

The Third International Timetabling Competition

Gerhard Post · Luca Di Gaspero · Jeffrey H. Kingston ·
Barry McCollum · Andrea Schaerf

© Springer Science+Business Media New York 2013

Abstract This paper is the organizers' report on the Third International Timetabling Competition (ITC2011), run during 2011–2012 with the aim of raising the profile of automated high school timetabling. Its participants tackled 35 instances of the high school timetabling problem, taken from schools in 10 countries. The paper describes the data model used, the competition, and the results.

Keywords High school timetabling · International timetabling competition

1 Introduction

High school timetabling is a long-established area of timetabling, but it has received less scientific attention than some other areas, for example examination timetabling. Its research community has been fragmented, with virtually no sharing of data.

Over the last few years, a group of researchers has tackled these problems. There is now an XML data format in which unsimplified instances and solutions can be specified precisely (Post et al. 2011, 2012). A web site is available for evaluating solutions (Kingston 2010). Instances have been widely sought. At the time of writing, 35 instances from 10 countries are available for download (Post 2011).

Building on this work and on two previous competitions (Paechter et al. 2002; McCollum et al. 2010), and supported financially by the PATAT conference series, EventMAP, and the Centre for Telematics and Information Technology at the University of Twente in the Netherlands, the Third International Timetabling Competition (ITC2011) was run by the authors in 2011–2012, with the aim of raising the profile of automated high school timetabling.

This paper is the organizers' report on that competition. It describes the data model used, the competition itself, and the results.

J.H. Kingston (✉)
School of Information Technologies, The University of Sydney, Sydney, Australia
e-mail: jeff@it.usyd.edu.au

2 Modelling high school timetabling

The competition used a data format called XHSTT in which high school timetabling problems and solutions can be expressed. This format has been described previously (Post et al. 2011, 2012), and a full specification is available online (Kingston 2010). An overview, omitting syntactic details, is given here for completeness.

An XHSTT file is an XML file containing one *archive*, which consists of a set of instances of the high school timetabling problem, plus any number of *solution groups*. A solution group is a set of solutions to some or all of the archive's instances, typically produced by one solver. There may be several solutions to one instance in one solution group, for example solutions produced using different random seeds.

Each instance has four parts. The first part defines the instance's *times*, that is, the individual intervals of time, of unknown duration, during which events run. Taken in chronological order these times form a sequence called the instance's *cycle*, which is usually one week. Arbitrary sets of times, called *time groups*, may be defined, such as the Monday times or the afternoon times. A *day* is a time group holding the times of one day, and a *week* is a time group holding the times of one week. For the convenience of display software, some time groups may be labelled as days or weeks.

The second part defines the instance's *resources*: the entities that attend events. Resources are partitioned into *resource types*. The usual resource types are a Teacher type whose resources represent teachers, a Room type of rooms, a Class type of *classes* (sets of students who attend the same events), and a Student type of individual students. However, an instance may define any number of resource types. Arbitrary sets of resources of the same type, called *resource groups*, may be defined, such as the set of Science laboratories, or the set of senior classes.

The third part defines the instance's *events*: meetings between resources. An event has a *duration* (a positive integer), a *time*, and any number of *resources* (sometimes called *event resources*). The meaning is that the resources are occupied attending the meeting for *duration* consecutive times starting at *time*. The duration is a fixed constant. The time may be preassigned or left open to the solver to assign. Each resource may also be preassigned or left open to the solver to assign, although the type of resource to assign is fixed. Arbitrary sets of events, called *event groups*, may be defined. A *course* is an event group representing the events in which a particular class studies a particular subject. Some event groups may be labelled as courses.

For example, suppose class 7A meets teacher *Smith* in a Science laboratory for two consecutive times. This is represented by one event with duration 2 containing three resources: one preassigned Classes resource 7A, one preassigned Teachers resource *Smith*, and one open Rooms resource. Later, a constraint will specify that this room should be selected from the *ScienceLaboratories* resource group, and define the penalty imposed on solutions that do not satisfy that constraint.

If class 7A meets for Science several times each week, several events would be created and placed in an event group labelled as a course. However, it is common in high school timetabling for the total duration of the events of a course to be fixed, but for the way in which that duration is broken into events to be flexible. For example, class 7A might need to meet for Science for a total duration of 6 times per week, in events of duration 1 or 2, with at least one event of duration 2 during which the students carry out experiments. One acceptable outcome would be five *sub-events*, as these fragments are called, of durations 2, 1, 1, 1, and 1; another would be three sub-events, of durations 2, 2, and 2. This is modelled by giving a single event of duration 6. Later, constraints specify how this event may be

