

## Adaptive selection of heuristics for improving exam timetables

Edmund K. Burke · Rong Qu · Amr Soghier

Published online: 26 June 2012  
© Springer Science+Business Media, LLC 2012

**Abstract** This paper presents a hyper-heuristic approach which hybridises low-level heuristic moves to improve timetables. Exams which cause a soft-constraint violation in the timetable are ordered and rescheduled to produce a better timetable. It is observed that both the order in which exams are rescheduled and the heuristic moves used to reschedule the exams and improve the timetable affect the quality of the solution produced. After testing different combinations in a hybrid hyper-heuristic approach, the Kempe chain move heuristic and time-slot swapping heuristic proved to be the best heuristic moves to use in a hybridisation. Similarly, it was shown that ordering the exams using Saturation Degree and breaking any ties using Largest Weighted Degree produce the best results. Based on these observations, a methodology is developed to adaptively hybridise the Kempe chain move and timeslot swapping heuristics in two stages. In the first stage, random heuristic sequences are generated and automatically analysed. The heuristics repeated in the best sequences are fixed while the rest are kept empty. In the second stage, sequences are generated by randomly assigning heuristics to the empty positions in an attempt to find the best heuristic sequence. Finally, the generated sequences are applied to the problem. The approach is tested on the Toronto benchmark and the exam timetabling track of the second International Timetabling Competition, to evaluate its generality. The hyper-heuristic with low-level improvement heuristics approach was found to generalise well over the two different datasets and performed comparably to the state of the art approaches.

---

E.K. Burke  
Department of Computing Science and Mathematics, University of Sterling, Sterling, FK9 4LA, UK  
e-mail: [e.k.burke@stir.ac.uk](mailto:e.k.burke@stir.ac.uk)

R. Qu · A. Soghier (✉)  
School of Computer Science, University of Nottingham, Nottingham, NG8 1BB, UK  
e-mail: [azs@cs.nott.ac.uk](mailto:azs@cs.nott.ac.uk)

R. Qu  
e-mail: [rxq@cs.nott.ac.uk](mailto:rxq@cs.nott.ac.uk)

## 1 Introduction

For more than 40 years exam timetabling has been a particularly well studied domain across the Artificial Intelligence and Operational Research communities. This is due to its importance in many academic institutions worldwide. However, much of the research has been aimed at developing methodologies that produce the best quality timetables for a single problem (Qu et al. 2009b). A more recent direction in this field, represented by some work on hyper-heuristics, aims to raise the level of generality of search methodologies to create algorithms that operate effectively over a range of problems. A hyper-heuristic can be seen as a heuristic to choose heuristics (Burke et al. 2003). In this case, the low-level heuristics represent the search space. In timetabling, the low-level heuristics can be categorised as heuristics which construct a timetable or heuristics which perform certain moves to improve a constructed timetable. This paper presents a random iterative hyper-heuristic approach which uses improvement low-level heuristics. This approach was tested on the Toronto benchmark and the second International Timetabling Competition (ITC2007) exam timetabling problems. It was found to generalise well over the two datasets. Furthermore, very competitive results have been produced against other approaches in the literature.

The following section presents a brief description of the benchmarks. Sections 2.3 and 2.4 provide an overview on different approaches developed in the exam timetabling domain. Furthermore, Sect. 2.5 concentrates on hyper-heuristic approaches. A random iterative hyper-heuristic to improve timetables is proposed in Sect. 3. An adaptive methodology to select low-level heuristics and the results obtained are presented in Sect. 4. The future extensions of this work are summarised in Sect. 5.

## 2 Exam timetabling

### 2.1 The Toronto benchmark

An exam timetabling problem consists of the allocation of a set of exams to a given set of timeslots. The generated timetable must satisfy the hard constraints of the problem which are the requirements that cannot be violated. For example, no one student must be scheduled to sit two exams during the same period. A timetable which meets all the hard constraints given is called a feasible timetable. A timetabling problem can also have a set of soft constraints. Soft constraints are constraints that can be violated and the level of violation determines the quality of the timetable generated. An example of a soft constraint is that there must be a certain number of periods between two exams taken by the same student. Therefore, high quality timetables contain the least number of soft constraint violations. The Toronto benchmark problem is well known in the exam timetabling community since it was first introduced by Carter et al. (1996) in 1996. Over the years, a slightly different version was made available and used to test some approaches in the literature. The characteristics of the two versions of this dataset are presented in Table 1. The problem has one hard constraint. This says that conflicting exams (with common students) must not be assigned to the same time slot. In addition, there is one soft constraint. This says that conflicting exams should be spread throughout the timetable as much as possible. The objective is to minimise the sum of proximity costs given as:

$$\sum_{i=0}^4 (w_i \times n) / S$$

where

**Table 1** Characteristics of the two versions of the Toronto Benchmark datasets

Problem	Exams I/II/IIC	Students I/II	Enrolments I/II	Density	Time slots
car91	682	16925	56877	0.13	35
car92	543	18419	55522	0.14	32
ear83 I	190	1125	8109	0.27	24
ear83 II	189	1108	8057	0.27	24
hec92 I	81	2823	10632	0.42	18
hec92 II	80	2823	10625	0.42	18
kfu93 I	461	5349	25113	0.06	20
lse91	381	2726	10918	0.06	18
sta83 I	139	611	5751	0.14	13
sta83 II	138	549	5417	0.19	35
tre92	261	4360	14901	0.18	23
uta92 I	622	21266	58979	0.13	35
uta92 II	638	21329	59144	0.12	35
ute92	184	2750	11793	0.08	10
yor83	181	941	6034	0.29	21
yor83 II	180	919	6002	0.3	21

- $w_i = 2^{4-i}$  is the cost of assigning two exams with  $i$  time slots apart. Only conflicting exams with common students and which are four or less time slots apart are considered as violations
- $n$  is the number of students involved in the conflict
- $S$  is the total number of students in the problem

This problem has been used to test and compare many approaches in the literature. Recently, a more constrained set of benchmarks was made available as part of the International Timetabling Competition (ITC2007) (McCollum et al. 2010). The next section describes the ITC2007 dataset in detail.

