Proceedings of the 2007 IEEE Symposium on Computational
Intelligence in Scheduling (CI-Sched 2007)

Super 14 Rugby Fixture Scheduling Using a Multi-Objective Evolutionary Algorithm

Lyndon While and Luigi Barone
School of Computer Science & Software Engineering
The University of Western Australia
{lyndon,luigi}@csse.uwa.edu.au

***Abstract*— Super 14 Rugby is not only a popular game, but also a hugely profitable business. However, determining a schedule for games in the competition is very difficult, as a number of different, often conflicting, factors must be considered. We propose the use of a multi-objective evolutionary algorithm for deciding such a schedule. We detail the technical details needed to apply a multi-objective evolutionary algorithm to this problem and report on experiments that show the effectiveness of this approach. We compare solutions found by our approach with recent fixtures employed by the organising authority; our results showing significant improvements over the existing solutions.**

Keywords: Fixture Scheduling, Multi-Objective Evolutionary Algorithm, Super 14

I. INTRODUCTION

Rugby Union is a team sport in which two teams compete to score the most number of points. An individual game pits two teams of 15 players against each other in a competition lasting 80 minutes. The game is somewhat similar to American Football, although it is much quicker, with far fewer interruptions to the flow of the game. The game is renowned for its skill and physicality; the game is "full-contact", with players allowed to physically impede opponent players by tackling them using their arms and torso. Players wear little to no protection, so the risk of injury is relatively high.

Super 14 is the largest Rugby Union competition in the southern hemisphere [1]. The competition consists of 14 teams from three countries: Australia, New Zealand, and South Africa; four from Australia and five from the other two countries. Games are played in batches called *rounds*, played over a weekend. Teams are only required to play one game per round.

Super 14 follows a single round-robin format, meaning each team must play every other team once. The competition lasts 14 rounds; each round consisting of either 6 or 7 games. In the case of 6 games, two teams do not play in the round and instead are said to have a *bye*. Each team has only one bye. Byes provide teams with an opportunity to rest, recover from injuries, and make preparations for travel. In order to minimise the effects of long-distance travel, Australian and New Zealand teams will typically play several rounds in South Africa (and vice versa) without returning home. A schedule for all games is known as the *fixture*.

Due to the fierce rivalry and involvement of the world's best players, Super 14 is very popular, not only in the three competing countries, but also across the entire world. Indeed, the managing director and CEO of Australian Rugby Union believes "the quality and intensity of competition between the New Zealand, South African and Australian Super 14 teams is without peer in the world" [2]. Like most popular team-based sports, Super 14 is also big business. As stated on the official Australian Rugby Union website, the organising authority for the game (SANZAR) have recently announced a new media deal worth US$323 million over five years [3] that will see Super 14 broadcast "live into 41 countries and more than 10 million homes" [4].

Like most team sports, the location at which the game is played can significantly affect the outcome of the game. The *home* team (the team who hosts the game), typically has the advantage because they are most familiar with the venue, and have more fans supporting them. The *away* team is at a disadvantage — they must contend with increased travel, time zone changes, unfamiliarity with the venue, and a hostile fan base. Indeed, this so called *home ground advantage* can be significant, so balancing the number of games a team plays at home and away is important to ensure fairness.

In Super 14, each team must play either 6 or 7 home games and 7 or 6 away games (notionally, some teams "play" their bye game at home while some "play" their bye game away). For any given game, the choice of which team plays at home and when into the competition they play can significantly affect the outcome of the game. Additionally, the placement of when teams have their bye can also affect team morale and preparedness; in order to minimise the effects of travel, teams prefer to have their bye either at the beginning or at the end of a sequence of games that see them cross the Indian Ocean. In order to ensure fairness, fixture designers must consider these effects when determining the schedule of games in the competition.

Scheduling games for Super 14 is indeed difficult, as the organising authority needs to balance a number of different, often conflicting, considerations. Besides competition fairness, factors like revenue expectations, political considerations, and availability of venues can also affect the fixture. For example, scheduling games against traditional rivals in the largest venues may increase revenue (more fans will be able to attend), but this may reduce the fairness of the competition, especially if it requires teams to travel incessantly. The placement of byes in the fixture also adds an extra complicating factor.

Further, balancing the number of games in each country

Proceedings of the 2007 IEEE Symposium on Computational Intelligence in Scheduling (CI-Sched 2007)

is crucial to ensure maximal revenue for the Super 14 competition through sale of their broadcasting rights. With the competition spread across three countries (and five time-zones), ensuring an "even" spread of games across the different playing regions allows broadcasters to schedule games in each region's television "prime-time" viewing, yielding maximum advertising income. More revenue for the broadcasters translates to more money for the competition (broadcasting rights earn significantly more than gate receipts), but may mean other factors (like competition fairness) are sacrificed. Trade-offs result — changing the fixture to improve one objective may result in worsening other objective. This is the realm of multi-objective optimisation — finding a suitably good set of solutions that vary the trade-offs in the different objectives by differing amounts in order to produce a range of alternative solutions.

