared to play either their preceding or following match with a shorter break. Of course, the schedule should be made in such a way that it places these teams' preceding and following matches as far away from Anzac Day as possible.

There are nine different constraints that must be used to describe the problem framework. Only one of the constraints is a soft constraint. We follow Nurmi et al. (2010) in the specification of the constraints:

- C01. There are exactly 23 rounds available for the tournament.
- C02. Nine matches can be assigned to rounds 1-10 and 14-23, and 6 matches to rounds 11-13.
- C07. There should be at least 0 and at most 1 home matches for four different pairs of teams. The pairs are Adelaide Crows and Port Adelaide, Brisbane Lions and Gold Coast Suns, Fremantle and West Coast Eagles, Greater Western Sydney Giants and Sydney Swans. What this means is that these pairs of teams can never play at home in the same round.
- C10. There are 42 pre-assigned matches (this number can vary year on year).
- C13. Teams cannot have more than 3 consecutive home matches.
- C14. Teams cannot have more than 3 consecutive away matches.
- C19. There must be at least 6 rounds between two matches with the same opponents (soft constraint).
- C24. If two teams play against each other twice, the second match cannot be played before round 11.
- C25. If two teams play against each other only once, this match cannot be played after round 22.

In addition, we had to add 6 constraints that were not included in Nurmi et al. (2010). Three of these constraints are hard constraints. We prefix these constraints with the letter X to separate them from the phase one constraints:

- XH1. Geelong cats must play only 4 home matches before round 10 (only valid for 2013).
- XH2. There should be a minimum of 6 days break between each match.
- XH3. There should be no overlapping broadcasting slots.
- XS1. There must be at least 6 rounds between visits to WA / Qld.
- XS2. Friday matches should mostly be played in the Etihad or MCG stadiums.
- XS3. The total number of breaks should be minimized.
- XS4. Venue contractual requirements. There has to be a predefined number of matches played at each stadium.

2.3 Phase 3: Deciding the Weekdays and Venues

In the last phase the exact weekdays and venues of the matches are decided. These decisions are mostly constrained by various broadcasting and venue requirements. All of the matches are broadcast on television. Foxtel is an Australian pay television company that produce and broadcast five matches a week (in a standard round) – three on Saturday and two on Sunday. The Seven Network broadcasts the remaining four matches (in a standard round) on a free-to-air network. It is important that the broadcasts do not overlap on Saturday or Sunday in Western Australia, Southern Australia, Queensland and New South Wales. This is applicable to the Adelaide and Port Adelaide teams, the Fremantle and West Coast Eagles teams, the Brisbane Lions and Gold Coast Suns teams, and also to the Sydney Swans and GWS Giants teams. This is manageable because there are at least 3 days on which to play, and subsequently there are also varying timeslots to schedule within each day so they can be scheduled on the same day in a different timeslot (eg. afternoon and night).

Two stadiums, Etihad and MCG, host almost half of all the matches (93/198). There are 10 teams that play home matches at these stadiums. Two of these teams play the majority of their home matches at Etihad Stadium and the remaining eight play a varying number of home matches at both stadiums. Of the latter, two of these teams have to play a minimum number of away matches at Etihad Stadium. We cannot be specific about the number of matches each team should play in these stadiums due to confidentiality. These stadiums should mostly be used for Friday matches.

At first sight, the two stadiums might seem to cause an optimization problem. However, this is not the case. Because of the local rivals, there are always 4 matches played in stadiums other than Etihad and MCG (local rivals can never play at home in the same round). This leaves a maximum of five matches to be played at Etihad, MCG and/or Simonds Stadium.

3 The Solution Method

This section describes the PEAST algorithm, which was used to solve all of the three phases of the schedule. The usefulness of an algorithm depends on several criteria. The two most important ones are the quality of the generated solutions and the algorithmic power. Other important criteria include flexibility, extensibility and learning capabilities. A successful heuristic most likely uses mixed local search and population-based methods. A local search method is defined by:

1) A neighborhood structure, which is a mechanism to obtain a new set of neighbor solutions by applying a small perturbation to a given solution.
2) A method of moving from one solution to another.
3) The parameters of the method.

