

A Time-Dependent Metaheuristic Algorithm for Post Enrolment-based Course Timetabling

Rhyd Lewis
Cardiff School of Mathematics
lewisR9@cf.ac.uk

May 27, 2011

Abstract

A metaheuristic-based algorithm is presented for the post enrolment-based course timetabling problem used in track-2 of the Second International Timetabling Competition (ITC2007). The featured algorithm operates in three distinct stages – a constructive phase followed by two separate phases of simulated annealing – and is time dependent, due to the fact that various run-time parameters are calculated automatically according to the amount of computation time available. Overall, the method produces results in line with the official finalists to the timetabling competition, though experiments show that this algorithm also seems to find certain instances more difficult to solve than others. A number of reasons for this latter feature are discussed.

1 Introduction

During late 2007 and early 2008 the Second International Timetabling Competition (ITC2007) was organised and run by a group of timetabling researchers from five different European Universities. The overall aim of this competition was to help people interested in timetabling from various fields to compare and contrast their timetabling methods using a common set of benchmark instances in a fair and accurate way.

On August the 1st 2007, the competition was officially started by the release of a number of problem instances into the public domain. Entrants were invited to register with the competition and to use these instances to help design algorithms that produced solutions according to the competition evaluation criteria. On January 11th 2008, two weeks before the end of the competition, a second set of problem instances was then also released. Competitors were required to submit their results from both instance sets to the competition organisers by the 25th of January. The organisers then collected the executables from selected entrants and ran these on their own benchmark machines, together with a third “hidden” set of problem instances to officially rank the various algorithms and choose a winner. Further details can be found in (McCollum et al., 2009) and also on the competition webpage at <http://www.cs.qub.ac.uk/itc2007/>

Before holding ITC2007, it was decided that the competition would be split into three tracks, each reflecting a different type of university timetabling problem, namely: (1) Exam timetabling, (2) Post Enrolment-based Course timetabling, and (3) Curriculum-based timetabling. Full descriptions of each of these problems can be found in the various specification reports (McCollum et al., 2007; Lewis et al., 2007; Di Gaspero et al., 2007), available on the competition webpage. In this paper, we present an algorithm for the problem used in track-2 of the competition. This particular formulation simulates the real-world situation where students are given a choice of lectures that they wish to attend, with the timetable then being constructed according to these choices. It is also based on the timetabling problem used for the first international timetabling competition run in 2003, albeit with extra features, which we outline in Section 2. Research arising due to the first competition includes the ant colony optimisation approach of Socha and Samples (2003), the simulated annealing-based algorithm of Kostuch (2005), and the mixed metaheuristic approach of

Lewis (2006). A variety of approaches for this problem-version are also offered by Rossi-Doria et al. (2002). Further information concerning the first competition is also available on the official website at www.idsia.ch/files/ttcomp2002/

In the next section we give a specification of the timetabling problem-version considered here. Readers interested in discovering more about this problem, including the rationale of *why* this problem-version takes the form that it does, are invited to consult the problem’s official specification document (Lewis et al., 2007), available on the ITC2007 website. Following this, the proposed algorithm is then described in detail in Section 3. Section 4 then contains details on the implementation and an analysis of the final results gained by the algorithm. Finally, Section 5 concludes the paper.

It should be noted that this algorithm was not officially entered into ITC2007 due to the fact that this author was one of the competition organisers.

2 Problem Description and Cost Functions

A problem instance for track-2 of ITC2007 contains the following information:

- A set of n events that are to be scheduled into 45 timeslots (to be interpreted as five days of nine, 1-hour timeslots);
- A set of m rooms where the events are to take place, each which has a specific seating capacity;
- A set of room-features that are *required* by events and which are *satisfied* by rooms;
- A set of s students who attend various combinations of the events;
- A set of available timeslots for each of the n events (i.e. not all events are available to be scheduled in all of the timeslots);
- A set of precedence requirements stating that certain events should occur before/after others in the timetable.

Given this information, the aim is to assign all of the n events to a room and a timeslot, whilst obeying the following five hard constraints:¹

HC_1 : No student should be required to attend more than one event in a particular timeslot;

HC_2 : Each event should be assigned to a room that has enough seats for all of the attending students and which satisfies all of the room-features required by the event;

HC_3 : Only one event should be put into each room in any timeslot (i.e. no double-booking of rooms);

HC_4 : Events should only be assigned to timeslots that are designated as “available” for those events;

HC_5 : Where specified, events should be scheduled to occur in the correct order in the week

In addition to these five hard constraints, it is also desirable for the following three soft constraints to be satisfied:

SC_1 : Students should not be scheduled to attend an event in the last timeslot of a day (that is, timeslots 9, 18, 27, 36, or 45);

SC_2 : Students should not have to attend events in three or more successive timeslots occurring in the same day;

¹Note that hard constraints 1-3 are the same as those used in the first timetabling competition.

SC_3 : Students should not be required to attend just one event in a particular day

As is typical in timetabling research, for this problem the satisfaction of the hard constraints is considered to be more important than the satisfaction of the soft constraints. Because of this, the competition rules state that a candidate solution should be evaluated according to two separate values: the *Distance to Feasibility* and the *Soft Cost*. The Distance to Feasibility is used because – in contrast to the first competition – it is necessary to allow for the fact that an algorithm may not be able to assign all of the n events into the timetable whilst obeying the hard constraints. For this competition, it is therefore permissible to allow some events to remain *unplaced* in order to ensure that none of the hard constraints are violated. (The ITC2007 rules stipulate that if an algorithm produces a solution that contains any violations of the hard constraints, then it is considered invalid and should be disqualified from the competition.) The Distance to Feasibility is thus calculated by considering the events that are *not* placed in the timetable. However, it is not the number of events that are considered here; rather it is the total number of students that *attend* each of these unplaced events, reflecting the real-world situation where we are interested in satisfying as many people’s needs as possible within a timetable. Thus if, for example, a particular solution timetable has three events that are unplaced, and the number of students attending each of these is 10, 5, and 8, then the Distance to Feasibility is simply $(10 + 8 + 5) = 25$. Note that if all events are inserted into a timetable (whilst obeying the hard constraints), then its Distance to Feasibility is zero.

The second value used for timetable evaluation is the *Soft Cost*, which is calculated by simply counting the number of soft constraint violations in a timetable. For SC_1 , if a student is scheduled to attend an event in the last timeslot of a day, then this results in one penalty point. (Naturally, if there are x students in this class, we consider this as x penalty points.) For SC_2 if one student has three events in a row we count this as one penalty point. If a student has four events in a row we count this as two, and so on. Note that adjacent events occurring over two separate days are not counted as a violation. Finally, for SC_3 each time a student attends just one event on a day, this results in one penalty point. The Soft Cost is simply the total of these three values. (Further information on this cost measure can be found in Lewis et al. (2007) or on either of the competition webpages.)

Given these two values, the following procedure is used to compare solutions. First, the solutions’ Distances to Feasibility are considered, and the solution with the lowest value is deemed the winner. However, when two or more solutions are equal in this respect, the winner is deemed the solution among these that has the lowest Soft Cost.

