

# **SCHOOL TIMETABLING USING HEURISTIC SEARCH**

by

**MIKE WRIGHT**

**Department of Management Science, Lancaster University, UK**

This paper concerns a computer system which produces the bulk of the timetable for a large comprehensive school in England. The complexities of the school's lesson structure are discussed and the various constraints and objectives described. The timetable thus produced was successfully implemented for the academic year starting in September 1994 and was considered by the school to represent a marked improvement on previous timetables as well as being achieved much more swiftly.

The solution method involves four phases of heuristic search with little or no manual intervention necessary. In contrast with other timetabling systems, the system completes all the difficult parts of the process to a high-quality standard, with only the final straightforward stages being left to the timetabler. Details of the solution method are outlined and discussed in further detail in an appendix, especially the more innovative parts which involve a form of tabu search with influential diversification guided by the values of the subcosts as well as the overall cost. The system could be generalised so as to be applied to other schools.

**Key Words:**       timetabling, tabu search, influential diversification, subcost-directed search, schools, heuristics, education.

## **INTRODUCTION**

The problem of constructing timetables for educational institutions is a classical OR problem. The academic literature on the subject has been prolific: de Werra <sup>1</sup> wrote a full review of the area and papers have since been published by many other authors including Abramson <sup>2</sup>, Hertz <sup>3</sup>, Balakrishnan *et al.*<sup>4</sup>, Ferland and Lavoie <sup>5</sup>, Dige *et al.*<sup>6</sup>, Johnson <sup>7</sup>, Ferland and Fleurent <sup>8</sup> and Thompson and Dowsland <sup>9</sup>. Many papers describe not only theory but also details of successful applications.

Timetabling is a necessity for all educational institutions including schools, universities and colleges. Moreover, even for a single institution such as a secondary school, there are many different types of timetabling problem: scheduling lessons, rooms, examinations, etc. The application described here involves the timetabling of lessons for a large and complex secondary school (for pupils aged between 11 and 18) in England.

Lesson timetabling for secondary schools covering the entire age range is probably the most difficult form of timetabling widely undertaken. The main reasons for this are not related primarily to the size of problem, although the problems usually are very large, but are as follows.

- (1)    The lesson structure can vary widely between year groups, involving "blocks" of lessons formed using different principles.
- (2)    Almost every school has its own particular rules, preferences and idiosyncrasies regarding the way in which its lessons are timetabled.
- (3)    There are always multiple objectives to be considered. Some of these are of a wholly practical nature but others, equally important, are more qualitative issues related to important educational requirements and preferences.
- (4)    There is very little temporal flexibility; most of the schoolchildren must have a lesson timetabled at all times. This can make it very difficult to achieve feasibility, let alone solutions of high quality.

For these reasons, most computer systems used in practice for school lesson timetabling are merely

decision aids. While such systems can indeed be very useful in checking, summarising and printing and can save a great deal of time, they do not do the timetabling themselves. All decisions are made by the user of the system.

Other systems go further and construct partial timetables. Although such systems may claim a "95% insertion rate" (i.e. 95% of lessons timetabled), they often do so without proper regard for the more qualitative issues. Moreover, the most difficult part of the work still remains to be done, since the final 5% is always much the hardest. It has to be accommodated manually, which can usually only be done by means of substantial restructuring of the given timetable with, often, consequent reduction in quality.

The system described in this paper, however, is much more ambitious than this. It produces the full lesson timetable for a particular school automatically with almost no need for manual intervention. As it stands it cannot be simply transferred without amendment to other schools, but the school concerned is large and has a complex and rigid lesson structure. The approaches and techniques described here are applicable to the production of lesson timetables for almost any school and could form the basis of a more general system in the future. The system described therefore represents a major advance in the field of school timetabling by computer.

## **THE SCHOOL**

The school for which the system was created is a comprehensive school in Lancashire, England, which contains over 1400 pupils of all levels of ability. It is the second largest school in Lancashire and one of the largest in the country.

There are over 80 teachers, nearly all of whom work full-time. The subjects taught include the full range of the UK National Curriculum (for a good summary of the requirements of the National Curriculum see, for example, Moon <sup>10</sup>). The school operates for five days a week, about forty weeks a year. There are five hour-long periods each day (three in the morning and two in the afternoon) and hence 25 time-slots (or, simply, "slots") per week.

The first three years (years 7 to 9) are known as the "lower school" and each of these year-groups is divided into eight classes. The rest of the school (years 10 to 13) forms the "upper school"; years 12 and 13 are also collectively known as the "sixth form". Pupils in years 7 to 11 have a lesson every period of every week. Each sixth-former has a few free periods each week because of the option structure, but there are sixth-form lessons timetabled in every period of every week.

