The Teaching Space Allocation Problem with Splitting

Camille Beyrouthy[1], Edmund K. Burke[1], Dario Landa-Silva[1], Barry McCollum[2,3], Paul McMullan[2], and Andrew J. Parkes[1]

[1] School of Computer Science and IT,
University of Nottingham, Nottingham NG8 1BB, UK
{cbb,ekb,jds,ajp}@cs.nott.ac.uk
[2] Queen's University of Belfast, Belfast, BT7 1NN, UK
{b.mccollum,p.p.mcmullan}@qub.ac.uk
[3] Realtime Solutions Ltd, 21 Stranmillis Road, Belfast, BT9 5AF
b.mccollum@realtimesolutions-uk.com

Abstract. A standard problem within universities is that of teaching space allocation which can be thought of as the assignment of rooms and times to various teaching activities. The focus is usually on courses that are expected to fit into one room. However, it can also happen that the course will need to be broken up, or 'split', into multiple sections. A lecture might be too large to fit into any one room. Another common example is that of seminars or tutorials. Although hundreds of students may be enrolled on a course, it is often subdivided into particular types and sizes of events dependent on the pedagogic requirements of that particular course.

Typically, decisions as to how to split courses need to be made within the context of limited space requirements. Institutions do not have an unlimited number of teaching rooms, and need to effectively use those that they do have. The efficiency of space usage is usually measured by the overall 'utilisation' which is basically the fraction of the available seat-hours that are actually used. A multi-objective optimisation problem naturally arises; with a trade-off between satisfying preferences on splitting, a desire to increase utilisation, and also to satisfy other constraints such as those based on event location and timetabling conflicts. In this paper, we explore such trade-offs. The explorations themselves are based on a local search method that attempts to optimise the space utilisation by means of a 'dynamic splitting' strategy. The local moves are designed to improve utilisation and satisfy the other constraints, but are also allowed to split, and un-split, courses so as to simultaneously meet the splitting objectives.

1 Introduction

An important issue in the management of university teaching space is that of planning for future needs. Support for such decision-making is generally divided into two broad, and sometimes overlapping, areas:

E.K. Burke and H. Rudová (Eds.): PATAT 2006, LNCS 3867, pp. 228–247, 2007.

– *space management:* near-term planning,
– *space planning:* long-term planning, including capacity planning.

A fundamental stage of capacity planning aims to estimate the projected student enrollments, and multiply by the expected weekly student contact hours to obtain the total demand for 'seat-hours'. Similarly, for the rooms we could just sum up the room capacities and multiply by the number of hours they are available in order to determine the 'seat-hours supply'. A naive way to perform capacity planning, based on such seat-hours estimates, would be simply to ensure that the supply exceeds the demand. However, it is very rare that it is possible to use all of the seats. The efficiency of space usage is usually measured by giving a figure for the 'utilisation': i.e., the fraction (or percentage) of available seat-hours that actually end up being used. In real institutions, the utilisation can be surprisingly low, perhaps only 20–50%. To compensate for this, when planning the amount of teaching space to supply, we need to build in excess capacity [13,14].

Naturally, such excess capacity is expensive, because it entails planning for seats to be underused. Good planning should reduce the excess capacity without increasing the risks that expected activities will not find a space. However, this is difficult because there is little fundamental understanding of why the utilisation is so low in the first place, or of the interaction of various constraints and objectives with the utilisation.

A study of this issue was initiated in [5,6]. However, that work, like the majority of work on (university) course timetabling research was concerned with unsplittable 'events' (or 'courses' or 'classes'). Such courses are 'atomic': i.e. they are not to be subdivided but need to be assigned to a single room and timeslot. However, in some circumstances, courses cannot be taken to be atomic, but must instead be subdivided, or 'split', before allocating them to rooms and timeslots. In this paper, we extend the work of [5,6] to the case of courses that require considerable splitting.

Our ongoing investigation into space management and space planning is closely related to research into automated timetabling but we emphasise that there is a crucial difference between the two. In automated timetabling, the set of events that should be accommodated into timeslots and rooms is usually fixed. This means that the space utilisation, in terms of seat-hours demand and offer, is also fixed from the outset. However, in this paper we want to study those factors that have an impact upon space utilisation (even before constructing the timetable). For this, we investigate a scenario in which the seat-hours demand (events to accommodate) is much larger that the seat-hours offer (available rooms). This allows us to vary the utilisation by selecting those events that will be accommodated and those that will be not. We note that although the algorithms presented here allocate events into rooms and timeslots, we are not proposing a timetabling approach. We are presenting a study that helps us to understand the interactions between space utilisation and aspects such as timetabling constraints and others.