Table 1 The 15 types of constraints, with informal explanations of their meaning

Name	Meaning
Assign Resource constraint	Event resource should be assigned a resource
Assign Time constraint	Event should be assigned a time
Split Events constraint	Event should split into a constrained number of sub-events
Distribute Split Events constraint	Event should split into sub-events of constrained durations
Prefer Resources constraint	Event resource assignment should come from resource group
Prefer Times constraint	Event time assignment should come from time group
Avoid Split Assignments constraint	Set of event resources should be assigned the same resource
Spread Events constraint	Set of events should be spread evenly through the cycle
Link Events constraint	Set of events should be assigned the same time
Avoid Clashes constraint	Resource's timetable should not have clashes
Avoid Unavailable Times constraint	Resource should not be busy at unavailable times
Limit Idle Times constraint	Resource's timetable should not have idle times
Cluster Busy Times constraint	Resource should be busy on a limited number of days
Limit Busy Times constraint	Resource should be busy a limited number of times each day
Limit Workload constraint	Resource's total workload should be limited

split into sub-events, and define the penalty imposed on solutions that do not satisfy those constraints.

Finally, an instance contains any number of *constraints*: conditions that an ideal solution would satisfy. There are 15 types of constraints, stating that events should be assigned times, prohibiting clashes, etc. (Table 1). Each has its own specific attributes. For example, a Prefer Times constraint lists the events whose time it constrains, and the preferred times for those events. Each also has attributes common to all, including a Boolean value saying whether the constraint is hard or soft, and an integer weight.

The *infeasibility value* of a solution is the sum over the hard constraints of the number of violations of the constraint times its weight. The *objective value* is similar, but using the soft constraints. One solution is better than another if it has a smaller infeasibility value, or an equal infeasibility value and a smaller objective value.

As mentioned earlier, solutions are stored separately from instances, in solution groups within the archive file. A solution is a list of sub-events, each containing a duration, a time assignment, and some resource assignments. The HSEval web site (Kingston 2010) calculates the infeasibility and objective values of the solutions of an archive, and displays comparative tables, lists of violations, and so on.

3 The competition

The competition attracted 17 registrations, 8 of whom are known to have worked for a significant amount of time, although only 5 submitted solutions in the end. This is many fewer than the over 100 registrations and about 40 active teams of ITC2007, even allowing for the fact that ITC2007 had three tracks. Why fewer teams registered is not known, but it could be due to high school timetabling's lower profile, or to the complexity of the format, which was remarked upon by several participants.

The competition had three independent parts, called Rounds 1, 2, and 3. In Round 1, participants were invited to submit solutions to 21 published instances. No restrictions were

Table 2 The source country, instance name, number of times (Ti), teachers (Te), rooms (Ro), classes (groups of students, Cl), individual students (St), and events (Ev), of the 21 instances of Round 1 (archive XHSTT-2012.xml). The last column shows the results of Round 1: where a new best solution was found, the team that found it and its cost, either a single soft cost or an ordered pair of hard cost and soft cost

Country	Instance	Ti	Te	Ro	Cl	St	Ev	New best
Australia	BGHS98	40	56	45	30		387	HySST (3, 494)
Australia	SAHS96	60	43	36	20		296	HySST (8, 52)
Australia	TES99	30	37	26	13		308	HySST (1, 140)
Brazil	Instance1	25	8		3		21	VAGOS (11)
Brazil	Instance4	25	23		12		127	VAGOS (44)
Brazil	Instance5	25	31		13		119	VAGOS (43)
Brazil	Instance6	25	30		14		140	VAGOS (77)
Brazil	Instance7	25	33		20		205	VAGOS (122)
UK	StPaul	27	68	67	67		1227	Lectio (136)
Finland	Artificial	20	22	12	13		169	
Finland	College	40	46	34	31		387	
Finland	HighSchool	35	18	13	10		172	
Finland	Secondary	35	25	25	14		280	Lectio (88)
Greece	HighSchool1	35	29		66		372	
Greece	Patras2010	35	29		84		178	
Greece	Preveza2008	35	29		68		164	
Italy	Instance1	36	13		3		42	VAGOS (12)
Netherlands	GEPRO	44	132	80	44	846	2675	GOAL (1, 382)
Netherlands	Kottenpark2003	38	75	41	18	453	1156	Lectio (532)
Netherlands	Kottenpark2005	37	78	42	26	498	1235	Lectio (533)
South Africa	Lewitt2009	148	19	2	16		185	VAGOS (0)

placed on how the solutions could be obtained. For each instance, a small prize was awarded to the participant who submitted the best solution to that instance, if it improved on the best solution previously known to the organizers.

The 21 published instances are listed in Table 2. The table shows that they differ greatly in size. For example, the number of times varies between 25 and 148, and the number of classes (groups of students) varies between 3 and 84. They also differ in structure: some require most events to be split into sub-events, others do not; some require teachers to be assigned as well as rooms, although most do not; and so on.