The relatively long periods (compared with most other UK schools) and the rigidity of the National Curriculum unfortunately mean that the timetable cannot currently operate under a weekly cycle. For every subject to be given at least its minimum amount of time for every year group every week would require more than twenty-five periods. A fortnightly cycle is therefore in operation.

## **LESSON BLOCKS**

The basic timetabling problem is to allocate each lesson to one of the fifty slots in the fortnight. However, the situation is considerably complicated by the fact that many lessons are grouped together in "blocks" of two or more lessons which must all take place simultaneously.

This is the case for the whole of the upper school, for which all pupils in a year group are timetabled together. These blocks are of various types. Some subjects in years 10 and 11 are taken by the whole year simultaneously to allow streaming by ability (English, Mathematics, Science) or rotation of activities (Technology); some blocks are option blocks, where pupils choose one subject from a given list; years 12 and 13 do games together; and the whole school has simultaneous "tutorial" lessons once a fortnight.

In the lower school there is blocking for Mathematics, Games, Technology, Art, Rural Science, Physical Education, Tutorial, some English and some French. These blocks may involve anywhere between 2 and 8 classes in the year group. The other lessons (Biology, Chemistry, Physics, French, German, Drama, Religious Education, Music, History, Geography and the rest of English and French) are not blocked. These are known as "freestanding" lessons and can be timetabled independently of all other lessons.

Most of the lessons in a block have a prespecified set of teachers who must teach a lesson in that block. However, there is sometimes flexibility here, especially for years 12 and 13.

For example, consider "block 1" for year 13. This block must occur ten times per fortnight and comprises lessons in Mathematics, Home Economics, Geography and History. However, for each subject there is more than one teacher involved. For example, teacher A teaches geography in Block 1 four times per fortnight; teacher B teaches geography in Block 1 for four different lessons in a fortnight; and teacher C teaches the other two geography lessons per fortnight. However, teachers A, B and C do not have to be timetabled in any particular order within the fortnight, so they can be interchanged among the ten "block 1" lessons. Similar arrangements apply for the other subjects in the block.

In practice there may be further flexibility in some cases; some prespecification of teachers represents preferences rather than rigid constraints. When the timetable was compiled manually many of these preferences were overridden in practice in order that a feasible and satisfactory timetable could be achieved, but the aim here is that this should occur as infrequently as possible, if at all.

The blocking arrangements do cause great difficulties to the timetabler, mainly because there are two types of block. For one type (e.g. in years 10 and 11) similar subjects are grouped together, whereas for the other type (in years 12 and 13) they are intentionally kept apart. This means that most individual teachers will find themselves blocked with a wide range of other teachers.

Henceforth, for simplicity, all lessons and blocks of lessons are referred to as blocks. Freestanding lessons are simply blocks of size 1. The block is the unit that must be timetabled.

## **CONSTRAINTS AND OBJECTIVES**

Some of the most basic constraints are treated as hard constraints. These include the requirement that every block must be timetabled and that there are only five lessons per day and ten days per fortnight. However, many other constraints are treated merely as important objectives. This approach has great flexibility and is in many ways much simpler. By treating any conceivable arrangement as feasible, if not actually acceptable, it is easy to find a starting solution which can then be improved by means of iterative heuristic search procedures. Thus it may be possible to remove all the unacceptable features as the process continues.

A notional "subcost" is associated with each objective - the worse a timetable is with regard to that objective, the higher the subcost. Summing weighted subcosts for all the different objectives gives a total "cost" for the timetable. These subcosts and the total cost are of course not financial, but the implication is that one timetable is better than another if its total cost is lower.

The weights were chosen to give an accurate reflection of the relative importance of the objectives. Of course, in practice, it is very difficult, if not impossible, to find weights which do this accurately, partly because the quality of a timetable is to a large extent subjective. However, it is not necessary to be absolutely precise in the determination of subcosts and weights, as long as the low-cost timetables achieved by the system are acceptable in practice.