## 2.2 The international timetabling competition (ITC2007) dataset

The ITC2007 exam timetabling track could be considered to be a complex and a more practical dataset in comparison to the Toronto benchmark. This is due to the larger number of constraints it contains. A full description of the problem and the evaluation function can be found in McCollum et al. (2010). In addition, the characteristics which define the instances are summarised in Table 2. The problem consists of the following:

- A set of timeslots covering a specified length of time. The number of timeslots and their durations are provided.
- A set of exams which should be allocated to the timeslots.
- A list of students enrolled in each exam.
- A set of rooms with different capacities.
- A set of additional hard constraints.
- A set of soft constraints and their associated penalties.

In comparison to the Toronto benchmark, the ITC2007 dataset has more than one hard constraint. The hard constraints can be outlined as follows:

**Table 2** Characteristics of the ITC2007 dataset

Instance	Conflict density	Exams	Students	Periods	Rooms	No. of hard constraints
exam 1	5.05	607	7891	54	7	12
exam 2	1.17	870	12743	40	49	14
exam 3	2.62	934	16439	36	48	185
exam 4	15.0	273	5045	21	1	40
exam 5	0.87	1018	9253	42	3	27
exam 6	6.16	242	7909	16	8	23
exam 7	1.93	1096	14676	80	15	28
exam 8	4.55	598	7718	80	8	21

- No student sits more than one exam at the same time.
- The capacity for each individual room should not be exceeded at a given period.
- Period lengths should not be violated.
- All period related hard constraints need to be satisfied e.g. exam A after exam B.
- All room related hard constraints need to be satisfied e.g. exam A must use room X.

The soft constraint violations can be summarised as follows:

- **Two Exams in a Row** The number of occurrences where a student sits two exams in a row on the same day.
- **Two Exams in a Day** The number of occurrences where a student sits two exams on the same day. If the exams are back to back then this is considered as a Two Exams in a Row violation to avoid duplication.
- **Period Spread** The exams have to be spread a certain number of timeslots apart.
- **Mixed Durations** The number of occurrences where exams of different durations are assigned to the same room.
- **Larger Exams Constraint** The number of occurrences where the largest exams are scheduled near the end of the exam session. The number of the largest exams and the distance from the end of the exam session are specified in the problem description.
- **Room Penalty** The number of times where certain rooms, which have an associated penalty, are used.
- **Period Penalty** The number of times where certain timeslots, which have an associated penalty, are used.

### 2.3 Exam timetabling approaches for the ITC2007 dataset

A three phased approach was developed by Muller (2008) to solve the problems in the ITC2007 exam timetabling track. The first phase consists of an Iterative Forward Search algorithm to find a feasible solution. Hill climbing is used to find the local optima in the second phase. Finally, a Great Deluge Algorithm is applied to further explore the search space.

Gogos et al. (2008) proposed a method which used a GRASP (Greedy Randomised Adaptive Search Procedure) approach. In the construction phase, five orderings of exams based on various criteria are generated. Tournament selection is used to select exams until they are all scheduled. A backtracking strategy using a tabu list is employed as required. In the improvement phase, Simulated Annealing is used. Finally, room allocations are arranged using integer programming in the third phase.

Atsuta et al. (2008) used a constraint satisfaction solver which employed tabu search and iterated local search. The solver differentiates between the constraints and their corresponding weights during the computation to improve performance. De Smet (2008) also incorporated local search techniques in a solver called Drools, an Open-Source Business Rule Management System (<http://www.jboss.org/drools/>).

Pillay (2008) introduced a biologically inspired approach which is analogous to cell behaviour. The exams are initially ordered using the saturation degree heuristic and scheduled sequentially in the available “cells” i.e. timeslots. If more than one timeslot is available, the slot which causes the least overall constraint violation is chosen. Rooms are chosen using the best fit heuristic. If a conflict occurs before all the exams are scheduled, the timetable is rearranged to reduce the soft constraint violation. This is described as cell division. If the overall soft constraint violation is not improved without breaking the hard constraints, cell interaction occurs. The timeslots are swapped in this process to remove hard constraint violations. The process continues until a feasible solution is achieved. Finally, the contents of cells having equal durations are swapped to improve the solution. This is called cell migration.

McCollum et al. (2009) proposed a two phased approach where an adaptive heuristic is used to achieve feasibility during the first phase. The second phase improves the solution through the employment of a variant of the Great Deluge Algorithm.

Finally, Pillay (2008) presented an evolutionary algorithm based hyper-heuristic using three different chromosome representations. An initial population is created and iteratively improved by applying the processes of evaluation, selection and recreation.

