Maximizing workforce diversity in project teams: a network flow approach

By: Joyendu Bhadury, E. Joy Mighty, and Hario Damar

Bhadury, J., E.J. Mighty and H. Damar "Maximizing Workforce Diversity In Project Teams: A Network Flow Approach." Published in *OMEGA: International Journal of Management Science*, 28, pp 143-153 (2000).

Made available courtesy of ELSEVIER:
http://www.elsevier.com/wps/find/journaldescription.cws_home/375/description

*****Note: Figures may be missing from this format of the document**

Abstract:
Several social, economic and political factors have contributed to the increasing diversity of today's workforce. In addition, in an era when organizations are continuously redesigning their work and restructuring their operations to achieve their goals with fewer resources, performing work in teams has become commonplace. These trends have increased the need for managing diverse work teams effectively. There are several existing models in the management science literature that help managers to assign employees to work groups in order to maximize the groups' diversity and hence, facilitate their effectiveness. This paper introduces a new model that recasts the problem of managing diversity in a different way: it is assumed that the population comes partitioned into `families' with a high degree of intra-familial similarity and inter-familial dissimilarity. The objective of the assignment then is to disperse these family members as evenly into the workgroups as possible. A little known network flow problem, known as the dining problem, is used to develop an efficient algorithm to produce solutions to this new model. This is followed by a report on an experimental application of the developed model to assign Master of Business Administration students in a business school to different projects in a course. As a part of this empirical report, an attractive feature of this model is also demonstrated; namely, how to conduct sensitivity analysis to determine the optimal levels of diversity in the presence of resource constraints. Finally, the paper concludes by discussing limitations of this new model and how they may be addressed in future research on this topic. Keywords: Decision making; Personnel/human resources; Manpower planning; Set partitioning; Optimization; Dining problem

Article:
1. Introduction
Among the many environmental trends affecting organizations in the 1990s is the rapidly changing composition of the workforce, a phenomenon known as workforce diversity. The word `diversity' refers to differences in a range of human qualities among individuals [17]. The traditional view of an organization characterized by workforce diversity is one in which there are increasing numbers of nondominant or minority social groups based on gender, race, ethnicity or nationality, resulting in heterogeneity in socio-cultural perspectives, world views, life styles, language and behavior [17]. However, recent approaches have also attempted to extend the concept of diversity to include other factors besides race, gender and ethnicity. As Thomas, Jr. argues in [25], employees differ on a variety of other dimensions such as age, functional and

educational backgrounds, tenure with the organization, lifestyles, geographic origins and many others.

Increasing workforce diversity is a reflection of several social, economic and political changes. Global interdependence and multinational corporations are becoming commonplace and free trade agreements and joint ventures often bring individuals from very different races and cultures into close contact within organizations. Moreover, recent changes in the international political arena and in national boundaries have increased the movement of workers across countries. For example, in North America (USA and Canada), more women, people of color and immigrants are participating in the labor force and it is estimated that by the year 2000 they will account for over 80% of all new recruits into the labor force (Bureau of Labor Statistics, USA and Statistics Canada). Another factor influencing the increase in workforce diversity is a legislative framework at all levels of government. Affirmative action programs in USA, the Employment Equity Act in Canada and the Federal Contractors' Program in these two countries are examples of government policies that require the establishment of employment equity programs.
As for the effects of increasing diversity on organizational effectiveness, several studies have suggested that workforce diversity can have both positive and negative impacts on organizations (e.g. [4]). However, the nature of the impact depends to a large extent on the type of diversity climate that exists rather than the fact of diversity itself [7,11]. Proponents of workforce diversity have argued that its advantages include enhanced abilities to deal more sensitively with multi-cultural domestic and foreign customers, thereby increasing customer satisfaction, keeping and gaining market share and attracting and retaining the best personnel. In addition, diversity may increase organizational effectiveness through improved decision- making and problem-solving, enhanced creativity and innovation, better quality of management, an increased ability to adapt to environmental change more effectively and reduction in costs related to turnover, absenteeism and legal action [7,11]. Further, when diversity within organizations is extended to include academic backgrounds and work experiences, managing it has become particularly necessary in organizations that rely heavily on work teams. For example, consulting firms that offer a full range of professional services often structure their workforces into teams comprising representatives of various professions. Similarly, in the context of continuous downsizing and restructuring operations, many government departments which offer different services to a wide variety of different people have recognized the need for multi-skilled, self-managed work teams to achieve organizational goals while utilizing fewer resources.

However, although much has been written in the organizational theory literature about the need to change organizational structures and systems for the effective management of diversity there has been remarkably little delineation in that literature of specific methodologies for making such changes in a rational, objective and nondiscriminatory manner. One notable exception is the development in [13] of two statistical models for setting numerical goals for the employment of women and minorities and for monitoring compliance with employment equity hiring goals once they have been established. These models are useful for planning and evaluating employment equity policies, but they do not help managers with practical aspects of managing diversity such as making decisions about assignments to work groups.

