SOLVING MULTI-CRITERIA ALLOCATION PROBLEMS: A DECISION SUPPORT SYSTEM APPROACH

by

Rob R. Weitz
Information Systems Department
Leonard N. Stern School of Business
New York University
100 Trinity Place
New York, New York 10006
U.S.A.

and

M. Tawfik Jelassi
Technology Management Area
INSEAD
Boulevard de Constance
77305 Fontainebleau
France

March 1990

Center for Research on Information Systems
Information Systems Department
Leonard N. Stern School of Business
New York University

Working Paper Series

STERN IS-90-5

SOLVING MULTI-CRITERIA ALLOCATION PROBLEMS: A DECISION SUPPORT SYSTEM APPROACH

ABSTRACT:

MCADSS is a multi-criteria allocation decision support system for assisting in the task of partitioning a set of individuals into groups. Based upon multiple criteria, MCADSS's goal is to maximize the diversity of members within groups, while minimizing the average differences between groups. (The project may be viewed from several perspectives: as a multi-criteria decision making problem, as a "reverse" clustering problem, or as a personnel assignment problem). The system is currently being used to allocate MBA students into sections and study teams at INSEAD, a leading European business school. This paper describes the rationale for MCADSS, design criteria, system methodology, and application results. It also suggests how the approach outlined here might be used for further applications.

KEY WORDS AND EXPRESSIONS:

Multiple Criteria Decision Making; Decision Support Systems; Personnel Assignment Problems; Systems Analysis and Design; Computer Implementation.

Acknowledgement:

The authors gratefully acknowledge the contributions and comments of Julie DelVecchio Smith, Helen Henderson, and particularly Martin East.

BACKGROUND:

INSEAD, a leading European business school, graduates two MBA classes every year: one which starts in September, the other in January. Each promotion contains over 220 students. The student body is very international, typically composed of some dozen different nationalities. Aside from cultural and linguistic differences, entering students are drawn from a variety of fields of study and actual work experiences.

At the start of the program, the entering promotion is divided into three sections. Members of each section attend the same required classes together. The goal of the MBA office is to create sections which are as similar to each other as possible, based on a set of criteria. Concomitant with this strategy is maximum diversity within sections. The criteria of interest are: nationality, mother tongue, area of academic training, type of work experience, sex, university attended, previous employer, and age. For example, each section should have roughly the same number of lawyers, French nationals and native English speakers while the average age of the sections should be approximately equal. It should be clear that trade-offs are frequently required in order to develop an acceptable partition. (Perhaps in order to keep the number of engineers constant across sections, an unequal weighting of ages will be required, for example.) The partitioning process is further constrained by various required and forbidden circumstances. Married couples frequently want to be in the same section (in order to coordinate schedules), while recipients of certain scholarships must be spread across sections in order to minimize disruptions during portions of the academic period when they return to their sponsoring companies.

After satisfactory sections have been created, study groups are formed within each section. (Students do case analyses and project work in these groups.) Each section is composed of approximately ten groups of six or seven students each. As before, the idea is to minimize the differences between groups - they should be as similar as possible, based on the previously mentioned criteria - and to maximize differences between individuals within groups. Again, a set of required and forbidden circumstances are encountered: no more than two members of the same nationality may be in any group, no more than two members of any occupational background (profession) may be in any one group, individuals married to each other should not be in the same group, certain scholarship cases should be separated, no two members of the same company in any group, etc.

In maximizing differences within sections/groups and minimizing them between sections/groups, the intent is to provide the student with a challenging exposure to work in a truly international business environment. Students typically cite this experience as one of the highlights of their study at INSEAD.