**The subcosts are as follows.**

- (1) The number of "class clashes", i.e. the number of occasions when a class has two different lessons timetabled at the same time. Obviously it is crucially important that this number is zero in the final timetable, so the associated weighting is extremely high (a million units of cost for each occurrence).
- (2) The number of "teacher clashes", i.e. the number of occasions when a teacher has to teach two different classes at the same time. The associated weighting is 500 units, much less high than in (1) because many of the teacher specifications are preferences rather than being absolutely fixed.
- (3) The number of "department clashes", i.e. the number of occasions when more teachers are required from a particular department (e.g. English, Science, History etc.) than are available. In this case, changing the teacher specifications will not help, so the associated weighting is higher (5000 units).
- (4) The number of room shortages. Some blocks require particular rooms or types of room; in particular, most science lessons require a laboratory. Clearly if too many science lessons are timetabled in the same slot, not enough laboratories will be available. The weighting here is fairly high (100), but not prohibitively so, since in practice it may be possible to hold some science lessons in ordinary classrooms.
- (5) The number of "double" requirements not met. Each of the four main block types in the sixth form is required to be timetabled in two consecutive slots one day a week, preferably without a break in between. The weighting here is higher for some types of block, where double periods are essential, than for others, where they are merely desirable. Costs range from 200 units to 10000.
- (6) Costs relating to the pattern of subjects for each class. All educational experts agree that this is a very important consideration. For example, Simper <sup>11</sup> states as his primary criterion for the quality of a timetable: "Does each subject - maths, English, history - have a reasonable distribution over the week, so that its lessons do not all occur in one day or all during the last period in the afternoon?"

The quality of the subject pattern is complex to define and measure, but its supreme importance means that it must be modelled explicitly.

The main factors giving rise to costs are as follows.

- \*\*\* The same subject occurring more than once on the same day, except where a double period is required. The weighting is especially high if the two periods are consecutive.
- \*\*\* The same subject being the last period on one day and the first period on the next.
- \*\*\* The same subject occurring predominantly in one half of the week.
- \*\*\* The same subject occurring at the same time of day twice in a week. The weighting is very high if this is the last period of the day, fairly high if it is the first period of the day and low otherwise.

These costs also apply (though with a lower weighting) to subjects which are in some ways similar. For example, if a class has both French and German on the same day, or Chemistry followed immediately by Biology, this gives rise to a cost.

Depending on the severity of the subject imbalance, the cost ranges from almost zero (0.02)

to very large indeed (up to 50,000).

(7) Various other costs relating to specific lessons or blocks. Each requirement is entered in a special file and given its own weighting to reflect its importance. The types of these requirements include the following.

- \*\*\* A particular block must be timetabled in a particular slot. For example, most Games lessons must be at given times on given days, to allow coordination with other schools for matches.
- \*\*\* A particular block must not be timetabled in a particular slot. For example, the fortnightly tutorial lesson must not be the last lesson of a day.
- \*\*\* Two blocks must be timetabled together, i.e. in the same slot, For example, years 12 and 13 must have Games together.
- \*\*\* Two blocks must not be timetabled together. This may be needed if specialised equipment is involved. Requirements of this type are also generated automatically if there is no slot for which both blocks are feasible (this is necessary for the first two stages of the solution procedure - see below).
- \*\*\* Two blocks must be timetabled on the same day. For example, both may require the same part-time teacher who only works one day a week.
- \*\*\* Two blocks must not be timetabled on the same day. For example, no class should have Physical Education and Games on the same day.
- \*\*\* A particular teacher must be free immediately before a particular block, or before a particular double block. This may be for preparation purposes.
- \*\*\* A rather more complicated requirement involving a group of blocks and a group of slots. The number of blocks from the group which are timetabled in a slot from the group of slots must not be above (below) a given maximum (minimum).

For example, it is thought undesirable for children below sixth-form level to be taught Mathematics (or English or French) on more than one afternoon in a week, because they may be tired and therefore less able to produce the necessary levels of concentration for academic subjects. This particularly applies to the lower-ability groups in the lower school. So one group of blocks is for Year 8 Mathematics (lower ability), the associated group of slots consists of all afternoon slots and the specified maximum is 1 (no minimum).

Another example applies to blocks including Food Science. Suppose that a particular set of children takes Food Science in blocks X, Y and Z. Once a week each Food Science class must prepare and eat its own lunch. So X, Y and Z form a group of blocks, with the associated group of slots being all the immediate pre-lunch and post-lunch slots, i.e. slots 3, 4, 8, 9, 13, 14, 18, 19, 23 and 24. The associated minimum is 1 (no maximum).

The costs for these requirements vary considerably, between 1 and 5,000 units.

## **TRANSFORMING A FORTNIGHTLY PROBLEM INTO A WEEKLY ONE**

Although any timetable produced for this school must of necessity have a fortnightly cycle, the head teacher and governing body of the school consider it highly desirable that the timetable should come as close as possible to following a weekly cycle. So, for example, a block type that must

occur seven times in a fortnight would ideally occur in the same three slots every week with one extra occurrence once a fortnight.