In sharp contrast, there has been no dearth of quantitative models in the management science literature inthe past decade that are intended to be used by managers who wish to manage diversity. In all of these models, the fundamental behavioral assumption has been that requiring diverse people to work together as a group increases their understanding of each other's backgrounds and modus operandi, thereby leading to better communication between diverse groups. In other words, the best way to manage diversity is to let people with different backgrounds work in close quarters with each other. Support for this hypothesis is provided by the fact that heterogeneous groups are known to perform better than homogeneous groups on creative, decision-making and problem solving tasks[22,26]; thus maximizing the diversity of work teams is highly desirable. Notwithstanding the reasons for this assumption, its implication in terms of mathematical formulation is that these models are fashioned along the same lines as the well known equitable partitioning problem (see [21]). Perhaps the earliest such model is the one in [27] which used a heuristic to allocate students to groups at the well-known European business school INSEAD; an adaptation of this was also applied at New York University [16]. In [19], the researchers suggested a heuristic and a Goal Programming model do similar allocations at Warwick Business School in the United Kingdom. Other studies that have used Integer or Mixed Integer Programming models to allocate groups with diversity as a consideration, include [9,12,14] and [23]. Similar models have also been studied in seemingly different contexts, namely VLSI design in [10] using a graph theoretic approach and scheduling in [15] where a heuristic is used. [28] contrasts graph theoretic and heuristic approaches and [29] compares different heuristic approaches.

One difficulty with all the models above is that all of them formulate the problem of group allocation as difficult, i.e. NP Hard, combinatorial optimization problems. As a result, optimal solutions to large problems are difficult, if not impossible, to obtain. Hence, almost all these models present heuristics for their respective models. The disadvantage with heuristics, of course, is that optimal solutions cannot be guaranteed. This necessitates recasting the problem in a different framework, so that efficient (i.e. polynomial time) exact algorithms can be developed. That is the point of departure of our paper.

In an effort to do so, close inspection will reveal that the NP Hardness of the above models is due to the excessive detail in the input; each individual member of the population is classified in terms of diversity with respect to every other and the objective function of these models typically tend to optimize the collective sum of these individual pairwise contributions to the diversity. This, in turn, frequently leads to a quadratic objective function, which, in conjunction with the fact that the decision variables are binary, tends to make these problems difficult to solve. Computational issues aside, such numerical assignment of diversity scores to each pair in the population has other, practical problems; as noted in [27]. It must be said, in all fairness however, that every model, including ours, that attempts to work quantitatively with qualitative data, has similar problems. What is interesting however, is that despite their limitations, evidence suggests that they outperform manual or random assignments (see [27,29]).

Therefore, in an attempt to address some of these problems, we propose a new model for assigning workgroups/project teams which adopts a completely different approach. We begin with the assumption that the population has been divided into several `families' where there is little diversity within the members of any family but a significant amount of diversity between the different families themselves. The families are supposed to be different enough so that small differences in inter-familial diversities can be ignored. In other words, when viewed in relation to the similarities that exist within each of the families, any two families appear to be `equally' different. This a priori grouping of the individuals into `similar' families, which is reminiscent of data agglomeration techniques used in statistics, is the reason why we are able to develop efficient exact algorithms for our model. However, this comes at a price; it restricts the applicability of the model to cases where these families are easily distinguishable on the basis of very few (preferably one or two) objective characteristics. We illustrate this with three examples. First, in case of assignment of families to public housing in large urban centers in USA, these could (and frequently are) defined on the basis of race/ ethnicity: Black/White/Asian/Hispanic. Second, in multilingual countries such as Canada, Switzerland, India, etc, when selecting members of national committees that affect language issues, these families would be given by the different linguistic groups; for example, anglophones and francophones in Canada. Finally, in an academic context, consider the allocation of students to project groups in an MBA (Masters In Business Administration) curriculum; here, the academic background in the baccalaureate degree of a student (engineering/arts/business) would define the basis for the families.

Notwithstanding the basis for defining the families themselves, the objective of our model then is to assign workgroups so as to disperse the family members as much as possible. This is achieved by recasting the problem as a little-known problem in network flow theory, namely, the dining problem. We then illustrate how this problem can be solved efficiently (i.e. in polynomial time) by repeated applications of the well known max-flow problem. Finally, we demonstrate the working of our model with an experimental study on a simple dataset constructed from the MBA students at a business school; however, we caution that due to its academic setting, the stringent condition of strong inter-familial diversity and intra-familial similarity was not entirely met in our dataset; hence, the study is meant solely to illustrate an application of our model. Nonetheless, through this numerical illustration, we also exhibit another novel and attractive feature of our model; namely, that it lends itself easily to sensitivity analysis, thereby allowing the decision maker to define and choose the `optimal' level of diversity when resources are constrained.

The remainder of the paper is divided as follows. The next section discusses the theoretical concepts needed for the paper, the dining model itself and our proposed efficient solution methodology for it. Section three then details the experimental implementation of the model including a discussion on how to perform sensitivity analysis. Finally, the fourth section distils the conclusions of our study, its limitations and how to address them in future research.