Evolutionary-based multi-objective optimisation has been used on a variety of different problems including scheduling [5], but little, if any, has been done in sports fixture scheduling. In this paper, we describe a study which uses a multi-objective evolutionary algorithm to determine a fixture for Super 14 rugby, subject to a number of constraints and objectives desired by the organising authority for the game.

The rest of this paper is structured as follows. Section II presents a summary of related work on using optimisation-based approaches for sports fixture scheduling. Section III gives an overview of multi-objective optimisation, including the terminology used in this paper. Section IV describes our multi-objective approach, providing the technical details used to solve this problem. Section V presents results of experiments that demonstrate the effectiveness of our approach. In particular, we show that the multi-objective approach is able to produce solutions significantly superior to recent fixtures employed by the organising authority. Finally, Section VI concludes the paper.

II. Previous Approaches for Sports Fixture Scheduling

Many approaches have been proposed for solving sports scheduling problems. These problems are especially difficult to solve because each sports league has its own idiosyncratic requirements, constraints, and preferences. Perhaps for this reason, many previous approaches simply seek any solution that satisfies all the problem constraints — that is, the problem is often cast as a satisfaction problem rather than an optimisation problem.

For example, a special case is the Sports League Scheduling Problem (*Prob026* of CSPLib [6]), as posed by McAloon et al. [7]. In this problem:

- there are T teams, where T is even, and each team plays the other once;
- there are $T-1$ rounds;
- each team plays one game per round;
- there are $\frac{T}{2}$ periods in a round, with one game per period; and
- no team plays more than once in a given period.

This has been solved by Hamiez and Hao [8] in linear time for the case where $(T-1)\ mod\ 3 \neq 0$ using an exhaustive repair method, improving on earlier results using Tabu search [9] and various earlier approaches including integer linear programming, constraint programming, and randomised complete searches with and without heuristics.

Another special case is the Travelling Tournament Problem, which is concerned with minimising the total distance travelled in a double round robin tournament. Easton et al. [10] present a combined integer programming and constraint programming approach to this problem for a group of 8 teams. This is a much simpler problem than the scheduling problem investigated in this work, but the authors note that solving even this simplified problem for even small numbers of teams has proven to be extremely difficult. A third example is the break minimization problem (see e.g. [11]), which deals with finding a round-robin schedule that minimizes the number of consecutive home or away games for the teams.

There has been less success on more "realistic" problems. Carefully formulated methods that rely on the regularity of the problem are all too readily rendered useless when additional constraints or preferences are added. This makes evolutionary algorithms an attractive option, as they can deal with complex objective functions, and there are good techniques available to handle constraints and preferences.

Although evolutionary algorithms have been much used for timetabling and scheduling problems (see e.g. [12]), there are only a few examples of evolutionary algorithms being used for sports scheduling. Some examples are: Schönberger et al. [13] used a genetic algorithm to schedule the rounds of a table-tennis competition; Yang et al. [14] used an evolution strategy algorithm to solve sports scheduling problems; Schönberger et al. [15] describes the use of a memetic algorithm for sports league scheduling; Costa [16] used a hybrid evolutionary Tabu search to schedule hockey leagues; and Yang et al. [14] used a genetic algorithm to schedule games for a baseball league. All these studies report good results compared to previously employed methods.

However, there appears to be little to no previous work in applying multi-objective evolutionary algorithms to sports scheduling. In previous work [17], we applied a multi-objective evolutionary algorithm to the fixture scheduling for Australian Rules Football. Results showed that the multi-objective approach was able to produce solutions superior to the existing fixture while still maintaining important constraints imposed by the organising authority, yielding significant improvements in a number of metrics. This works focuses on applying the same techniques to Super 14 Rugby.

III. Multi-Objective Optimisation

Multi-objective optimisation is the task of finding appropriately "good" solutions to a problem in which candidate solutions are judged according to multiple criteria that conflict with each other to some degree. With multiple conflicting criteria, a good solution can be improved on one criterion only by accepting worse performance in at least one other criterion. The aim in multi-objective optimisation is to generate a set of solutions that compromise the different criteria to varying degrees — the solution to be used in any given situation can be selected according to the particular needs of that situation.

Without loss of generality, consider a multi-objective optimisation problem defined in terms of a search space of allowed values consisting of parameters and a vector of objective functions mapping parameter vectors into fitness space. Given two vectors $\vec{a}$ and $\vec{b}$, $\vec{a}$ is said to *dominate* $\vec{b}$ iff $\vec{a}$ is at least as good as $\vec{b}$ in all objectives and better in at least one. A vector $\vec{a}$ is *non-dominated* with respect to a set X iff there is no vector in X that dominates $\vec{a}$. A set X is a *non-dominated set* iff all vectors from X are mutually non-dominating. The set of corresponding objective vectors is called the *non-dominated front*.

A vector $\vec{a}$ is *Pareto optimal* iff $\vec{a}$ is non-dominated with respect to the set of all possible vectors. Such a vector is characterised by the fact that improvement in any one objective necessarily means a worsening in at least one other objective. The *Pareto optimal set* is the set of all possible Pareto optimal vectors. The goal of multi-objective optimisation is hence to find this Pareto optimal set, although for continuous problems a representative subset suffices.

