
An XML Format for Benchmarks in High School Timetabling II

Gerhard Post · Jeffrey H. Kingston · Samad Ahmadi · Sophia Daskalaki · Christos Gogos · Jari Kyngas · Cimmo Nurmi · Haroldo Santos · Ben Rorije · Andrea Schaerf

Abstract We present the progress on the benchmarking project for high school timetabling that was introduced at PATAT 2008. In particular, we announce the High School Timetabling Archive HSTT2010 with 15 instances from 7 countries and an evaluator capable of checking the syntax of instances and evaluating the solutions.

Keywords. Timetabling, high school, benchmark, XML, scheduling.

1 Introduction

“It is surprising that no standard format for exchanging datasets in the field of high school timetabling has emerged until now.” This sentence was the motivation for a

Gerhard Post

University of Twente, P.O. Box 217, 7500 AE Enschede, The Netherlands and
ORTEC, Groningenweg 6k, 2803 PV Gouda, The Netherlands

Jeffrey H. Kingston

School of Information Technologies, The University of Sydney, Australia

Samad Ahmadi

Dept. of Informatics, De Montfort University, The Gateway, LE1 9BH, UK Leicester

Sophia Daskalaki

Engineering Sciences Department, University of Patras, 26500 Rio Patras, Greece

Christos Gogos

Engineering Sciences Department, University of Patras, 26500 Rio Patras, Greece

Jari Kyngas

Satakunta University of Applied Sciences, Tiedepuisto 3, 28600 Pori, Finland

Cimmo Nurmi

Satakunta University of Applied Sciences, Tiedepuisto 3, 28600 Pori, Finland

Haroldo Santos

Computing Department, Universidade Federal de Ouro Preto, Brazil

Ben Rorije

University of Twente, P.O. Box 217, 7500 AE Enschede, The Netherlands

Andrea Schaerf

DIEGM, University of Udine, via delle Scienze 206, 33100 Udine, Italy

group of researchers to define a format capable of expressing high school timetabling instances from all over the world, see (Post 2010).

The high school organization is different around the world, consequently the problems in high school timetabling that arise from real cases in various countries differ as well. As a specific example, one of the main differences that has emerged from our research is related to allowing idle times for students during school hours versus the cases where this is not allowed. In the first case teachers are usually not preassigned to the lessons, as this may lead to infeasibilities. In the second case, teachers are mostly preassigned, leading to the problem of eliminating idle times for students and minimizing them for teachers.

Another important difference is related to the granularity of the scheduling process: sometimes it is performed at the level of an entire class, whereas in other cases of a single student. In the latter case the problem usually becomes harder, since the schedule of each individual student has to be evaluated during the solving process, thus making the process computationally more expensive.

The purpose of this paper is to report on the progress of this collaborative research and reflect on the current situation; we will give a short overview and motivation of the XML format in Section 2, discuss the current archive in Section 3, the evaluator in Section 4, and give an outlook to the future in Section 5.

2 The format

Differences in the organization of high schools in different countries imply that the definition of a unified format for high school timetabling is not a trivial task. The current format has emerged after many iterations; indeed the format discussed in (Post 2010) differs considerably from the original version presented at PATAT 2008.

The format of our benchmark is mapped out by an XML schema which defines the compulsory and optional elements that need to be present in the XML files holding the instances and solutions. The basic elements of an instance are the *times*, *resources*, and *events*, complemented by the *constraints*, which are imposed on them. We believe that the structure of the format as it is now will essentially remain the same over time. The reason lies in the principal choice to embed the “business logic” in the constraints, and not in the basic elements. At present, the format contains (only) 15 constraint types, including “obvious” ones, like *AssignTimeConstraint* (assign a start time to selected events) and *AvoidClashesConstraint* (a resource may be involved with at most one event at a time). The modular nature of the schema assures that new constraints can be added without having to change its structure. Indeed, we believe that the set of constraints will probably be extended further to incorporate new specialized constraints to deal with unforeseen problems in other countries.

A fragment of an instance file is shown in Figure 1. All objects have the attribute *Id* for referencing, and the child *Name* for displaying. *Times*, *Resources* and *Events* can be grouped in *TimeGroups*, *ResourceGroups* and *EventGroups*, respectively. For demonstration purposes the *Times* section is expanded with some more detail.