#### 2.4 Exam timetabling approaches for the Toronto benchmark

Abdullah et al. (2009) presented a hybridisation of an electromagnetic-like mechanism (EM) and the Great Deluge algorithm. The technique estimates the quality of the solution and calculates a decay rate at every iteration of the search. These values depend on a force value calculation using the EM approach. This approach produced the best quality solution for four of the instances. An approach which uses a sequential construction method, employed by Caramia et al. (2008), to assign exams in the least number of timeslots was able to produce the best quality timetables for another four of the Toronto benchmark instances. It uses a greedy scheduler to obtain a feasible solution. A penalty decreaser and trader are then applied to improve the quality of the constructed solution. Burke et al. (2010) introduced an approach which combines a variable neighbourhood search with a genetic algorithm which produced the best quality solution for one of the Toronto instances. In addition, Burke and Bykov (2008) proposed a method where a hill-climber compares the candidate solution with a solution produced a couple of iterations back instead of the current solution. This was called the “late acceptance criteria” and it produced the best quality solutions for another two instances. The results obtained by these approaches are presented in Sect. 4.1.

Recently, some new methods were investigated to automatically find the best heuristic to solve a set of instances. This has led to the introduction of hyper-heuristics. The next section summarises some of the hyper-heuristic methods applied to the exam timetabling domain.

#### 2.5 Hyper-heuristics in exam timetabling

A hyper-heuristic can be seen as a method to choose low-level heuristics depending on the problems in hand. Furthermore, it could be used to adapt or tune heuristics and meta-heuristics. Hyper-heuristics in exam timetabling can be categorised, according to the low-level heuristics they use, into two types as follows:

1. Hyper-heuristics with constructive low-level heuristics
2. Hyper-heuristics with improvement low-level heuristics

A Tabu search approach was developed by Burke et al. (2007) to optimise a search space of heuristic sequences comprised of two or more low-level heuristics. This work was extended in later research by Qu et al. (2009a) to construct heuristic sequences which produce feasible timetables. The combinations are then analysed to find distribution patterns of low-level heuristics. This leads to the generation of sequences which are adaptively adjusted to construct better timetables. In addition, hybridisations of the graph based hyper-heuristic with local search methods were investigated in Qu and Burke (2009). The formal definition of the search problem of hyper-heuristics, where low-level heuristics represent the search space, has also been provided in Qu and Burke (2009).

Asmuni et al. (2004) used fuzzy logic to combine two out of three graph colouring heuristics. The idea was to combine the two heuristics into a single value which calculates the difficulty of allocating an exam to a timeslot. The exams are ordered using this value and are scheduled in order. Furthermore, the approach was extended to tune the fuzzy rules instead of keeping them fixed (Asmuni et al. 2009).

Ersoy et al. (2007) developed an approach called the hyperhill-climber where a hyper-heuristic is embedded in a memetic algorithm. The aim of this hyper-heuristic was to select the best hill-climber to apply or decide the best order in which hill-climbers are executed. In addition, Pillay and Banzhaf (2007) created another approach where genetic programming was used to evolve hyper-heuristics.

Biligan et al. (2007) presented different heuristic selection methods and acceptance criteria for hyper-heuristics in exam timetabling. Finally, a different method of combining heuristics was presented by Pillay and Banzhaf (2009). The low-level heuristics were combined hierarchically and applied simultaneously instead of being applied sequentially.

### 3 A hyper-heuristic with low-level improvement heuristics

Several low-level heuristics can be used to improve a timetable with varying quality. The different low-level heuristics used could be considered as different methods for escaping from local optima. However, the order in which exams are moved and the type of moves performed play an important role in finding the best quality solution. An initial feasible solution is constructed using the Largest Degree heuristic where the exams in the ordering are assigned to the timeslot causing the least penalty. In the case where there is more than one timeslot with the lowest penalty, one of them is randomly chosen. Our objective is to analyse the performance of the different low-level heuristics used to minimise the penalty incurred from a constructed solution. In addition, we test the effect of using different orderings for the exams causing penalties in the solution. Finally we develop an adaptive approach which orders the exams causing violations and automatically selects the best heuristic to use for each exam to produce an improvement.

#### 3.1 The low-level heuristics

In this paper we investigate the effect of using different low-level heuristics or neighbourhood operators to improve timetables. A combination of two improvement low-level heuristics is used in our approach. The following is a list of the heuristics investigated:

1. Move Exam (ME): This heuristic selects an exam and reassigns it to the timeslot causing the least penalty.

---

**Algorithm 1** The pseudo-code of the random iterative hyper-heuristic with low-level improvement heuristics

---

```

order the exams using the LD heuristic
construct a feasible solution by assigning the exams to the time slot causing the least
penalty
create a randomly ordered list of the exams which contribute to the overall penalty in-
curred
for  $i = 1 \rightarrow i = 50$  do
  for  $n = 1 \rightarrow n = e // e$ : number of exams causing penalty do
    initialise heuristic sequence  $h = [\text{KCM KCM KCM KCM}] // h$  has a length of  $e$ 
     $h =$  randomly change  $n$  heuristics in  $h$  to ME, SE or ST
    construct a solution  $s_c$  using  $h$ 
    if solution  $s_c$  is feasible then
      save  $h$  and the penalty of its corresponding solution  $s_c$ 
    end if
  end for
end for

```

---

2. Swap Exam (SE): This heuristic selects an exam and tries to swap it with a scheduled exam leading to the least penalty timetable.
3. Kempe Chain Move (KCM): This is similar to the SE heuristic but is more complex as it involves swapping a subset of conflicting exams in two distinct timeslots. This neighbourhood operator was successful in some earlier studies (Burke et al. 2010; Thompson and Dowsland 1996).
4. Swap Timeslot (ST): This heuristic selects an exam and swaps all the exams in the same timeslot with another set of exams in a different timeslot. After testing all the timeslots, the swap producing the least penalty timetable is applied.