For this competition track sixteen problem instances were made publicly available during the competition: eight “early” instances, which were released on the start date, and eight “late” instances, released two weeks before the close of the competition. Since completion of the competition in 2008 the eight “hidden” instances used by the competition officials have also been released into the public domain. In all cases these problems were created using an automated generator and all are known to feature at least one perfect solution – that is, a solution where all of the n events are assigned to the timetable without any constraint violations, hard or soft. For the competition, a benchmark timing program was also released which entrants were required to execute on their own machines. This program specifies a time limit for each machine and ensures that entrants use approximately the same amount of computational effort when testing their algorithms.

3 Algorithm Description

In this approach our strategy is to tackle the problem in three distinct stages, each with a strict time limit, T_1 , T_2 , and T_3 respectively (such that the full time limit $T = T_1 + T_2 + T_3$). If a particular stage completes before reaching its time limit, then the remaining time is passed on to the next stage. If Stage 3 completes early, then the algorithm also halts early. For guidance, a description of the main objectives of each stage is given in fig. 1. As this demonstrates, the idea is to arrange the constraints into three different levels of importance. At each successive stage, violations of constraints satisfied in previous stages are then disallowed. In the first stage, attempts are made




Figure 1: High level description of the three stage algorithm. Here, T represents the time limit defined by the competition benchmarking program

to try and satisfy hard constraints 1 to 4 using specialised procedures proposed in some of our earlier work (Lewis, 2006). Although these methods are generally effective, they do not, however, seem immediately applicable to the remaining hard constraint HC_5 ; thus the second stage of the algorithm is concerned with the removal of violations of this constraint. Finally, in Stage 3, the algorithm concentrates on removing as many soft constraint violations as possible from the current timetable. At this point any unplaced events are ignored.

In Stages 2 and 3 of this algorithm, optimisation is carried out using simulated annealing (Kirkpatrick et al., 1983). Because this algorithm runs according to time limits, when applying this metaheuristic it is useful to calculate cooling schedules that take the amount of available computation time into account. The aim is to therefore allow the algorithm to perform a slower cooling (and therefore a wider, more global search) when presented with a generous amount of run time, and a quicker cooling (with a more intensive, greedy search) when only small amounts of time are available. The method for calculating these parameters is described in Section 3.3.

Note that in all experiments described here, $1/3$ of the available time limit T is allocated to each stage. In preliminary tests we also experimented with a scheme that allowed each stage to be run until all of its objectives were met (or until T was reached), but this often turned out to give worse performance in instances where full feasibility was difficult to achieve as this would result in no time being allocated to Stage 3. We also experimented with different time allocations, such as granting Stages 1 and 2 a total of $9/10$ of the available run time, but on the whole we found that our $1/3$ allocation gave each stage of the algorithm sufficient time to complete their objectives (or to stagnate in the process).

In the next subsection, we now discuss some encoding and preprocessing issues that are relevant in the design of this algorithm. Subsections 3.2, 3.3, and 3.4 then describe Stages 1, 2, and 3 of the algorithm respectively.

3.1 Encoding and Preprocessing Issues

For this approach a timetable is encoded using a two-dimensional ($r \times 45$) matrix (i.e. grid) in which rows represent rooms and columns represent timeslots. Throughout this paper we refer to this timetable matrix as tt and use the notation $tt[i, j]$ to denote the contents at location (i, j) . Each cell in this grid (i.e. *place* in the timetable) can be *blank* or will be *occupied* by exactly one event. Note that this latter feature means that it is impossible to double-book a room, allowing us to disregard HC_3 .

It is also useful to carry out some preprocessing steps before executing the main body of the algorithm. First, two additional matrices are constructed: the *event-room* matrix, and the *conflicts matrix* (these were also used in the first timetabling competition by Kostuch (2005)). The event-room matrix is of dimensions $n \times m$ and is used to indicate which rooms are suitable for which events. This can be easily calculated for an event i by simply identifying which rooms satisfy both the seating capacity and the features required by i . The $n \times n$ conflicts matrix, meanwhile, is very

Table 1: Heuristic rules used in Stage 1.

Heuristic	Description
h_1	Choose the event with the smallest number of suitable places in tt to which it can be assigned.
h_2	Choose the unplaced event that conflicts with the most other events.
h_3	Choose an event randomly.
h_4	Choose the place that is suitable for the least number of other unplaced events in U .
h_5	Choose the place in the timeslot with the fewest events in.
h_6	Choose a place randomly.
h_7	Choose the event with the least number of students.

much like the standard adjacency matrix used for representing graphs and indicates which pairs of events *conflict* (i.e. cannot be assigned to the same timeslot). Thus, if two events i and j have one or more students in common, or if both i and j can only be assigned to the same single room r , then it is obvious that events can never be feasibly assigned to the same timeslot, and so elements (i, j) and (j, i) in the conflicts matrix can be marked as true.

Our final act of preprocessing considers the hard constraint HC_4 . First of all, note that if we have a constraint “event i must occur before event j ”, then this will automatically imply the constraint “event j must occur *after* event i ”. For this approach, this means that all occurrences of HC_5 can be conveniently stored in a compact way using an array A of n lists, where each list $A[i]$ contains the events that need to be scheduled *after* event i in the timetable. (The “before” constraints do not need to be considered). Second, we can also make further additions to A by noting that hard constraint HC_5 is *transitive* (i.e. if event i must occur before event j , and event j must occur before event k , then this implies that event i must also occur before event k). In some of the competition instances, not all of the implied constraints due to this transitivity are present in the given problem files, and so it makes sense to calculate these in order to gain a better understanding of the number of constraints that are being considered when trying to solve the problem.

3.2 Algorithm Description: Stage 1

In Stage 1, the objective is to insert as many of the n events into the timetable as possible without violating the first four hard constraints. A precise pseudo-code description is presented fig. 2. As is shown, this stage takes as arguments the empty timetable tt , an iteration-limit I , and a list of currently unplaced events U (to begin with, $|U| = n$). Events are then taken one-by-one from U and are inserted into suitable places in tt .² In order to try and maximise the number of events that are inserted, heuristic rule h_1 is used to select the next event, with ties being broken using h_2 , and further ties with h_3 (refer to Table 1). Note that these particular heuristics are akin to those used in the Dsatur algorithm for graph colouring (Brelaz, 1979), though in this case h_1 also takes the issue of room allocation into account. Rule h_1 therefore selects events based on the state of the current partial timetable tt , and prioritises those with the least remaining feasible options. Meanwhile, rule h_2 prioritises those events that have the highest number of conflicts which, as a rule of thumb, are often the more problematic events to insert. Note that events with no remaining place-options are ignored at this point. Finally, to select a place for each event, rule h_4 is used, which chooses the place whose occupation will have the least effect on the place-options of the remaining unplaced events in U . Ties of this rule are broken using h_5 and further ties with h_6 .