Since evolutionary algorithms are population based, the partial order imposed on the search space necessitates the need for an appropriate ranking scheme. Two schemes are commonly employed. Both schemes employ the concept of domination to assign a *Pareto rank* to individuals — a lower rank implies a superior candidate. In Goldberg's [18] ranking procedure, non-dominated vectors are assigned a rank of 0 while any dominated vector $\vec{a}$ in the population X is assigned a rank equal to one plus that of the highest-ranked vector from X that dominates $\vec{a}$. In contrast, Fonseca and Fleming propose a scheme [19] in which a dominated vector $\vec{a}$ in the population X is assigned a rank equal to the number of vectors in X that dominate $\vec{a}$. It is this Pareto rank, rather than some (weighted) combination of the objectives, that is used as the basis for selection in a multi-objective evolutionary algorithm.

IV. Our Multi-Objective Approach

In this section, we describe the technical details needed to equip a multi-objective evolutionary algorithm to address the problem of fixture determination in the Super 14 competition. The multi-objective evolutionary algorithm we use in this work is a form of a $(n + n)$ evolutionary strategy.

A. Representation

Recall that the Super 14 competition follows a single round-robin format [20], meaning every team must play each other team precisely once. Each team has one bye, so the competition consists of 14 rounds with 6 or 7 games each.

At first thought, representing fixtures for a round-robin tournament can be achieved by simply enumerating the teams who play each other in every round of the competition (i.e. instantiating a "round matrix" that lists the round that each team plays each other). However, if we choose this representation, exploration of alternative solutions is extremely difficult as most modifications of the round matrix are likely to generate fixtures that violate the round-robin constraint. When this occurs, the search flounders in the large infeasible regions of the search space. Instead, we chose an alternative representation that still allows the exploration of different solutions, but which better maintains the round-robin constraint imposed by the competition structure.

A round-robin tournament for an even number of teams can be constructed using the *polygon construction method* [20], or via instantiating a schedule from a known valid design (after all, the polygon method simply generates a valid design). In this work, we take the latter approach and use an existing fixture as a *template* from which candidate solutions in our evolving population use to generate their corresponding schedule of games. As described below, candidate solutions are allowed to vary: which teams will play at home for any given game pairing, which order the rounds of the template fixture are played, and which "actual" teams correspond to which "logical" teams in the template fixture, hence allowing the exploration of different alternative fixtures of interest. Candidate solutions capture these variations to the template fixture, and then using the template, are converted to a "real" fixture that is assessed by the evolutionary algorithm.

To determine home teams, we use a $14{\times}7$ Boolean *home team matrix* that indicates which team for each game pairing in a round plays at home. Each row represents one round in the competition (recall, there are 14 rounds), and each column represents one of the games in the round (there are 7 games in each round, noting that some may represent byes). Evolutionary selection pressure will then drive the algorithm to locate solutions with good home-and-away sequences.

To avoid bias in the search, instead of using the actual team names and round numbers in the fixture template, we use logical team names and round numbers, and then convert these logical values to actual values via two maps: a *logical-to-actual team map* for converting team names, and a *logical-to-actual round map* for round numbers. This allows the evolutionary algorithm to test different scenarios in which the rounds and teams making up the template fixture are permuted, offering more exploration of the search space in order to find better solutions.

As described in Section IV-B, since there is no difference in teams from the same "region" of the competition, the actual team names used in the logical-to-actual team map are further placeholders for the real team names from the Super 14 competition. These placeholders must distinguish teams from different regions, but need not distinguish teams from within the same region (e.g. Eastern Australia Team 1). When finally required, these regional placeholders can be randomly instantiated with real team names from the Super 14 competition.

It is these three matrices (the home team matrix, the logical-to-actual team map, and the logical-to-actual round map) that form the genotype of a candidate solution in our evolving population and are what are optimised by the multi-objective evolutionary algorithm.

B. Objectives

We assess each candidate Super 14 according to three objectives. For consistency, we cast each objective in terms of a minimisation problem.

1) Home Games: As indicated previously, the location of where a game is played can have a significant impact on the outcome of the game. To ensure equity (fairness of the competition), a fixture should try and balance, as far as possible, the number of home games each team has. However, unlike other team based sports, continuous sequences of home (or away) games are allowed in Super 14 (even though they have some negative psychological impacts) in order to minimise the amount of foreign travel required (Australasian teams much prefer having a sequence of away games in South Africa instead of constantly having to make long journeys across the Indian Ocean, and vice versa).

Our first objective captures this equity measure:

$$Equity = \sum_{t \in T} \max(0, 6 - H_t, H_t - 7)$$

where T is the set of all teams, and H_t is the number of games played at home by team t.

The structure of this expression captures the fact that the desired number of home games for any team is 6 or 7 (recall, some teams "play" their bye game at home, while some "play" their bye game away). Should a team play less than 6 or more than 7 home games, one of the terms $6 - H_t$ or $H_t - 7$ for that team will be positive, and hence the expression $\max(0, 6 - H_t, H_t - 7)$ will be positive. Summing for all teams, $equity$ is optimal when all teams play either 6 of 7 home games, resulting in a final $equity$ score of zero.

