Multidisciplinary International Conference on Scheduling : Theory and Applications (MISTA 2009)
10-12 August 2009, Dublin, Ireland

MISTA 2009

Recent Work on Planning and Management of University Teaching Space

Andrew J. Parkes · Barry McCollum · Paul McMullan · Edmund K. Burke

1 Motivations and Overview of Existing Works

Universities have to invest considerable financial and human resources in the provision of space for teaching activities such as lectures, seminars, tutorials and workshops. Naturally, they would like such investments to be made wisely and efficiently. However, there is considerable evidence that in many Universities the resulting space is considerably underused. In a report by the Higher Education Funding Council for England (HEFCE), roughly speaking, it was found that often space was only used half the time, and then only half filled [9]. At least on the face of it this seems like an inefficient use of resources; it is natural to expect or hope that better planning for the space capacity would improve this. However, a fundamental problem has been that there was no real foundation for deciding whether or not such usage levels are in fact an inevitable result of meeting the timetabling and teaching resources, or alternatively, to provide methods to improve the situation. In this abstract we briefly overview our work towards remedying this situation. Here, we cannot hope to cover all work in the topic; instead the aim is to discuss our various strands of research and how they are woven together.

The space planning process needs to decide upon what space resources need to be provided for the envisaged teaching events. For this, an essential first step is to decide upon a measure of the quality of the provision. In reality, the true measure is likely to be a complex financial issue and highly specific to each institution. Hence, to make general progress, we follow standard practice and use the simpler measure of "utilisation", which we take to simply be the average fraction (or percentage) of the seats used within the teaching space [5,6]. That is, it is a direct measure of the demand (by students) for a given supply (or space). So the (unfortunately common) situation that 'half of the time rooms are empty, and are only half full when used' corresponds to a utilisation of a paltry 25%.

Andrew J. Parkes (contact author), Edmund K. Burke,
School of Computer Science, University of Nottingham, Nottingham, NG8 1BB, UK.
E-mail: { ajp, ekb }@cs.nott.ac.uk

Barry McCollum, Paul McMullan
School of Computer Science, Queen's University, Belfast, BT7 1NN, UK.
E-mail: { b.mccollum, p.p.mcmullan }@qub.ac.uk

Firstly, it is important to observe that in a standard timetabling problem (see for example [11]) both the set of events and the rooms are fixed, hence, the utilisation is immediately and trivially derived from the problem description itself. The utilisation has the same value for all (feasible) solutions – standard timetabling software works to find feasible solutions and improve various quality measures, however, utilisation is not one of those measures. In particular, in order to be able to study the factors that influence utilisation we have to consider changes to at least one of the set of rooms (the supply) or the set of events (the demand): Our papers consider both of these changes.

Also, it is crucial to realise that space planning is inevitably performed without access to complete information about future teaching events and student enrolments. Space planning is an imperfect information problem; that is, a question of decision support (or optimisation) under uncertainty. This places it in stark contrast with the (vast majority) of timetabling research which works on classic perfect-information combinatorial optimisation problems. Hence, whilst space planning (in the sense we use) naturally includes timetabling issues, it is is separate problem in its own right; and substantially less developed than the timetabling arena. Note in the literature the term "space planning" might also be used, for example, finding the minimum space needed for a given set of events, and then this is a packing problem: Such planning is useful (and hard), however does not have the planning under uncertainty issues that inevitably arise in longer-term campus planning.

Achievement Curves:
The first step [5,6] was to define what we call "Achievement Curves" which are based on sampling different sets of events from a pool, and then attempting to assign as many as possible to rooms and timeslots. Each selected set of events corresponds to a "Requested Utilisation (U_R)" and the successfully assigned events gives a "Achieved Utilisation (U_A)". Empirically, we observed that the statistical behaviour of U_A depends upon U_R in a fashion that exhibits a threshold (or phase transition) [8]. There was a "Critical Utilisation (U_C)" that separates two 'phases'. In the under-constrained phase, $U_R < U_C$, with high probability all events can be scheduled. In the over-constrained phase, $U_R > U_C$, with high probability some events will inevitable remain unscheduled. Naturally, the actual value of of the critical utilisation U_C is not universal, but depends on the particular room resources, characteristics of the set of events, details of the timetabling constraints, etc. However, we have essentially taken the value of U_C as a measure of the effectiveness of the resources compared to the demands, within the context of the specified constraints. Hence, improving utilisation is done on the basis of planning or modifying the set of resources so as to improve the critical utilisation. That is, the achievement curves themselves provide the basis of evaluation of a given set of room resources over a set of candidate events.

This initial work was done using only lecture events which are (generally) assigned as a single event. However, other teaching activities such as tutorials will often require splitting into smaller events. After such splitting into multiple sections, "student sectioning" then assigns students to sections so as to minimise timetable clashes (for recent work on this see [10]). The work on achievement curves was also applied to such cases [3,1].

Computational Hardness Peaks:
Such thresholds are well known (at least within the field of artificial intelligence) to often be associated with strong peaks, and 'easy-hard-easy' patterns, in the computational difficulty of solving the typical problem instances [12]. In [4] we studied this behaviour and the resulting dramatic increase in computational effort needed as approaching the critical utilisation. We observed that this will have a practical impact; as potential utilisation improves then it is likely to need ever better timetabling software in order to fully exploit the potential. Poorer

algorithms might well be usable at low levels of utilisation but are likely to rapidly become impractical on improving space planning.