The system constrains this to happen, though manual override is possible. In the example given above, there would be three "weekly" blocks and one "fortnightly" block. When a weekly block is allocated to a slot, this applies to both weeks of the cycle.

This transforms the original fortnightly timetabling problem into a weekly one, with the added complication that two fortnightly blocks may be allocated to the same slot for a given class. This does, however, give rise to possible "hidden" infeasibilities.

Consider the following example, where all the blocks concerned are fortnightly blocks.

Block 1, for year 10, needs Teacher A and Teacher B.

Block 2, for year 10, needs Teacher A and Teacher C.

Block 3, for year 11, needs Teacher B and Teacher C.

On the face of it, it would appear possible to timetable all three of these fortnightly blocks in the same slot, since no teacher is required more than twice in a fortnight. However, if this is done, two of the three blocks must occur together, in the same week of the cycle, which is impossible. The system therefore generates such sets of fortnightly blocks and adds high costs if they are all timetabled in the same slot.

Apart from this complication, the transformation of the problem from a fortnightly one to a weekly one makes the problem more manageable, since the number of blocks to be timetabled is reduced by about 40% and the number of slots is halved. It does, however, mean that the problem is now more constrained than it need be.

## **SOLUTION METHOD**

The solution procedure has six stages, the first four of which form the computer system. Between stages the partial solution needs to be agreed to be acceptable by the school, possibly with manual alterations. The first two stages put blocks "together" (i.e. in the same slot as each other) without specifying the particular slots. Stage 1 is mainly for Upper School blocks, which are more difficult to timetable. Stage 2 adds blocks of types which occur more than once a week in the lower school. Stage 3 then determines the precise slots and Stage 4 includes the remaining non-freestanding blocks in the Lower School.

Stages 5 and 6 are manual stages. In Stage 5 the freestanding lessons are timetabled and in Stage 6 rooms are allocated. There are advantages to keeping these stages manual as a greater knowledge and understanding of the school is necessary and there is much more flexibility in practice. Anyway, these stages are not difficult or very time-consuming to carry out manually.

All four computer stages proceed by means of local search, using perturbations as described below.

## **PERTURBATIONS**

**The perturbations are of the following types:**

- (a) Teacher-only swaps. For example, suppose two blocks of type 1 for year 13 are timetabled in slots 1 and 2 (i.e. the first two periods on Mondays). Both blocks involve geography (as well as several other subjects), but they could have different geography teachers. Then a teacher-only swap might involve simply swapping the two geography teachers between the

two blocks.

- (b) Whole-year swaps between two slots. For example, all the lessons for Year 9 are swapped between slots 1 and 2.
- (c) "Kempe Chain" perturbations between two slots. For example, a block is moved from slot A to slot B; then any block already timetabled in slot B which has a class in common with the first block is moved to slot A; then any block already timetabled in slot A which has a class in common with any of these blocks is moved to slot B; etc. Whole-year swaps are often particular cases of Kempe Chain perturbations; care is taken to avoid duplication. Kempe Chain perturbations were suggested by Morgenstern and Shapiro <sup>12</sup> and have been used by Thompson and Dowsland <sup>9</sup> for examination scheduling, a problem closely related to the school timetabling problem.
- (d) Fortnightly block moves. For example, a fortnightly block is moved from slot A to slot B.
- (e) Fortnightly block swaps. For example, two fortnightly blocks are swapped between slots A and B.
- (f) Full swaps between slots. Everything for all years is swapped between slots A and B.
- (g) Full morning swaps. Everything for all years is swapped between two mornings (the first three periods of the day are in the morning).
- (h) Full late-morning swaps. Everything for all years is swapped between two late mornings (i.e. the second and third periods of the day).
- (i) Full afternoon swaps. Everything for all years is swapped between two afternoons (the last two periods of the day are in the afternoon).

The last four types of perturbation are not used in Stages 1 and 2; they would have no effect because the identities of the slots are not established until Stage 3.

## STAGE 1

Stage 1 timetables the whole of the Upper School plus Tutorial and Technology for the Lower School. The Tutorial lessons (once a fortnight for each class) need to be simultaneous throughout the school, so must all be timetabled together. Lower School Technology is included because it is all blocked, often the whole year together, with at least six teachers in each block (the first version of the system did not include Lower School Technology in Stage 1 and this led to severe difficulties in Stage 2).