The organising authority of the Super 14 competition deems this measure to be very important. Analysis shows that the organising authority always attempts to use fixtures in which all teams play, as far as possible, the same number of home games, effectively viewing this objective as a constraint that must be true in order for a fixture to be accepted.

In this work, we still choose to include this measure as a separate objective in our evolutionary algorithm for two reasons. Firstly, using this measure as an objective creates a gradient for which the evolutionary algorithm is able to make regular progress in the search (without a gradient, the search would be punctuated — many fixture variations would be rejected because they do not satisfy the constraint). Additionally, interesting solutions in other objectives may be possible via transitioning through solutions that violate the constraint. This is especially desirable in the earlier parts of the search (when selection pressure will not immediately remove these non-zero $equity$ solutions), as we wish the evolutionary algorithm to explore (diversify) the entire search space of potential solutions in order to find better alternative solutions to the seed fixture.

2) Travel: As described previously, teams in the Super 14 competition are located in three countries: Australia, New Zealand, and South Africa; countries that are spread across half of the southern hemisphere. Travel is hence unavoidable, and coupled with significant changes in time-zones, is generally considered detrimental to team performance. This effect is also believed to increase when travelling for sequential weeks. To produce a fair competition, we desire to reduce the effects of such travel.

Teams in Australia are located on the two coasts of the continent. As Australia is a large country (approximately the same size as continental U.S.), travel between these locations is also quite significant. To capture a measure of travel, we first classify each team as belonging to a particular region: New Zealand, South Africa, Eastern Australia, and Western Australia. As the journey between Australasian (Australia and New Zealand) and South Africa is particularly long (requiring teams to cross the Indian Ocean), we weight these journeys significantly more than other journeys.

Travel between the other regions (e.g. Australia to New Zealand, Eastern Australia to Western Australia, etc) have less of an impact than the Indian Ocean crossings, but due to time-zone changes, can still adversely affect team performance. We hence weight these journeys less than those of crossing the Indian Ocean. Note that while the journey from Western Australia to New Zealand is significantly longer than from Eastern Australia to New Zealand, we weight these trips equally as both involve time-zone changes, immigration concerns, and similar psychological effects on the players. We consider travel within a region as relatively insignificant, as flights are not very long and do not involve changes in time-zones. We hence ignore these journeys in our calculation of the travel objective.

Recall that the location of a bye within a fixture can significantly affect the performance of a team. In order to minimise the effects of long journeys, teams prefer to have their bye either at the beginning or at the end of a sequence of games that takes them across the Indian Ocean. Assuming resting for one round is sufficient to recover from travelling, we ignore any travel which immediately proceeds or follows a bye. Similarly, travel before the start of the competition (i.e. travel involved to get to the location of a game for the first round) and travel after the competition has ended (i.e. travel to return home after the last round) are ignored as teams can obtain sufficient time to prepare or recover for/from the game.

For some given fixture, a measure of the travel undertaken is captured as:

$$Travel = 100L + \sum_{t \in T} S_t$$

where T is the set of all teams, L is the maximum number of long journeys (crossing the Indian Ocean) by any team in the competition not immediately proceeding or following a bye, or the start or end of the competition, and S_t is the number of other inter-regional journeys undertaken by team t.

The $travel$ objective consists of two distinct sub-components: one capturing long journeys across the Indian Ocean, and the other capturing the number of other regional crossings. We use the sum of all short journeys in order to minimise the total number of flights across regions. For long journeys, we use the maximum number of trips made by any team, thus simultaneously creating selection pressure to both minimise and equalise the number of crossings of the Indian Ocean. Note that if we simply used the sum of the number of long journeys, the distribution of long journeys may not be equalised. Similarly, if we used the difference between the maximum and minimum numbers of long journeys, the total number of these journeys may not be minimised.

The weighting of 100 for long journeys effectively places an ordering on the two sub-components of the $travel$ objective — long journeys are more important than short journeys. We could use two separate objectives to capture these sub-measures, but this increases the difficulty of the search for little to no benefit. Experience shows that the organising authority regards the number of Indian Ocean crossings much more highly, and thus recasting this measure into two different objectives allows for the creation of "undesirable" (even if they are Pareto optimal) solutions. In essence, the two sub-measures are not equally important, and hence a linear combination is best in this case.

3) Game Distribution: The Super 14 competition consists of 14 teams based in 4 regions, distributed as follows: 3 in Eastern Australia, 1 in Western Australia, 5 in New Zealand, and 5 in South Africa. Distributing the games equally, as far as possible, across the three countries is important for a number of reasons, including political reasons (e.g. contractual obligations), promoting interest (fans prefer a regular distribution of games), venue availability, and maintenance concerns (venues can be over-used, degrading the quality of the playing surface). But of more importance, revenue is maximised by equalising the distribution of games.

With an "even" spread of games across the different countries, broadcasters are able to schedule games in each country's "prime-time" viewing, thus maximising revenue from advertising income. More revenue for the broadcasters translates to more money for the competition (broadcasting rights earn significantly more than gate receipts), thus maximising the profitability of the competition as a whole. Contributing another effect, equalising the distribution of games ensures regular interest from fans, not only ensuring maximal gate receipts, but also interest in the television broadcasts (via television ratings).