### 3.2 The random iterative hyper-heuristic

The study presented in this paper takes a similar approach to that presented in Qu et al. (2009a) where a random iterative hyper-heuristic generates heuristic sequences of different quality to solve the benchmark problem mentioned in Sect. 2.1. Instead of being used to construct solutions, the heuristic sequences are used here to improve constructed feasible solutions by rescheduling exams causing penalties. Algorithm 1 presents the pseudo-code of this random iterative hyper-heuristic. The process starts by constructing an initial feasible solution. Since the initial solution constructed affects the improvement process, a random largest degree graph colouring heuristic which orders exams according to the number of conflicts each exam has with others is used (Burke et al. 2010). This allows us to compare our approach to other approaches in the literature which use a similar method in construction. At every iteration, the exams causing violations in the constructed solution are identified and a random sequence of moves is generated. A move is the application of one of the low-level heuristics described in Sect. 3.1. The sequence of moves is then applied to the sequence of exams as they are unscheduled one by one. Only moves that improve the current solution are accepted. If a move does not improve the solution, it is skipped and the exam stays in its current position. A sequence is discarded if an improvement is not obtained after the whole sequence is applied.

**Table 3** Results using KCM without a hybridisation and with several different moves

	hec92 I	yor83 I	sta83 I	tre92
KCM without hybridisation Best	13.50	43.84	160.43	8.99
KCM with ME Best	12.03	43.84	157.48	8.91
KCM with SE Best	12.03	42.37	157.75	8.75
KCM with ST Best	<b>11.30</b>	<b>41.79</b>	<b>157.27</b>	<b>8.57</b>

This approach was applied to four instances (hec92 I, sta83 I, yor83 I and tre92) of the Toronto benchmark exam timetabling problems described in Sect. 2.1 for off-line learning of the best heuristic hybridisations and the order of execution leading to the best improvement. These instances were chosen as they vary in size and cover a range of conflict densities (see Table 1). After running this process for ( $e \times 50$ ) times, where  $e$  is the number of exams causing soft constraint violations in the constructed solution, a set of sequences and the penalties of their corresponding solutions are obtained for further investigation on the effectiveness of the different heuristics used. Note that, only 50 samples were collected for each rate of hybridisation as we found that using more samples does not improve the quality of the final outcome at this point. Finally, an adaptive approach was developed and applied to the Toronto benchmark. Furthermore, to test the generality of the approach, it was applied to the ITC2007 exam timetabling track. The approach is presented in Sect. 4.

### 3.3 Analysis of hybridising improvement low-level heuristics

In order to clearly observe the effect of the different low-level heuristics in improving solutions, the heuristic sequences generated consist of two heuristics. We use the Kempe chain move heuristic as the basic heuristic in the sequences as it has proved to be successful in previous work (Burke et al. 2010; Thompson and Dowsland 1996). The Kempe chain move involves swapping a subset of exams in two distinct timeslots making sure that a hard constraint violation does not occur. The rest of the heuristics (ME, SE and ST) are randomly hybridised into the list of KCM.

The random sequences are generated with different percentages of hybridisation by inserting  $n$  ME, SE or ST,  $n = [1, \dots, e]$  in the sequences. For each hybridisation of KCM with either ME, SE or ST, 50 samples are obtained for each amount of hybridisation. Duplicate sequences are discarded and another sequence is generated instead. The sequences are re-initialised in each iteration to lead the search to explore a wider area of the search space at this stage.

We applied this approach to four instances of the Toronto benchmark exam timetabling problems (Carter et al. 1996). Table 3 presents the results obtained using ME, SE and ST in a hybridisation with KCM as well as a comparison against using KCM only.

It is observed that using a Kempe chain only produces the worst results. After introducing other heuristics in a hybridisation with the Kempe chain moves, better results were obtained. Another observation from Table 3 is that swapping timeslots and performing Kempe chain moves produces the best improvement for all the problems. One possible reason may be that swapping timeslots allows the search to be more diverse and to sample different areas of the search space to find good solutions more quickly. In addition, no obvious trends could be obtained on the amount of ST hybridisation within the best heuristic sequences. However, it is observed in all the sequences leading to the best timetables, the ST heuristic is randomly distributed within the sequence and the percentage of hybridisation is less than 50 %.



**Table 4** Results of hybridising KCM with ST using different orderings of the exams causing a soft constraint violation. The notation  $X$  tb  $Y$  means heuristic  $Y$  is used to break ties in heuristic  $X$ 

	hec92 I	yor83 I	sta83 I	tre92
KCM with ST + RO Average	11.99	42.63	159.74	9.02
KCM with ST + RO Best	11.60	41.33	158.46	8.66
KCM with ST + LD Average	12.69	42.10	163.32	9.00
KCM with ST + LD Best	12.50	39.69	159.50	8.66
KCM with ST + LWD Average	12.06	42.08	159.74	8.91
KCM with ST + LWD Best	11.39	39.69	157.76	8.64
KCM with ST + LP Average	12.15	42.09	159.52	8.85
KCM with ST + LP Best	11.32	39.69	157.49	8.56
KCM with ST + SD tb LWD Average	11.45	41.96	159.39	8.74
KCM with ST + SD tb LWD Best	<b>11.25</b>	<b>39.56</b>	<b>157.37</b>	<b>8.54</b>

### 3.4 Variations of orderings of the exams causing a penalty

To analyse the effect of ordering the unscheduled exams causing a soft constraint violation in a previous solution, we decided to test different orderings while using the Kempe Chain and swapping timeslot hybridisation stated in the previous section. After the exams causing violations are identified, they are ordered first before being reassigned to a timeslot. Several orderings can be used to guide the search as follows:

- Largest Degree (LD): The exams are ordered decreasingly according to the number of conflicts each exam has with others.
- Largest Weighted Degree (LWD): The exams are ordered similarly to LD but weighted according to the number of students involved in the conflict.
- Saturation Degree (SD): The exams are ordered increasingly according to the number of remaining timeslots available to assign them without causing conflicts. In the case where ties occur, LWD is used as a tie breaker. From our previous work it was shown that SD produces the best results when LWD is used to break ties in the ordering (Burke et al. 2011).
- Largest Penalty (LP): The exams are ordered decreasingly according to the penalty they incur in the current solution.
- Random Ordering (RO): The exams are ordered randomly.