However, the specific slots are not identified. Stage 1 will determine, for example, that there is a slot for which Year 13 have option block 3, Year 12 have option block 2, Year 11 have Games and Year 10 have English, but it will not specify which slot this is. This is done in Stage 3.

An initial solution is created by a quick semi-random construction procedure whose main aim is to avoid class clashes (though it does not guarantee to be successful even in this modest aim). Alternatively an initial solution can be used which is the result of a previous run of Stage 1.

The cost of the initial solution is calculated, excluding those subcosts which cannot be determined until the identities of the slots are known. So the subcosts at this stage therefore reflect class clashes, teacher clashes, department clashes, room shortages and blocks being timetabled (or not being timetabled) together.

The first stage of the solution procedure is simple local improvement (see Eglese <sup>13</sup>). Perturbations as described earlier are made systematically to the solution and accepted if and only if the resulting solution has lower cost. An accepted solution then becomes the new current solution and perturbations are made to this solution, etc. This continues until a local optimum is reached.

Since the local optimum may still not represent a very good solution, a form of tabu search is then used in which the subcosts guide the search. The precise method is discussed in greater detail in Appendix 1. The approach is very similar to that used successfully by Wright <sup>14</sup> for timetabling cricket fixtures and is an example of what Hübcher and Glover <sup>15</sup> call "influential diversification". Just as "standard" tabu search (see Glover <sup>16</sup>) averts the possibility of the search process being trapped in a single unsatisfactory local optimum, influential diversification may be needed for complex problems in order that the process does not become trapped in a *region* of unsatisfactory local optima.

Involving subcosts in the search process not only helps to produce a low total cost, but also aims to ensure that no single subcost is much higher than it need be. This mitigates the problem mentioned earlier of the precise determination of the weights allotted to the objectives.

The program is left to run overnight or longer on a 486 PC, so that several solutions may be produced, each starting from a different initial solution. These are then inspected; usually the one with lowest total cost is the preferred alternative. There may be feasible solutions among them, or there may not. If not, the school's deputy head teacher needs to be consulted to see whether it is possible to change one or more of the teacher specifications, or make some other change, such that the best solution (or one of the other solutions) becomes feasible. In any case, the deputy head is consulted for approval.

## **STAGES 2, 3 AND 4**

In Stage 2 further blocks are introduced. These are Lower School blocks which require more than one teacher and/or occur more than once a week. The subjects involved are Games, Mathematics, English and French. Stage 2 starts with a solution from Stage 1 and timetables the new blocks randomly to form an initial solution. Alternatively it is possible to start from a previous Stage 2 solution. As in Stage 1, the specific slots are not identified. The solution method is exactly as for Stage 1, followed again by possible manual intervention until the deputy head is satisfied. It is possible that this stage may introduce some minor alterations to the Stage 1 solution, but these would not be such as to introduce any new teacher clashes.

Stage 3 does not introduce any further blocks, but it identifies the slots. Therefore the remaining subcosts are included. These relate to double lessons, the pattern of subjects for each class, pairs of blocks being timetabled on (or not on) the same day, teachers being free immediately before particular blocks and any complex requirements involving groups of blocks and groups of slots (see above). Again the method is broadly similar to that of Stage 1, except that a greater variety of perturbations is allowed.

Stage 4 introduces the remaining blocked subjects: Rural Science, Art and Physical Education. These blocks need to be timetabled once a week. The philosophy of this stage is slightly different from the others; instead of a few long runs involving tabu search, there are several short runs of about 30 seconds duration, each ending at the first local optimum.

The reason for the different approach used by Stage 4 is that the timetable is very full by this time. This has the effect that every initial solution is very close to a local optimum. The methods used for the first three stages could not be expected to work very well under these circumstances. The "sledgehammer" approach of many random initial solutions, however, in practice very quickly produced an initial solution with no major costs which was modified only very slightly before a



local optimum was reached (if no good solution had been reached, it would have been necessary to use a different Stage 3 solution or to make manual changes).

## **RESULTS**

The system was used in practice for the 1994-5 academic year. The results were very successful. A timetable was produced which required almost no manual alteration between Stages, apart from one or two very minor amendments caused by changes in teaching staff.

The contrast with the previous year's timetable was extremely marked. Whereas there had been an enormous difference between the two weeks of the cycle in 1993-4, there was now almost no difference. Moreover, in 1993-4 a large number of the preferred teacher specifications had had to be changed to fit in with the timetable; but in 1994-5 this only happened once, for the Year 11 German option (and this may well have been inevitable - every single Stage 1 solution exhibited this problem).