Recall, each round of the Super 14 competition contains 6 or 7 games. Therefore, a distribution in which at least two games are played in each country each round is desirable. Indeed, experience shows that the organising authority for the competition prefers this distribution where possible.

We capture the distribution of games by:

$$Distribution = \sum_{r=1}^{R} \sum_{c \in C} \max(0, 2 - G_{c,r})$$

where R is total number of rounds, C is the set of countries in the competition, and $G_{c,r}$ is the number of games played in country c in round r.

This expression is minimal when all rounds contain at least two games in each country (the term $2-G_{c,r}$ will be positive if a country contains less than two games, resulting in the expression $\max(0, 2 - G_{c,r})$ being positive). By summing over all countries in all rounds, the total variation from "optimal" is captured, allowing the evolutionary algorithm a means of minimising the mis-match.

C. Mutation

As we indicated in Section IV-A, we represent candidate solutions in our evolving population by three matrices (the home team matrix, the logical-to-actual team map, and the logical-to-actual round map) that control how to interpret entries in the template fixture. In order to explore the search space but retain constraint-preserving solutions (fixtures that preserve the round-robin structure), we require a constraint-preserving mutation operator that modifies these matrices.

When a candidate solution is mutated, one of the three matrices is selected uniformly randomly; the mutation operator modifying the contents of the selected matrix. Mutation of the home team matrix is achieved by randomly selecting between one and four games inclusive and reversing the home team for these games. Mutation of the logical-to-actual maps occurs by simply swapping two co-domain entries.

Note that under this scheme, mutation may yield different genotypes that produce the same phenotype and hence have the same net performance. This occurs when the mutation operator switches two teams from the same region in the logical-to-actual team map; the re-ordering not affecting the performance of the mutated fixture as no objective differentiates between different teams from the same region. When this occurs, the mutation is rejected and another mutation is generated as its replacement.

D. Selection

As is the case in most multi-objective evolutionary algorithms, determination of which candidate solutions survive and reproduce in our multi-objective evolutionary algorithm is primarily based on Pareto rank. We use the Pareto ranking scheme proposed by Fonseca and Fleming [19].

Inevitably though, at some stage in the algorithm, the evolutionary algorithm will be forced to select between candidate solutions with the same rank (i.e., choose a subset of all candidate solutions with the same rank). When this occurs, a scheme for resolving ties is needed. In this work, we choose between equally Pareto ranked solutions based on performance of the $equity$ objective. Recall that the equity objective assesses the fairness of the competition by ensuring all teams play, as far as possible, an equal number of home games. By using this metric as the means of resolving ties, evolutionary selection pressure will drive the population to find solutions with a minimum $equity$ score. As this objective captures a constraint imposed by the organising authority, this is indeed desirable, forcing unwanted constraint-violating solutions from the population. Any subsequent ties are resolved by random choice of the remaining "equal" solutions.

Proceedings of the 2007 IEEE Symposium on Computational Intelligence in Scheduling (CI-Sched 2007)

V. EXPERIMENTAL RESULTS

The multi-objective evolutionary algorithm we use in this work is a form of a $(n+n)$ evolutionary strategy. Selection is determined using the ranking scheme detailed in Section IV-D. We use an elitism rate of 50%, thus meaning we preserve the best 50% of the population from one generation to generation. Initial experiments demonstrate that a population of 300 yields good results in a reasonable amount of time. All experiments were run for 1000 generations.

In the experiments below, the starting populations were initialised to mutated versions of either the 2006 or 2007 Super 14 fixture. The corresponding fixture was used as the template fixture (see Section IV-A) in the generation of a phenotype for a candidate solution.

A. Pareto Front Evolution

Fig. 1 plots the Pareto front at different stages (generations) during a single run of our multi-objective evolutionary algorithm. Noting that the *equity* objective acts more as a constraint (the organising authority strongly prefers "fair" fixtures) than a varying objective, we only plot the non-dominated solutions with an *equity* measure of zero.

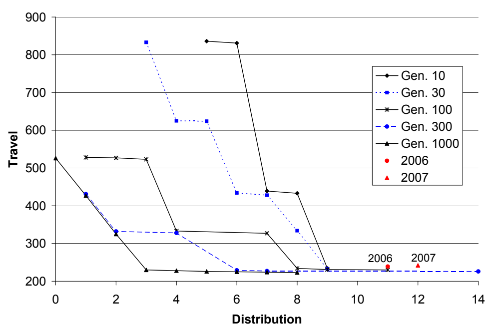

Fig. 1. Pareto front evolution for a run of our multi-objective evolutionary algorithm. Only non-dominated solutions with $equity = 0$ are plotted.

Fig. 1 shows that our multi-objective evolutionary algorithm is able to simultaneously optimise both the travel and distribution objectives at the same time, while maintaining the constraint captured by the equity objective. Fig. 1 also shows that the evolutionary algorithm is able to generate a good range of different solutions, each trading-off the different objectives by varying amounts. During the course of the run we see good coverage across both objectives, highlighting the evolutionary algorithm's ability to explore different parts of the search space. Indeed, by the end of the run, solutions with a wide variety of *distribution* scores have been discovered. The narrowing of the Pareto front towards the end of run highlights that the algorithm is able to locate a few solutions good in both objectives.