Table 4 presents the results of applying different orderings to the unscheduled exams, then running a random heuristic sequence of KCM and ST to assign them in better timeslots.

As shown in Table 4, we found that using SD and breaking any ties in the ordering using LWD produced the best results. This is because SD orders the unscheduled exams according to the number of timeslots available to assign them without causing conflicts. Therefore, the chances of moving exams at the top of the SD list and finding better timeslots for them become higher. Ordering the exams according to the penalty they incur was seen to be the second best ordering followed by LWD. The effect of using LD and RO was the same as randomly choosing an exam to reschedule.

A t-test is also carried out to give an indication if the results using SD, LP and LWD are significantly different. Tables 5, 6 and 7 summarise the p-values of the t-tests carried out between the results of different orderings, which are significantly different in all the cases.

**Table 5** t-test on the results from ordering exams causing violations using SD and LP

	hec92 I	yor83 I	sta83 I	tre92
p-value	1.3E-05	5.27E-21	1.9E-18	1.33E-04
t Stat	5.96	74.66	50.35	4.73

**Table 6** t-test on the results from ordering exams causing violations using SD and LWD

	hec92 I	yor83 I	sta83 I	tre92
p-value	1.58E-07	5.3E-18	6.85E-16	5.56E-08
t Stat	8.67	47	33.89	9.41

**Table 7** t-test on the results from ordering exams causing violations using LP and LWD

	hec92 I	yor83 I	sta83 I	tre92
p-value	0.18	7.51E-04	1.58E-11	3.36E-09
t Stat	0.93	3.87	17.04	11.62

#### 4 Adaptive selection of low-level heuristics for improving exam timetables

Algorithm 2 presents the initialisation stage of the adaptive approach. The exams causing a penalty are first identified and are unscheduled. They are then put in a list and ordered using SD. Random heuristic sequences are generated using KCM and ST to reschedule the exams. The sequences are then applied to the ordered exams and the corresponding solutions are saved. Note that, only 10 sequences are generated for each rate of hybridisation to be able to adhere to the time limitation and compare our results with the best in the literature for the ITC2007 dataset. Furthermore, limiting the number of sequences generated in this stage makes it easier to analyse and observe any trends in the sequences generating the best results.

The above observations indicate that the best solutions were obtained when ordering the exams causing violations using SD, and rescheduling them using either a Kempe-chain move or swapping timeslots. It was also observed that the heuristic sequences producing the top 5 % results used the same move for the majority of the exams (i.e. the same heuristic appears at the same position in more than 75 % of the sequences). Therefore, we developed an intelligent approach that performs an analysis of the best 5 % of the sequences produced to generate a new set of sequences. The new set of sequences obtained better results for all the problem instances. The adaptive approach was tested and showed to be effective and comparable with the best approaches in the literature.

Algorithm 3 presents the pseudo-code of the approach which hybridises ST with KCM in two stages. The process is presented as follows:

1. In the first stage, the best 5 % of heuristic sequences are collected and analysed. If the same heuristic is used at the same position for more than 75 % of the heuristic sequences, then it is stored. Otherwise the position is kept empty. Note that, we also tested collecting more than 5 % of the best sequences to analyse them. However, the behaviour was random since no trends were seen. Therefore, to guarantee the effectiveness of the ap-

---

**Algorithm 2** The pseudo-code of the initialisation stage of the adaptive hyper-heuristic with low-level improvement heuristics

---

```

order the exams using the LD heuristic
construct a feasible solution by assigning the exams to the time slot causing the least
penalty
create a list of the exams which contribute to the overall penalty incurred, ordered by SD
for  $i = 1 \rightarrow i = (e \times 10)$  //  $e$ : number of exams causing penalty do
  for  $n = 1 \rightarrow n = e$  do
    initialise heuristic sequence  $h = \{\text{KCM KCM} \dots \text{KCM KCM}\}$ 
     $h =$  randomly change  $n$  heuristics in  $h$  to ST
    construct a solution using  $h$ 
    if solution  $S_c$  is feasible then
      save  $h$  and the penalty of its corresponding solution  $S_c$ 
    end if
  end for
end for

```

---

**Algorithm 3** Adaptive generation of heuristic sequences hybridising KCM and ST

---

```

construct the initial heuristic sequences // see Algorithm 2
collect the best 5 % of the heuristic sequences
for  $i = 0 \rightarrow i = (\text{number of exams causing penalty} - 1)$  do
  count = 0
  for  $j = 1 \rightarrow j = (\text{number of sequences} - 1)$  do
    if heuristicSequence[ $i$ ][ $j$ ] = KCM then
      count ++
    end if
  end for
  if count >  $(0.75 \times \text{number of sequences})$  then
    finalHeuristicSequence[ $i$ ] = KCM
  else if count <  $(0.25 \times \text{number of sequences})$  then
    finalHeuristicSequence[ $i$ ] = ST
  else
    finalHeuristicSequence[ $i$ ] = empty
  end if
end for
 $n = (\text{number of empty positions} \times 5)$  for large problems or  $(\text{number of empty positions} \times 10)$  for small problems
for  $i = 0 \rightarrow i = (n - 1)$  do
  for  $j = 0 \rightarrow j = (\text{number of sequences} - 1)$  do
    if finalHeuristicSequence[ $j$ ] = empty then
      finalHeuristicSequence[ $j$ ] = KCM or ST
    end if
  end for
  construct a solution using finalHeuristicSequence[ $j$ ]
  if a better solution is obtained then
    save finalHeuristicSequence[ $j$ ] and the penalty of its corresponding solution
  end if
end for

```

---

proach, the heuristic was chosen for a certain position if it appears in 75 % of the best 5 % heuristic sequences collected.

