Scheduling Research Grant Proposal Evaluation Meetings and the Range Colouring Problem

Patrick Healy

CS Department, University of Limerick, Limerick, Ireland
patrick.healy@ul.ie

Abstract. In many funding agencies a model is adopted whereby a fixed panel of evaluators evaluate the set of applications. This is then followed by a general meeting where each proposal is discussed by those evaluators assigned to it with a view to agreeing on a consensus score for that proposal. It is not uncommon for some evaluators to be unavailable for the entire duration of the meeting; constraints of this nature, and others, complicate the search for a solution and take it outside the realm of the classical graph colouring problem. In this paper we (a) report on a system developed to ensure the smooth running of such meetings and (b) compare two different ILP formulations of a sub-problem at its core, the list-colouring problem.

1 Background

The process of evaluating research grant proposals presents some interesting opportunities for operations research practitioners and researchers. The model we discuss in this paper assumes that a fixed pool of evaluators exists and a set of grant proposals is distributed amongst them subject to the evaluators' stated abilities to evaluate each. Several constraints affect the allocation of proposals to evaluators. For example, insofar as it is reasonable, all evaluators should be allocated the same load, although some evaluators, such as vice-chairs, may be expected to take a reduced load of evaluations due to other duties; upper limits may be enforced on the number of evaluators employed; in the interests of (particularly, novice) evaluators not being intimidated at the associated meetings that occur later, no proposal should have a majority of vice-chairs evaluating it; for each proposal, one evaluator should be appointed as proposal reporter amongst the – usually 3, although requests for larger financial sums necessitate more – evaluators assigned to it with this extra duty evenly allocated amongst all evaluators.

These complications, and other side constraints, make the allocation problem an interesting study in its own right. Nonetheless, for present purposes, we will assume that evaluations have been completed and the evaluators' recommendations must now be unified in to an ordered *ranking list*. From this list grant applications will be offered funding subject to budget and other administrative considerations. Although other procedures for arriving at a ranking list are possible [5] the most commonly used procedure appears to be that each

E.K. Burke and H. Rudová (Eds.): PATAT 2006, LNCS 3867, pp. 119–131, 2007.

evaluator initially assigns a numerical mark to each proposal allocated to them. At a follow-up meeting of the (usually three) evaluators associated with each proposal a consensus mark is arrived at. The ranking list is then based on the consensus marks for each proposal.

Our interest begins when a general panel meeting (a panel may be thought of as a general research area, e.g. computer science) brings all of the evaluators together. Each proposal is discussed face to face by the assigned evaluators for the purpose of agreeing a consensus position and category scores, from which the final evaluation report may be written; the consensus meeting runs for a fixed length of time. Following the entry of all scores in a database, a ranking list is generated which forms the basis for funding decisions by the grant agency. So that the entire panel of evaluators may agree to the ranking list it is desirable that all consensus meetings be completed as quickly as possible allowing time for the inevitable clean-up before the final ranking list acceptance process.

In its most restricted form the problem may be expressed as, given an assignment matrix of proposals to evaluators, where each proposal has been read by some subset of evaluators, generate a schedule of consensus meetings that uses the fewest number of time periods where the meetings can be held. A maximum of T time periods exist.

The constraints that must be respected, then, are:

1. A consensus meeting can only take place during one time slot;
2. An evaluator can only be in one consensus meeting during a time slot;
3. If a consensus meeting for a proposal takes place then all of the evaluators assigned to it must be present;
4. No more than T time slots may be used.

The goal is to minimize the number of time slots used. While the problem in its purest form is similar to the exam timetabling problem, we will see some additional constraints that impinge on it later.

In the following section we describe our first approach to solve the problem exactly by modelling it as an ILP. This will act as a reference point for what will follow later. We then describe a two-phase heuristic that combines tabu search to find an initial colouring with a weighted bipartite graph matching problem on the colouring, and the deficiencies of this model. Section 4 proposes a return to the search for an exact solution to a fundamental sub-problem by comparing two competing ILP formulations. Section 5 concludes the paper.

2 An ILP Model

Given $E = e_{ij}$, a binary assignment matrix that indicates what evaluators have been assigned to read proposal j, we can view its transpose $E^{\mathrm{T}} = P = p_{ij}$ as the matrix that indicates what proposals have been assigned to evaluator j.

We introduce binary variables x_{ij} that indicate that the consensus meeting for proposal i will take place in time slot j, and y_{ij} that indicate that evaluator i is in some meeting during time slot $j, 1 \leq j \leq T$.