At the end of this assignment stage, the list U will be empty (in which case all of the events have been assigned to the timetable), or U will only contain events that have no suitable places in

²In this section the term “suitable” is used to indicate a place in the timetable that will not result in the violation of hard constraints 1 to 4.

```

STAGE-1( $tt, U, I$ )
(1)   while ( $\exists$  events in  $U$  with suitable places in  $tt$ )
(2)       Select an event  $e \in U$  that has suitable places in  $tt$ ;
(3)       Choose a suitable place  $p$  for  $e$ ;
(4)       Move  $e$  from  $U$  into  $tt$  at place  $p$ ;
(5)       ITERATED-HEURISTIC-SEARCH( $tt, U, I$ );

ITERATED-HEURISTIC-SEARCH( $tt, U, I$ )
(1)   while ( $U \neq \emptyset$  and (timelimit  $T_1$  not reached))
(2)       HEURISTIC-SEARCH( $tt, U, I$ );
(3)       if ( $U \neq \emptyset$ )
(4)            $V \leftarrow$  EXTRACT-SOME-EVENTS( $tt, |U|$ );
(5)           HEURISTIC-SEARCH( $tt, U, I$ );
(6)            $U \leftarrow U \cup V$ ;

HEURISTIC-SEARCH( $tt, U, I$ )
(1)   Make a list  $P$  of all the unoccupied places in  $tt$ ;
(2)    $i \leftarrow 0$ ;
(3)   while ( $U \neq \emptyset$  and  $P \neq \emptyset$  and  $i < I$ )
(4)       foreach ( $u \in U$  and  $p \in P$ )
(5)           if ( $p$  is a suitable place in  $tt$  to assign  $u$ )
(6)               Put  $u$  into  $p$  in  $tt$ ;
(7)               Remove  $u$  from  $U$  and  $p$  from  $P$ ;
(8)       if ( $U \neq \emptyset$  and  $P \neq \emptyset$ )
(9)           repeat
(10)              Choose a random event  $e$  in  $tt$  and  $p \in P$ ;
(11)              if ( $p$  is suitable place in  $tt$  to assign  $e$ )
(12)                  Move  $e$  from its current place to  $p$ ;
(13)                  Update  $P$  to reflect the changes;
(14)               $i \leftarrow i + 1$ ;
(15)           until ( $i = I$  or ( $e$  has been moved to  $p$ ))

EXTRACT-SOME-EVENTS( $tt, q$ )
(1)    $V \leftarrow \emptyset$ ;
(2)   for ( $i \leftarrow 1$  to  $q$ )
(3)       Randomly choose two events  $e$  and  $g$  in  $tt$ ;
(4)       Move either  $e$  or  $g$  (according to  $h_7$ ) from  $tt$  to  $V$ ;

```

Figure 2: Pseudo-code description of Stage 1. Here, tt represents the $(r \times 45)$ timetable matrix and U and V are lists of unplaced events. (When STAGE-1 is first called, $|U| = n$.) I defines the iteration limit of the HEURISTIC-SEARCH procedure.

tt . In the latter case, the procedure ITERATED-HEURISTIC-SEARCH is called, which is used to try and transfer further events from U into tt , ensuring that hard constraints 1-4 are not violated in the process. To start, the sub-procedure HEURISTIC-SEARCH is called, which operates by repeatedly attempting to move events from U into free (i.e. blank) places in tt (lines (3)-(7)). While doing this, however, HEURISTIC-SEARCH also shuffles the events *within* tt so that the free places change position (lines 9-15). The rationale of this latter action is that it offers the possibility of further events in U being added to tt when we loop back to line (1) of HEURISTIC-SEARCH. However, although the HEURISTIC-SEARCH procedure is quite effective in reducing the number of unplaced events, in initial experiments it was also noticed that it is only able to do this for a fairly short period of time, after which the process stagnated, with no further events being transferred from U into tt . To counter this, ITERATED-HEURISTIC-SEARCH therefore includes a mechanism intended for re-invigorating the process, which is achieved by the procedure EXTRACT-SOME-EVENTS, which removes other events from tt and puts these into a second list V . Of course, by removing events

Table 2: Percentage of events inserted into the timetable before and after applying ITERATED-HEURISTIC-SEARCH. Presented figures are the means of 51 runs on each instance together with the standard deviation. Instances 1-8 are the “early” instances, 9-16 the “late” instances, and 17-24 the “hidden” instances.

Instance	n	m	s	Before (%)	After (%)
comp-2007-2-1	400	10	500	86.3 ± 1.2	100 ± 0.0
comp-2007-2-2	400	10	500	83.9 ± 1.2	100 ± 0.0
comp-2007-2-3	200	20	1000	95.3 ± 1.2	100 ± 0.0
comp-2007-2-4	200	20	1000	92.1 ± 1.3	100 ± 0.0
comp-2007-2-5	400	20	300	92.7 ± 1.0	100 ± 0.0
comp-2007-2-6	400	20	300	92.2 ± 0.9	100 ± 0.0
comp-2007-2-7	200	20	500	93.2 ± 1.1	100 ± 0.0
comp-2007-2-8	200	20	500	93.4 ± 1.1	100 ± 0.0
comp-2007-2-9	400	10	500	85.5 ± 1.4	100 ± 0.0
comp-2007-2-10	400	10	500	81.2 ± 1.4	100 ± 0.1
comp-2007-2-11	200	10	1000	94.5 ± 1.3	100 ± 0.0
comp-2007-2-12	200	10	1000	91.3 ± 1.2	100 ± 0.0
comp-2007-2-13	400	20	300	90.0 ± 1.2	100 ± 0.0
comp-2007-2-14	400	20	300	90.7 ± 0.9	100 ± 0.0
comp-2007-2-15	200	10	500	91.5 ± 1.3	100 ± 0.0
comp-2007-2-16	200	10	500	98.0 ± 0.8	100 ± 0.0
comp-2007-2-17	100	10	500	99.9 ± 0.4	100 ± 0.0
comp-2007-2-18	200	10	500	89.4 ± 2.2	100 ± 0.0
comp-2007-2-19	300	10	1000	86.9 ± 1.0	99.4 ± 0.6
comp-2007-2-20	400	10	1000	98.6 ± 0.7	100 ± 0.0
comp-2007-2-21	500	20	300	94.3 ± 1.0	100 ± 0.0
comp-2007-2-22	600	20	500	83.9 ± 0.8	95.0 ± 1.0
comp-2007-2-23	400	20	1000	85.6 ± 1.4	99.8 ± 0.3
comp-2007-2-24	400	20	1000	97.6 ± 0.6	100 ± 0.0

from tt , extra free places are created that can be used by some of the events in U . Thus the events in V are put to one side temporarily, and HEURISTIC-SEARCH is again applied using U and the new, emptier timetable. Finally, upon completion of this second phase of heuristic search, the events in V are added to the events (if any) that still reside in U and the entire ITERATED-HEURISTIC-SEARCH process is repeated.