Round 2 compared solvers under uniform conditions. Solvers were allowed to use only freely available software libraries, and a time limit was imposed: 1000 seconds of single-threaded execution per instance on the organizers' computer. Participants used a benchmark program to estimate how much time on their own computer was equivalent to the time limit on the organizers' computer.

For initial testing, the Round 1 instances were used, except that the Australian instances from Table 2 were omitted. This was because these are the only instances so far to use Avoid Split Assignments and Limit Workload constraints, and omitting them reduced the implementation burden for the participants.

The original plan was to select up to 5 finalists based on this initial testing, but since there were only 5 active teams, all were invited to become finalists. One team declined owing to problems with their solver. This left 4 finalists, who submitted their solvers, which

Table 3 Summary of results of Round 2. For each finalist, the average over 10 runs with different random seeds of the ranks of the solutions produced. The entries marked ‘–’ were excluded because the authors of the solver supplied the instances. The asterisk indicates failure to produce valid solutions

Instance	GOAL	HySST	Lectio	HFT
BrazilInstance2	–	2.0	1.0	3.0
BrazilInstance3	–	1.0	2.0	3.0
BrazilInstance4	–	1.9	1.1	3.0
BrazilInstance6	–	2.0	1.0	3.0
FinlandElementarySchool	1.9	3.0	1.1	4.0
FinlandSecondarySchool2	1.0	2.4	2.6	4.0
Aigio 1st High School 2010–2011	1.0	2.8	2.2	4.0
Italy Instance4	1.1	3.0	1.9	4.0
KosovaInstance1	1.0	4.0	2.0	3.0
Kottenpark2003	1.4	1.6	3.0	4.0
Kottenpark2005A	1.4	1.6	3.4	3.6
Kottenpark2008	1.0	2.1	2.9	4.0*
Kottenpark2009	1.0	2.0	3.0	4.0
Woodlands2009	1.7	2.8	1.5	4.0
Spanish school	1.0	2.0	3.0	4.0
WesternGreeceUniversityInstance3	1.0	2.0	4.0	3.0
WesternGreeceUniversityInstance4	1.0	2.0	3.0	4.0
WesternGreeceUniversityInstance5	1.0	2.0	3.0	4.0
Average rank	1.18	2.23	2.32	3.64

the organizers ran on their own computer. For each of 18 hidden (unpublished) instances, and for each of 10 random seeds, each solver was run on that instance and seed for at most 1000 seconds, and the solutions for that instance and seed were ranked. If a solver failed to produce a solution for any instance and seed, it received a high rank for that instance and seed; and in the few cases where the authors of a solver were also the creators of an instance, their results on that instance were not ranked. (No organizers were participants.) Then each solver’s ranks were averaged, and the solver with the lowest average rank was declared the winner.

The 18 hidden instances were 14 new ones plus 4 small corrections or variants of published ones, and were not very different from the published ones. The Round 2 finalists have described their solvers in invited papers for the PATAT 2012 conference, written before results were announced (Domrös and Homberger 2012; Fonseca et al. 2012; Kheiri et al. 2012; Sørensen et al. 2012). Two other related papers appear in the PATAT 2012 proceedings (Sørensen and Stidsen 2012; Valouxis et al. 2012).

After Round 2, the hidden instances were published and Round 3 was conducted on them. As for Round 1, no restrictions were placed on how the solutions could be obtained. The participant with the lowest average rank over the hidden instances was declared the winner, under the same detailed rules as in Round 2.

After the competition ended, the solutions were published on the competition web site, as XHSTT archive files. The results of Rounds 2 and 3 were calculated by passing these files to the appropriate HSEval operation, so can be publicly verified. Unfortunately the public version of HSEval cannot evaluate the large archive holding the Round 2 results in

Table 4 Results of Round 3: for each finalist, the rank of the best solution supplied for each instance. The entries marked ‘—’ were excluded because the authors of the solver supplied the instances. Blank entries indicate that no solution was supplied; maximal ranks were assigned in those cases

Instance	GOAL	HySST	Lectio	HFT	VAGOS
BrazilInstance2	—	3.0	1.0	4.0	2.0
BrazilInstance3	—	3.0	2.0	4.0	1.0
BrazilInstance4	—	3.0	2.0	4.0	1.0
BrazilInstance6	—	3.0	1.0	4.0	2.0
FinlandElementarySchool	2.5	2.5	2.5	2.5	
FinlandSecondarySchool2	2.0	2.0	2.0	4.0	
Aigio 1st High School 2010–2011	1.0	3.0	2.0	4.0	
Italy Instance4	2.0	1.0	3.0	4.0	
KosovaInstance1	1.0	2.0	4.0	3.0	
Kottenpark2003	3.0	2.0	1.0	4.0	
Kottenpark2005A	2.0	1.0	4.0	3.0	
Kottenpark2008	1.0	2.0	3.0	4.0	
Kottenpark2009	2.0	1.0	3.0	4.0	
Woodlands2009	1.0	2.0	3.0	4.0	
Spanish school	1.0	3.0	2.0	4.0	
WesternGreeceUniversityInstance3	1.5	3.0	5.0	4.0	1.5
WesternGreeceUniversityInstance4	1.0	2.0	4.0	3.0	
WesternGreeceUniversityInstance5	2.0	2.0	5.0	4.0	2.0
Average rank	1.64	2.25	2.75	3.75	3.86