2. In the second stage, the empty positions are randomly assigned as KCM or ST. ( $n \times 5$ ) sequences for the large problems (uta92 I, uta92 II, car91 and car92) and ( $n \times 10$ ) sequences for the small problems are generated, respectively. The generated sequences are then applied to the problem.

#### 4.1 The Toronto benchmark results

We tested this approach on the Toronto benchmark exam timetabling problems and present the results in Tables 8 and 9. The average computational time for each stage across the instances is also presented for 30 runs on a Pentium IV machine with a 1 GB memory. In addition, the number of exams causing a penalty is presented in the tables.

The best results stated in the literature are presented in Table 10. These include the hybridisation of an electromagnetic-like mechanism (EM) and the Great Deluge algorithm employed by Abdullah et al. (2009), the hill-climbing with a late acceptance strategy implemented by Burke and Bykov (2008), the variable neighbourhood search incorporating the use of genetic algorithms used by Burke et al. (2010) and the sequential construction method developed by Caramia et al. (2008). These algorithms are briefly discussed in Sect. 2.4.

The results obtained indicate the generality of our approach to different constructed timetables regardless of the size. We also make a comparison with other hyper-heuristics which produced the best results in the literature in Table 11. In comparison with the graph-based hyper-heuristic in Burke et al. (2007), our approach performs better in all the cases reported. In addition, it performs better in 8 out of 11 cases in comparison with the hyper-heuristics investigated in Pillay and Banzhaf (2009) and Qu and Burke (2009). Finally, it performs better in 10 out of 11 cases compared to the tabu search hyper-heuristic investigated in Kendall and Mohd Hussin (2005). Only the problems presented in Table 11 were compared to other results since the results for the other instances in Table 1 were not reported in the literature.

#### 4.2 The international timetabling competition (ITC2007) results

To test the generality of our approach, we applied it to the ITC2007 exam timetabling dataset. The initial solution is constructed by ordering the exams according to their saturation degree. The exams are assigned a random timeslot in the situation where more than one timeslot is available. After a feasible solution is constructed the Adaptive Improvement Hyper-heuristic in Algorithm 3 was applied to the constructed solution. To allow a fair comparison with the reported competition results, the approach was run for the same amount of time using 11 distinct seeds for each instance. Table 12 presents the results we obtained in comparison with the best in the literature. The description of the approaches used for comparison is presented in Sect. 2.3. We do emphasise that the objective here is not to beat the best reported results but to demonstrate the generality of our approach to different problems with different constraints. A dash in the table means that no feasible solution was obtained.

The Extended Great Deluge in McCollum et al. (2009) obtained the best results for 5 out of the 8 instances. However, the approach was run for a longer time as it was developed after the competition. In the competition, the best results for all the 8 instances were reported in Muller (2008) using a three phased approach. The GRASP used in Gogos et al. (2008) produced the second best results.

**Table 8** Results from the adaptive improvement Hyper-heuristic (AIH) approach on the Toronto Benchmark dataset.

	hec92 I	yor83 I	ear83 I	sta83 I	car92	car91	uta92 I	ute92	lse91	tre92	kfu93
AIH Average	12.69	41.74	38.98	159.21	4.49	5.39	3.56	27.97	11.45	8.90	15.54
AIH Best	11.19	39.47	35.79	157.18	4.31	5.19	3.44	26.70	10.92	8.49	14.51
Time spent in first stage (s)	254	1227	1139	576	10286	74973	35149	178	1189	2065	2153
Time spent in second stage (s)	143	456	553	183	41954	22988	26135	641	277	2228	592
Number of exams causing violations	18	22	26	22	42	49	40	27	39	29	34

**Table 9** Contd. Results from the adaptive improvement Hyper-heuristic (AIH) approach on the Toronto Benchmark dataset

	hec92 II	yor83 II	ear83 II	sta83 II	uta92 II
AIH Average	12.43	50.49	41.98	35.00	3.54
AIH Best	11.35	49.72	39.60	32.57	3.45
Time taken in first stage (s)	352	861	1517	1204	53966
Time taken in second stage (s)	146	513	1275	684	27350
Number of exams causing violations	23	19	30	17	53

**Table 10** Best results obtained by the Adaptive Improvement Hyper-heuristic (Algorithm 3) compared to the best approaches in the literature on the Toronto Benchmark

Problems	AIH Best	Abdullah et al. (2009) Best	Burke and Bykov (2008) Best	Burke et al. (2010) Best	Caramia et al. (2008) Best
hec92 I	11.19	9.73	10.06	10.00	<b>9.20</b>
sta83 I	157.18	156.94	157.03	<b>156.90</b>	158.20
yor83 I	39.47	34.95	<b>34.78</b>	34.90	36.20
ute92	26.70	24.90	24.79	24.80	<b>24.40</b>
ear83 I	35.79	36.00	32.65	32.80	<b>29.30</b>
tre92	8.49	8.5	<b>7.72</b>	7.90	9.40
lse91	10.92	10.03	9.86	10.00	<b>9.60</b>
kfu93	14.51	<b>12.62</b>	12.81	13.00	13.80
car92	4.31	<b>3.76</b>	3.81	3.90	6.00
uta92 I	3.44	<b>2.99</b>	3.16	3.20	3.50
car91	5.19	<b>4.42</b>	4.58	4.60	6.60

In comparison with other hyper-heuristic techniques, our approach was able to produce better results in only one instance when compared to the evolutionary algorithm based hyper-heuristic presented in Pillay (2010). However, it was stated that these experiments did not adhere to the time limitation imposed by the competition.