BEFORE MCADSS:

The section/group making procedure has been carried out manually for over 20 years. With the growth of the MBA program, the task became increasingly time consuming and onerous. As of the start of MCADSS development, two high-level administrators in the MBA office spent two days twice each year in creating the sections and groups. Working from sheets of paper generated for each student during the admission proces[illegible] procedure required keeping track of a large amount of information, and undergoing a trial and error process until a satisfactory solution was reached. The administrators involved in the task, although unhappy with the manual process, were disinclined to pass the task off to their subordinates for two main reasons. First, much of the data on the admission sheets was confidential (salary levels and interview results, for example). Second, a detailed knowledge of the incoming students was required in order to do a thorough job as admission data might be technically correct, but less than helpful. For example, a student may carry the citizenship of one country, but have been born, raised and educated in another. This eliminated anyone who was not intimately familiar with the profiles of the incoming students from doing the allocation task. Moreover, it called into question the usefulness of a completely automated procedure, though a full spectrum of approaches for addressing the problem was considered.

ALTERNATIVE APPROACHES:

The personnel assignment problem may be characterized by the requirement of assigning a limited number of individuals to a limited number of entities, under a specified set of constraints, with the intent of maximizing some overall utility function. Examples include assigning workers to jobs, students to schools, and salespeople to sales territories. Hill et al.[7], and Beheshtian-Ardekani and Mahmood [2], in surveying the personnel assignment problem area, cite four categories of solution approaches: random assignment, free market assignment, optimal algorithm (i.e., mathematical programming), and heuristic assignment.

A random procedure is appropriate when assignments are so unimportant as to not warrant data collection and model development efforts. In the free market approach, individuals engage in auction style bidding for assignments, based upon their preferences. Clearly, neither of these approaches is suited to the section/group assignment application described here.

Optimal algorithm approaches use mathematical, that is integer or goal, programming models. Some existing applications involve allocating teachers to private schools [8], faculty to class schedules [5, 6], students to projects, and students to company job interviews [7]. There are a number of problems with this approach for the section/group application.

First, as mentioned in the previous section, the student data does not always fully capture reality. The major advantage of a mathematical programming model vis-a-vis a heuristic is its capability, at least in theory, of arriving at an optimal solution. However, this advantage is negated if the data upon which the model is based is not precise.

Second, given the number of constraints, individuals, criteria, and values for each criterion, the mathematical programming model would be extremely complex. Aside from the difficulties of formulating such a model (for example, specifying an objective function), given the integer nature of the problem, there is no guarantee of arriving at a solution within a reasonable time limit. Dhar and Ranganathan [4] use an integer programming formulation, in addition to an expert system approach for assigning faculty to course schedules. In contrasting the results they note (p. 18), "Like many integer programs, we found the solution time to be highly unpredictable, varying from a few minutes to a few days. ...In effect, if a feasible solution was not found quickly, it was not found at all." Essentially, the problem must be well structured in order to ensure solution times.

Third, the mathematical programming approach is rigid. For example, a single objective function may not reflect the richness of the decision maker's goals. The alternative, a multiple objective model is yet more complex and may burden the decision maker with a large set of solutions to analyze.

Finally, mathematical programs are not transparent to non-technical users. It was felt that a system which operated in a way the users could intuitively grasp, and which interacted with them to enhance their abilities would more likely be accepted.

Ultimately, given the nature of the problem, the environment, and the alternatives, assistance to the administrators in the form of a decision support system (DSS) seemed most appropriate. Taking into account the complexity of the task, the DSS would have to incorporate some sort of heuristic procedure for making "reasonably good" assignments of students to sections and groups.

Heuristic assignments are sequential procedures which fall into two basic categories: people-sequential and utility-sequential. These procedures are described by Hill et al. ([7], p. 10).

> The people-sequential heuristic ranks all people according to some rule (e.g., seniority, grade-point average, etc.) and then gives the first person on the list their highest (feasible) preference assignment. This process is continued sequentially for all people. ...The utility-sequential heuristic assigns all people to their first choice assignments without regard to a constraint on the number of people assigned to a task. For the tasks that have an overload of people, a person is randomly selected and reassigned to a second choice assignment (as long as that assignment is not overloaded). The process is continued until a feasible set of assignments is found.

These heuristic methods are easily comprehensible and computationally quite simple. The drawback is that they may provide a poor solution, or in fact, no solution at all. As both heuristics are based upon user preferences, neither is directly appropriate for the section/group assignment problem.