Examining the pseudo-code of ITERATED-HEURISTIC-SEARCH, two important features become apparent. First, it is noticeable that if we choose to transfer some events from tt into V (line (4)) and then subsequently add the contents of V to U (line (6)), then in some cases the overall number of unplaced events may actually increase, thus conflicting with the objectives of Stage 1. However, in practice if such a situation arises, it is usually only temporary because the number of events in U is generally seen to decrease again when the algorithm loops back to the start of the procedure. The second issue, meanwhile, concerns the strategy of event extraction used in EXTRACT-SOME-EVENTS. One choice available here is to simply choose events randomly for removal. However, in this particular case it seems sensible to bias the choice towards removing smaller events from the timetable, due to the fact that unplaced events containing less students will attract a lower Distance to Feasibility measure when the timetable is evaluated. Thus, heuristic rule h_7 (Table 1) is used here. Note also that $|U|$ events are currently extracted here, though a different value could be used here in theory.

Table 2 summarises the effectiveness of Stage 1 on the twenty-four problem instances. Details on the sizes of each instance are also included here: the number of events n , rooms m and students s . With sixteen of the instances, $\geq 90\%$ of events are inserted into the timetable by the heuristic assignment rules on average. With the remaining eight instances this figure averages in the 80s. We also see that after the subsequent application of ITERATED-HEURISTIC-SEARCH, $> 99\%$ of events are inserted into the timetable in all cases, with the exception of instance 22, the largest problem in the set, where an average of only 95% of events are assigned.

3.3 Algorithm Description: Stage 2

In Stage 2 of the algorithm, attention is turned towards eliminating violations of the remaining hard constraint HC_5 . As mentioned earlier, this is done using simulated annealing (SA). The cost function $C(tt)$ used in this phase is:

$$C(tt) = \sum_{i=1}^n \sum_{j=1}^{|A[i]|} f(i, j), \quad (1)$$

where

$$f(i, j) = \begin{cases} \frac{\text{size}(i) + \text{size}(A[i]_j)}{B} & \text{if } (\text{slot}(i) \geq \text{slot}(A[i]_j)) \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

Here, A refers to the array of lists described in Section 3.1, $A[i]_j$ indicates the j th element in the list $A[i]$, and $\text{slot}(i)$ indicates the timeslot that event i is assigned to in timetable tt . In addition, $\text{size}(i)$ indicates the number of students attending event i , and B is a constant that reflects the total number of students in the two largest events. Cost function C thus reflects the number of violations of HC_5 in tt , with each violation incurring a penalty value of between zero and one. In a previous version of this algorithm (Lewis, 2008) all H_5 violations incurred a penalty value of one, however this current method is an improvement as it prioritises the satisfaction of HC_5 with larger events, thus paying heed to the ultimate measure of quality: the Distance to Feasibility. Obviously, the aim in this stage is to try and produce a solution tt with $C(tt) = 0$.

The neighbourhood operator used in this phase operates by randomly selecting a cell $tt[a, b]$ in the timetable. A second cell $tt[c, d]$ is then also randomly selected, and the contents of the two cells are swapped. If such a move causes a violation of hard constraints 1-4, then it is immediately rejected and reset; otherwise it is accepted, and tt is re-evaluated using C . Note that in choosing the second cell $tt[c, d]$ it is necessary to ensure that $tt[a, b] \neq tt[c, d]$. This guarantees (a) that the two selected cells are distinct, and (b) that at least one of the selected cells is occupied (obviously, if either (a) or (b) were not true then the corresponding move would result in an identical timetable).

Note that an application of this neighbourhood operator results in one of two actions. The first occurs when two occupied cells are selected, which causes the places of the two associated events to be *swapped* in the timetable. The second occurs when one occupied and one blank cell are selected, causing just one event to be *moved* to a different place in the timetable. Due to the way in which cells are selected in this case, the probability of each of these actions occurring is directly related to the proportion of occupied cells in the timetable: assuming (without loss of generality) that n events are present in a timetable with $p = 45m$ places, then the probability of a swap occurring is:

$$P(\text{swap}) = \frac{n}{p} \times \frac{n-1}{p-1}, \quad (3)$$

(i.e. the conditional probability of selecting one occupied cell in the grid, followed by another, different occupied cell). The probability of performing a *move*, meanwhile, is:

$$P(\text{move}) = \left(\frac{p-n}{p} \times 1.0 \right) + \left(\frac{n}{p} \times \frac{p-n}{p-1} \right), \quad (4)$$

(i.e. the sum of the probability of selecting a blank cell followed by an occupied cell, and the probability of selecting an occupied cell followed by a blank). Note that these probabilities arise naturally due to our chosen method of cell selection and will be fixed in each run of the algorithm. We highlight them here in order to aid our analysis of this algorithm and we note that in practice it would be easy to attach different probabilities to the *swap* and *move* actions if it were desired.

Given the above cost function and neighbourhood operator, a straightforward application of SA is now used: starting at an initial ‘‘temperature’’ t_0 , during execution the temperature variable is slowly reduced according to a temperature update rule $t_{i+1} = \alpha t_i$, where α ($0 < \alpha < 1$) is a variable known as the ‘‘cooling rate’’. At each temperature t_i , a Markov chain of length n^2 is then

Table 3: Cost of the timetable (using cost function C) at the start and end of Stage 2, and the number of extra events that are removed in order to free the timetables of remaining hard constraint violations. Presented figures are the means of 51 runs on each instance together with the standard deviation.

Instance	Start	End	Removed
comp-2007-2-1	14.9 ± 3.0	0.5 ± 0.7	0.5 ± 0.8
comp-2007-2-2	13.8 ± 3.1	1.6 ± 1.6	1.8 ± 1.8
comp-2007-2-3	7.4 ± 1.4	0.0 ± 0.0	0.0 ± 0.0
comp-2007-2-4	8.1 ± 1.7	0.0 ± 0.0	0.0 ± 0.0
comp-2007-2-5	57.1 ± 7.6	0.0 ± 0.0	0.0 ± 0.0
comp-2007-2-6	52.9 ± 4.1	0.0 ± 0.0	0.0 ± 0.0
comp-2007-2-7	6.4 ± 1.8	0.0 ± 0.0	0.0 ± 0.0
comp-2007-2-8	7.7 ± 2.1	0.0 ± 0.0	0.0 ± 0.0
comp-2007-2-9	16.9 ± 2.6	2.8 ± 2.1	3.0 ± 2.1
comp-2007-2-10	16.0 ± 2.4	4.3 ± 2.6	4.6 ± 2.8
comp-2007-2-11	6.0 ± 1.6	0.0 ± 0.0	0.0 ± 0.0
comp-2007-2-12	7.4 ± 2.5	0.0 ± 0.0	0.0 ± 0.0
comp-2007-2-13	52.0 ± 5.9	0.0 ± 0.0	0.0 ± 0.0
comp-2007-2-14	52.6 ± 5.2	0.0 ± 0.2	0.0 ± 0.2
comp-2007-2-15	7.7 ± 2.0	0.0 ± 0.0	0.0 ± 0.0
comp-2007-2-16	10.9 ± 2.1	0.0 ± 0.0	0.0 ± 0.0
comp-2007-2-17	3.1 ± 1.0	0.0 ± 0.0	0.0 ± 0.0
comp-2007-2-18	7.3 ± 1.8	0.0 ± 0.0	0.0 ± 0.0
comp-2007-2-19	11.9 ± 2.3	4.0 ± 2.1	4.4 ± 2.3
comp-2007-2-20	16.7 ± 3.0	0.0 ± 0.0	0.0 ± 0.0
comp-2007-2-21	70.3 ± 6.1	0.0 ± 0.0	0.0 ± 0.0
comp-2007-2-22	77.0 ± 8.5	32.4 ± 10.4	31.1 ± 9.0
comp-2007-2-23	18.0 ± 3.5	4.3 ± 2.4	4.5 ± 2.3
comp-2007-2-24	116.2 ± 10.6	1.8 ± 2.4	1.8 ± 2.1