2. The quantitative model

The development of the quantitative model presented in this paper requires some concepts from the area of network flows in management science, (see [2]). We therefore discuss these relevant concepts before presenting the model.

2.1. Max-flow problem

This is one of the most well known and extensively studied models in management science. The max-flow problem is posed as follows. Consider a graph G=(V, A) with V={v1, v2, v3, . . ., $v_{|V|}$} being the set of vertices of G and A representing the set of a total of $|A|$ arcs that connect them, where we will assume that (i ,j) represents the arc going from node i to j. Assume further that one of the nodes, say v1, is designated as a source node from where a certain commodity has to be shipped to another node, say vn, which is referred to as the destination node. Associated with each arc (i, j) is a positive number c_{ij} that represents the maximum capacity of that arc; in other words, a maximum of c_{ij} units of the commodity can be shipped through the arc (i, j) at any point in time. As an example, the nodes could represent the cities in a certain geographical region with the arcs denoting the set of roads connecting them. The arc capacities can then be supposed to represent the maximum amount of traffic that can be accommodated on each road. Given these capacities however, the max-flow problem is to find the maximum number of units that can be shipped from the source to the destination, without exceeding any of the arc capacities. As mentioned before, the max-flow problem is a classical one that has been extensively studied for quite some time now and very fast, polynomial time, algorithms exist for solving it. The two fastest algorithms to date are those in [9] and [3]. The former runs in $O(|V|^3 /\log|V|)$ time and the latter in $O(|V|\,|A|)$ time when $|A|$ r $|V|^{5/3}\log|V|$ and $O(|V|\,|A|\log|V|)$otherwise.

2.2. Dining problem and its application to assigning diverse workgroups

Our use of the max-flow problem will be indirect and only in the context of a lesser known network flow problem that is referred to as the dining problem [2], p. 198]. To the best of the authors' knowledge, the dining problem has no known prior application; however, it belongs to a celebrated class of problems that go by the name of set partitioning problems, that are very well studied combinatorial optimization problems (see for example [20]). In particular, other problems similar to the dining problem, that deal with the issue of `equitable' partitioning (see, for example [21]) have been studied in a wide variety of contexts such as: political districting [18], crime-based districting [24], highway patrol districting [1] etc. However, since the dining problem is crucial to our proposed model and is not well studied in the literature, we will discuss it in detail and also propose a solution procedure for it. Having done that, we will also show how the dining problem can be applied in a straightforward way to develop a model that helps a manager allocate workgroups in his/her organization.

In the dining problem, we are given n families, that are going out to dinner. It is assumed that family i has a(i) members, with A being the total number of family members to be seated (i.e., $A = a(1) + a(2) + ... + a(n)$). The restaurant that they wish to go to has m tables, 1, 2, 3,...m and the seating capacity of table j is given by $c(j)$. To eliminate trivial instances of the problem, it is assumed that the total seating capacity of the restaurant is at least as large as the total number of family members, i.e.,

$$a(1) + a(2) + \ldots + a(n) \leq c(1) + c(2) + \ldots + c(m).$$

In order to maximize social interaction, the families wish to develop a seating arrangement where it is ensured that at each table there is representation from as many different families as possible. To that end, the dining problem seeks to find a seating arrangement with the property that the maximum number of members from the same family that are seated at any one of the T tables in this arrangement is as small as possible. To get an alternate definition, define the *diversity index* of any given seating arrangement, denoted by k, as the maximum number of members from any one family at any one of the tables in the given seating arrangement. Then the objective of the dining problem is to find a seating arrangement with the minimum diversity index, i.e., smallest value of *k* — and we denote this minimum diversity index as *k**. Obviously, for any given instance of this problem, *k** is at least unity and any seating arrangement that has a diversity index of unity ensures maximum possible representation at each table as it guarantees that no two members from the same family are seated at the same table.

Although the dining problem is a hypothetical one, it can be applied in a corporation to devise a quantitative model that enables a manager to allocate work- groups with diversity into consideration. An effective way to accomplish this in any organization is to assign people to different teams or projects in a manner that ensures everybody an opportunity to work with as many `different' kinds of people as possible. Thus, if an organization has different `groups' based on any criteria of diversity such as gender, ethnicity, educational backgrounds or technical capabilities, that are to be assigned to different teams, workgroups or projects, the dining problem can be effectively used by a manager to assign people from these groups to the various teams/projects/workgroups in a manner that maximizes the diversity within each team. For that purpose, the manager need only solve the dining problem discussed above with the different groups representing the families and the teams/projects/workgroups denoting the tables in that setting.

2.3. Solution procedure for dining problem

We will now propose a simple solution procedure for the dining problem. To do so, note that although the two problems mentioned may seem different, the dining problem can be formulated and solved as a max-flow problem and this is the solution procedure that we use. In order to understand this formulation, assume that we have been given the same instance of the dining problem as above and seek to find a seating arrangement whose diversity index is at least *k*, where k is an integer that is at least unity. Then we will solve this problem by first constructing the graph G shown in Fig. 1. The set of nodes in G are constructed so as to fall into the following four groups:

1. The source node *S*.
2. The destination node *D*.
3. *n* family nodes, *F(1), F(2), ... , F(n)*, where each node *F(i)* is supposed to represent family i.
4. *m* table nodes *T(1), T(2), ..., T(m)),* where each node *T(j)* is supposed to represent table *j*.