Room size profiles:
It is generally expected that to obtain good utilisation the distribution of sizes of the rooms needs, in some sense, to match the distribution of sizes of events. Conversely, a poor 'room size profile' will lead to poor utilisation. In practice, this can arise because of changes in teaching practices. A teaching facility might have been designed for the case of a few large lectures, but teaching practices changed to use many small lectures and tutorials. In [2], we provided methods to quantify this effect. This was then extended to provide some methods, based upon stochastic programming ideas, to design robust room size profiles. That is, to provide sets of target room capacities that would lead to improved critical utilisation values.

Space Type Mixing:
In teaching there are many different types of teaching activity (lectures, seminar, computer laboratories, etc), and each such event-type will have preferred space-type; that is, a type of room desinged for a particular activity. In most work the event and space-types are treated individually and separately. However, in the real-world there are many interactions between space-types and often events will be placed in different space-type. In recent work [7] we have looked the effect of allowing such space-type mixing. A common problem is that the set of space-types becomes a poor match to the mix of event-types, and so also provided methods to select a better mix of space-types; that is, to provide a better 'space-type profile'.

2 Planned Work

Our existing work has looked at the many facets of space planning, but has generally treated them separately. Naturally, the most immediate goal is to combine all the work into an integrated approach. Also, of course, any such approach will need to be tested on a wide range of problems.

Finally, we note that there is a general and interesting issue about whether the timetabling model used within space planning needs to be a complete model or whether it suffices to use a simpler approximate model. We will investigate the question of whether some of the details needed when doing actual timetabling can be neglected when using timetabling within some system for space planning. That is, we will look at the extent to which some approximate timetable model can be used for space planning without causing a significant loss in the quality of the resulting space plans when presented with the complete timetabling model.

References

1. Camille Beyrouthy. *Models, Solution Methods and Threshold Behaviour for the Teaching Space Allocation Problem*. PhD thesis, School of Computer Science, University of Nottingham, 2008. (Submitted).
2. Camille Beyrouthy, Edmund K. Burke, Dario Landa-Silva, Barry McCollum, Paul McMullan, and Andrew J. Parkes. Improving the room-size profiles of university teaching space. Technical report, School of Computer Science, University of Nottingham, 2007. Presented at the Dagstuhl Workshop on 'Cutting, Packing and Space Allocation', March, 2007.
3. Camille Beyrouthy, Edmund K. Burke, Dario Landa-Silva, Barry McCollum, Paul McMullan, and Andrew J. Parkes. The teaching space allocation problem with splitting. In Edmund K. Burke and Hana Rudova, editors, *Revised Selected papers from the 6th International Conference on the Practice and Theory of Automated Timetabling (PATAT 2006), Brno, Czech Republic*, volume 3867 of *Lecture Notes in Computer Science*, pages 228–247. Springer-Verlag, 2007.

4. Camille Beyrouthy, Edmund K. Burke, Dario Landa-Silva, Barry McCollum, Paul McMullan, and Andrew J. Parkes. Threshold effects in the teaching space allocation problem with splitting. *European Journal of Operational Research (EJOR)*, 2008. To appear.
5. Camille Beyrouthy, Edmund K. Burke, J. Dario Landa-Silva, Barry McCollum, Paul McMullan, and Andrew J. Parkes. Understanding the role of UFOs within space exploitation. In *Proceedings of the 6th International Conference on the Practice and Theory of Automated Timetabling (PATAT 2006)*, pages 359–362, Brno, Czech Republic, 2006.
6. Camille Beyrouthy, Edmund K. Burke, J. Dario Landa-Silva, Barry McCollum, Paul McMullan, and Andrew J. Parkes. Towards improving the utilisation of university teaching space. *J Oper Res Soc*, 60(1):130–143, 2008.
7. Camille Beyrouthy, Edmund K. Burke, Barry McCollum, Paul McMullan, and Andrew J. Parkes. University space planning and space-type profiles. Submitted to Journal of Scheduling, 2007.
8. B. Bollobas. *Random Graphs*. Academic Press, London, England, 1985.
9. HEFCE. Estates management statistics project. Technical report, Higher Education Funding Council for England, March 1999. Report 99/18. http://www.hefce.ac.uk/pubs/hefce/1999/99_18.htm.
10. Keith Murray and Tomáš Müller. Real-time student sectioning. In Philippe Baptiste, Graham Kendall, Alix Murnier-Kordon, and Francis Sourd, editors, *Proceedings of the 3rd Multidisciplinary International Conference on Scheduling: Theory and Application (MISTA 2007), Paris, France*, pages 598–600, 2007.
11. Andrea Schaerf. A survey of automated timetabling. *Artif. Intell. Rev.*, 13(2):87–127, 1999.
12. Bart Selman, David G. Mitchell, and Hector J. Levesque. Generating hard satisfiability problems. *Artificial Intelligence*, 81(1-2):17–29, 1996.