B. Hypervolume

The *hypervolume* metric for non-dominated front comparison [21], [22] measures the ratio of the hypervolume dominated by a front to the hypervolume dominated by the idealised minimum. It provides a numerical measure that rewards both closeness to the Pareto optimal front and the extent of the obtained non-dominated front. Importantly, the hypervolume metric is more robust than other unary numerical metrics [23].

Fig. 2 plots the hypervolume of the Pareto front (including solutions that violate the equal-home-games constraint) at different stages (generations) during a single run of our multi-objective evolutionary algorithm. Fig. 2 also plots the hypervolume contributed by just the fixtures that observe this constraint (i.e. those with $equity = 0$). After dividing the travel objective by 100, the reference point for the hypervolume calculations was (20,20,20).

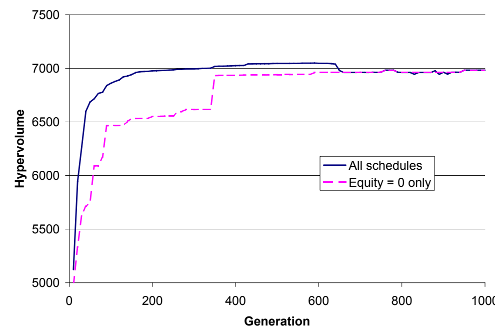

Fig. 2. Progression of the hypervolume of the Pareto front for a run of our multi-objective evolutionary algorithm

Fig. 2 shows the benefits of using an objective to capture the equal-home-games "constraint" desired by the organising authority. We observe that initially, much of the Pareto front consists of non-preferred solutions that have an "unequal" number of home games. Over time, we see an increase in the proportion of the hypervolume contributed by preferred solutions (solutions with $equity = 0$), suggesting an emergence of these solutions in the population. This gives us the effect we seek — some exploration of non-preferred solutions early in the run, with a general convergence to preferred solutions towards the end of the run.

Fig. 2 also confirms that our algorithm is able to maintain a good rate of progression during the entire run — we observe a general increase in the hypervolume of the Pareto front over time. Occasional drops in the hypervolume of the whole population (including a significant drop around generation 650) correspond to generations where solutions with $equity = 0$ are selected over other solutions. This leads to a loss in the spread of the front and consequently a drop in hypervolume. However this is the preferred outcome late in the evolution, as only solutions with $equity = 0$ would be considered viable fixtures anyway.

C. Comparison with Existing Super 14 Fixtures

Also marked on Fig. 1 are the fixtures employed by the Super 14 organising authority for the 2006 and 2007 seasons. Examination of this figure shows that our evolutionary algorithm is quickly able to produce solutions that dominate both of these fixtures.

Table I compares the final Pareto optimal solutions found from twenty runs of our multi-objective algorithm, listing the performance of each solution using the different objectives described in Section IV-B.

	Objective		
Solution	Equity	Distribution	Travel
Fixture A	0	0	326
Fixture B	0	1	321
Fixture C	0	2	231
Fixture D	0	3	224
Fixture E	0	4	222
Fixture F	0	5	221
Fixture G	0	7	220
Fixture H	0	10	219
Fixture I	0	12	218
2006 Super 14 Fixture	0	11	239
2007 Super 14 Fixture	0	12	242

TABLE I

PERFORMANCE OF THE PARETO OPTIMAL SOLUTIONS TAKEN FROM TWENTY RUNS OF OUR MULTI-OBJECTIVE EVOLUTIONARY ALGORITHM

We see from Table I that our multi-objective evolutionary algorithm is able to produce a good range of solutions that trade-off the different objectives by varying amounts, while still maintaining the equal-home-games constraint (i.e. a $equity$ score of zero). We also see from Table I that most of the evolved Pareto solutions strictly dominate both the 2006 and 2007 Super 14 fixtures.

Table II lists the details of the different Pareto optimal solutions listed in Table I. For each different solution, the break-down in travel for each team is listed using the format X-Y, where X represents the number of long journeys (crossings of the Indian Ocean) without significant rest (proceeding or following a bye, or at the start or end of the season) and Y the number of other inter-regional trips made without significant rest.

Table II shows that both of the recent fixtures used by the Super 14 competition have a maximum of two long journeys for any single team in the competition without significant rest. The 2006 fixture has 9 teams with two long journeys, while the 2007 fixture has 7 teams with two long journeys. Note however that the 2007 fixture is somewhat more inequitable — one team (Eastern Australia Team 3) has zero long journeys compared to 7 teams with two!

We see from Tables I and II that the optimisation of the distribution objective has a significant negative impact on the travel objective. Table II shows that the best distribution solution (Fixture A) represents a fixture in which four teams cross the Indian Ocean three times without significant rest. Fixture B, which represents a solution with $distribution = 1$ contains a better total number of short journeys (21 versus the 26 of Fixture A), but has 7 teams crossing the Indian Ocean three times without significant rest. This large number of long journeys means Fixture B is unlikely to be a solution that the Super 14 organising authority would employ.