a reasonable time. This archive was evaluated using an identical private copy of HSEval running on a faster computer.

4 Results

In Round 1, new best solutions were found for 15 of the 21 instances (Table 2). In most cases the improvement was significant. Four teams received prizes in this round.

A summary of the results of Round 2 appears in Table 3. The clear winner here was the GOAL team (Fonseca et al. 2012), with an average rank of 1.18. Their solver produced the best solutions to most of the tests, and was never worse than equal second on any test. GOAL would also have won if the excluded tests had been included.

A summary of the results of Round 3 appears in Table 4. Here the results were closer, and as for Round 1 new best solutions were found by four teams, but once again GOAL produced the best average result.

5 Conclusion

The competition’s aim was to raise the profile of automated high school timetabling. This it has undoubtedly done.

Judging by previous competitions, the ITC2011 instances are likely to continue to attract researchers for years to come. These are unsimplified instances from real high schools around the world—a great foundation to build on. The authors will shortly be placing a selection of these instances on the project web site (Post 2011) under archive name XHSTT-2013, together with a suggested experimental protocol. Hopefully, a new era of research into automated high school timetabling will then begin.

References

- Domrös, J., & Homberger, J. (2012). An evolutionary algorithm for high school timetabling. In *Proceedings of the ninth international conference on the practice and theory of automated timetabling*, PATAT 2012, Son, Norway, August 2012.
- Fonseca, G. H. G., Santos, H. G., Toffolo, T. A. M., Brito, S. S., & Souza, M. J. F. (2012). A SA-ILS approach for the high school timetabling problem. In *Proceedings of the ninth international conference on the practice and theory of automated timetabling*, PATAT 2012, Son, Norway, August 2012.
- Kheiri, A., Özcan, E., & Parkes, A. J. (2012). HySST: hyper-heuristic search strategies and timetabling. In *Proceedings of the ninth international conference on the practice and theory of automated timetabling*, PATAT 2012, Son, Norway, August 2012.
- Kingston, J. H. (2010). The HSEval high school timetable evaluator. <http://www.it.usyd.edu.au/~jeff/hseval.cgi>.
- McCollum, B., Schaerf, A., Paechter, B., McMullan, P., Lewis, R., Parkes, A., Di Gaspero, L., Qu, R., & Burke, E. (2010). Setting the research agenda in automated timetabling: the second international timetabling competition. *INFORMS Journal on Computing*, 22, 120–130.
- Paechter, B., Gambardella, L. M., & Rossi-Doria, O. (2002). The first international timetabling competition. <http://www.idsia.ch/Files/ttcomp2002/>.
- Post, G. (2011). Benchmarking project for high school timetabling. <http://www.utwente.nl/ctit/hstt/>.
- Post, G., Kingston, J. H., Ahmadi, S., Daskalaki, S., Gogos, C., Kyngäs, J., Nurmi, C., Musliu, N., Pillay, N., Santos, H., & Schaerf, A. (2011). XHSTT: an XML archive for high school timetabling problems in different countries. *Annals of Operations Research*. doi:10.1007/s10479-011-1012-2.
- Post, G., Ahmadi, S., Daskalaki, S., Kingston, J. H., Kyngäs, J., Nurmi, C., & Ranson, D. (2012). An XML format for benchmarks in high school timetabling. *Annals of Operations Research*, 194, 385–397.
- Sørensen, M., & Stidsen, T. R. (2012). High school timetabling: modeling and solving a large number of cases in Denmark. In *Proceedings of the ninth international conference on the practice and theory of automated timetabling*, PATAT 2012, Son, Norway, August 2012.
- Sørensen, M., Kristiansen, S., & Stidsen, T. R. (2012). International timetabling competition 2011: an adaptive large neighborhood search algorithm. In *Proceedings of the ninth international conference on the practice and theory of automated timetabling*, PATAT 2012, Son, Norway, August 2012.
- Valouxis, C., Gogos, C., Alefragis, P., & Housos, E. (2012). Decomposing the high school timetable problem. In *Proceedings of the ninth international conference on the practice and theory of automated timetabling*, PATAT 2012, Son, Norway, August 2012.