In comparison to the Constraint Based Solver developed in Atsuta et al. (2008), our approach performed better in 3 out of the 8 instances. The approach using the Drools solver in De Smet (2008) obtained feasibility for only 5 instances. Our approach outperformed it as we were able to gain feasibility for all the 8 instances. This demonstrates the generality of our approach to solving exam timetabling problems. Finally, our approach performed better on 6 of the 8 instances in comparison with the biologically inspired approach proposed in Pillay (2008).

## 5 Conclusions

The study presented in this paper implements a hyper-heuristic approach which adaptively adjusts heuristic combinations to achieve the best improvement on constructed timetables. An investigation is made of the low-level heuristics used and the order in which exams causing soft constraint violations are rescheduled. The analysis is performed on a set of four

**Table 11** Best results obtained by the Adaptive Improvement Hyper-heuristic (AIH) compared to other hyper-heuristics approaches in the literature on the Toronto Benchmark

Problems	AIH Best	Kendall and Mohd Hussin (2005) Best	Burke et al. (2007) Best	Pillay and Banzhaf (2009) Best	Qu and Burke (2009) Best
hec92 I	<b>11.19</b>	11.86	12.72	11.85	11.94
sta83 I	<b>157.18</b>	157.38	158.19	158.33	159.00
yor83 I	<b>39.47</b>	–	40.13	40.74	40.24
ute92	<b>26.70</b>	27.60	31.65	28.88	28.30
ear83 I	<b>35.79</b>	40.18	38.19	36.86	35.86
tre92	8.49	<b>8.39</b>	8.85	8.48	8.60
lse91	<b>10.92</b>	–	13.15	11.14	11.15
kfu93	<b>14.51</b>	15.84	15.76	14.62	14.79
car92	4.31	4.67	4.84	4.28	<b>4.16</b>
uta92 I	3.44	–	3.88	<b>3.40</b>	3.42
car91	5.19	5.37	5.41	<b>4.97</b>	5.16

**Table 12** Best results obtained by the Adaptive Improvement Hyper-heuristic (AIH) compared to the best approaches in the literature on the ITC2007 dataset

Instances	AIH Best	McCollum et al. (2009) Best	Muller (2008) Best	Gogos et al. (2008) Best	Atsuta et al. (2008) Best	De Smet (2008) Best	Pillay (2008) Best	Pillay (2010) Best
Exam 1	6235	4633	<b>4370</b>	5905	8006	6670	12035	8559
Exam 2	2974	405	<b>400</b>	1008	3470	623	3074	830
Exam 3	15832	<b>9064</b>	10049	13862	18622	–	15917	11576
Exam 4	35106	<b>15663</b>	18141	18674	22559	–	23582	21901
Exam 5	4873	3042	<b>2988</b>	4139	4714	3847	6860	3969
Exam 6	31756	<b>25880</b>	26950	27640	29155	27815	32250	28340
Exam 7	11562	<b>4037</b>	4213	6683	10473	5420	17666	8167
Exam 8	20994	<b>7461</b>	7861	10521	14317	–	16184	12658

benchmark instances of differing difficulty in an off-line learning process. It is shown that, of the heuristics tried, the timeslot swapping heuristic is the best to combine with Kempe chains. In addition, better solutions are produced when ordering the exams causing a soft constraint violation using Saturation Degree and breaking any ties with Largest Weighted Degree. Based on the output of the learning process, an adaptive approach which analyses and adjusts some randomly generated sequences is implemented and applied to the rest of the instances. Furthermore, the hyper-heuristic approach is applied to a more constrained dataset, and showed to produce very competitive results compared to other approaches in the literature on both datasets.

Future research directions include performing improvements during the timetable construction stage instead of performing the improvements at the end of the construction. Using hybridisations of more than two low-level heuristics could also be investigated. Finally, the approach investigated in this paper can be applied to course timetabling.