Beheshtian-Ardekani and Mahmood [2] use a variation of the people-sequential heuristic in assigning students to project groups. Their goal is to provide groups balanced by the criteria of "experience". To do this, each student is asked to respond to a questionnaire which measures the student's experience in the field. The questionnaire produces a single composite score directly correlated with experience. The scores are then ranked in decreasing order. To create n groups, the n most experienced students are "dealt out", one to a group. The process continues, n students at a time; at each iteration, the student with the lowest experience score is added to the thus far created group with the highest total score. The process is complete when no students are left. While conceptually very simple, this procedure hints at the heuristic assignment mechanism developed for the group/section assignment task (which is described below).

DESIGN CRITERIA:

MCADSS was designed to support the MBA program administrators in the task of partitioning students into sections and groups. The task is difficult due to the large amount of loosely organized data, the inherent complexity of sorting and classifying upwards of over 220 individuals based on some half-dozen variably weighted criteria, and due to the judgment involved. The administrators felt that the system should initially provide them with a flexible means of examining the data and useful summary statistics, so that they could get a good "feel" for the data. Next the DSS should suggest a partition of the promotion into sections, based on the objectives of maximum diversity within sections and minimum diversity across sections, and on the required/forbidden constraints. Well designed output should provide full and summary listings of each section, as well as measures indicating the quality of the partition. The system should then facilitate manual adjustment of the proposed sections as required.

After a successful section partition has been accomplished, the system should propose a set of groups for each section, again based upon diversity objectives as well as required and forbidden constraints. Here also, useful visual output is critical, as are quality measures and the ability to manually adjust the suggested solution. Overall the system must be easy to use, preferably menu driven, and designed for use on existing hardware and software platforms.

SYSTEM METHODOLOGY:

Data Handling and Presentation :

Input to MCADSS is a flat file comprised of data gleaned from the admission sheets. (The fields comprising the student file are presented in Figure 1.) MCADSS's first contribution to the process of section and group partition lies simply in its data handling and presentation capabilities. Information about the students, or selected subsets of students, may be sorted in a variety of ways, and printed out in an abbreviated or complete format. Summary statistics on the entire promotion are provided, i.e., number of British, number of engineers, number of French mother tongue, etc. (See Figures 2 and 3.) This summary provides the users with a starting sense of what a good partition will look like; for example, a preponderance of one subset of students, or a scarcity of another will suggest limitations to within section/group diversity.

Student Number
Name
Age (at entry to program)
Sex
Nationality
Background (one of: Science/Engineering, Economics, Law, Other)
Mother Tongue(s)
Universities Attended
Previous Employers/Locations

Figure 1: Data for each student provided as input to MCADSS

Total Number of Students in Promotion	
Average Age	
Number of: Females	French Mother Tongue
British Nationals	English Mother Tongue
French Nationals	German Mother Tongue
German Nationals	
Scandanavians	

Provided at User's Request:

Number of Other Mother Tongues
Number of Other Nationalities
Job Locations (number of students who have worked in each country)
Companies (number of students who have worked for each company)
Universities (number of students who have attended each university)

Figure 2: Summary data of the entire promotion provided by MCADSS at the start of the consultation

Students	231	English Mother Tongue	85
Average Age	28	French Mother Tongue	58
Females	37	German Mother Tongue	12
British Nationals	44	Scandinavian Mother Tongue	16
French Nationals	47	Science Background	130
German Nationals	5	Economics Background	61
Lawyers	15		
Others	25		

Figure 3: Sample Summary Statistics for Entire Promotion

Sections

At this point, partitioning into sections may commence. The first step is to specify individuals who must, or must not, be in the same section. This is done on a pairwise basis via the "Rules/Sections" menu option. First, individual pairs of students required to be in the same section are entered, then pairs of students forbidden to be in the same section are keyed in. The process is guided by system prompts. A listing of the complete set of forbidden/required pairs that have been provided may be obtained via a menu option.