The PEAST algorithm obtains a new neighbor solution by applying the GHCM operator. It explores promising areas in the search space by extending the basic hill-climbing step to generate a sequence of moves in one step, leading from one solution to another. The operator is based on ideas similar to the Lin-Kernighan procedures (Lin and Kernighan, 1973) and ejection chains (Glover, 1992). It moves an object, o_1, from its old position, p_1, to a new position, p_2, and then moves another object, o_2, from position p_2 to a new position, p_3, and so on, ending up with a sequence of moves. A detailed discussion of the GHCM operator can be found in Kyngäs et al. (2012).

```
Input  the population size n, the iteration limit t, the cloning interval c,
           the shuffling interval s and the ADAGEN update interval a
    Generate a random initial population of schedules S_i for i = 1, …, n
    Set best_S = null and iteration = 1
WHILE iteration ≤ t
      k = 1
      WHILE k ≤ n
     (explore promising areas in the search space)
         Apply GHCM to schedule S_k to get a new schedule
         IF Cost(S_k) < Cost(best_S) THEN Set best_S = S_k
         k = k + 1
      END REPEAT
      (avoid staying stuck in the promising search areas too long)
      Update the simulated annealing framework
      IF iteration ≡ 0 (mod c) THEN
         (favor the best schedule, i.e. use elitism)
         Replace the worst schedule with the best one
      IF iteration  ≡ 0 (mod s) THEN
         (escape from the local optimum)
         Apply shuffling operators
      IF iteration ≡ 0 (mod a) THEN
         Update the ADAGEN framework
      iteration = iteration + 1
    END WHILE
    Output best_S
```

Figure 1: The pseudo-code of the PEAST algorithm.

The algorithm avoids staying stuck (i.e., the objective function value does not improve for some predefined number of generations) in the same areas of the search space using tabu search and the refined simulated annealing method. A tabu list (Glover et al., 1985) is used to prevent reverse order moves in a single application of the GHCM operator. The simulated annealing refinement is used to decide whether or not to commit to a sequence of moves in the GHCM operator. This refinement is different from the standard simulated annealing (Kirkpatrick et al., 1983). It is used in a three-fold manner: 1) when choosing an object to be moved, 2) when choosing the destination of the object, and 3) when the sequence of moves is cut short (a worsening move is made, and it worsens the solution more than the previous worsening move did). A detailed discussion of the tabu search and simulated annealing refinement can be found in Kyngäs et al. (2012).

The pseudo-code of the algorithm is given in Figure 1. The algorithm uses a population of solutions that enables it to explore a wide range of promising areas in the search space. In every *c* iteration, the least fit schedule is replaced with a clone of the fittest individual. This operation is completely irrespective of the globally fittest schedule (*best_S* in Figure 1) found by that time in the search process.

The PEAST algorithm applies a number of shuffling operators to perturb a solution into a potentially worse solution in order to escape from local optima. The operators are called according to a rule. The idea of shuffling is the same as in hyper-heuristics (Burke et al., 2013) but the other way around. Hyper-heuristic is a mechanism that chooses a heuristic from a set of simple heuristics, applies it to the current solution to get a better solution, then chooses another heuristic and applies it, and continues this iterative cycle until the termination criterion is satisfied. We introduce a number of simple heuristics that are used to worsen the current solution instead of improving it. Examples of shuffling operators can be found in Kyngäs et al. (2012).

The PEAST algorithm is used to solve multi-objective optimization problems - i.e., problems where multiple objective functions have to be optimized simultaneously. The objective functions usually compete in such a way that improving one objective function value most likely improves the other objective function values. The PEAST algorithm uses the adaptive genetic penalty method (ADAGEN) (Nurmi, 1998) to solve the multi-objective problems. A traditional penalty method assigns positive weights (penalties) to the soft constraints and sums the violation scores to the hard constraint values to get a single value to be optimized. The ADAGEN method assigns dynamic weights to the hard constraints based on the search trajectory and the constant weights assigned to the soft constraints. The soft constraints are assigned fixed weights according to their significance. The significance is given by the problem owner (end-user).

The PEAST algorithm uses random initial solutions. In our extensive test runs we have found no evidence that a sophisticated initial solution improves results. On the contrary, random initial solutions tend to yield superior or at least as good results (Kyngäs et al., 2012).