Apart from this, the total "cost" of the 1994-5 timetable was 414.5 units, which was very small compared with the major costs (e.g. 1000 for a single teacher clash once per week or 10000 for a double lesson not being provided). Full details of all the factors giving rise to cost are listed in Appendix 2 - none of these factors was considered by the school to be a serious problem.

In contrast, one single facet of the 1993-4 timetable would have incurred a cost of over 10000 units: Year 12 pupils had "Block 2" for the last two periods on Tuesday and the first two periods on Wednesday once a fortnight. There were many other major problems which would have given rise to significant costs.

There was also an enormous improvement in timescale. Timetabling for 1993-4 started in February 1993 and was not completed until the end of June 1993. In contrast, timetabling for 1994-5 did not start until mid-May 1994 and was finished by early June 1994.

The author received a letter from the head teacher in September 1994 which stated: "From long and bitter experience, I'd concluded that the traditional chaotic start to a school year was inevitable. You've proved that can be avoided. I haven't known a year begin so smoothly, and given the size of our school and the complexity of our curriculum and setting arrangements, that is saying a great deal ... By moving us from the complicated 10-day timetable to the present 5-day one (with only a few amendments), you have greatly improved the quality of learning of our pupils and students."

## **POSSIBLE VARIATIONS OF THE METHOD USED**

Many aspects of the formulation of the problem and the solution method are fairly arbitrary and could be varied. These include:

- \*\*\* The precise formulation of the subcosts - for example, how the cost of a poor spread of subjects is defined and quantified.
- \*\*\* The weights for the subcosts - it is not known how sensitive the final outcome is to these weights.
- \*\*\* The division of the problem into four stages - is this really sensible? Have the lesson blocks been categorised in the best possible way?
- \*\*\* The particular types of perturbation used - it could be worth experimenting with more complex types.

- \*\*\* The precise form and frequency of diversification - this is no more than a "reasonable try", as is the choice of perturbation in each of the four cases.
- \*\*\* The length and number of discrete runs of each stage of the program - what is the optimal balance between the two?

In each case, the choice made was guided by "common sense" and was to a large degree arbitrary, though some variations were tried. This suggests that, with further detailed experimentation and analysis, the program could be made even more successful.

## **IMPLICATIONS FOR FURTHER RESEARCH AND PRACTICE**

This application and its results to date conclusively demonstrate that there is great scope for improved computer timetabling systems which can produce timetables automatically for large, complex secondary schools. While the current system cannot be immediately transferred to other schools without amendment, the approaches and techniques have very wide applicability. Further work should therefore concentrate on making these methods more robust and generally applicable, so that many schools may reap the benefits.

This work also has important implications for the solution of a variety of other complex combinatorial problems. Influential diversification as described by Hübscher and Glover <sup>15</sup> is an approach which is likely to be very valuable for many types of problem. The special type of influential diversification used here (subcost-directed search) has been used before by Wright <sup>14</sup> with marked success, but needs further investigation as suggested by Marett and Wright <sup>17</sup> in order that it may be more widely applied.

## APPENDIX 1 - The subcost-guided diversification method

Let us use notation as follows.

- T,** the total cost at the current local optimum;  
**P,** the total cost at the previous local optimum;  
**Q,** the total cost two local optima ago;  
**Z,** the best cost found to date since the last creation of an initial solution;
- C,** the total cost increase caused by a particular perturbation;  
**B,** the highest cost decrease for an individual subcost caused by the perturbation;  
**S,** the specific subobjective for which B applies;  
**D,** the diversification parameter (see below - this indicates the degree of diversification required) at the current local optimum;  
**E,** the diversification parameter at the previous local optimum;  
**F,** the diversification parameter two local optima ago;  
**L,** the tabu list size;  
**X,** the "extra" list size (see below - sometimes needed to avoid the possibility of cycling).

There are four cases:

- (1)  $\frac{T < P \text{ and } T < 1.1 * Z}{: \text{ set } D = 1.}$
- (2)  $\frac{T < P \text{ and } T > 1.1 * Z}{: \text{ set } D = 2.}$
- (3)  $\frac{T > P \text{ and } T < 1.1 * Z}{: \text{ set } D = 3.}$
- (4)  $\frac{T > P \text{ and } T > 1.1 * Z}{: \text{ set } D = 4.}$

However, if the resulting value of D is greater than either (E+1) or (F+1), this may result in premature diversification; so, in this case, reset  $D = \min \{(E+1), (F+1)\}$ .