Thus there are a total of ($m+n$ + 2) nodes in *G*. The arcs in *G* are then constructed so as to be classified into the following three groups.

1. The *family arcs* that go from the source node S to each of the family nodes *F(i*). The capacity of the arc (*S*,*F(i*)) is assumed to be a(*i*), i.e., the total number of members in family *i.* Thus there are a total of n family arcs.
2. The table arcs that go from each of the table nodes *T(j)* to the destination node *D*. The capacity of each such arc (*T(j), D*) is assumed to be c(*j*), the total capacity of table *j*. Thus there are a total of m table arcs.
3. The *seating arcs* that run from each family node *F(i)* to each table node T(j). The capacity of all the seating arcs is fixed at *k*. Therefore, there are a total of *mn* seating arcs.

Hence there are a total of ($n+m+nm$) arcs in G. In this formulation, the `commodity' to be shipped is represented by the people who need to be seated at the restaurant and every unit of flow on any arc represents a person from a given family. Given this interpretation, note that (i) fixing the capacity of each seating arc at k, ensures that no more than k members from any given family are seated at the same table. (ii) The upper bound of a(i) on each family arc (S, F(i)) symbolizes the fact that family i has no more than a(i)

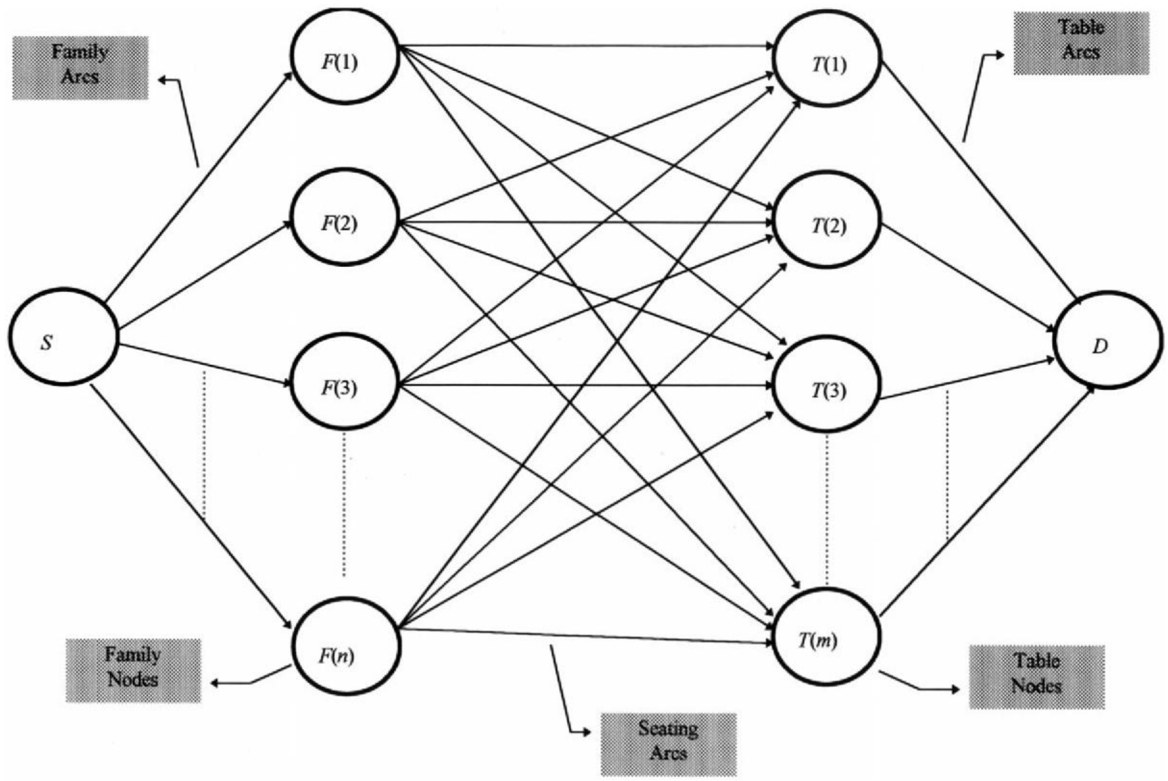

Fig. 1. The graph *G*.

members to seat and (iii) the upper bound of c(j) on each table arc (T(j),D) ensures that the seating capacity of that table is not exceeded. Thus the maximum amount of the `commodity' that can be shipped from S to D in G represents the maximum number of family members that can be seated if we require that no more than k members of the same family be seated at the same table.

