Scheduling Cricket Umpires Using Neighbourhood Search The Dramatic Impact of a Simple Change in Neighbourhood Definition

Mark Wright

One widespread application of employee timetabling is scheduling officials for sports leagues. This can be complex, as there are frequently many objectives, preferences and constraints (hard and soft) to accommodate. While one can sometimes use optimising software for such problems, frequently heuristics are necessary. This often involves neighbourhood search.

In a cricket league, two umpires must be appointed to each match during the cricket season, which typically lasts between 16 and 20 weeks. There may be several divisions at different levels.

A computer system was developed to schedule umpires for the Devon League in south-west England, with four hierarchically organised divisions of ten teams each.

The hard constraints are:

- every match requires two umpires
- no umpire may be used on a date for which he has declared himself unavailable
- no umpire may be used for two matches on the same date
- some specific rules: for example, some umpire-team combinations are considered infeasible

Soft constraints concern the *target* number of matches for each umpire in each division over the season – it is usually possible to satisfy all targets exactly if they add up correctly for each division, but the system accepts, during the search, solutions which break soft constraints, using penalty costs.

Objectives concern travel distances and spread of umpires between teams, home grounds and each other. Further adaptations weight travel more heavily at certain times of year and encourage situations where two umpires can travel together, thus reducing costs.

The original neighbourhood definition was simple: replace one umpire by another for a specific match (a *replacement* perturbation), or exchange two umpires between two matches (a *swap* perturbation). The search technique was the variant of Simulated Annealing described in Wright (2001), using subcost information, and considering perturbations in a predefined order rather than at random. For precise details, see Wright (2007).

Recently, the system was adapted for the Home Counties League in South-East England. The structure is slightly different, with one top division and two others of equal status, organised geographically. The computer system was amended so that two divisions could be considered at an equal *level*; thus soft constraints concerned the targets for each umpire at each level. This was then implemented in practice.

While the client was happy with the results and implemented them, they contained disappointing features. In particular, on several occasions an umpire/team combination occurred twice within three weekends, which rarely happened for the Devon League, and there were other unexpectedly high costs.

Further analysis uncovered a possible reason, concerning the umpires' targets. The Devon League deliberately allows some slack for most umpires, enabling flexibility when schedules must be changed, e.g. because of illness. Thus, for example, an umpire available for 17 of the 18 weekends might be given an overall target of 15 matches. Thus there was plenty of scope for swapping umpires between dates without breaking any hard constraints.

However, mainly because of a shortage of suitable umpires, the Home Counties League uses most umpires as often as possible. Thus, once the umpires' targets have been met, swaps involving such umpires are feasible only between two matches taking place on the same date. This greatly restricts the search; moreover, if an umpire currently meets his targets there are few perturbations available which do not break these, i.e. only those between matches at the same divisional level and on the same date.

Variable Neighbourhood Search (Hansen & Mladenovic, 2001) was considered, but not implemented since it can increase neighbourhood size enormously. Instead, a new type of perturbation was defined, involving two umpires, two dates and four matches, such that the umpires swap matches on both dates (a *double* perturbation). This allows umpires to change division levels on each date without breaking any soft constraint.

Tests on Home Counties data showed that including such perturbations reduced the total cost by about 22%, averaged over 100 runs. Surprisingly, given that it had not been apparent that there was a problem, almost as large an effect was demonstrated for the Devon League, where the total cost was reduced by about 20%.

Other interesting features were apparent. About 73% of perturbations were doubles, compared with 11% for replacements and 16% for swaps. Replacement perturbations were accepted 0.27% of the time; swaps 1.06% of the time; and doubles only 0.02% of the time, which is interesting given the great improvements seen from including such perturbations. Overall, only 0.21% of perturbations were accepted, compared with 0.64% for the original system without double perturbations.

Moreover, 18.3% of accepted replacements were accepted in the first 0.05% of the run, compared with 0.61% of swaps and 0.65% of doubles; and no replacements were ever accepted in the last 1% of the run, compared with 0.053% for swaps and 0.016% for doubles. (These figures are for the Home Counties League – similar patterns were evident for the Devon League.)

This suggests that the new system's success may owe much to the presence of three distinct "meta-types" of perturbation:

1. "Get rich quick", making strong early gains, eliminating most unnecessary major costs such as soft constraint violation penalties, but of little value when fine-tuning later in the search;
2. "Steady as she goes", valuable throughout the process, especially when fine-tuning at the end;
3. "Occasional gem", only rarely useful, but capable of producing dramatic improvements.

It is hypothesised that any neighbourhood search process is more likely to produce high-quality solutions to complex problems if it contains perturbations of all three meta-types.

References

Hansen, P. and Mladenovic, N. (2001) "Variable Neighborhood Search: principles and applications". *European Journal of Operational Research* **130(3)**: 449-467.

Wright, M.B. (2001) "Subcost-guided simulated annealing", in: "Essays and surveys in metaheuristics", eds. C.C. Ribeiro and P. Hansen, chapter 28, 631-639 (Kluwer Academic Publishers, Boston).

Wright, M.B. (2007) "Case study: problem formulation and solution for a real-world sports scheduling problem", *Journal of the Operational Research Society* **58 (6)**, 439–445.