If  $T = P$  or  $T = Q$ , the process may well be repeating itself. To escape from this cycle, set X to a random number between 2 and 4; otherwise set  $X = 0$ . X affects the tabu list size L as described below.

If  $D = 1$ : there are good prospects of an improved solution in this vicinity. Therefore intensification is indicated. The list size (L) is set equal to (X+2) or (X+3) (chosen randomly) and perturbations are chosen for which  $B > 0$  and C is lowest. This has the effect of ensuring that the total cost will not be increased by very much.

If  $D = 2$ : the cost is tending to decrease, but is still some way from the best yet found. Some slight diversification may be necessary. L is set to (X+2) or (X+3) (at random) and perturbations are chosen for which  $B \geq 200$  and (B-C) is highest. Thus there must be at least a reasonable decrease in at least one subcost (200 represents a fairly important change in cost) and the influence of the change in total cost is reduced.

If  $D = 3$ : the cost is not tending to decrease, though it is still not too far away from the best yet found. Rather more diversification is required. L is set to (X+3) or (X+4) (at random) and perturbations are chosen for which  $B \geq 500$  and (5B-C) is highest. Thus a significant subcost decrease is required and the effect on total cost is a less influential factor than in the previous cases.

If  $D = 4$ : the cost is a long way from the best yet found and is not tending to decrease. Severe diversification is called for. L is set to (X+4) or (X+5) (at random) and perturbations are chosen for which  $B \geq 5000$  and (100B-C) is highest. Thus there must be a very substantial decrease in at least one subcost and the influence of the change in total cost is very small indeed.

The random elements in X and L represent a diversificatory device whose purpose is to reduce the risk of cycling between local optima. Taillard<sup>18</sup> used a similar approach in his work on quadratic assignment problems.

A further rule is applied which is also intended to reduce the risk of cycling. If a perturbation has been chosen by this procedure in iteration  $n$ , the same perturbation cannot also be so chosen in iteration  $(n+1)$ . This could be viewed as the reverse of the standard tabu search philosophy; move repetitions are banned instead of move reversals. However, since most of the perturbations are simple swaps, to repeat a move is the same as reversing it, so this is merely a way of increasing the length of time for which tabu status applies.

Whether intensifying or diversifying, the list sizes used here are small compared with those used in most applications (though not all - Widmer<sup>19</sup> used a list size of 1 for a job shop problem) described in the literature. However, as Glover<sup>16</sup> points out, the ideal list size depends on the number of perturbations made tabu by each entry on the list. Here this number may be fairly large; for example, if a perturbation chosen at a local optimum involves three blocks, then any other perturbation causing any of these three to return to its previous slot is made tabu, which is quite a severe restriction on the neighbourhood.

## APPENDIX 2 - Costs of final timetable for 1994/5

The factors giving rise to cost in the final timetable for 1994-5 are given below, together with their costs. Unless otherwise stated, occurrences are every week. The number of classes within the year group for which the condition applies is given for Years 7 to 9; for Years 10 to 13 everything applies for the whole year group.

Some of the figures below may look anomalous, but the costs are calculated according to quite complicated formulae which take account of various different factors. For example, Year 11 science occurs five times per week, so the fact that it occurs twice on one day every week is not a major drawback. However, Year 10 Humanities occurs only five times per fortnight, so its occurrence twice on the same day is less acceptable and thus incurs a far higher cost. As another example, it is much more important for the bottom two classes in Years 7 to 9 to avoid having "academic" subjects in the afternoon than it is for the top six classes.

Year 7 -	Technology twice on same day once per fortnight, all classes	56.8	
	One Mathematics lesson in afternoon for bottom two classes	5.0	
	One English lesson in afternoon for bottom two classes		5.0
	One French lesson in afternoon for bottom two classes		5.0
	Two Mathematics lessons in afternoon for two of the top six classes		4.0
Year 8 -	Technology twice in period 2 once per fortnight, four classes	0.4	
	French biased towards start of week, three classes	1.5	
	Mathematics twice in period 4, six classes	4.2	
	Mathematics twice in period 2, two classes	1.4	
	One Mathematics lesson in afternoon for bottom two classes	5.0	
	One English lesson in afternoon for bottom two classes		5.0
	Two Mathematics lessons in afternoon for all the top six classes		12.0
			Two French lessons in afternoon for three of the top six classes
Year 9 -	Mathematics twice in period 1, six classes	12.0	
	English biased towards end of week, four classes	2.0	
	Technology twice in period 2 once per fortnight, four classes	0.4	
	Rural Science biased towards start of week, two classes		1.0
	One Mathematics lesson in afternoon for bottom two classes	5.0	
	One English lesson in afternoon for bottom two classes		5.0
	One French lesson in afternoon for bottom two classes		5.0
	One Music/Drama lesson in morning for bottom two classes	5.0	