The heuristic by which students are actually allocated to sections is based upon the manual method. The program takes one student and puts him/her into the first section. It then selects the student who is the most similar to the first one and assigns him/her to the second section. The program continues in this fashion, at each iteration taking the student it deems most similar to the previous student and placing him/her in the next section. (After a student is allocated to the last section, for example section three of a three section promotion, the "next" section is section one.) Before placing a student into a section, MCADSS checks the appropriate required and forbidden constraints to make sure the placement is permitted. If a constraint is violated, MCADSS attempts to place the student in the next section.

The measure of similarity between students is based upon the seven criteria provided in Figure 1: Age, Sex, Nationality, Mother Tongue, Background, University, and Company. Two instances of students are provided as an example in Figure 4.

	Student A	Student B
Age	30	25
Sex	Male	Female
Nationality	F	GB
Mother Tongue	French	English
Background	Science	Other
University	Ecole Polytechnique	Oxford
Company	Schlumberger	Bain & Co.

Figure 4: An example of two students data

Aside from age, each of these characteristics is assumed to be measured on a nominal scale. That is for example, being French is as different from being British as it is from being Belgian. MCADSS treats age as a variable; i.e., the difference of five years is five times the difference of one year. It would consider Student A and Student B to be 100% different as they differ on each criterion. How different would these two students be if they were of the same sex? The relative importance of each of the criteria in determining the difference (and hence the similarity) between students is determined by a user specified weighting scheme. If the criterion Sex had a weighting of 10%, then two students with the characteristics of A and B, but of the same sex, would be 90% different. If the weighting of Sex was 25%, then the students would be 75% different.

The program's default weightings, which have proved to give good results for the partitioning into sections, are given below.

Age:	0%
Sex:	5%
Nationality	25%
Mother Tongue	10%
Background	35%
University	20%
Company	5%

Arriving at these weights was an empirical process. Initial values were provided by the administrators, based on their experience and goals in partitioning the class into sections. These initial values were adjusted through a series of trials on historical data. (Discussions of the weighting scheme and of testing are provided in later sections.) The default weights perform well, given the current criteria for section partitioning, though they may be easily adjusted by the user.

After selection of the weightings, the user is prompted for the number of sections (currently three at INSEAD), and the starting section number ("1" for the September promotion, "4" for the January promotion). Lastly, as the system starts section allocation by placing a student in the first section, the user may choose the student with which to begin, or enter "0" to start the process with a randomly selected student. (Assuming that the distribution of students is clustered around a "typical" profile, the allocation process based on selecting a student with the most common characteristics provides a "better" partition. This is due to the fact that students with such characteristics can be well spread out across sections which (a) makes the allocation task easier to do than with other students profiles, and which (b) improves the quality of the obtained partition. It should be noted that the system is usually run several times, until the administrator is satisfied with the proposed solution.

When the sections have been partitioned, the user may select a menu option for a listing of the promotion by section, and sorted within section as desired (by name, nationality, etc.) For each section, summary statistics, as per Figure 2, are provided. Included as well is the average difference between all pairs of students in the section; this average is used as an overall measure of the partition quality.

The sorted output and summary statistics provide a basis for the administrators to evaluate the quality of the computer-suggested allocation. The menu allows manual adjustments to be made to the partition, by moving students from one section to another, or swapping students. Updated statistics are then provided. If a poor initial allocation is suggested, the process may be rerun with different weights, or by using a different starting student.

Section I

Students	75	English Mother Tongue	29
Average Age	28	French Mother Tongue	18
Females	12	German Mother Tongue	3
British Nationals	15	Scandinavian Mother Tongue	5
French Nationals	15	Science Background	42
German Nationals	1	Economics Background	20
		Lawyers	5
		Others	8

Average Difference = 70%

Section II

Students	75	English Mother Tongue	29
Average Age	28	French Mother Tongue	18
Females	10	German Mother Tongue	4
British Nationals	14	Scandinavian Mother Tongue	5
French Nationals	15	Science Background	43
German Nationals	2	Economics Background	21
		Lawyers	4
		Others	7