Hence, the next step after constructing G is to solve for the maximum flow from S to D and check if it is A units. If this is the case then the desired seating arrangement can be easily obtained from the optimal flow in the following way. Consider any family node, say F(1) — since the maximum flow is found to be A, all the a(1) members must be seated, i.e., the total flow on the family arc (S, F(1)), which also represents the total flow emanating out of node F(1), is equal to a(1). Choose the first seating arc from this node, i.e., (F(1), T(1)) — if the flow on this arc is l units, then it implies that l members from the first family should be seated at table 1. Now proceed to arc (F(1), T(2)) to find out how many family members from the first family should be seated at table 2 and so on, until a table assignment is found for each member in the family. Proceeding in this manner for each of the families, we can find the desired seating arrangement for all the family members. Thus it can be claimed that each feasible flow in G, that transports A units from S to D, represents a seating arrangement for the families with a diversity index of k.

Remember however, that in order to solve the dining problem we need to find the minimum diversity index k*. In order to do so, note that k" is at least unity. Further, it is also trivial to see that it is no larger than the number of members in the largest family, say p. Thus we can find the value of k" by iterating on all values of k between 1 to p. In other words, a simple solution procedure for the dining problem is as follows. Begin with a value of k = 1 and solve the max-flow problem described above on G, with this value, i.e., with the capacities of all seating arcs fixed at unity. If doing so still allows A units to be transported from S to D, then we stop with k* =1. Else, we know that k* is at least two and we repeat this process with k = 2. The smallest value of k that ensures a maximum flow of A units from S to D is then the optimal value k*. It is this idea that is summarized in the form of the following algorithm that we propose to solve the Dining problem.

Algorithm (Dining problem).

begin

Step 1. Set k equal to 1.

Step 2. Construct G as discussed and set the capacity of all seating arcs in G to k.

Step 3. Solve Max Flow Problem on G.

Step 4. if (maximum flow in Step 3 above is found to be A units) then

{

Step 4.1 The optimal seating arrangement has been found. Construct it from the optimal flow the manner discussed.

Step 4.2 Set the Optimal Diversity Index k" equal to the present value of k.

Step 4.3 Stop the algorithm.

}

else

{

Step 4.4 Increment k to the next higher integer.

Step 4.5 Go to Step 2.

}

End

Note that the algorithm described above can be considerably embellished by introducing more complex techniques for the theory of algorithms, such as binary search, or by using sophisticated data structures to store and retrieve the data. Since the focus of this paper is on showing an application of the dining problem in assigning workgroups rather than developing algorithms for it, we do not discuss these refinements here. As an aside however, we note that by modifying the above algorithm to perform a binary search on k, no more than O(logp) Max-Flow problems would have to be solved. Hence, using a fast algorithm for the max flow problem, such as the one in [9], an instance of the dining problem could be solved in $O((n^3 m^3 \log(p))/\log(n + m))$ time. Thus, from the computational standpoint, the Algorithm Dining-problem described above is efficient, i.e., its running time is polynomial in the size of the input.

3. An experimental implementation

One of the potential applications of the model introduced in the previous section can be found in allocating Master of Business Administration (MBA) students in a business school to different projects in a course. MBA students are typically characterized by a tremendous diversity in terms of their prior educational background. It is not uncommon to have students with undergraduate degrees in engineering, arts, pure sciences and other fields, all within the same graduating class. Further, team work, in terms of class projects and group seminars, is also very important in the MBA curriculum. As one of the fundamental aims of any MBA program is to produce managers capable of working with people of different backgrounds, it is particularly desirable to assign students to groups in a manner such that there is sufficient heterogeneity within each group.

Therefore our model was applied to a batch of 90 MBA students in the Faculty of Administration at the University of New Brunswick, Fredericton, Canada. It was decided to divide the students into different groups based on (i) their educational backgrounds (arts, sciences, engineering, business, nursing, kinesiology and other) and (ii) their gender. In keeping with the original notation of the dining problem, these groups were referred to as the `families'. Thus this batch was divisible into 14 such `families', based on the heterogeneity of their backgrounds and their genders. However, a cautionary note here; due to its academic setting, our dataset did not fully meet the stringent condition of strong inter-familial diversity and intrafamilial similarity that is required by our model. Thus the study is intended primarily to serve as a platform to demonstrate the working of the model.

The objective of our model in this case was to enable a professor to develop an assignment of students to respective class projects such that there is a desirable mix of a variety of students that work on a project. Examination of the notation employed in the original Dining Problem easily reveals that in this case, the `tables' at which the families are seated are represented by the class projects to which individual students are assigned. Thus, in this application the decision maker is faced with an optimal choice of three decision variables.

