Linear programming brings marital bliss

Jay Sethuraman * Chung-Piaw Teo †

1 Motivation

By all means marry; if you get a good wife, you'll be happy. If not, you'll become a philosopher.
–Socrates

Alan, Bob, Carl and Dan, the only four bachelors in oddtown, finally contemplate marriage. They approach Marx, the matchmaker, who introduces them to Alice, Brenda, Cindy and Debbie. After the meeting, each person ranks all of the members of the opposite sex, and hands this list to Marx.

Alan	Cindy	Alice	Debbie	Brenda
Bob	Alice	Cindy	Debbie	Brenda
Carl	Debbie	Cindy	Alice	Brenda
Dan	Cindy	Brenda	Alice	Debbie

Men's preference lists

Alice	Dan	Carl	Alan	Bob
Brenda	Carl	Dan	Alan	Bob
Cindy	Carl	Bob	Alan	Dan
Debbie	Dan	Bob	Carl	Alan

Women's preference lists

For example, Alan likes Cindy best and Brenda least; Cindy, on the other hand, likes Carl and Bob better than Alan. Marx's job is to find a match for each man, and his reputation depends on the number of successful marriages performed. What should Marx do?

It is easier to answer what Marx *should not* do. Suppose he matches Alice & Alan, Brenda & Bob, Cindy & Carl, and Debbie & Dan. Observe that Dan likes Brenda better than Debbie (his current partner), and Brenda likes Dan better than Bob (her current partner); so, the proposed marriage would "break-up," and Dan & Brenda would "elope." Generalizing this observation, we see that a proposed marriage would be "unstable" if there is a man and a woman who *both* benefit by marrying each other. Obviously, such a situation is undesirable for Marx, and so the least he should look for is a *stable marriage.* Can he always find a stable marriage? How many stable marriages are there? Are some stable marriages better than the others? Is it possible for Debbie to lie about her preferences and get a better partner? The rest of this essay seeks to answer these questions and more.

*Operations Research Center, MIT, Cambridge, MA 02139.
†Faculty of Business Administration, Decision Sciences Department, NUS, Singapore.

2 Preliminaries

The stable marriage problem has held a fascination for computer scientists, mathematicians and economists ever since its introduction in the pioneering paper of Gale and Shapley[6]. Research conducted during the past thirty-five years has helped us understand and appreciate its connections to a variety of problems arising in combinatorics, operations research and economics. This essay is a modest attempt to explain a recent approach to this problem; the path we follow will take us through some beautiful mathematics, interesting applications, and hopefully expose the reader to some of the latest tools and techniques used in discrete optimization.

We first fix notation and define some of the terms used in the rest of this paper. This is followed by a brief overview of the best-known application of the stable marriage problem—matching interns to hospitals.

2.1 Notation and terminology

An instance of the stable marriage problem involves two disjoint sets of size n, the men and the women. Each person has a *strictly* ordered preference list of *all* of the members of the opposite sex. For convenience, we say that $w_1 >_m w_2$ whenever man m likes woman w_1 better than woman w_2. A marriage M is a one-to-one correspondence between the men and the women. A man-woman pair (m, w) is said to *block* the marriage M if m and w are not married in M, but prefer each other to their partners under M; such a pair is a *blocking pair* for the marriage M. A stable marriage M is one with no blocking pairs.

2.2 Matching interns to hospitals

The best-known application of the stable marriage problem is the mechanism used to match graduating medical students to hospitals. One glance at the history of the labor market for interns is enough to convince the skeptic of the importance of the stability condition. We provide a very brief overview of this history—for further details, see Roth[16].

The National Resident Matching Program (NRMP) assigns internships at hospitals to graduating medical students. Prior to 1952, such appointments were made on an ad-hoc basis, which resulted in a lot of dissatisfaction to both hospitals and students. Each hospital would make offers of appointments to applicants independently of its rivals, and the applicants were given a strict deadline by which they had to make their decisions. An applicant receiving an offer from a not-so-good hospital had to either reject this offer (in the hope that a better offer might be forthcoming), or accept it. An applicant who accepted such an offer would be unhappy if a better offer came along subsequently. Therefore an applicant would not reject any offer until the last moment, which caused considerable difficulty to hospitals. A hospital receiving a late rejection would be unhappy if many of its preferred alternatives ("wait-listed" candidates) had accepted other offers (rather, were "forced" to accept other offers). Typically, this process resulted in an "unstable" matching, and caused regret and anguish to both interns and hospitals.