generated by performing n^2 applications of the neighbourhood operator. Any move that increases the cost of the timetable is then accepted with a probability $\exp(-\delta/t_i)$, where δ is the cost change that this move causes. Any move that reduces or leaves the cost unchanged, meanwhile, is accepted automatically.

Because this algorithm aims to be time dependent, it is a good idea to choose a cooling rate α that allows Markov chains to be generated at as many temperatures as possible between the initial temperature t_0 and some end temperature. To calculate such a cooling rate the SA algorithm is first run for 5% of Stage 2's allocated time, and the number of Markov chains generated is recorded. This figure is then used to predict the number μ of Markov chains that will be generated in the remaining 95% of time. Using μ , we can then calculate a value for α that ensures that the temperature will be reduced from t_0 to a specific end temperature t_μ in exactly μ steps as:

$$\alpha = (t_\mu/t_0)^{1/\mu}. \quad (5)$$

In our case, following other works (van Laarhoven and Aarts, 1987; Abramson et al., 1996) the initial temperature t_0 is determined automatically by performing a small sample of neighbourhood moves and then calculating the variance of the cost over these moves. The end temperature t_μ , meanwhile, needs to be set by the user (see Section 4).

Finally, Stage 2 completes either when a timetable tt with a cost $C(tt) = 0$ is found, or when the time limit T_2 is reached. If the latter occurs, then the best solution found during this stage is taken, and events that are seen to be causing a violation of HC_5 are removed one-by-one from tt until it is completely free of any hard constraint violations. This resultant timetable is then passed on to Stage 3.

Table 3 summarises the effects of Stage 2 on the twenty-four competition instances. Note that there is no significant correlation between the initial and final costs of the timetables. In total, sixteen of the instances feature final costs with a means and standard deviations close to zero, indicating that the procedure is able to successfully complete its objectives in the majority of runs. We see that the instances that caused difficulties with the assignment heuristics in Stage 1, namely instances 1, 2, 9, 10, 19, 22, and 23 also prove to be troublesome here, and events usually need to

be removed from these timetables to free them of remaining violations of HC_5 . This is also the case with instance 24, where an average of 1.8 events need to be removed. One possible reason for the lack of performance in these eight cases is that candidate solutions to some of these problems have high proportions of their cells occupied ($> 66\%$ of available places with instances 1, 2, 9, 10, 19, and 22), thus according to eq. (3) a greater proportion of proposed neighbourhood moves will be swaps, which are more likely to be rejected since both events involved need to be assigned to places that are suitable. There are exceptions to this observation however: instances 23 and 24, for example, feature relatively low place occupancy rates of 44% but still show a marked lack in performance in this stage. However, as we will see in fig. 3 later, all of the eight instances causing difficulties here also tend to feature quite a high rejection rate of proposed neighbourhood moves, which seems to make movements in the search space more restricted, making exploration and resultant improvements in cost harder to achieve. We will return to this topic in Section 4.

3.4 Algorithm Description: Stage 3

In Stage 3, attention is turned towards satisfying the soft constraints of the problem. When this part of the algorithm is invoked, one of two situations will have occurred: either a valid timetable with a Distance to Feasibility of zero will have been produced or, after having spent 2/3's of the available run time trying to deal with the hard constraints, we will have settled for a timetable that has some unplaced events. In the latter case, these unplaced events are not considered any further – i.e. they are effectively eliminated from the problem.

The application of SA used here is again straightforward, with the cooling scheme being calculated in the same way as Stage 2. The neighbourhood operator is also the same as Stage 2, though in this case, moves that cause a violation of hard constraint HC_5 (in addition to the previous four) are also immediately reset. Finally, the cost function used here is simply the Soft Cost (see Section 2), which is appropriate due to the fact that no hard constraint violations are permitted in the timetable from this point onwards.

4 Implementation and Analysis of Final Results

The algorithm was implemented using C++ under Linux using the g++ 4.1.1 compiler under the `-O3` optimisation option. All experiments were run on a 1.8GHz machine with 256MB RAM, which was granted a time limit 636 seconds by the competition benchmarking program. Because many of this algorithm's parameters are calculated automatically, only two values need to be chosen for these experiments: I , the iteration limit used for ITERATED-HEURISTIC-SEARCH in Stage 1, and the end temperature t_μ used in the two annealing phases. In practice, the algorithm did not seem to be particularly sensitive to variations in I , providing that values of around $100n$ or more were used.³ Values of I less than this tended to cause the process to stagnate too readily. Considering the end temperature t_μ , if this was set too high then it tended to mean that too many increases in cost were permitted throughout the run, making the search more of a random walk. On the other hand, if t_μ was too low, then the algorithm spent too much time at low temperatures, making it more greedy and increasing its probability of getting stuck in a local minimum. On the whole, however, the algorithm did not seem too sensitive to variations in t_μ , provided that values of around 0.000005 to 0.0001 were used. For all experiments here, parameter settings of $I = 1000n$ and $t_\mu = 0.00001$ (for both annealing phases), were used. Note that no fine tuning of these values was conducted.

Table 4 summarises the final results achieved by this algorithm after performing 51 runs (from different random seeds) on each instance. Because timetable quality is ranked according to a pair of values (see Section 2), results are summarised using the best, worst, and median, and lower and upper quartiles (we have used fifty-one runs so that these statistics can be calculated without the need for interpolation, which would be inappropriate here). From the table we can see that the algorithm has achieved feasibility for all instances at least once, with the exception of instance 22,

³The value n was used to allow the setting to scale with instance size.

Table 4: Summary of the final results obtained from 51 runs on each problem instance. In each case the Distance to Feasibility is presented together with the associated Soft Cost (in brackets).