One of the main discussions during the design of this format was about the “domain specific” structure and the “solver needed” structure. The “domain specific” structure reflects how a timetabler at a school has structured the data; for example, a timetabler will distinguish days of the week, will think in terms of students, school classes or teachers (certainly not of general resources), and will consider courses and subjects

```

<Instance Id="Example">
  <Times>
    <TimeGroups>
      <Day Id="Day1"/> <Name>Monday</Name> </Day>
      ...
      <Day Id="Day5"/> <Name>Friday</Name> </Day>
      <TimeGroup Id="AllTimes"/> <Name>AllTimes</Name> </TimeGroup>
      ...
    </TimeGroups>
    <Time Id="Day1_1"> <Name>Monday 1</Name>
      <TimeGroups>
        <Day Reference="Day1"/>
        <TimeGroup Reference="AllTimes"/>
      </TimeGroups>
    </Time>
    ...
  </Times>
  <Resources>
    ...
  </Resources>
  <Events>
    ...
  </Events>
  <Constraints>
    ...
  </Constraints>
</Instance>

```

Fig. 1 A problem instance in the XML format

as principal objects for scheduling. The solver instead requires a structure organized in terms of variables which represent units of lessons or resources, while conceptual entities are not important.

In the format, although some of the domain specific structural elements (for example “days” and “courses”) are supported and their use is recommended, it is not obligatory to use them. We mention the two most important ones. A *Day* is introduced as a time group, that almost certainly will be needed, because in high school timetabling there are many constraints at a daily level. For example, constraints about “working hours”, “idle times”, and “number of days present”. In addition, by introducing days we are able to display daily schedules for the resources. Another element is the *Course*, which is introduced as an *EventGroup* for a subject and a student group combination. This is important in order to control the spreading of the individual lessons (events) of a course over the week days. *Courses* also allow control on events with similar properties; if certain events are identical, they can be clustered to one event and allow the lessons to be sub-events. Moreover, constraints like the *SplitEventsConstraint* can prescribe how such an event should be split into sub-events.

Our view on inclusion of new constraints has changed during the past two years. Originally we tried to include all the constraints that we encountered in the literature, or that we could imagine to be useful. On the contrary, the current set of constraints in the format reflect only the constraints needed by the contributors, and no more. The reason for this is that complicated constraints usually need to be clarified by an expert.

3 The archive HSTT2010

At the website (Post 2008) the archive HSTT2010 with 15 instances from 7 different countries is available. These instances have appeared previously in the literature, but were not available for download. Apart from the instances, solutions are also available. We intend to keep record of the best found solutions for researchers to be able to validate their solvers. In Table 1 we present the instances that have been contributed. In the columns are given

- the country (Country);
- name of the instance (Name);
- total duration of all events (EvD);
- the number of teachers (T);
- the number of school classes (SC);
- the number of students (St);
- the number of rooms (Ro).

Country	Name	EvD	T	SC	St	Ro
Australia	BGHS98	1564	56	30		45
	SAHS6	1876	43	20		36
	TES99	806	37	13		26
Brazil	Instance 1	75	8	3		
Brazil	Instance 4	300	23	12		
England	St Paul	1227	68	67		67
Finland	Artificial	200	22	13		12
	College	854	46	31		33
	High school	319	18	102		13
	Secondary school	306	25	14		25
Greece	High School 1	372	29	66		
Italy	Instance 1	133	13	3		
Netherlands	GEPRO	2675	132	44	846	80
	Kottenpark 2003	1203	75	18	453	41
	Kottenpark 2005	1272	78	26	498	42

Table 1 The instances in the archive HSTT2010.

Note that the instances vary significantly in size. Most instances are described at the level of school classes, which might split further to form sub groups. The Dutch instances, however, carry information at the level of individual students as well. For the lower grades the groups of students (school classes) are fixed and all students of a group attend most lessons together. Conversely, for the higher grades the timetable of each student is mostly personal, since the compulsory lessons constitute only $\frac{1}{3}$ of the lessons. For the Australian instances, the teachers have to be assigned as well in the timetabling process and in such case split assignments should be avoided.

4 The evaluator

Apart from the instances, the format also models solutions. A solution is presented by describing the duration of all events (as mentioned previously it is possible to define

a “course” event of duration 3, which can be split into, for example, three sub-events of duration 1), the time slots assigned to each event, and (in some cases) the resources assigned to events. Once a solution is provided we can then evaluate it. The evaluation leads to two integers: the infeasibility value (i.e. the total cost of the hard constraints violations) and the objective value (the total cost of the soft constraints). The total cost is generated from two different constraint types: if the constraint is “hard”, then the cost is added to the infeasibility value, otherwise to the objective value. Depending on the type of the constraint, the cost is attributed to: an event (for example: “is there a time assigned to the event?”), to an event group (for example: “is the course well-spread?”) or to a resource (for example: “are the idle times within the given limits?”). The cost of the schedule is the sum of all separate costs. The cost value V is then calculated from the deviation D of the constraint C with weight λ by the formula:

$$V = \lambda \cdot f_C(D)$$

Here f_C is a cost function and its type is specified in the constraint. The cost functions supported currently are: step function, linear function, and quadratic function.