It was agreed in 1950 to switch from a competitive market to a centralized system involving voluntary participation. Each applicant would submit a rank-ordered list of hospitals; likewise, each hospital would submit a rank-ordered list of its applicants. A matching would be arranged based on these preferences. The initial attempts to provide a solution were not satisfactory. Furthermore, it was recognized that some of the participants could benefit by deliberately presenting false preference lists. To get rid of these limitations, the algorithm was modified. The final algorithm was essentially a version of the Gale-Shapley algorithm, and was discovered by a trial-and-error process, which is quite remarkable! Even more remarkable is the fact that this algorithm has been in use from 1952, predating the work of Gale and Shapley by a decade.

Since participation in this matching program is entirely voluntary, one of the measures of the success of this matching mechanism is the participation rate. The fact that the participation rate has remained high throughout is a testimony to the effectiveness of the Gale-Shapley algorithm. In the early years, more than 95% of the applicants and hospitals participated in this scheme.

Any good application gives rise to a host of interesting theoretical problems, and this is no exception. Of particular relevance in this context is the problem of finding a stable matching with couples. In its present form, the algorithm cannot cater to the needs of married couples (who would like to be assigned to hospitals in the same city). Indeed, this is a shortcoming of the algorithm, and has become a big factor in the last two decades. As a consequence, the participation rate in the voluntary matching program has dropped to 80% percent. Modifications and improvements to the algorithm are being made to this day.

Note: The interns/hospitals problem departs from our definition of the stable marriage problem in two important ways. Firstly, the number of "positions" in hospitals need not be the same as the total number of applicants. Secondly, an applicant is not required to rank order *all* of the hospitals. However, a simple transformation reduces this problem to the standard stable marriage problem, and the same methods can be used.

2.3 Organization of the rest of the paper

Gusfield and Irving, in their book on the stable marriage problem[8], say:

> ... the paper of Gale and Shapley, with its neat and, in a sense, conclusive result, might have been expected to mark the end of the matter. On the contrary, however, it was only the beginning. Such is the elegance of the fundamental result and of the Gale-Shapley algorithm that it has, over the years, sparked much interest in stable matching problems from a variety of viewpoints.

Indeed, this paper is a survey of the stable marriage problem from an *algebraic* viewpoint. This point of view, in addition to unifying most of the known results, offers many new insights and provides us with a better understanding of the structure of solutions of the stable marriage problem. Preliminary experiments also suggest that the this approach might prove useful computationally in some (more difficult) stable matching problems, where the underlying graph is not *bipartite*.

In the next section, we provide an algebraic formulation of the example introduced in section 1. All of the results are first discussed in this specialized setting. In section 4, we present an algebraic formulation of the general stable marriage problem, discuss solution methods, and address issues such as fairness and deceit. For convenience, this section is patterned exactly like section 3; so, whenever a concept seems too abstract or fuzzy, the reader can refer to section 3, and try to understand it on the special case discussed there. Extensions to the stable marriage problem and some interesting open problems are discussed in section 5.

Since the main purpose of this paper is expository, we avoid making any bibliographic references to literature in the rest of this paper; a brief discussion of some of the important papers in the area is provided in the "notes and sources" section.

3 Marx's dilemma revisited

In this section we express Marx's problem in an algebraic language. We define decision variables which are allowed to be either zero or one; this point of view will help us translate Marx's problem to finding values for decision variables satisfying certain linear equations and inequalities. Finally, we find all solutions to Mark's problem and describe different ways of evaluating the quality of stable marriages.

3.1 Formulation

For each man-woman pair (m, w), we introduce a boolean variable $x_{m,w}$: thus, $x_{m,w}$ is either zero or one. For (mathematical) convenience, we rename the players involved. We use 1, 2, 3, 4 for Alan, Bob, Carl and Dan respectively; similarly, we use 1, 2, 3, 4 to represent Alice, Brenda, Cindy and Debbie respectively. This will cause no confusion because we will always use ordered pairs, where the first component denotes a man, and the second denotes a woman. For example, the pair (Bob, Cindy) would be denoted by (2,3). Altogether, we have sixteen decision variables in this problem.

The key idea is to express all of the constraints of the problem as linear relations involving these decision variables. So, what are the constraints? A little thought convinces us that there are two classes of constraints.

- **Matching Constraints**: Each person has exactly one partner (of the opposite sex).
- **Stability Constraints**: If man m and woman w are not partners, then at least one of them is better off. In other words, if m and w are not married to each other, either m is married to somebody he prefers to w, or w is married to someone she prefers to m. (If this is not true, clearly, (m, w) blocks the current marriage.)

