Next Steps for the Examination Timetabling Format and Competition

Barry McCollum · Paul McMullan · Tomáš Müller · Andrew J. Parkes*

June 2012

1 Background and Motivations

The second International Timetabling Competition, ITC2007[1], [4] included a track on examination timetabling with results presented during PATAT-2008. It provided a well-defined representation [3] with many features not appearing in previous benchmarks. For example, the 'Toronto benchmarks'[2] were highly influential and well-studied (e.g. see [5]) but provided only basic enrolment data and had only a limited mechanism to spread out the exams in time; however, these mechanisms were greatly extended in ITC-2007. A follow-up competition is being arranged, and here we briefly list the most important intended changes (at the time of writing, the exact list and syntax is not yet fixed, and so relatively minor changes are still likely). A challenge in the design of benchmark problems and competitions is to select a compromise between ease of implementation, and the ability to represent all problems that might be encountered in real problems. Accordingly, extensions and changes were selected with the aim they are relatively straightforward and direct to implement, but will make an important contribution to the usability of the format and solvers. Note that changes are inspired by the practical experiences both at Purdue[3] (e.g. see [6]), and also of EventMAP Limited Ltd[4] based at Queen's University and producing software for many institutions.

2 Direct Extensions

By direct extensions we mean those that are closely related to constraints and objectives already used in the previous ITC2007 format; note that the weights assigned to various penalties were collected together into an "Institutional Model (IM)" and many of the proposed changes will simply correspond to new declarations within the IM.

Barry McCollum & Paul McMullan
Queen's University, Belfast, U.K. E-mail: {B.McCollum,P.P.McMullan}@qub.ac.uk
Tomáš Müller
Purdue University, West Lafayette, IN 47907, USA E-mail: muller@unitime.org
Contact author: Andrew J. Parkes
University of Nottingham, NG8 1BB, U.K. E-mail: ajp@cs.nott.ac.uk

[1] http://www.cs.qub.ac.uk/itc2007/
[2] ftp://ftp.mie.utoronto.ca/pub/carter/testprob/
[3] http://unitime.org/
[4] http://www.eventmap-uk.com/

Relaxed Conflicts: Occasionally, some students take a very unusual mix of modules, and the resulting induced conflicts between exams can have a large negative effect on the overall problem, possibly rendering it infeasible or leading to many other students having worse timetables. In such circumstances, it can be better to allow the student to have two of their modules placed into the same time-slot, though with a (high) penalty, and so we will allow a declaration in the IM of the form:

```
TwoExamsAtSameTime <W>
```

giving a penalty of W for each student affected and that will hence need special arrangements. Absence of such a declaration will correspond to an infinite penalty as (implicitly) in the ITC2007 format.

Extended Exam Pattern Penalties: Multiple copies of the "FrontLoad" and "PeriodSpread" declarations will be permitted. This is an easy extension to implement but will allow a much finer-grained control; permitting dividing exams into multiple size classes rather than just the 'large' and 'small' of ITC2007. As explained in [3], multiple copies of the period spread will also capture the system used with the Toronto set of benchmark problems. The IM will also allow "NInARow" and/or "NInADay" penalties to extend the current "TwoInARow" and "TwoInADay" penalties in a natural way to N=3, etc. In combination with TwoExamsAtSameTime, this will also allow modelling of cases when a student can demand a special session when they are assigned to multiple exams in the same day.

Individual Penalties by Pairs of Exam, Room or Period: Extra sections in the data will specify for each (exam,room) or (exam,period) pair a corresponding specific individual penalty of assigning the exam to that room or period. This is useful for exams having special requirements; they might need facilities only present in given rooms, or alternatively might need a special set of times during which a facility or staff member is available. It will also be possible to penalise pairs of (room,period) to account for rooms being unavailable. Penalties be infinite for cases when a particular combinations is prohibited. Separate penalties for individual "(exam,period,room)" triples, were considered but not thought to be useful enough to warrant the extra complexity.