Year 10 - Humanities twice on same day once per fortnight		100.0
The two Block C lessons not on same day as preferred	50.0	
English twice in period 2	6.0	
English biased towards start of week	4.0	
Year 11 - Humanities twice in period 4 once per fortnight		1.5
Humanities biased towards start of week		1.0
	Technology twice in period 4	
Block B biased towards end of week	1.0	
Science twice on same day	1.1	
Mathematics twice in period 1		16.0
English twice in period 3	6.0	
Year 12 - Block 1 twice on same day as well as double period	20.0	
Block 1 twice in period 1	2.4	
Block 1 biased towards end of week		3.2
Block 2 twice in period 1	2.4	
Block 2 twice in period 2		0.4
Block 2 - combined effect of the previous two points	0.4	
Block 3 twice in period 5	4.4	
Block 4 twice in period 3	0.4	
Block 4 twice in period 5	4.4	
Block 4 - combined effect of the two previous points		0.4
Year 13 - Block 1 twice on same day as well as double period	20.0	
Block 1 twice in period 5	4.4	
Block 1 biased towards start of week		3.2
Block 2 twice in period 3	0.4	
Block 2 twice in period 5	4.4	
Block 2 - combined effect of the two previous points	0.4	
Block 3 twice in period 1	2.4	
Block 4 twice in period 2	0.4	
General Studies twice on same day once per fortnight	1.0	
General Studies biased towards end of week	0.2	
<b>TOTAL COST</b>		<b>414.5</b>

## REFERENCES

1. D. de WERRA (1985) An introduction to timetabling. *European Journal of Operational Research* 19(2), 151-162.
2. D. ABRAMSON (1991) Constructing school timetables using simulated annealing: sequential and parallel algorithms. *Management Science* 37, 98-113.
3. A. HERTZ (1991) Tabu search for large-scale timetabling problems. *European Journal of Operational Research* 54, 39-47.
4. N. BALAKRISHNAN, A. LUCENA and R. T. WONG (1992) Scheduling examinations to reduce second-order conflicts. *Computers and Operations Research* 19(5), 353-361.
5. J. A. FERLAND and A. LAVOIE (1992) Exchange procedures for timetabling problems. *Discrete Applied Mathematics* 35, 237-253.
6. P. DIGE, C. LUND and H. F. RAVN (1993), "Timetabling by simulated annealing", *Lecture Notes in Economics and Mathematical Systems* 396, 151-174.
7. D. JOHNSON (1993) A database approach to course timetabling. *Journal of the Operational Research Society* 44, 425-433.
8. J. A. FERLAND and C. FLEURENT (1994) SAPHIR: a decision support system for course scheduling. *Interfaces* 24(2), 105-115.
9. J. M. THOMPSON and K. A. DOWSLAND (1995) Variants of simulated annealing for the examination timetabling problem. To appear in *Annals of Operations Research*.
10. R. MOON (1991) *A guide to the National Curriculum*. Oxford University Press, Oxford.
11. R. SIMPER (1980) *A practical guide to timetabling*. Ward Lock Educational, London.
12. C. MORGENSTERN and H. SHAPIRO (1989) Chromatic number approximation using simulated annealing. Technical Report CS86-1, Department of Computer Science, University of New Mexico.
13. R. W. EGLESE (1986) Heuristics in operational research. In *Recent developments in operational research* (V. BELTON and R. W. O'KEEFE Eds.), pp. 49-67, Pergamon Press, Oxford.
14. M. B. WRIGHT (1994) Timetabling county cricket fixtures using a form of tabu search. *Journal of the Operational Research Society* 45, 758-770.
15. R. HÜBSCHER and F. GLOVER (1994) Applying tabu search with influential diversification to multiprocessor scheduling. *Computers and Operations Research* 21(8), 877-884.
16. F. GLOVER (1990) Tabu search: a tutorial. *Interfaces* 20(4), 74-94.
17. R. C. MARETT and M. B. WRIGHT (1995) A comparison of neighbourhood search techniques for multi-objective combinatorial problems. To appear in *Computers and Operations Research*.
18. E. TAILLARD (1991) Robust tabu search for the quadratic assignment problem. *Parallel Computing* 17, 443-455.
19. M. WIDMER (1991) Job shop scheduling with tooling constraints: a tabu search approach. *Journal of the Operational Research Society* 42, 75-82.