1. k, the diversity index: this represents the maximum number of members from each family that are permitted to be seated at the same table. The minimal value of this decision variable is unity and the higher the diversity index of any assignment, the less desirable it is from the viewpoint of maintaining diversity in the assignments. However, as our results show, attempting to achieve a high level of diversity within each project group, without the availability of sufficient resources, may lead to infeasible and/or impractical assignments. For example, stipulating that the diversity index always be one, may lead to highly undesirable values of decision variables 2 and 3 (say, an unacceptably high or low number of class projects or numbers of students in each project) leading to impractical assignments. Thus, it may be desirable to compromise on this value slightly to get better assignments that are more easily implementable. We will demonstrate a method to make such compromises rationally.
2. Number of `tables' (i.e. class projects): for experimental purposes, it was decided that the value of this variable should be no fewer than 4 and no larger than 15. However, for a real application, this number would be decided by the instructor(s) teaching the course, depending on factors such as the total enrollment, the complexity of the projects and the method of grading them. Note that this would imply that in the case of our data set, the

number of tables would have to be at least 14, the total number of families, in order to be able to achieve the minimum diversity index of unity; this fact is also borne out by our results. This provides a quick check for the minimum resources needed in order to achieve maximum diversity.

3. Maximum number of `family members' allowed at a `table' (i.e. the maximum number of students allowed in each class project): again, it was decided that for the purpose of our experiment, we would (i) assume the same upper limit for each class project and (ii) vary this upper limit from five to 25. Once again, for a real application, these assumptions and the final range would also be decided by the individual professor depending on course and class specific factors.

Our model was formulated and solved using the standard optimization software, LINDO. The hardware platform used was a 33 MHz, Intel 486 based PC. Despite the fact that this platform is outmoded and slow by present standards and that we used Linear programming formulations to solve all of our max- flow problems, we found that the execution times of all the runs was under 5 s; a testimony to the efficiency of our model. In light of this fact that the execution times were so small, CPU times were not recorded. A representative sample of our results is summarized in Table 1.

Consider the first row of Table 1 when the total number of class projects is only 4. In this case, requiring an optimal diversity index of unity would lead to only 47 of the 90 students being assigned with an allowable maximum of 23 students in each project. This is clearly an infeasible solution, as not all students are assigned. In fact, for the same number of projects, the diversity index would have to be four, with an upper limit of 23 in each project, for all students to be assigned. Thus, we observe that, Observation 1. The maximal level of diversity is not always achievable if constraints are put on the resources available such as the total number of class projects or the maximum allowable enrollment in each project. Hence, given a fixed number of resources, it might be necessary to trade off on the diversity index in order to achieve feasible and/or implementable schedules. For example, if only 8 class projects are allowed (row five of Table 1), then a diversity index of 2 is the best that

Table 1
Sample results from the experimental study

Total number of projects available	Maximum permissible enrollment in any project	Maximum number of students that can be assigned to projects			
		diversity index = 1	diversity index = 2	diversity index = 3	diversity index = 4
4	23	47	73	86	90
5	18	55	80	90	90
6	15	61	86	90	90
7	13	69	89	90	90
8	12	73	90	90	90
9	10	77	90	90	90
10	9	80	90	90	90
11	9	83	90	90	90
12	8	86	90	90	90
13	7	88	90	90	90
14	7	89	90	90	90
15	6	90	90	90	90

is achievable with a maximum of 12 students being assigned to any class project.

It is clear that if we need to assign all students to the projects (which would be realistic in any real-world application) then one will have to compromise on larger values of the Diversity Index if the number of projects available is held constant. On the other hand, if the maximum enrollment in any project is held constant, then the minimum Diversity Index necessary to accommodate every student can be decreased with an increase in the number of projects. Sensitivity analysis with our model is capable of producing such a tradeoff curve. For example, Fig. 2 expresses this trade off in the form of a plot of the number of projects available vs. minimum diversity index necessary to assign all students. As is evident from the curve, with 4 projects, the minimum achievable diversity index is 4. Between 5 and 7 projects, the best possible diversity index is 3. When the total number of projects is between 8 and 14, the best possible diversity index achievable is 2 and above 15 projects, it is possible to achieve the optimal diversity index of unity, with a maximum of 6 students being allowed in any project group. As expected, the maximum number of students that are allowed to be enrolled in the same class project decreases with the increasing number of class projects and larger values of diversity index. Thus, we conclude that,

Observation 2. Using the sensitivity analysis procedure described above and the resulting trade-off curve, it is therefore possible for a decision maker, in this case, the individual instructor, to choose an assignment that has an implementable number of class projects and enrollment limits and at the same time achieves a desirable level of diversity. If there is a constraint on the availability of resources, (for example, if the instructor is allowed only a fixed number of class projects), the same curve can also be used to decide the `optimal' assignment that maximizes diversity subject to these resource constraints. Such analysis is of particular use in industry applications, where resource constraint is an everyday reality.