Instance	Best	Q_1	Median	Q_3	Worst
comp-2007-2-1	0 (1166)	0 (1571)	0 (1819)	32 (1535)	60 (1790)
comp-2007-2-2	0 (1665)	12 (1871)	42 (1866)	59 (1652)	168 (1734)
comp-2007-2-3	0 (251)	0 (361)	0 (436)	0 (472)	0 (673)
comp-2007-2-4	0 (424)	0 (488)	0 (552)	0 (617)	0 (666)
comp-2007-2-5	0 (47)	0 (121)	0 (190)	0 (276)	0 (533)
comp-2007-2-6	0 (412)	0 (574)	0 (621)	0 (701)	0 (794)
comp-2007-2-7	0 (6)	0 (320)	0 (383)	0 (461)	0 (719)
comp-2007-2-8	0 (85)	0 (147)	0 (215)	0 (252)	0 (346)
comp-2007-2-9	0 (1819)	48 (1800)	67 (2095)	106 (1805)	232 (1916)
comp-2007-2-10	0 (2091)	73 (2197)	110 (2293)	173 (2033)	264 (2120)
comp-2007-2-11	0 (288)	0 (433)	0 (496)	0 (572)	0 (681)
comp-2007-2-12	0 (474)	0 (663)	0 (744)	0 (865)	0 (1045)
comp-2007-2-13	0 (298)	0 (515)	0 (592)	0 (700)	0 (892)
comp-2007-2-14	0 (127)	0 (563)	0 (690)	0 (778)	12 (571)
comp-2007-2-15	0 (108)	0 (284)	0 (319)	0 (351)	0 (423)
comp-2007-2-16	0 (138)	0 (167)	0 (192)	0 (210)	0 (247)
comp-2007-2-17	0 (0)	0 (17)	0 (31)	0 (61)	0 (104)
comp-2007-2-18	0 (25)	0 (187)	0 (273)	0 (332)	0 (515)
comp-2007-2-19	0 (2146)	177 (2075)	275 (1916)	398 (1969)	549 (1989)
comp-2007-2-20	0 (625)	0 (811)	0 (844)	0 (884)	0 (1048)
comp-2007-2-21	0 (308)	0 (442)	0 (500)	0 (547)	0 (646)
comp-2007-2-22	636 (1808)	937 (1472)	1024 (1401)	1169 (1670)	1312 (1110)
comp-2007-2-23	0 (3101)	151 (3665)	262(4523)	396 (3383)	690 (3998)
comp-2007-2-24	0 (841)	0 (1212)	28 (1244)	89 (1018)	199 (1529)

which proved the most troublesome in all stages of the algorithm. Feasibility has also been achieved in all runs with fifteen of the twenty-four instances. For instance-17 a perfect solution has also been found in one run. Following on from observations made in previous sections, once again instances 1, 2, 9, 10, 19, 22, 23, and 24 prove to be the hardest problems to solve, with feasibility being found in just 59%, 22%, 4%, 6%, 2%, 0%, 6%, and 35% of runs respectively.⁴ Note that the algorithm is able to find feasibility in $> 95\%$ of runs with all remaining instances, however.

Earlier, in Section 3.3 it was suggested that when considering instances in which a large proportion of neighbourhood moves were rejected, movements in the search space would be more restricted, possibly making improvements to the candidate solution more difficult to achieve. At this point, a relevant question to now ask is whether the same feature also applies when performing optimisation according to the Soft Cost. To investigate this, in trials where a distance to feasibility of zero was achieved, we recorded the proportion of all moves that were accepted during Stage 3, together with the proportion by which the Soft Cost was ultimately reduced.⁵

The statistics collected here are displayed in fig. 3 where a weak positive, though statistically significant, correlation between the two variables can be observed (a correlation of $r = 0.53$, using a two-tailed test at the 1% level). Curiously, one particular group of smaller instances – namely instances 7, 8, 15 and 16, which all have $n = 200$ and $s = 500$ – seems to go against this trend, with the algorithm experiencing a low proportion of accepted moves and yet still achieving large reductions in the Soft Cost (removing these four instances causes r to increase to 0.64). From these admittedly limited results, we propose that there is *some* relationship between the restrictiveness of the search space and the improvements in cost that are achievable, but there are certainly other factors that will also have an effect here, including the shape of the cost landscape, the amount of computation that is required for each application of the evaluation function, and, perhaps most importantly, the amount of time that is ultimately allocated to this stage of the algorithm (obviously this stage can be allocated more than the original $T/3$ seconds if the previous stages of the algorithm are completed early).

Considering this latter point, in fig. 4 we illustrate the effects that differing time limits have on

⁴These percentages improve upon those reported in an earlier version of this paper, reported in (Lewis, 2008).

⁵We remember that for Stage 3 a move is “accepted” if the alteration to the timetable does not cause a violation of *any* of the hard constraints.

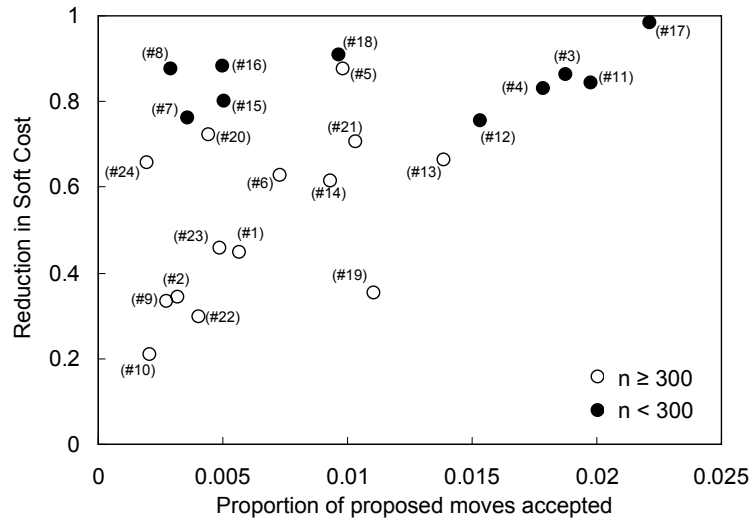


Figure 3: Scatter diagram showing the relationship between the proportion of accepted neighbourhood moves in Stage 3, and the resultant reduction in the Soft Cost (expressed as a proportion). Each point in the graph is averaged across all runs with a particular instance where a Distance to Feasibility of zero was achieved, with the exception of instance 22 where the run with the lowest Distance to Feasibility of the 51 runs was used.

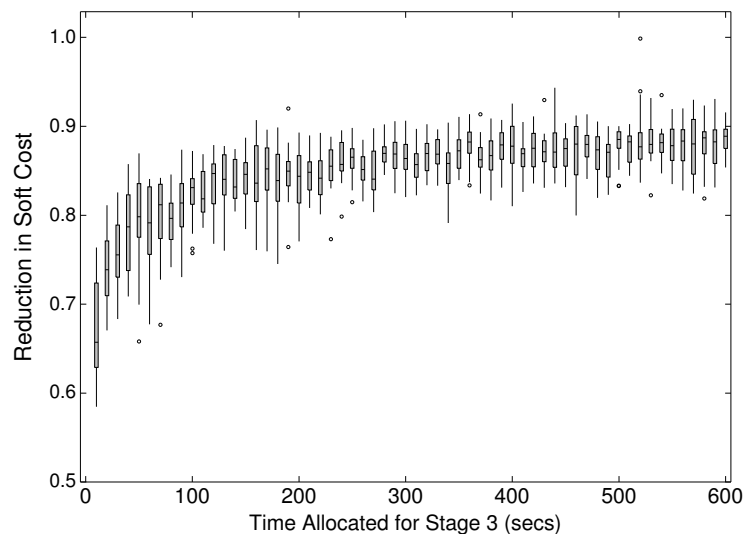


Figure 4: Box plot demonstrating the relationship between the time allocated to Stage 3, and the resultant reduction in the Soft Cost (expressed as a proportion). Results were gained from twenty runs on instance comp-2007-2-16 for each time threshold 10 to 600 seconds, incrementing in steps of 10.