The PEAST algorithm has been used to solve several types of real-world scheduling problems and is in industrial use. The first version of the algorithm was used to solve school timetabling problems (Nurmi, 1998). The later versions of the algorithm have been used to solve sports scheduling problems (see, e.g., Kyngäs et al., (2014) and Kyngäs and Nurmi (2009)) and workforce scheduling problems (see, e.g., Kyngäs et al. (2012) and Kyngäs et al. (2013)). Furthermore, we have used the algorithm to solve several artificial and benchmark problems, including school timetabling (Post et al., 2012), balanced incomplete block design(Nurmi et al., 2011), single round robin tournaments with balanced home-away assignments and pre-assignments (Nurmi et al., 2011), days-off scheduling (Nurmi and Kyngäs, 2011) and constraint minimum break problems (Nurmi et al., 2010).

The industrial use of the PEAST algorithm, as well our experiences in solving different benchmark problems, can be summarized as follows:

1) The crucial components of the algorithm are random initial solutions, making moves in sequences (the GHCM operator) and using a population of solutions.
2) Either the simulated annealing refinement or shuffling operators should be used. Both should be used in real-world instances to ensure good quality solutions.
3) A tabu list improves the efficiency of the GHCM operator.
4) The initial object in a move sequence should be chosen using tournament selection.

Even though the best parameter values vary depending on the problem and the instance, our extensive test runs over several years have shown that the following values can safely be used in different real-world problems and instances:

– The population size is 10.
– The cloning interval is 500.
– The shuffling interval is 5,000.
– The maximum length of the move sequence in the GHCM operator is 10.
– The size of the tournament selection is 7.

- The length of the tabu list is 10 (equals the length of the move sequence).
- In the simulated annealing framework we stop the cooling at some predefined temperature. Therefore, after a certain number of iterations, *m*, we continue to accept an increase in the cost function with some constant probability. We choose *m* equal to *t* /2, where *t* is the maximum number of iterations and *p* is equal to 0.0001.

We are aware of the fact that we have used many different heuristic methods in the PEAST algorithm. The acronym PEAST stems from the methods used: Population, Ejection, Annealing, Shuffling and Tabu. One might think that the outcome is nothing more than a collection of old ideas. However, to the best of our knowledge, the heart of the algorithm, the GHCM operator, is one of a kind. The same applies to our implementation of the shuffling operators, simulated annealing and penalty method.

We can steadily note that the PEAST algorithm realizes the criteria given at the beginning of the section. Our industrial partners from different lines of business have stated that the algorithm constantly produces good-quality solutions in acceptable time.

4 Computational Results

We compared our schedule to the running schedule for the 2013 season. Table 1 summarizes the comparison. The optimization model was created in cooperation with one of the league authorities. However, the running schedule was not generated using exactly the same model. For example, the topics H11 and H12 may not have been in effect in the model behind the running schedule. We still believe the comparison is fair enough. The comparison shows that in our view our solution is better than the running schedule. Our solution is at least as good as the running schedule on every topic.

The biggest differences are highlighted. The most significant difference is the sum of differences of the total traveling of Non-Victorian teams (H12). We are not sure whether this was optimized in the running schedule, but we think it should have been. Balancing the traveling distance of local teams most certainly increases overall satisfaction. The second most significant difference is that the top 8 teams should play against each other a minimum of three times (S1). We chose to use a more difficult constraint in our model - that is, “exactly three times” - because we think it would be unfair if a team had to meet the top eight teams more than three times. We were able to find such a tightened solution. Note also that our solution has clearly fewer breaks.

5 Conclusions and Future Work

The format of the AFL fixture was something we had not run into before. One of the most interesting features is that it includes only a single round robin with 18 teams. In addition, each team has 5 extra matches, which increases the total number of matches played during the season to 22. The 5 extra matches are selected in a way that balances the fixtures between the teams.

The traveling is a big issue for the teams. The traveling distance is from 12,000 kilometers up to 70,000 kilometers per season per team. These are big numbers considering only 11 away matches are played during the season.

We showed that the PEAST algorithm is capable of finding good solutions to the AFL fixture. It was quite easy to apply the algorithm to handle all the three phases needed to generate the schedule. One thing to do in the future could be to merge these two phases.

Table 1: Comparison between the running schedule and our schedule. The topics are ordered as follows: phase one, phase two and miscellaneous.