However, Fixture A, with far fewer short journeys than either of the 2006 and 2007 fixtures, and a perfect distribution of games (contrast Fixture A's $distribution$ score of zero with the $distribution$ scores of 11 and 12 for the 2006 and 2007 fixtures) makes an interesting viable alternative for the Super 14 organising authority to consider. With a $distribution$ score of zero, Fixture A promises a better distribution of games across the season, allowing for increased revenue (through better broadcasting opportunities). This improvement comes at the cost of two teams who now have to complete one additional long journey. Note that the Western Australian team is somewhat harshly done by in this fixture, needing to travel seven times in total without significant rest — the cost of achieving much better performance in the distribution objective.

Fixtures G through I minimise the total travel required by the teams in the competition, sacrificing performance on the distribution objective to achieve the better travel performance. While these solutions have comparable distribution scores to the 2006 and 2007 fixtures, their performance relative to Fixtures C and D probably mean the lost revenue associated with the uneven distributions does not justify employing them over either Fixture C or D. However, they do make interesting alternatives that might be considered should external forces significantly change the trade-off between travel and distribution (e.g. should travel become restrictively more difficult). This is one advantage of the multi-objective approach over a single-objective approach — the multi-objective approach is able to present a range of solutions that trade-off the different objectives by varying amounts, leaving the choice of which solution to use to the expert decision maker who can consider external variables/factors impossible to capture in the model.

Fixtures C and D (with $distribution$ scores of 2 and 3 respectively) offer strong alternatives to the recent fixtures employed by the Super 14 organising authority. Their far better performance in both the distribution and travel objectives relative to the 2006 and 2007 fixtures mean significant improvements over these existing solutions. These solutions are able to achieve better performance in the travel objective by positioning the timing of byes in the fixture to ensure teams have sufficient time to recover from travelling.

Note that the more equitable distribution of travel in Fixture D relative to Fixture C may give this fixture the edge, offering a more palatable solution for the organising authority to "sell" to the teams in the competition (the organising authority does not want to be seen to be playing favourites). Note that both Fixtures C and D arguably have a more equitable travel distribution than both of the 2006 and 2007 fixtures. It is also interesting that these similarly performing solutions were derived from different seedings of the initial population, further showing the robustness of this approach.

Solutions E and F offer further trade-offs between the travel and distribution objectives that may be of interest to the organising authority. In summary, comparison of the different fixtures shows that the evolutionary algorithm is able to "tweak" general solutions in order to optimise performance in any single objective. The good spread of resultant solutions offers significantly different solutions that the organising authority may consider as alternatives to the recent fixtures they have used.

	2006	2007	Fix. A	Fix. B	Fix. C	Fix. D	Fix. E	Fix. F	Fix. G	Fix. H	Fix. I
Eastern Australia Team 1	2-3	2-4	2-4	2-1	2-2	2-1	2-1	2-1	2-2	2-2	2-2
Eastern Australia Team 2	1-4	1-7	2-0	2-0	2-3	1-2	1-2	1-2	1-2	1-2	1-2
Eastern Australia Team 3	2-5	0-3	2-2	3-1	2-1	2-4	2-4	2-4	2-4	1-4	1-4
Western Australia Team 1	1-5	2-5	3-4	2-2	2-2	2-3	2-2	2-2	2-3	1-1	2-0
New Zealand Team 1	2-4	2-2	1-1	1-2	2-3	2-1	2-1	2-1	2-1	2-3	2-3
New Zealand Team 2	1-2	1-4	3-2	2-2	1-3	1-4	1-4	1-4	1-0	2-2	2-2
New Zealand Team 3	2-3	1-4	2-1	3-1	1-5	2-1	2-1	2-1	2-0	2-1	2-1
New Zealand Team 4	2-1	2-2	2-0	2-2	2-2	1-1	1-1	1-1	1-1	2-3	2-2
New Zealand Team 5	2-3	2-2	2-3	3-3	1-5	2-2	2-2	2-2	2-2	2-0	2-0
South Africa Team 1	2-2	1-1	3-1	3-1	2-0	2-0	2-0	2-0	2-0	2-0	2-0
South Africa Team 2	1-2	2-2	3-2	3-2	1-3	2-1	2-0	2-0	2-0	1-0	1-0
South Africa Team 3	2-1	1-2	2-3	3-1	2-1	1-1	1-1	1-1	1-2	2-1	2-2
South Africa Team 4	2-2	2-2	2-2	2-2	2-1	2-1	2-1	2-1	2-1	2-0	2-0
South Africa Team 5	1-2	1-2	2-1	3-1	1-0	1-2	1-2	1-1	1-2	1-0	1-0
Maximum Indian Ocean journeys	2	2	3	3	2	2	2	2	2	2	2
Sum of other inter-regional journeys	39	42	26	21	31	24	22	21	20	19	18
Seed fixture	-	-	2006	2006	2006	2007	2007	2007	2007	2006	2006

TABLE II

THE DIFFERENT PARETO OPTIMAL SOLUTIONS FROM TABLE I

VI. CONCLUSIONS

In this paper, we have presented a multi-objective evolutionary algorithm for fixture determination for the Super 14 Rugby competition. Like many team sports that involve teams spread over significant distances, fixture designers for the Super 14 competition face the difficult problem of balancing a number of different, often conflicting, factors like competition fairness, amount of travel, availability and distribution of games, and of course revenue.