the ability of the Stage 3 annealing process to reduce the Soft Cost. In order to eliminate noise, only instance 16 is considered here – an instance that seemed to be quite easy for finding feasibility, and yet still proved to be quite troublesome when trying to reduce the Soft Cost down to zero. It is obvious from fig. 4 that as the time limit is increased (therefore providing a slower cooling) the resultant reduction in the Soft Cost also increases. Note, however, that these improvements seem to “tail off” and are only marginal beyond 200-or-so seconds. Note also that improvements

of $> 90\%$ are quite rare on the whole, even when very generous amounts of run time are granted, suggesting that this algorithm seems to encounter some difficulty in achieving perfect solutions in this case. We will reflect on potential reasons for this in our discussions in Section 5.

5 Conclusions

In this paper, a three stage metaheuristic-based algorithm has been presented for the post enrolment-based course timetabling problem used in track-2 of the Second International Timetabling Competition. Crucially this algorithm is time-dependent, meaning that it is able to alter the intensity of the search conducted in accordance with the amount of available run time granted. In practical circumstances such a feature will be advantageous as users are not required to manually assign values to parameters *a priori*, nor are they required to perform any sort of tuning procedure – a strategy used by some other entrants in the competition.

We performed a comparison between this algorithm and the five official finalists of the competition (Cambazard et al., 2008; Chiarandini and Stützle, 2002; Atsuta et al., 2008; Mayer et al., 2008; Muller, 2008). This was done by taking a random sample of ten runs on all instances and incorporating the results into the ranking process used for choosing the competition winner. The comparison reveals that our algorithm comes in 6th place with a rank-average of 38.9.⁶ For reference, the resultant rank-averages of the five finalists with these results included are, in order, 1st place = 14.8; 2nd = 28.2; 3rd = 31.9; 4th = 33.6; and 5th = 35.6. Full results and further details of these experiments are detailed in Appendix A.

One issue with this algorithm is that for Stages 2 and 3, in order to ensure that the constraints satisfied in previous stages are not re-violated the proposed neighbourhood operators have been restricted so that moves causing such violations are automatically rejected. In Stage 3, for example, this means that the algorithm only searches in the space of feasible solutions. However, in this case, by using a restricted neighbourhood operator there is no guarantee that all feasible solutions will communicate with one another. In other words, the feasible-only search space may, in effect, be split into a number of disjoint subspaces, with areas of non-feasibility – that cannot be traversed using the current neighbourhood operator – occupying the space in-between. A practical implication of this, as we saw in fig. 4, is that even if the algorithm is granted excess computation time, there will be no guarantee of finding an optimal (i.e. perfect) solution, despite the fact that these problem instances are all known to feature at least one. One way of increasing the freedom to move about the feasible-only search space is obviously to try and increase the number of neighbourhood moves that are accepted by the algorithm. This might be achieved, for example, by making use of a maximum matching algorithm in order to dynamically assign events to rooms, thus maximising timeslot occupancy (see, for example, the competition entries of Cambazard et al. (2008) and Chiarandini et al. (2008)). Another strategy along similar lines would be to use a variety of different, albeit more expensive, neighbourhood operators such as the Hungarian operator (Cambazard et al., 2008) or the Kempe chain neighbourhood operator (Thompson and Dowsland, 1998; Chiarandini et al., 2008). We suggest that it is these omissions that are the main contributors to the differences in performance between this approach and the leading entries to the competition. Nevertheless, certain aspects of the approach outlined in this paper have shown to be very successful, and the combination of such aspects with features used in the other competition entries, particularly those mentioned above, seems to offer a promising future research path.

References

Abramson, D., Krishnamoorthy, H., Dang, H., 1996. Simulated annealing cooling schedules for the school timetabling problem. *Asia-Pacific Journal of Operational Research* 16, 1–22.

⁶Refer to the ITC2007 website for full listings of the other algorithms' results and for details of this ranking process.

- Atsuta, M., Nonobe, K., Ibaraki, T., 2008. Itc-2007 track 2: An approach using general csp solver.
URL www.cs.qub.ac.uk/itc2007/winner/bestcoursesolutions/Atsuta_et_al.pdf
- Brelaz, D., 1979. New methods to color the vertices of a graph. *Commun. ACM* 22 (4), 251–256.
- Burke, E., Gendreau, M. (Eds.), 2008. *Proceedings of the Seventh International Conference for the Practice and Theory of Automated Timetabling*. Universite de Montreal, Canada.
URL <http://www.asap.cs.nott.ac.uk/patat/patat08>
- Cambazard, H., Hebrard, E., Osullivan, B., A., P., 2008. Local search and constraint programming for the post-enrolment-based course timetabling problem. In: Burke and Gendreau (2008).
URL <http://www.asap.cs.nott.ac.uk/patat/patat08>
- Chiarandini, M., Fawcett, C., Hoos, H., 2008. A modular multiphase heuristic solver for post enrolment course timetabling. In: Burke and Gendreau (2008).
URL <http://www.asap.cs.nott.ac.uk/patat/patat08>
- Chiarandini, M., Stützle, T., 2002. An application of iterated local search to graph coloring. In: Johnson, I. D., Mehrotra, A., Trick, M. (Eds.), *Proceedings of the Computational Symposium on Graph Coloring and its Generalizations*. New York, USA, pp. 112–125.
- Di Gaspero, L., McCollum, B., Schaerf, A., August 2007. The second international timetabling competition (itc-2007): Curriculum-based course timetabling (track 3). Tech. Rep. QUB/IEEE/Tech/ITC2007/CurriculumCTT/v1.0/1, School of Computing, Queens University, Belfast.
- Kirkpatrick, S., Gelatt, C., Vecchi, M., 1983. Optimization by simulated annealing. *Science* 4598, 671–680.
- Kostuch, P., 2005. The university course timetabling problem with a 3-phase approach. In: Burke, E., Trick, M. (Eds.), *Practice and Theory of Automated Timetabling (PATAT) V*. Vol. 3616. Springer-Verlag, Berlin, pp. 109–125.
- Lewis, R., 2006. *Metaheuristics for university course timetabling*. Ph.D. thesis, School of Computing, Napier University, Edinburgh.
- Lewis, R., 2008. A time-dependent metaheuristic algorithm for post enrolment-based course timetabling. In: Burke and Gendreau (2008).
URL <http://www.asap.cs.nott.ac.uk/patat/patat08>
- Lewis, R., Paechter, B., McCollum, B., August 2007. Post enrolment based course timetabling: A description of the problem model used for track two of the second international timetabling competition. *Cardiff Working Papers in Accounting and Finance A2007-3*, Cardiff Business School, Cardiff University, ISSN: 1750-6658.
- Mayer, A., Nothegger, C., Chwatal, A., Raidl, G., 2008. Solving the post enrolment course timetabling problem by ant colony optimization. In: Burke and Gendreau (2008).
URL <http://www.asap.cs.nott.ac.uk/patat/patat08>
- McCollum, B., McMullan, P., Burke, E., Parkes, A., Qu, R., September 2007. The second international timetabling competition: Examination track. Tech. Rep. QUB/IEEE/Tech/ITC2007/Exam/v4.0/17, School of Computing, Queens University, Belfast.
- McCollum, B., Schaerf, A., Paechter, B., McMullan, P., Lewis, R., Parkes, A., de Gaspero, L., Qu, R., Burke, E., 2009. Setting the research agenda in automated timetabling: The second international timetabling competition. *INFORMS Journal on Computing* 10.1287/ijoc.1090.0320.
- Muller, T., 2008. ITC2007: Solver description. In: Burke and Gendreau (2008).
URL <http://www.asap.cs.nott.ac.uk/patat/patat08>