We conclude this section with the observation that

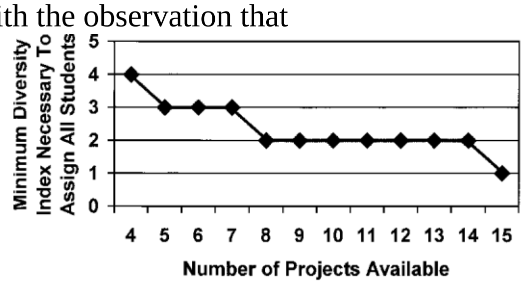

Fig. 2. Trade-off curve.

the inability to achieve optimal diversity may actually be beneficial for organizations. As pointed out in [8, p. 51], ``decision quality is best when neither excessive diversity nor excessive homogeneity are present''. Although diversity among group members reduces the probability of `groupthink', brings different perspectives to bear on problem solving and ultimately increases team creativity and innovation, it is desirable to have diversity balanced by some core of similarity among group members. Having a team characterized by diversity share some core organizational value reduces the potential for conflict associated with excessive diversity and increases the probability of a coherent focus on organizational or group goals. This is consistent with research findings in [6] that workgroups assigned on the basis of sociometric criteria (i.e.

some perceived similarity in background, values or opinions) report higher levels of communication, coordination, peer ratings, group cohesion and job satisfaction than ability-based workgroups (i.e. assigned on the basis of ability to perform a task). This implies that in considering the diversity criteria for assigning employees to groups, managers should ensure that some commonality, based on a core organizational value, exists. Thus, given the capacity for making trade offs between optimal diversity and feasibility, another value of our model lies in its ability to allow managers to create work groups with sufficient diversity to reap the inherent benefits while reducing the problems of both excessive diversity and excessive homogeneity.

4. Conclusions, limitations and future research

Diversity of workforce, in terms of ethnic backgrounds, cultures, work experience etc, is an increasingly common phenomenon in most of organizations today. One way of dealing with this heterogeneity is to make people of different backgrounds work on common projects so as to facilitate understanding and communication between them. In this paper we have introduced a new model to accomplish this. In contrast to the existing models in the literature, our model assumes that the population is classified into `families' where individuals within a family are `similar' with respect to the diversity criterion being used but are very different from individuals in other families. This enables us to formulate our problem as a network flow problem, for which and an efficient exact algorithm can be developed; this is also in contrast to the existing methodologies in the literature, which are mostly based on heuristic methods that do not guarantee optimality. An illustrative study with this algorithm was also conducted on a dataset constructed from MBA students at a business school. Finally, we also demonstrated another useful feature of the model; namely, that it can be used to conduct sensitivity analysis that helps a manager choose the `optimal' level of diversity in the presence of resource constraints.

Almost every model that attempts to deal quantitatively with qualitative problems has limitations; ours is no exception. In the following we mention the two major ones. First and foremost is that the performance of this model rests strongly on the definition of the families. In order for the model to be applied usefully, the population of individuals should lend itself meaningfully to division into different groups with the property that there is a high degree of similarity within groups and large dissimilarities between the groups themselves. Further, since our model merely disperses the family members, each pair of families is implicitly assumed to be `equally' different. In practice this would imply that the `families' are so different with respect to each other that when viewed in comparison to similarities that exist within each family, differences in inter-familial dissimilarities can be ignored. Whereas this assumption about the existence of these families allows the development of an efficient algorithm to solve our model, it comes at a price. As we mentioned in the introduction, such requirements automatically restrict the applicability of the model. Further, it makes our model difficult to compare directly to the existing models for benchmarking.

A second drawback of our model is that, like most other existing models, it has also been tested only in the context of student groups rather than in `real' organizations. We must therefore be cautious in generalizing the usefulness of our model to the workplace before more extensive tests are done on real life data- sets obtained from actual organizations. Nonetheless, the model does appear applicable to any type of work group, in contrast to some other models that exist in the management literature; for example, the study in [5] only investigates the impact of diversity on the performance of specialized (new product) groups.

Finally, another word of caution regarding the applicability of our model (and to some extent, any of the other ones existent in the literature): despite its obvious usefulness in helping managers to assign employees to work groups, the dining model should not be regarded as a panacea for the effective management of diversity. Managing diversity effectively involves a wide range of procedures for creating an organizational culture and building systems and practices that unite different people in a common pursuit without undermining their individual differences. While it is beyond the scope of this paper to explore fully other procedures and processes for managing diversity, our model is not to be used in isolation, but as one tool that can help in the complex process of successfully managing diversity.

Future research on this topic should address the two limitations above. In particular, one interesting strand of future research would be address the issue of how to define families in any given population and then, test to see if they are sufficiently homogenous with respect to their own members and heterogeneous with respect to those in other families. Perhaps statistical tools such as cluster analysis or modifications thereof will be of use to us in this regard. This also ties into another interesting problem for future research: devise experimental setups that will be able to compare our model with the existing ones for benchmarking. The primary obstacle there is to be able to adapt the data- sets available for the existing models for use by our model. We suggest one potential, albeit convoluted, approach in that regard. Typically, in these data sets each individual is characterized by several, measurable attributes of diversity. Thus families of `similar' individuals could be identified through repeated application of cluster analysis techniques. At each step, it needs to be verified if the families identified are `sufficiently different' from each other and if such `differences' are `approximately equal'; perhaps statistical test similar to analysis of variance (ANOVA) may be of assistance to us in that regard. Notwithstanding the actual methods used to identify these differences, if `sufficient differences' are not found between the families identified, data agglomeration would have to be done and the process repeated on the agglomerated data set. However, it is not difficult to see the potential difficulties of such an approach.

