Properties of Yeditepe Examination Timetabling Benchmark Instances

Andrew J. Parkes · **Ender Özcan**

May 2010

1 Introduction

Examination timetabling is a type of educational timetabling which is a highly challenging field for the researchers and practitioners. Examination timetabling problems require a search for the best assignment of examinations into a fixed number of time-slots possibly along with other resources, such as, a set of rooms with certain capacities, subject to a set of constraints. There are two common types of constraints: *hard* and *soft*. The hard constraints must not be violated, while the soft constraints represent preferences that can be infringed. Examination timetabling problems are proven to be NP-complete (Even et al 1976). A recent survey on exam examination timetabling can be found in Qu et al (2009).

There are many variants of examination timetabling problems due to the fact that each educational institution have their own rules, regulations and expectations resulting with various constraints. This situation also makes it extremely difficult to compare different solution methods. Not only comparability but also reproducibility of the results is vital for the research community, as pointed out in Schaerf and Gaspero (2006). McCollum (2006) discusses real world issues in examination and course timetabling. Although practitioners and researchers have to deal with different aspects of examination timetabling, it has been always of interest for both communities to design robust and flexible approaches that can solve new problem instances. ITC2007 (http://www.cs.qub.ac.uk/itc2007/) competition is organised considering the real world examination timetabling complexities and capturing them within the problem instances. The state of the art method for examination timetabling turned out to be a hybrid multistage approach combining Iterated Forward Search (IFS) for feasible initial solution construction and great deluge for improvement as described in Müller (2009). The source code of the solver is available from http://www.unitime.org/itc2007.

Yeditepe University (Faculty of Engineering) data set contains real problem instances from a total of eight semesters in three consecutive years. Bilgin et al (2007) modified the initial data set provided in Özcan and Ersoy (2005) with new properties and also generated a variant of Toronto benchmarks (Carter et al 1996) that fits into the problem formulation which will be referred to as modified Toronto benchmark. This problem is a capacitated variant of examination timetabling. There is a maximum *capacity* of seating available during

School of Computer Science, University of Nottingham, NG8 1JX, UK
http://www.cs.nott.ac.uk/∼{ajp,exo}/

Table 1 Characteristics of the modified Toronto benchmark dataset.

Instance	No. of Exams	No. of Students	No. of Enrolments	Conflict Density	Days	Capacity
car91 I	682	16925	56877	0.13	17	1550
car92 I	543	18419	55522	0.14	12	2000
ear83 I	190	1125	8109	0.27	8	350
hecs92 I	81	2823	10632	0.42	6	650
kfu93	461	5349	25118	0.06	7	1955
lse91	381	2726	10918	0.06	6	635
pur93 I	2419	30029	120681	0.03	10	5000
rye92	486	11483	45051	0.07	8	2055
sta83 I	139	611	5751	0.14	4	3024
tre92	261	4360	14901	0.18	10	655
uta92 I	622	21266	58979	0.13	12	2800
ute92	184	2749	11793	0.08	3	1240
yor83 I	181	941	6034	0.29	7	300

Table 2 Characteristics of the Yeditepe benchmark dataset.

Instance	No. of Exams	No. of Students	No. of Enrolments	Conflict Density	Days	Capacity
yue20011	126	559	3486	0.18	6	450
yue20012	141	591	3708	0.18	6	450
yue20013	26	234	447	0.25	2	150
yue20021	162	826	5755	0.18	7	550
yue20022	182	869	5687	0.17	7	550
yue20023	38	420	790	0.2	2	150
yue20031	174	1125	6714	0.15	6	550
yue20032	210	1185	6833	0.14	6	550

exams at each time slot. The timetable size is fixed with three examination slots per day for a given number of days. The characteristics of each problem instance of modified Toronto and Yeditepe benchmark problem instances are summarised in Table 1 and 2, respectively.

Yeditepe examination timetabling problem has the usual hard constraints:

– Examination conflict (C_1): A student must not sit for more than one examination at any given time.
– Capacity (C_2): At a given period, the overall number of students seated for all examinations should not exceed the fixed capacity.

and the soft constraint

– Examination spread (C_3): Examinations of a student in the same day should not be scheduled consecutively.

As yet, to our knowledge, optimality has not been proven for any solutions of the examination timetabling problem instances in the Toronto and ITC2007 benchmarks, even including the smallest problem instances. This study focuses on the smallest Yeditepe instances which can be solved exactly, and so allows us to test and compare the optimal solutions and the state of the art approach of Müller. Additionally, a multi-objective formulation of the problem based on the trade-off between the room size (capacity) and solution quality is analysed.

Table 3 The results for `yue20023`.

