Modeling and Solution of a Complex University Course Timetabling Problem

Keith Murray[1], Tomáš Müller[1], and Hana Rudová[2]

[1] Space Management and Academic Scheduling, Purdue University
400 Centennial Mall Drive, West Lafayette, IN 47907-2016, USA
{kmurray,muller}@purdue.edu
[2] Faculty of Informatics, Masaryk University, Botanická 68a,
Brno 602 00, Czech Republic
hanka@fi.muni.cz

Abstract. The modeling and solution approaches being used to automate construction of course timetables at a large university are discussed. A course structure model is presented that allows this complex real-world problem to be described using a classical formulation. The problem is then tackled utilizing a course timetabling solver model that transforms it into a constraint satisfaction and optimization problem. The tiered structure of this approach provides flexibility that is helpful in solving the multiple subproblems that arise from decomposition of the university-wide problem. A production system has been partially implemented and results of early use are presented. Practical issues raised during the implementation of the automated timetabling system are also discussed.

1 Introduction

Timetabling is a widely studied area and many potentially useful algorithms have been offered for solving the university course timetabling problem, as evidenced by several recent surveys [7,16,19]. Unfortunately, much of the work in this area has been conducted using artificial data sets or based on actual problems that have been greatly simplified. Methods developed have also rarely been extended to the solution of actual university problems of any large scale. McCollum offers a good review of this situation in [11].

The major differences between many of the problems studied and their real-life counterparts are the additional complexity imposed by course structures, the variety of constraints imposed, and the distributed responsibility for information needed to solve such problems at a university-wide level. University timetabling problems may also involve the solution of multiple subproblems with very different characteristics. In practice, therefore, the solution process should not be specifically tailored to a single problem type.

The work described here has been motivated by the need to create and modify course timetables at Purdue University that better meet student course demand and allow students to be assigned to the constituent course sections in a way

E.K. Burke and H. Rudová (Eds.): PATAT 2006, LNCS 3867, pp. 189–209, 2007.

that minimizes conflicts. Purdue is a large (39,000 students) public university with a broad spectrum of programs at the undergraduate and graduate levels. In a typical term there are 9,000 classes offered using 570 teaching spaces. Approximately 259,000 individual student class requests must be satisfied. The complete university timetabling problem is decomposed into a series of subproblems to be solved at the academic department level, where the resources required to provide instruction are controlled. Several other special problems, where shared resources or student interactions are of critical importance, are solved institution wide. A major consideration in designing the system has been supporting distributed construction of departmental timetables while providing central coordination of the overall problem. This reflects the distributed management of instructional resources across multiple departments at the University. The general definition of the university-wide timetabling problem described in this paper is similar to the problem studied by Carter [6] at the University of Waterloo, and the influence of that work can be seen here, though the solution methods used differ significantly. It is hoped that the results of the present work will, likewise, be beneficial to other institutions seeking to improve their ability to construct course timetables for their students.

This paper discusses the approach used for modeling and solving the additional complexities involved in developing automated solution techniques for a real-life course timetabling problem on the scale of a large university. Although a specific example is discussed, many of the methods used should be applicable to modeling and solving other complex problems.

The complexity of the university course timetabling problem studied here has been broken down by developing a logical data model allowing all courses to be represented as hierarchical groupings of classes with additional parent–child relationships and constraints governing their placement. This allows use of one standard class-oriented problem formulation rather than having to develop different models and solution methods to work with the wide variety of ways that departments organize the instruction given in their courses.

A flexible and general solution technique has been developed for solving course timetabling problems and applied to all of the departmental and special subproblem. Rather than applying multiple solution methods, each optimized around the characteristics of a specific problem, a single solution approach allows the outcomes of all of these subproblems to be easily combined into a complete solution and facilitates optimization of the sectioning of students across the complete timetable.

Each of these topics will be explored in greater depth in the next three sections of this paper, followed by observations on creating a general framework that is very useful in developing a practical solver. Several practical issues faced while implementing a real-world system are then discussed, including competitive behavior among users, making changes to solutions, and managing data consistency. Some results from actual use of the system to solve departmental and campus problems are also presented. In addition, links are provided to the problem data used in this work to promote further study by other researchers using real data instances.