Average Difference = 70%

Section III

Students	75	English Mother Tongue	26
Average Age	28	French Mother Tongue	21
Females	13	German Mother Tongue	5
British Nationals	14	Scandinavian Mother Tongue	6
French Nationals	16	Science Background	42
German Nationals	2	Economics Background	19
		Lawyers	6
		Others	8

Average Difference = 71%

Figure 5: Sample Summary Statistics for Section Allocation

Groups

The process for allocating groups is similar to the one determining sections. The user specifies the weightings, and for each section the number of groups, the initial group number, and the starting student. There is however a major difference between the section and group allocation processes at INSEAD: diversity of the groups' composition is more rigidly controlled. As the number of students in each group is significantly smaller than the number in each section, slight deviations towards homogeneity within a group have a greater impact. Moreover, while the administrators could only present general guidelines for section diversity, they provided and insisted upon strict specifications for groups. These specifications entailed a set of forbidden conditions, such as "no more than three students of any one nationality in a group", and "no two students from the same university in any one group". MCADSS therefore incoporates these rules, some ten in all, into the group partitioning process. These restrictions in fact operate in an analagous manner to the "forbidden" pairs constraints in the section allocation process.

It should be noted here that if a specification forbids two students sharing the same characteristic from being in the same group, the weighting of that characteristic should be set to zero. Clearly, the system ensures maximum diversity based on that characteristic by means of the rule.

Again, MCADSS provides for an assortment of useful listings of the groups, as well as summary statistics. Figure 6 is a sample report of the summary values for two groups in section I.

Group 1

Students	7	English Mother Tongue	2
Average Age	27	French Mother Tongue	1
Females	1	German Mother Tongue	1
British Nationals	2	Scandinavian Mother Tongue	0
French Nationals	1	Science Background	3
German Nationals	0	Economics Background	2
		Lawyers	0
		Others	2

Average Difference = 63%

Group 2

Students	7	English Mother Tongue	2
Average Age	29	French Mother Tongue	1
Females	2	German Mother Tongue	1
British Nationals	2	Scandinavian Mother Tongue	1
French Nationals	1	Science Background	2
German Nationals	1	Economics Background	3
		Lawyers	0
		Others	2

Average Difference = 62%

Figure 6: Sample Summary Statistics for Group Allocation

After examining the listings and summary data, individual students may be moved, or pairs of students switched, in order to incrementally improve the groups allocation.

Testing and Evaluation

The system was developed and tried using students data of the two past years. Formal testing occurred prior to the beginning of the 1989 - 1990 academic year, using the September entering class data. The criteria for success, agreed upon at the outset of system development, were: 1) The partitions developed through the use of the system are as good or better than those obtained via the traditional manual process; 2) The use of the system results in significant time savings; and 3) Administrators/users of MCADSS are satisfied with the system.

Complete partitioning of the September entering class was completed by the two administrators and MCADSS in several hours - a significant savings over the four person-day manual procedure. Both for section and group partitioning, the administrators "tweaked" the computer-based solution suggested by MCADSS by moving few individual students. Their conclusion was that the quality of the solution achieved with the use of the DSS was at least as good as the one manually obtained. The administrators found the use of the system to be intuitive; while some instruction and guidance during the test were required, any further training was deemed unnecessary. The weighting scheme was the only source of some confusion (see the next section). Given the large multivariate data set and enormous solution space, MCADSS was seen as making the problem manageable through its data handling and display facilities, good proposed solutions along with measures of solution quality, and capability to modify those solutions according to the administrators' judgment. Overall, the administrators were pleased and impressed with the power and flexibility they achieved through the use of MCADSS. For the following January, only a single administrator working with the system satisfactorily completed the partitioning task in a matter of hours.

Weighting

The administrators were initially uneasy at assigning numerical scores to the characteristics. While in fact they implicitly weighted each criterion when manually partitioning the class, they found the explicit formulation of weights difficult. Part of this had to do with the well known problem of attempting to express qualitative judgments quantitatively ([3], pp. 209-292; [9]). This was compounded by their unfamiliarity with mathematics; for example, the idea of normalizing the weights, if necessary, to 100% seemed to them a peculiar notion.