## References

- Abdullah, S., Turabieh, H., & McCollum, B. (2009). A hybridization of electromagnetic-like mechanism and great deluge for examination timetabling problems. In *Lecture notes in computer science: Vol. 5818. Proceedings of the 6th international workshop on hybrid metaheuristics*, HM2009 (pp. 60–72). Berlin: Springer.
- Asmuni, H., Burke, E. K., Garibaldi, J., & McCollum, B. (2004). Fuzzy multiple ordering criteria for examination timetabling. In E. K. Burke & M. Trick (Eds.), *Lecture notes in computer science: Vol. 3616. Selected papers from the 5th international conference on the practice and theory of automated timetabling* (pp. 334–353). Berlin: Springer.
- Asmuni, H., Burke, E. K., Garibaldi, J., McCollum, B., & Parkes, A. J. (2009). An investigation of fuzzy multiple heuristic orderings in the construction of university examination timetables. *Computers & Operations Research*, *36*(4), 981–1001.
- Atsuta, M., Nonobe, K., & Ibaraki, T. (2008). ITC2007 track 1: an approach using general CSP solver. In *Practice and theory of automated timetabling*, PATAT 2008, August 2008 (pp. 19–22).
- Biligan, B., Ozcan, E., & Korkmaz, E. E. (2007). An experimental study on hyper-heuristics and exam timetabling. In E. Burke & H. Rudova (Eds.), *Lecture notes in computer science: Vol. 3867. Practice and theory of automated timetabling VI: selected papers from the 6th international conference*, PATAT 2006 (pp. 394–412).
- Burke, E. K., & Bykov, Y. (2008). A late acceptance strategy in hill-climbing for examination timetabling problems. In *Practice and theory of automated timetabling*, PATAT 2008, August 2008.
- Burke, E. K., Kendall, G., Newall, J., Hart, E., Ross, P., & Schulenburg, S. (2003). Hyper-heuristics: an emerging direction in modern search technology. In F. Glover & G. Kochenberger (Eds.), *Handbook of meta-heuristics* (pp. 457–474). Norwell: Kluwer Academic.
- Burke, E. K., McCollum, B., Meisels, A., Petrovic, S., & Qu, R. (2007). A graph-based hyper-heuristic for educational timetabling problems. *European Journal of Operational Research*, *176*, 177–192.
- Burke, E. K., Eckersley, A., McCollum, B., Petrovic, S., & Qu, R. (2010). Hybrid variable neighbourhood approaches to university exam timetabling. *European Journal of Operational Research*, *206*, 46–53.
- Burke, E. K., Qu, R., & Soghier, A. (2011). An adaptive tie breaking and hybridisation hyper-heuristic for exam timetabling problems. In D.A. Pelta et al. (Eds.), *Nature inspired cooperative strategies for optimization*, NISCO 2011 (pp. 205–223)
- Caramia, M., Dell Olmo, P., & Italiano, G. F. (2008). Novel local-search-based approaches to university examination timetabling. *INFORMS Journal on Computing*, *20*(1), 86–99.
- Carter, M. W., Laporte, G., & Lee, S. Y. (1996). Examination timetabling: algorithmic strategies and applications. *Journal of the Operational Research Society*, *74*, 373–383.
- De Smet, G. (2008). ITC2007—examination track. In *Practice and theory of automated timetabling*, PATAT 2008, August 2008 (pp. 19–22).
- Ersoy, E., Ozcan, E., & Uyar, S. (2007). Memetic algorithms and hill-climbers. In P. Baptiste, G. Kendall, A. M. Kordon, & F. Sourd (Eds.), *Proceedings of the 3rd multidisciplinary international conference on scheduling: theory and applications conference*, MISTA 2007 (pp. 159–166).
- Gogos, C., Alefragis, P., & Housos, E. (2008). A multi-staged algorithmic process for the solution of the examination timetabling problem. In *Practice and theory of automated timetabling*, PATAT 2008 (pp. 19–22).
- Kendall, G., & Mohd Hussin, N. (2005). An investigation of a tabu search based hyper-heuristic for examination timetabling. In G. Kendall, E. Burke, S. Petrovic, & M. Gendreau (Eds.), *Selected papers from MISTA 2005* (pp. 309–328). Berlin: Springer.
- McCollum, B., McMullan, P., Parkes, A. J., Burke, E. K., & Abdullah, S. (2009). An extended great deluge approach to the examination timetabling problem. In *Proceedings of the 4th multidisciplinary international scheduling: theory and applications 2009*, MISTA 2009, Dublin, Ireland, 10–12 August (pp. 424–434).
- McCollum, B., Schaerf, A., Paechter, B., McMullan, P., Lewis, R., Di Gaspero, L., Parkes, A. J., Qu, R., & Burke, E. K. (2010). Setting the research agenda in automated timetabling: the second international timetabling competition. *INFORMS Journal of Computing*, *22*(1), 120–130.
- Muller, T. (2008). ITC 2007 solver description: a hybrid approach. In *Practice and theory of automated timetabling*, PATAT 2008, August 2008 (pp. 19–22).
- Pillay, N. (2008). A developmental approach to the examination timetabling problem. In *Practice and theory of automated timetabling*, PATAT 2008, August 2008 (pp. 19–22).
- Pillay, N. (2010). Evolving hyper-heuristics for a highly constrained examination timetabling problem. In *Practice and theory of automated timetabling*, PATAT'10 (pp. 336–346).
- Pillay, N., & Banzhaf, W. (2007). A genetic programming approach to the generation of hyper-heuristic systems for the uncapacitated examination timetabling problem. In J. Neves, et al. (Eds.), *Lecture notes in artificial intelligence: Vol. 4874. Progress in artificial intelligence* (pp. 223–234).



- Pillay, N., & Banzhaf, W. (2009). A study of heuristic combinations for hyper-heuristic systems for the uncapacitated examination timetabling problem. *European Journal of Operational Research*, 197, 482–491.
- Qu, R., & Burke, E. K. (2009). Hybridisations within a graph-based hyper-heuristic framework for university timetabling problems. *Journal of the Operational Research Society*, 60, 1273–1285.
- Qu, R., Burke, E. K., & McCollum, B. (2009a). Adaptive automated construction of hybrid heuristics for exam timetabling and graph colouring problems. *European Journal of Operational Research*, 198(2), 392–404.
- Qu, R., Burke, E. K., McCollum, B., Merlot, L. T. G., & Lee, S. Y. (2009b). A survey of search methodologies and automated approaches for examination timetabling. *Journal of Scheduling*, 12(1), 55–89.
- Thompson, J. M., & Dowsland, K. A. (1996). Variants of simulated annealing for the examination timetabling problem. *Annals of Operations Research*, 63, 105–128.