Our multi-objective approach to this problem produces a range of different fixtures, each varying the trade-offs in the objectives by differing amounts. This provides the organising authority the ability to explore different options, allowing them to choose the option that best suits the requirements of the day. Our experiments show that this multi-objective approach is able to evolve solutions that strictly dominate recently employed fixtures, promising better returns in every measure of success.

REFERENCES

[1] "Super 14," November 2006, URL: `http://en.wikipedia.org/wiki/Super_14`.

[2] "Investec new naming rights partner of rugby's Super 14," October 2006, URL: `http://www.rugby.com.au/news/investec_new_naming_rights_partner_of_rugbys_s14,49808.html/news/archive/section/40074`.

[3] "More for players in new SANZAR deal," December 2004, URL: `http://www.worldcupweb.com/WCrugby/item.asp?NewsID=845`.

[4] "SANZAR announces ABC Asia Pacific as broadcast partner," January 2006, URL: `http://aru.com.au/news/sanzar_announces_abc_asia_pacific_as_broadcast_par,40373.html`.

[5] C. A. Coello Coello and G. B. Lamont, Eds., *Applications of Multi-objective Evolutionary Algorithms*. World Scientific Publishing, 2004.

[6] B. Huich, I. Miguel, I. P. Gent, and T. Walsh, "CSPLib: A problem library for constraints," January 2006, URL: `http://csplib.org`.

[7] K. McAloon, C. Tretkoff, and G. Wetzel, "Sports league scheduling," in *Proceedings of the Third ILOG Optimization Suite International User's Conference*, 1997.

[8] J.-P. Hamiez and J.-K. Hao, "A linear time algorithm to solve the sports league scheduling problem (prob026 of CSPLib)," *Discrete Applied Mathematics*, vol. 143, pp. 252–265, 2004.

[9] ——, "Solving the sports league scheduling problem with tabu search," *Lecture Notes in Artificial Intelligence*, vol. 2148, pp. 24–36, 2001.

[10] K. Easton, G. Nemhauser, and M. Trick, "Solving the travelling tournament problem: A combined integer programming and constraint programming approach," *Lecture Notes in Computer Science*, vol. 2740, pp. 100–109, 2003.

[11] R. Miyashiro and T. Matsui, "Round-robin tournaments with a small number of breaks," Department of Mathematical Informatics, The University of Tokyo, Mathematical Engineering Technical Reports METR 2003-29, 2003.

[12] E. Burke and P. Ross, Eds., *The Practice and Theory of Automated Timetabling*. Springer Verlag, 1996.

[13] J. Schönberger, D. C. Mattfeld, and H. Kopfer, "Automated timetable generation for rounds of a table-tennis league," in *Proceedings of the 2000 Congress on Evolutionary Computation (CEC'00)*. IEEE Publications, 2000, pp. 277–284.

[14] J. T. Yang, H.-D. Huang, and J.-T. Horng, "Devising a cost effective baseball scheduling by evolutionary algorithms," *Proceedings of the 2002 Congress on Evolutionary Computation (CEC'02)*, pp. 1660–1665, 2002.

[15] J. Schönberger, D. C. Mattfeld, and H. Kopfer, "Memetic algorithm timetabling for non-commercial sport leagues," *European Journal of Operational Research*, vol. 153, no. 1, pp. 102–116, 2004.

[16] D. Costa, "An evolutionary tabu search algorithm and the NHL scheduling problem," Department of Mathematics, Swiss Federal Institute of Technology, Tech. Rep. ORWP 92/11, 1992.

[17] L. Barone, L. While, P. Hughes, and P. Hingston, "Fixture-scheduling for the australian football league using a multi-objective evolutionary algorithm," in *Proceedings of the 2006 Congress on Evolutionary Computation (CEC'06)*. IEEE Publications, 2006, pp. 954–961.

[18] D. E. Goldberg, *Genetic Algorithms in Search, Optimization & Machine Learning*. Addison-Wesley, 1989.

[19] C. M. Fonseca and P. J. Fleming, "Genetic algorithms for multiobjective optimization: formulation, discussion and generalization," in *5th International Conference on Genetic Algorithms*, 1993, pp. 416–423.

[20] J. Dinitz, E. Lamken, and W. Wallis, "Scheduling a tournament," in *Handbook of Combinatorial Designs*, J. Dinitz and C. Colbourn, Eds. CRC Press, 1995, pp. 578–584.

[21] E. Zitzler, "Evolutionary algorithms for multiobjective optimization: methods and applications," Ph.D. dissertation, Swiss Federal Institute of Technology, Zurich, Switzerland, 1999.

[22] L. While, P. Hingston, L. Barone, and S. Huband, "A faster algorithm for calculating hypervolume," *IEEE Transactions on Evolutionary Computation*, vol. 10, no. 1, pp. 29–38, January 2006.

[23] E. Zitzler, L. Thiele, M. Laumanns, C. M. Fonseca, and V. G. da Fonseca, "Performance assessment of multiobjective optimizers: An analysis and review," *IEEE Transactions on Evolutionary Computation*, vol. 7, no. 2, pp. 117–132, April 2003.