However, part of their confusion lies with the weighting scheme itself. The following questions must be posed:

> What does a weight of, say, 20% for University really mean? That is, what does it mean to say that two students are 20% different?
>
> Is 20% different really twice 10% different?
>
> Is the 20% difference between two students of different universities the same as the 20% difference between two students who differ on Mother Tongue (10%), Company (5%), and Sex (5%)?
>
> How robust are the weights? (How sensitive are good solutions to variations in weights?)

Compounding the problem, the criteria are not independent. (For example, Nationality and Mother Tongue are clearly correlated.)

It should be recalled however, that related methodologies pose similar problems. Statistical clustering techniques, for example, rely on an arguably ad hoc choice of similarity measure and, "in most cases are not supported by an extensive body of statistical reasoning" ([1], p.14). The justification for MCADSS's weighting scheme is identical to that for the use of the "certainty factor" approach for representing judgment in expert systems ([3], p. 233-262). Though not firmly rooted in mathematical theory, the approach is both intuitively appealing and empirically sound, while superior alternatives do not exist. While cognizant of the limitations of the weighting scheme, it reflects the goal of the project - to develop an appropriate, heuristic-based DSS for the problem at hand.

PAPER SUMMARY:

MCADSS is a decision support system designed to aid administrators in an MBA program in allocating students to sections, and groups within those sections. The goal of the allocation process is to maximize diversity within groups and sections, and minimize diversity between groups and sections. MCADSS is menu driven, implemented in C, and runs on PC compatible microcomputers. The section allocation problem requires several minutes (real time) on an AT class machine; partitioning each section into groups takes about several minutes per section.

MCADSS works from a database of students and their salient characteristics; the partitioning heuristic it uses is a novel approach, inspired from the manual procedure. The system has been adopted for regular use at INSEAD. It could clearly be adapted for use in other academic or training (i.e., executive education) programs. Other areas for potential applications include portfolio planning and marketing strategy.

REFERENCES

[1] Aldenderfer, M.S. and Blashfield, R.K., Cluster Analysis, Sage Publications, 1984.

[2] Beheshtian-Ardekani, M. and Mahmood, M.A., "Development and Validation of a Tool for Assigning Students to Groups for Class Projects," Decision Sciences, vol. 17, pp. 92-113, 1986.

[3] Buchanan, B.G. and Shortliffe, E.H. (Eds.), Rule-Based Expert Systems, Addison-Wesley, 1984.

[4] Dhar, V. and Ranganathan, N., "Experiments with an Integer Programming Formulation of an Expert System," Working Paper CRIS #205, STERN #89-39, Center for Research on Information Systems, Stern School of Business, New York University, March 1989. (Also MCC AI Laboratory Technical Report ACA-AI-022-89.)

[5] Dyer, J.S. and Mulvey, J.M., "An Integrated Optimization/Information System for Academic Department Planning," Management Science, vol. 22, pp. 1332-1341, 1976.

[6] Harwood, G.B. and Lawless, R.W., "Optimizing Organizational Goals in Assigning Faculty Teaching Schedules," Decision Sciences, vol. 6, pp. 513-524, 1975.

[7] Hill, A.V., Naumann, J.D. and Chervany, N.L., "SCAT and SPAT: Large-Scale Computer Based Optimization Systems for the Personnel Assignment Problem," Decision Sciences, vol. 14, pp. 207-220, 1983.

[8] Lee, S.M. and Schniederjans, M.J., "A Multicriteria Assignment Problem: A Goal Programming Approach," Interfaces, vol. 13, pp. 75-81, August 1983.

[9] Zadeh, L., "A Theory of Approximate Reasoning", in R.R. Yager, S. Ovchinnikov, R.M. Tong and H.T. Nguyen (editors): Fuzzy Sets and Applications, New York: John Wiley & Sons, pp. 367-412, 1987.