Topic	Running schedule	Our schedule
All clubs to play a minimum of five matches in Victoria (H1)	5-7	6-7
Victoria teams travel maximum of six times (H2)	ok	ok
Minimum one home match against Essendon or Collingwood for each club (H3)	ok	ok
Each team should visit Perth once (H4)	ok	ok
All teams to play at least one match at MCG (H5)	ok	ok
Blockbuster matches (H6)	ok	ok
Local derbies (H7)	ok	ok
Top 4 teams not to play against bottom 4 teams twice (H8)	Western Bulldogs vs. Adelaide	ok
No top 8 team (other than Syd v GWS x 2) to play either Gold Coast Suns or GWS Giants twice (H9)	ok	ok
Round 23 – alternate travel between non-Vic teams (H10)	ok	ok
Each team should visit Queensland once (H11)	no visit for 4 teams	ok
Total visits to Perth and Queensland between 2-3 (H11)	2 teams once 2 teams four times	ok
The sum of differences of the total traveling of non-Victoria teams (H12)	9474	632
Top 8 teams to play against each other twice a minimum of three times (S1)	Once 2 times Once 4 times Once 5 times	ok
Bottom 10 teams to play bottom 10 teams twice a minimum of three times (S2)	ok	ok
Non-fixed (not pre-assigned) matches played outside regular slots (C01)	1 time	0 times
One bye per team in rounds 11-13 (C02)	ok	ok
20 rounds of nine matches and 3 rounds of six matches (C02)	ok	ok
Number of 3-breaks at home (C13)	2	2
Number of 3-breaks away (C14)	3	0
Local rivals never playing home in same round (C07)	ok	ok
Pre-assigned matches (C10)	ok	ok
Must be a minimum of six weeks between playing a team for the first and second time (C19)	ok	ok
No teams to play for the second time until after round 10 (C24)	ok	ok
All teams must play each other once by round 22 (C25)	ok	ok
No home matches for Geelong Cats at Simonds Stadium until Round 10 (XH1).	ok	ok
Back-to-back Perth minimum of 6 rounds gap (XS1)	3 times 5 rounds gap	1 time 5 rounds gap
Back-to-back Queensland minimum of 6 rounds gap (XS1)	1 time 2 rounds gap 1 time 5 rounds gap	2 times 4 rounds gap
Minimum six-day break between each match, with exceptions for Anzac Day (XH2)	ok	ok
Never play in the same or overlapping timeslot, so that all local matches can be broadcast on free-to-air in each market (XH3)	ok	ok
Number of Friday matches at MCG or Etihad (XS2)	15	17
Total number of breaks (XS3)	94	74
Venue contractual requirements (XS4)	ok	ok
Total traveling of Gold Coast and Brisbane Lions (appr. using Google Maps kilometers)	GC 18,347 BL 21,614	GC 18,347 BL 18,347
Total traveling of Adelaide and Port Adelaide	PA 11,176 A 13,126	PA 11,776 A 11,850
Total traveling of West Coast Eagles and Fremantle	WCE 33,601 F 34,801	WCE 34,101 F 34,101
Total traveling of Sydney Swans and GWS	SS 12,243 GWS 15,300	GWS 12,243 SS 12,800
Total traveling of Victoria teams	6,718-10,992	6,711-10,987
Total traveling of non-Victoria teams	160,208	153,558
Total traveling of Victoria teams	82,917	86,607
Minimum of 45 matches in MCG	ok	ok
Minimum of 46 or 48 matches in Etihad	ok	ok
Other contractual matches	ok	ok
No day or twilight matches at TIO Stadium	ok	ok
No Sunday early or Saturday afternoon matches at Patersons Stadium	ok	ok
Key Features 1-17	ok	ok
Number of breaks for one team	8 (four times) 7 (once) 6 (four times)	8 (once) 7 (once)

References

G. Blainey, "A Game of Our Own – The Origins of Australian Football", Black Inc., 2010, pp. 7.

E,K. Burke, M. Gendreau, M. Hyde, G. Kendall, G. Ochoa E. Özcan and R. Qu, "Hyper-heuristics: a survey of the state of the art", in *Journal of the Operational Research Society* 64, 2003, pp. 1695–1724.

K. Easton, G. Nemhauser, and M. Trick "The traveling tournament problem: description and benchmarks" in Proc of the 7th. International Conference on Principles and Practice of Constraint Programming, Paphos, pp. 580–584, 2001.

K. Easton, G. Nemhauser and M. Trick, "Sports scheduling", in Handbook of Scheduling, edited by Leung, Florida, USA: CRC Press, pp 52.1-52.19, 2004.