Table 5: Results used in the comparison with the competition entries. For each instance, results are ordered from best (1) to worst (10).

Instance	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
comp-2007-2-1	0 (1235)	0 (1459)	0 (1580)	0 (1661)	0 (1729)	0 (1819)	32 (1535)	32 (1766)	53 (1532)	57 (1799)
comp-2007-2-2	0 (1718)	0 (1751)	0 (2018)	42 (1473)	42 (1866)	52 (2037)	58 (2195)	73 (1904)	98 (1774)	168 (1734)
comp-2007-2-3	0 (327)	0 (349)	0 (355)	0 (391)	0 (415)	0 (461)	0 (472)	0 (487)	0 (614)	0 (663)
comp-2007-2-4	0 (424)	0 (475)	0 (485)	0 (496)	0 (544)	0 (586)	0 (595)	0 (612)	0 (622)	0 (627)
comp-2007-2-5	0 (106)	0 (175)	0 (176)	0 (189)	0 (192)	0 (271)	0 (276)	0 (276)	0 (420)	0 (533)
comp-2007-2-6	0 (479)	0 (562)	0 (567)	0 (583)	0 (597)	0 (612)	0 (644)	0 (708)	0 (753)	0 (774)
comp-2007-2-7	0 (107)	0 (327)	0 (361)	0 (375)	0 (381)	0 (390)	0 (391)	0 (432)	0 (577)	0 (658)
comp-2007-2-8	0 (119)	0 (164)	0 (164)	0 (184)	0 (226)	0 (229)	0 (244)	0 (249)	0 (301)	0 (379)
comp-2007-2-9	45 (1858)	49 (1914)	54 (2029)	60 (1687)	60 (2067)	65 (1964)	80 (2063)	95 (1850)	117 (1768)	232 (1916)
comp-2007-2-10	0 (2143)	31 (2184)	51 (2434)	57 (2017)	87 (2236)	87 (2396)	107 (2149)	116 (2410)	178 (2222)	189 (2072)
comp-2007-2-11	0 (373)	0 (399)	0 (424)	0 (436)	0 (451)	0 (501)	0 (550)	0 (572)	0 (607)	0 (681)
comp-2007-2-12	0 (649)	0 (681)	0 (743)	0 (767)	0 (771)	0 (797)	0 (820)	0 (823)	0 (905)	0 (906)
comp-2007-2-13	0 (482)	0 (520)	0 (582)	0 (587)	0 (601)	0 (639)	0 (645)	0 (659)	0 (674)	0 (811)
comp-2007-2-14	0 (265)	0 (560)	0 (584)	0 (591)	0 (625)	0 (683)	0 (700)	0 (739)	0 (785)	0 (788)
comp-2007-2-15	0 (265)	0 (274)	0 (299)	0 (313)	0 (321)	0 (338)	0 (340)	0 (374)	0 (381)	0 (423)
comp-2007-2-16	0 (151)	0 (157)	0 (161)	0 (162)	0 (180)	0 (202)	0 (208)	0 (214)	0 (231)	0 (238)
comp-2007-2-17	0 (0)	0 (9)	0 (15)	0 (17)	0 (17)	0 (28)	0 (32)	0 (33)	0 (66)	0 (69)
comp-2007-2-18	0 (25)	0 (73)	0 (159)	0 (213)	0 (270)	0 (273)	0 (318)	0 (349)	0 (365)	0 (374)
comp-2007-2-19	139 (1802)	141 (1770)	152 (2029)	190 (1868)	198 (2046)	249 (1964)	262 (1923)	341 (1780)	489 (2187)	549 (1989)
comp-2007-2-20	0 (771)	0 (811)	0 (824)	0 (844)	0 (849)	0 (862)	0 (874)	0 (884)	0 (947)	0 (963)
comp-2007-2-21	0 (429)	0 (474)	0 (490)	0 (498)	0 (500)	0 (524)	0 (539)	0 (552)	0 (562)	0 (646)
comp-2007-2-22	878 (1748)	901 (1608)	937 (1472)	986 (1498)	994 (1651)	1007 (1689)	1021 (1481)	1030 (1506)	1182 (1226)	1252 (1366)
comp-2007-2-23	112 (3883)	219 (3828)	232 (3468)	236 (3837)	246 (2905)	324 (4041)	360 (3949)	372 (4083)	404 (3797)	515 (3994)
comp-2007-2-24	0 (841)	0 (1077)	0 (1343)	26 (1030)	28 (1244)	40 (1239)	45 (1232)	66 (1310)	89 (1018)	93 (1226)

Rossi-Doria, O., Samples, M., Birattari, M., Chiarandini, M., Knowles, J., Manfrin, M., Mastrolilli, M., Paquete, L., Paechter, B., Stutzle, T., 2002. A comparison of the performance of different metaheuristics on the timetabling problem. In: Burke, E., De Causmaecker, P. (Eds.), Practice and Theory of Automated Timetabling (PATAT) IV. Vol. 2740. Springer-Verlag, Berlin, pp. 329–351.

Socha, K., Samples, M., 2003. Ant algorithms for the university course timetabling problem with regard to the state-of-the-art. In: Evolutionary Computation in Combinatorial Optimization (EvoCOP 2003). Vol. 2611. Springer-Verlag, Berlin, pp. 334–345.

Thompson, J., Dowsland, K., 1998. A robust simulated annealing based examination timetabling system. Computers and Operations Research 25 (7/8), 637–648.

van Laarhoven, P., Aarts, E., 1987. Simulated Annealing: Theory and Applications. Kluwer Academic Publishers, Reidel, The Netherlands.

A Comparison with Competition Results

Table 5 shows the results achieved by our algorithm within the specified time limit in a random sample of ten runs on each problem instance. Note that the algorithm was executed “blindly” on the hidden instances – that is, all refinements were made to our implementation *before* performing runs on the hidden instances. This is consistent with the process used for officially ranking the finalists of the competition.