The format supports multiple instances and multiple groups of solutions to these instances:

```
<HighSchoolTimetableArchive>
  <Instances>
    <Instance Id="Instance1">
      ...
    </Instance>
  </Instances>
  <SolutionGroups>
    <SolutionGroup>
      <Solution Reference="Instance1">
        ...
      </Solution>
    </SolutionGroup>
  </SolutionGroups>
</HighSchoolTimetableArchive>
```

In benchmarking an indisputable interpretation of the data and constraints is essential. The documentation for the current constraints is given on the website devoted to this project (Post 2008). Although we believe that this effort so far was very important, we still thought it was not enough. For this reason an evaluator was additionally developed and can be accessed from the Internet, see (Kingston 2009). The task of the evaluator is three-fold: first it checks if the provided instances and solutions satisfy the syntax rules. This includes checking consistency of the used Ids and whether a solution respects the *preassignments*. The second task is to provide the infeasibility value and the objective value of the solution and, if indicated, a full report on the deviations for all constraints. Finally, if several solutions of the same instance are included, a comparison table is presented. The first two parts are very useful for the implementation as they provide the user with checks on the generated format and implementation of the constraints. In this sense the evaluator is the ultimate documentation: either the result of the evaluator is accepted, or the behaviour of the evaluator is marked as “bug”.

Some cases led to discussions of the interpretation. One example is the constraint *LimitBusyTimesConstraint*, which limits the busy times of a resource on a day between a minimum and a maximum. Initially this constraint generated the cost for the days without work too. In the revised implementation these days are skipped. This is

reasonable, since by using another constraint one can describe the number of days a resource should be busy.

5 The future

After the past and the present let us make some speculations about the future, based on our experiences till now. First of all we have noted that there is a great interest in this format; interest from researchers in high school timetabling, but also from other areas of timetabling. In our opinion this shows that many researchers feel the urge for exchangeable datasets. Our vision in building this data format is our belief that efforts to define very general formats right from the start have a great chance to fail. We believe that the current format keeps a good balance between tangibility and abstraction.

Though several researchers have expressed their interest, converting their formats to the proposed format requires time. In view of this we are proud that we can present an archive with 15 datasets from 7 contributors. By active acquisition and support we hope to extend this in the near future. As an example we can refer to the development of a automated repository for High School timetabling. This repository will have facilities to convert data sets to the standard data format, uploading new data sets, download of existing data and use of the evaluator. A work in progress version of this site is available (Ahmadi 2010).

If new contributors appear, new constraints or variants of the current constraints may appear. One type of constraint needed is a sequencing constraint. Looking at a resource the sequence of events can be important. An example is when the events take place at different locations: in such cases one would like to minimize the number of location changes, or have breaks or idle times in between. Another example could be a sequence of the subjects like Mathematics, Physics, Chemistry, and Biology for a group of students.

Though new constraints will be needed, one should keep in mind that an instance is usually just an approximation of practice. The timetabler at school will have a clear view of the schedules, but to formalize this in constraints is not always easy (or interesting). In practice hard constraints can turn out to be soft, if necessary, while giving weights to the soft constraints can be difficult. The violation of some important soft constraints can turn out to be unacceptable to the timetabler. The timetabler will rather change the data, and try again to find a good solution. On one hand this might be discouraging for the researchers, but on the other hand it is always a challenge to cope with the reality.

References

- [(Ahmadi 2010)] Samad Ahmadi and Ben Rorije, ‘High School Timetabling Problem Repository’, <http://opt-kd.cse.dmu.ac.uk/www/>, (2010).
- [(Kingston 2009)] , Jeffrey H. Kingston, ‘The HSEval High School Timetable Evaluator’, <http://www.it.usyd.edu.au/~jeff/hseval.cgi>, (2009).
- [(Post 2008)] Gerhard Post, ‘High school timetabling web site’, <http://wwwhome.math.utwente.nl/~postgf/BenchmarkSchoolTimetabling/>, (2008).
- [(Post 2010)] Gerhard Post et al, ‘An XML Format for Benchmarks in High School Timetabling’, *Annals of Operations Research*, (2010) [DOI 10.1007/s10479-010-0699-9].