Let us consider the (easy) matching constraints first. How do we ensure that Alan has a partner? Well, Alan could only be married to Alice, or Brenda, or Cindy, or Debbie, and so he has a partner if and only if he is married to exactly one of these four women. The corresponding decision variables are $x_{1,1}, x_{1,2}, x_{1,3}, x_{1,4}$ respectively, and so mathematically, this constraint can be written as

$$x_{1,1} + x_{1,2} + x_{1,3} + x_{1,4} = 1. \tag{1}$$

Similarly, we can write down the matching constraints for each of the other people. The remaining seven constraints are:

$$x_{2,1} + x_{2,2} + x_{2,3} + x_{2,4} = 1, \tag{2}$$

$$x_{3,1} + x_{3,2} + x_{3,3} + x_{3,4} = 1, \tag{3}$$

$$x_{4,1} + x_{4,2} + x_{4,3} + x_{4,4} = 1, \tag{4}$$

$$x_{1,1} + x_{2,1} + x_{3,1} + x_{4,1} = 1, \tag{5}$$
$$x_{1,2} + x_{2,2} + x_{3,2} + x_{4,2} = 1, \tag{6}$$
$$x_{1,3} + x_{2,3} + x_{3,3} + x_{4,3} = 1, \tag{7}$$
$$x_{1,4} + x_{2,4} + x_{3,4} + x_{4,4} = 1. \tag{8}$$

To write down the stability constraints we use a similar reasoning. Notice that we have sixteen stability constraints—one for each (man, woman) pair. For starters, let us write the stability constraint corresponding to the pair (Bob, Debbie). The stability constraint, in English, says that if Bob and Debbie are not married to each other, then at least one of them has a *better* partner. In our example, Bob prefers Alice and Cindy to Debbie, and Debbie prefers Dan to Bob. In a stable marriage, either Bob and Debbie are married, or Bob is married to Alice or Cindy, or Debbie is married to Dan. Another way to express this is that any marriage in which Bob is married to Brenda, and Debbie is married to either Alan or Carl, is unstable. A mathematical translation of the last statement is

$$x_{2,4} + x_{2,2} + x_{1,4} + x_{3,4} \leq 1. \tag{9}$$

Let us convince ourselves that equation (9) does indeed capture the stability condition for the pair (2,4). When is equation (9) violated? If $x_{2,4} = 1$, all the other decision variables appearing in equation (9) are forced to be zero by the matching constraints (2) and (8), and so (9) is satisfied. So, the only way in which equation (9) is violated is if $x_{2,4} = 0$, $x_{2,2} = 1$, and $x_{1,4} + x_{3,4} = 1$. (Since $x_{1,4} + x_{3,4} \leq 1$, by the matching constraint.) Thus equation (9) is violated if and only if Bob is married to Brenda, and Debbie is married to Alan or Carl, which is exactly the stability condition for the pair (Bob, Debbie). The remaining fifteen stability constraints are:

$$x_{1,1} + x_{1,2} + x_{1,4} + x_{2,1} \leq 1, \tag{10}$$
$$x_{1,2} + x_{2,2} \leq 1, \tag{11}$$
$$x_{1,3} + x_{1,1} + x_{1,2} + x_{1,4} + x_{4,3} \leq 1, \tag{12}$$
$$x_{1,4} + x_{1,2} \leq 1, \tag{13}$$
$$x_{2,1} + x_{2,2} + x_{2,3} + x_{2,4} \leq 1, \tag{14}$$
$$x_{2,2} \leq 1, \tag{15}$$
$$x_{2,3} + x_{2,2} + x_{2,4} + x_{1,3} + x_{4,3} \leq 1, \tag{16}$$
$$x_{3,1} + x_{3,2} + x_{1,1} + x_{2,1} \leq 1, \tag{17}$$
$$x_{3,2} + x_{1,2} + x_{2,2} + x_{4,2} \leq 1, \tag{18}$$
$$x_{3,3} + x_{3,1} + x_{3,2} + x_{1,3} + x_{2,3} + x_{4,3} \leq 1, \tag{19}$$
$$x_{3,4} + x_{3,1} + x_{3,2} + x_{3,3} + x_{1,4} \leq 1, \tag{20}$$
$$x_{4,1} + x_{4,4} + x_{1,1} + x_{2,1} + x_{3,1} \leq 1, \tag{21}$$
$$x_{4,2} + x_{4,1} + x_{4,4} + x_{1,2} + x_{2,2} \leq 1, \tag{22}$$
$$x_{4,3} + x_{4,1} + x_{4,2} + x_{4,4} \leq 1, \tag{23}$$
$$x_{4,4} + x_{1,4} + x_{2,4} + x_{3,4} \leq 1. \tag{24}$$

The problem of finding a stable marriage now reduces to finding a 0-1 vector $\mathbf{x}$ satisfying constraints (1)-(24).

3.2 Solution

Let us now find all possible 0-1 vectors satisfying the constraints (1)-(24). We begin by identifying all the variables which are forced to be zero or one.