F. Glover, C. McMillan and B. Novick, "Interactive Decision Software and Computer Graphics for Architectural and Space Palnning", *Annals of Operations Research* 5, 1985, pp. 557-573.

F. Glover, "New ejection chain and alternating path methods for traveling salesman problems", in *Computer Science and Operations Research: New Developments in Their Interfaces*, edited by Sharda, Balci and Zenios, Elsevier, 1992, pp. 449–509.

G. Kendall, S. Knust, C.C. Ribeiro and S. Urrutia, "Scheduling in Sports: An annotated bibliography", Computers and Operations Research 37, pp. 1-19, 2010.

S. Kirkpatrick, C.D. Gelatt Jr. and M.P. Vecchi, "Optimization by Simulated Annealing", *Science* 220, 1983, pp. 671-680.

S. Knust, "Sports Scheduling Bibliography", [Online], Available: http://www.inf.uos.de/knust/sportssched/sportlit_class/, (Last update 12.05.2014).

J. Kyngäs and K. Nurmi, "Scheduling the Finnish 1st Division Ice Hockey League", in *Proc of the 22nd Florida Artificial Intelligence Research Society Conference*, Florida, USA, 2009.

N. Kyngäs, D. Goossens, K. Nurmi and J. Kyngäs, "Optimizing the Unlimited Shift Generation Problem", *Applications of Evolutionary Computation, Lecture Notes in Computer Science,* Vol. 7248, Springer, USA 2012.

N. Kyngäs, K. Nurmi and J. Kyngäs, "The Workforce Optimization Process Using the PEAST algorithm", *Lecture Notes in Engineering and Computer Science: Proc. of The International MultiConference of Engineers and Computer Scientists*, Hong Kong, 2013.

N. Kyngäs, K. Nurmi and J. Kyngäs, "Scheduling the Highly Constrained Finnish Major Ice Hockey League", *Lecture Notes in Engineering and Computer Science: Proc. of The International MultiConference of Engineers and Computer Scientists*, Hong Kong, 2014. (accepted)

S. Lin and B.W. Kernighan, "An effective heuristic for the traveling salesman problem", *Operations Research* 21, 1973, pp. 498–516.

K. Nurmi, "Genetic Algorithms for Timetabling and Traveling Salesman Problems", Ph.D. dissertation, Dept. of Applied Math., University of Turku, Finland, 1998. Available: http://www.bit.spt.fi/cimmo.nurmi/

K. Nurmi, D. Goossens, T. Bartsch, F. Bonomo, D. Briskorn, G. Duran, J. Kyngäs, J. Marenco, CC. Ribeiro, FCR. Spieksma, S. Urrutia and R. Wolf-Yadlin, "A Framework for Scheduling Professional Sports Leagues", in Ao, Sio-Iong (ed.): IAENG Transactions on Engineering Technologies Volume 5, Springer, USA, 2010.

K. Nurmi, D. Goossens and J. Kyngäs, "Scheduling a triple round robin tournament for the Finnish national ice hockey league for players under 20", in *Proc of the IEEE Symposium on Computational Intelligence in Scheduling*, Paris, France, 2011.

K. Nurmi and J. Kyngäs, "Days-off Scheduling for a Bus Transportation Staff", *International Journal of Innovative Computing and Applications* Vol. 3(1), Inderscience, UK, 2011.

OnlyMelbourne, "History of Australian Football", [Online], Available: http://www.onlymelbourne.com.au/melbourne_details.php?id=1543#.Ur6LMrRbHnY (Accessed 20.1.2014).

G. Post, J.H. Kingston, S. Ahmadi, S. Daskalaki, C. Gogos, J. Kyngas, K. Nurmi, N. Musliu, N. Pillay, H. Santos and A. Schaerf, "XHSTT: an XML archive for high school timetabling problems in different countries", *Annals of Operations Research*, Springer, USA, 2012.

P. Rasmussen and M. Trick, "Round robin scheduling - A survey", European Journal of Operational Research 188, pp. 617-636, 2008.

Sydney University, Australian National Football Club, "The Origin of Australian Rules", [Online], Available: http://www.suanfc.com/suanfc/recent-history/the-origin-of-australian-rules (Accessed 20.1.2014)

D. de Werra, "Scheduling in sports", in Studies on graphs and discrete programming, edited by Amsterdam and Hansen, pp. 381-395, 1981.