

Multicriteria Optimisation in (Large Scale) Real World Applications

by

M. Ehrgott

Department of Engineering Science

University of Auckland

m.ehrgott@auckland.ac.nz

When I read the Forum column in the last newsletter, I mentioned in an email to José Figueira, that I found it interesting, and that hopefully current projects I am involved with will be among objects D, E, or F according to Vincke's classification of applications. José invited me to write the next column. So here it is.

Discussion about applications of MCDM/MCDA technologies has been going on for a while in our community. What constitutes a "real" application? Which methods are the "best"? are some questions that might never really find a final answer. I don't want to discuss the former question. As for the latter: Probably it just depends on the situation. Often decision aiding techniques are the appropriate tools, especially in cases where a relatively small number of alternatives is considered and they certainly have had their well-deserved success. Interactive methods work well when the intermediate computational steps can be carried out quickly. But sometimes both are not applicable due to the structure of the decision problem underlying the application. Here I would like to put forward some arguments for using multicriteria optimisation methodology in real world projects

Let us consider the situation of a large scale application, where a very large number or a continuity of alternatives might exist. Solution of such a problem could involve multi-million dollar savings, or could make the difference between survival and death of a patient. I would like to report on two ongoing projects of that kind. The problems themselves are very well known and more or less routinely solved as single objective optimisation problems. Perhaps this familiarity with the models is the reason for the arising insight that what would really be needed in both applications is the consideration of more than just the cost criterion or another aggregate measure of performance.

Airline Crew Scheduling at Air New Zealand

In the competitive airline markets of today, major airlines can no longer rely solely on minimal cost solution of their scheduling and rostering problems. Robustness of solutions is of increasing importance to avoid the cascading effects of delayed flights and to gain an edge in the business. In a project in co-operation with Air New Zealand we set up a model of bicriteria optimisation of the classic crew scheduling problem, where we consider both the cost and robustness criterion as a linear objective function. The resulting mathematical model is a large scale set partitioning problem. Its solution requires sophisticated problem specific technologies in integer programming - such as constraint branching and column generation - as well as multiobjective (linear programming) techniques. The system will provide the users with a trade-off analysis of the two criteria and the possibility to navigate among solutions by specifying certain parameters, e.g. the increase in cost management is willing to cover in order to obtain a more robust schedule. In this case the decision makers (DM) involved, i.e. the network logistics department at Air New Zealand, are familiar with optimisation software for more than 15 years (previous work by Air New Zealand and the University of Auckland was presented in the finalist round of last years Edelman prize). Thus they are comfortable with OR technology and understand well what OR can do for them and what their role in the process can be. In fact, they prefer a system, in which they have the ability to steer the process -

and which shows them the range of possibilities they have - to one, in which the preferences would be considered from the start, as they felt that this would have them searching around "in darkness".

Radiation Therapy Planning in Cancer Treatment

The inverse planning problem in radiation therapy planning is a key issue to increase the effectiveness of therapy plans (in Germany, about one third of patients diagnosed with curable cancer die nevertheless). In this problem, the goal is to find optimal intensity profiles of radiation beams, given a desired dose distribution, provided by the physician. This turns out to be a precarious problem of keeping the balance between an ineffective underdosing of the tumour and a dangerous overdosing of healthy organs. Traditionally, the model involves minimising a weighted sum of deviations from prescribed doses with a trial and error approach for adjusting the weights. However, the more natural formulation is a multicriteria model, in which one objective occurs for the target volume (tumour) and each of the organs at risk. Using a discretisation of digitised CT or MRT images and beam heads, a large scale multiobjective linear program results. For this problem a good representation of efficient solutions can be computed and maintained in a database (what, exactly, constitutes a "good" representation is an ongoing topic of research). Note that (perhaps except from setting some initial parameters) this does not need the involvement of the physician. The DM (physician) can then consult the database, navigate among the solutions and pick the one most suited for the patient awaiting treatment, and he can do so on-line! Thus the involvement of the physician in the process is increased, the time consuming trial and error process needed earlier is avoided, we get rid of the problem of the well-known extreme sensitivity of the weighted sum approach to the weight values and the chances of obtaining good treatment plans are greatly increased. The advantages of a multicriteria optimisation approach in this application are self-evident.

From my point of view, these two applications clearly illustrate the need to develop appropriate multicriteria optimisation methodology. It is simply not sufficient to have sophisticated methods of aggregating - in which way whatsoever - the multiple objectives into one overall goal, too much valuable information will be lost. In addition, the direct application of interactive methods must be discarded out of hand, because the intermediate optimisation problems are too big to be solved on-line. We should also be aware of the fact that "real" projects will always involve two groups of people: the DMís (schedulers, physicians, ...), who are experts in the application or the real world problem, and consultants (us), who are experts in Operations Research and/or multicriteria methodology. Co-operation is needed and indispensable for success of the project in the development of the underlying mathematical model as well as in the validation of results. Apart from that, let the DM do what she does best: Judge which solution to choose by her expertise and experience. But let the consultant provide this possibility by the development of correct, efficient and easy to use systems of multicriteria optimisation (tailored to the needs of the DM, e.g. by visualising radiation dose values by colour coded pictures). With such an approach the interaction takes place at the appropriate stage and the DMís expertise is utilised to maximum benefit. Needless to say that ample possibilities for research open up if we want to make multicriteria optimisation a successful, respected, and well established technique in OR practice.