Generalised Room Usage Penalties: In ITC2007, for each room it is only possible to add a penalty for each specific period that it is used. However, for some rooms, their standard configuration might not be suitable for exam usage at all. Preparation of the room for usage at any time during the exam session would have a setup penalty, though the setup would need to be done only once. Hence, it is proposed that the information for each room be extended to:

```
<capacity> <per-usage-penalty> <setup-penalty>
```

With this, many potential rooms could be given and the solver allowed to select which rooms would best match the problem instance and hence should be prepared for usage. In the extreme case, then the setup penalty might be very high and be taken to correspond to building (or at least remodelling) the room itself. This would allow some (limited) usage for space planning [1] in that it could be used to give insight into what combination of room sizes and types are best suited for the examination timetabling problems.

3 Structural Extensions

These are not specific to examination timetabling but are changes to the way that the instances are presented to a solver, and the resources that the solver can use (and will also be retrofitted to previous benchmarks).

Multiple Time Limits and Cores: Firstly, the runtime on instances will not be fixed, but allowed to vary. This also has the practical advantage that the same data instance could be used, but with two (radically) different time limits, e.g. with both a 2 minute and a 120 minute limit. This is practically useful as finding suitable data instances can be difficult, and so it is better to make maximal use of them. Also, even desktop machines are rapidly moving to having many cores, and so it is important that future solvers can exploit this. To encourage this, the next competition will have a sub-track using 8 (or more) cores. This can be done trivially using multiple threads each running the same (randomized) solver, but we hope that solvers will make more interesting cooperative usage of the cores.

Multiple Institutional Models (IMs): Real-world examination timetabling problems have an objective that contains many terms, and (basically) with the IM providing the weights. A natural way extension to treat them as a multi-objective optimisation (MOO) problem, e.g. see [2]. It was indeed considered to try to make a competition with the evaluation being truly multi-objective. For example, solvers could be expected to produce their best approximation of the Pareto Front, and comparisons could be based on standard MOO evaluation methods such as volume dominance; however, we believed that this would lead to too drastic a change. Instead, a simpler version, "MOO-lite", is proposed here: that each data instances is simply associated with multiple different IMs. The solver will then be expected to produce a separate solution for each IM, and they will be used for ranking as if solved independently. However, the solver will be free to share the computational resources and intermediate solutions between IMs in any way that it feels appropriate. The hope is that participants will develop methods that are more efficient (and interesting) than simply solving each IM in turn.

4 Summary

We have listed the main changes intended for the examinations timetabling format and competition.

References

1. C. Beyrouthy, E. K. Burke, B. McCollum, P. McMullan, and A. J. Parkes. University space planning and space-type profiles. *Journal of Scheduling*, 13:363–374, August 2010.
2. E. K. Burke, B. McCollum, P. McMullan, and A. J. Parkes. Multi-objective aspects of the examination timetabling competition track. In *Proceedings of PATAT 2008*, pages 3119 –3126, 2008.
3. B. McCollum, P. McMullan, A. J. Parkes, E. K. Burke, and R. Qu. A new model for automated examination timetabling. *Annals of Operations Research*, 194:291–315, 2012.
4. B. McCollum, A. Schaerf, B. Paechter, P. McMullan, R. Lewis, A. J. Parkes, L. Di Gaspero, R. Qu, and E. K. Burke. Setting the research agenda in automated timetabling: The second international timetabling competition. *INFORMS Journal on Computing*, 22(1):120–130, Winter 2010.
5. R. Qu, E. K. Burke, B. Mccollum, L. T. Merlot, and S. Y. Lee. A survey of search methodologies and automated system development for examination timetabling. *J. of Scheduling*, 12:55–89, February 2009.
6. H. Rudová, T. Müller, and K. Murray. Complex university course timetabling. *Journal of Scheduling*, 14:187–207, 2011. 10.1007/s10951-010-0171-3.