RoomCap.	penalty	time(secs)	IFS-GD
132	70	1123	86, 94, 86
134	68	1468	100, 115, 105
135	65	935	72, 87, 87
136	64	1022	81, 74, 83
137	59	818	73, 73, 87
146	56	875	80, 65, 67
153	55	304	77, 76, 67
157	54	402	73, 73, 67
166	50	371	58, 65, 67
170	48	295	72, 72, 64
176	47	268	64, 76, 66
187	46	234	48, 63, 63

2 Some Exact Results

In this section, we report results of completely solving one of the smaller instances as a case study; yue20023 (chosen simply because it was the largest that we could solve exactly). It was solved using an encoding[1] of the exam timetabling problem as described in (McCollum et al 2008) into ILOG/IBM OPL and solved using CPLEX 11. The encoding is not optimised (e.g. there is no branch and cut) and so timing results are purely for comparison of the relative hardness of different cases. Also, in order to give a better insight into the problem, the size of the room used was varied, and the effect on the final penalty studied. The results given in Table 3 are all illustrated in Figure 1. The last column simply gives the quality obtained from three 3 separate runs of 1200 seconds each and using Müller's winning submission to the examination timetabling track of ITC2007.

There are two main observations:

Firstly, the Pareto front is not trivial, there are fairly wide ranges of unsupported solutions - that is, solutions that are Pareto optimal but not on the convex hull of the Pareto front, and so are not optimal with respect to any linear combination of the objectives. In this case, for example, the Pareto optimal solutions with room size of 153 or 157 will be missed if solving optimally using a linear combinations of the room size and penalty; as any linear combination will not be able to access the 'indented portion' of the Pareto front.

Secondly, even though the hybrid approach tested was the clear winner of ITC2007, it still did not manage to find optimal solutions. This suggests that even on these small instances there is still significant room for improvement in the performance of meta-heuristics.

3 Summary

We have made available some exam timetabling instances from Yeditepe. Despite their independent origin, they fit reasonably well into the format of the ITC2007 benchmarks, suggesting that this format captures real-world issues (McCollum et al 2008). On the smaller instances, we were able to solve them completely using integer programming. In terms of a multi-objective trade-off between the room size and solution quality, the Pareto fronts were found to be interesting with large unsupported regions. This suggests that weighted sum

---

[1] The website `http://www.cs.nott.ac.uk/∼ajp/timetabling/exam/` gives the encoding, the instances, and other supplementary material.
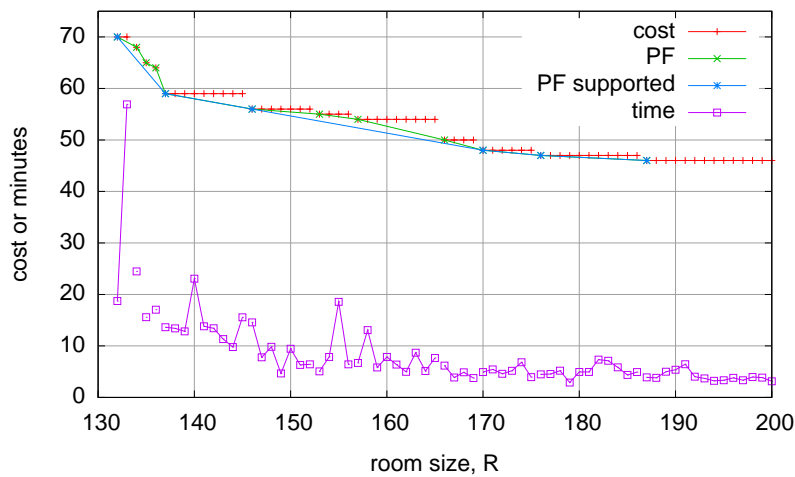
Fig. 1 Results for the instance 'yue20023' as given as a function of the room size. 'Cost' is the optimal (minimal) penalty. The 'PF' are those solutions that are non-dominated. The 'supported PF' are those on the convex hull of the PF line. The 'time' is minutes for CPLEX to solve the instance, which includes the proof of optimality.

methods would potentially miss many interesting solutions. It was also interesting that even the best meta-heuristic solver was consistently failing to find the optimal, suggesting that this research area still has room for considerable improvement.

References

Bilgin B, Özcan E, Korkmaz EE (2007) An experimental study on hyper-heuristics and final exam scheduling. In: Proc. of the International Conference on the Practice and Theory of Automated Timetabling (PATAT'06), Lecture Notes in Computer Science, vol 3867, pp 394–412

Carter MW, Laporte G, Lee S (1996) Examination timetabling: Algorithmic strategies and applications. Journal of the Operational Research Society 47(3):373–383

Even S, Itai A, Shamir A (1976) On the complexity of timetable and multicommodity flow problems. SIAM Journal on Computing 5(4):691–703

McCollum B (2006) University timetabling: Bridging the gap between research and practice. In: Proc. of the 5th International Conference on the Practice and Theory of Automated Timetabling, Springer, pp 15–35

McCollum B, McMullan P, Burke EK, Parkes AJ, Qu R (2008) A new model for automated examination timetabling (under review)

Müller T (2009) ITC2007 solver description: a hybrid approach. Annals of Operations Research 172:429–446

Özcan E, Ersoy E (2005) Final exam scheduler - FES. In: Proc. of the Congress on Evolutionary Computation, IEEE, pp 1356–1363

Qu R, Burke EK, McCollum B, Merlot L, Lee S (2009) A survey of search methodologies and automated system development for examination timetabling. Journal of Scheduling 12(1):55–89

Schaerf A, Gaspero LD (2006) Measurability and reproducibility in timetabling research: State-of-the-art and discussion. In: Proc. of the 6th Int. Conf. on the Practice and Theory of Automated Timetabling, pp 53–62