Finally, we mention that none of the models, including ours, have been tested on any industrial dataset. An interesting avenue for further research is to actually conduct such an empirical study and determine the effectiveness of the different models by measuring the effects of workgroups assigned through these methods on multiple outcomes, including performance, group coordination and cohesiveness, and job satisfaction.

Acknowledgements
The authors wish to gratefully acknowledge the efforts of two anonymous referees whose comments and suggestions have substantially improved the form and content of this paper.

References

1. Agnetis A, Pacifici A, Mirchandani PB. Partitioning of Bi-Weighted Trees. 1998. Presented at INFORMS Conference, Seattle.
2. Ahuja RK, Magnanti TL, Orlin JB. Network flows: theory, algorithms and applications. Englewood Cliffs, NJ: Prentice Hall, Inc, 1993.

3. Alen N. Generating pseudo-random numbers and the max-flow algorithm. Information Processing Letters 1990;35:201-4.
4. Alderfer CP, Alderfer CJ, Tucker L, Tucker R. Diagnosing race relations in management. Journal of Applied Behavioral Science 1980;16:135-65.
5. Ancona DG, Caldwell DF. Demography and design: predictors of new product team performance. Organization Science 1992;3(3):321-41.
6. Colarelli SM, Boos AL. Sociometric and ability-based assignment to workgroups: some implications for personnel selection. Journal of Organizational Behavior 1992;13:187-96.
7. Cox T. Cultural diversity in organizations: theory, research and practice. San Francisco, CA: BerrettKoehler Publishers, Inc, 1993.
8. Cox T, Blake S. Managing cultural diversity: implications for organizational competitiveness. Academy of Management Executive 1991;5(3):45-56.
9. Cheriyan J, Hagerup T, Melhorn K. Can a maximum flow be computed in o(nm) time? In: Proceedings of the 17th International Colloquium on Automate, Languages and Programming, 1990. p. 235±48.
10. Feo T, Khellaf M. A class of bounded approximation algorithms for graph partitioning. Networks 1990;20:181- 95.
11. Fernandez JP. Managing a diverse work force. Lexington, MA: Lexington Books, 1991.
12. Hill AV, Naumann, Chervany NL. SCAT and SPAT: large scale computer based optimization systems for the personnel assignment problem. Decision Sciences 1983;14:207-20.
13. Hopkins SP. Models for affirmative action planning and evaluation. Management Science 1980;26(10):994- 1006.
14. Kuo CC, Glover F, Dhir KS. Analyzing and modeling the maximum diversity problem by zero-one programming. Decision Sciences 1993;24:1171-85.
15. Lotfi V, Cerveny R. A final exam scheduling package. Journal of Operational Research Society 1991;42:205-16.
16. Lakshminarayanan S, Weitz R. Creating groups of maximum diversity: contrasting math programming and heuristic approaches. Working paper, Stillman School of Business, Seton Hall University, South Orange, NJ 07079, 1994 (revised January 1996).
17. Loden M, Rosener JB. Workforce America: managing employee diversity as a vital resource. Homewood, IL: Business One Irwin, 1991.
18. Mehrotra A, Johnson EL, Nemhauser GL. An optimization based heuristic for political districting. Management Science 1998;4(8):1100-14.
19. Mingers J, O'Brien FA. Creating student groups with similar characteristics: a heuristic approach. OMEGA 1995;23(3):313-21.
20. Nemhauser GL, Wolsey LA. Integer and combinatorial optimization. John Wiley and Sons, Inc, 1988.
21. O'Brien FA, Mingers J. A heuristic algorithm for equitable partitioning problem. OMEGA 1997;25(2):215± 23.
22. Shaw ME. Group dynamics: the psychology of small group behavior. New York: McGraw-Hill, Inc, 1981.
23. Reeves GR, Hickman EP. Assigning MBA students to filed study project teams: a multi-criteria approach. Interfaces 1992;22(5):52-8.

24. Sarac A, Batta R, Bhadury J, Rump CM. Reconfiguring the police reporting districts in the city of Buffalo. Working Paper, Dept. of Industrial Eng., SUNYBuffalo, NY 14260, 1998.
25. Thomas Jr. RR. Beyond race and gender: unleashing the power of your total work force by managing diversity. New York: Amacom, 1991.
26. Watson WE, Kumar K, Michaelson LK. Cultural diversity's impact on interaction processes and performance: comparing homogeneous and diverse task groups. Academy of Management Journal 1993;36:590- 602.
27. Weitz RR, Jelassi MT. Assigning students to groups: a multi-criteria decision support system approach. Decision Sciences 1992;23(3):746-57.
28. Weitz RR, Lakshminarayanan. An empirical comparison of heuristic and graph theoretic methods for creating maximally diverse groups, VLSI design and exam scheduling. OMEGA 1997;25:473-82.
29. Weitz RR, Lakshminarayanan. An empirical comparison of heuristic methods for creating maximally diverse groups. Journal of the Operational Research Society 1998;49:635-46.