- Constraints (12) and (1) force $x_{4,3} = 0$; Constraints (21) and (5) force $x_{4,4} = 0$.
- Constraints (19) and (7) force $x_{3,1} = x_{3,2} = 0$.
- Constraints (20) and (3) force $x_{1,4} = 0$.
- Constraints (22), (4), $x_{4,3} = 0$ and $x_{4,4} = 0$ force $x_{1,2} = x_{2,2} = 0$.
- Constraints (6), $x_{1,2} = 0$, $x_{2,2} = 0$ and $x_{3,2} = 0$ force $x_{4,2} = 1$.
- Constraints (4) and $x_{4,2} = 0$ force $x_{4,1} = 0$.

Fixing the values of the variables, and ignoring redundant constraints, we obtain the following simplified problem.

$$\begin{aligned} x_{1,1} + x_{1,3} &= 1, & (25)\\ x_{1,1} + x_{2,1} &= 1, & (26)\\ x_{1,3} + x_{2,3} + x_{2,4} &\leq 1, & (27)\\ x_{1,3} + x_{2,3} + x_{3,3} &= 1, & (28)\\ x_{2,1} + x_{2,3} + x_{2,4} &= 1, & (29)\\ x_{2,4} + x_{3,4} &= 1, & (30)\\ x_{3,3} + x_{3,4} &= 1, & (31)\\ x_{4,2} &= 1. & (32) \end{aligned}$$

Notice that there is a one-to-one correspondence between the 0-1 vectors satisfying equations (25)-(32) and the set of all stable marriages corresponding to Marx's problem.

- Suppose $x_{1,3} = 1$. Equation (25) forces $x_{1,1}$ to be zero, which in turn forces $x_{2,1} = 1$ (equation (26)). By equation (29), $x_{2,3} = x_{2,4} = 0$, and so equation (30) forces $x_{3,4} = 1$. This specifies a complete solution, which is $M_1 = \{(1,3),(2,1),(3,4),(4,2)\}$.
- Suppose $x_{2,3} = 1$. Equation (28) forces $x_{1,3}$ and $x_{2,4}$ to be zero, which in turn force $x_{1,1} = 1$ (equation (25)) and $x_{3,4} = 1$ (equation (30)). This specifies the complete solution, $M_2 = \{(1,1),(2,3),(3,4),(4,2)\}$.
- Suppose $x_{3,3} = 1$. Equation (28) forces $x_{1,3}$ and $x_{2,3}$ to be zero, which in turn force $x_{1,1} = 1$ (equation (25)). Since $x_{3,3} = 1$, $x_{3,4} = 0$ (by equation (31)), and so $x_{2,4} = 1$ by equation (30). The corresponding solution is $M_3 = \{(1,1),(2,4),(3,3),(4,2)\}$.

Since exactly one of $x_{1,3}$, $x_{2,3}$, $x_{3,3}$ is one (by equation (28)), the three cases discussed above are exhaustive, and so Marx's problem has exactly three stable marriages.

3.3 Fairness

There are three marriages which satisfy the stability condition. Are all stable marriages "equal" or are some marriages better than the others? The answer, of course, depends on how we measure the "goodness" of a marriage. For example, if Alan is allowed to choose among the three marriages, he would obviously pick M_1, because he gets his best partner. Would any other male disagree with Alan's choice? Remarkably, the answer is no! Indeed, it is very surprising that the men, who are essentially in competition with each other for the women, can agree on a marriage that is best for *all* of them. Another interesting property of M_1 is that it is the *worst* marriage for all the women: a good marriage from the male point of view is obtained at the expense of the women! Reversing the roles of men and women, we see that if women are given the power to choose, they would obviously pick M_3; again, M_3, the best marriage for women, is the worst marriage for the men. This observation prompts our next question: Are there any marriages "fair" to both sexes? Well, in this example, the only other marriage is a candidate. It is "in-between" the male optimal and the female optimal marriage. In general, there could be many more marriages, in which case we need to define an appropriate measure to guide us in the search of a good stable marriage.

Suppose m and w are partners in a marriage M. Let us define the *rank* of m in M to be the position of w in m's preference list. Similarly, the *rank* of w in M is the position of m in w's preference list. We emphasize that our definition of *rank* is with respect to a marriage M. For example, in marriage M_1, the matched pairs are $(1,3),(2,1),(3,4),(4,2)$, with the ranks of men being 1, 1, 1, 2 and the ranks of women being 3, 4, 3, 2. Define the *regret* of a marriage M to be the m-rank or w-rank of the person who is worst off in M. In other words,

$$\text{Regret(M)} = \max_{(m,w)\in M}\{\text{m-rank(m,w), w-rank(m,w)}\}.$$

Using the notions we can measure the goodness of a marriage in a variety of ways.

- **Men-optimal marriage**: A stable marriage which minimizes the sum of m-ranks.
- **Women-optimal marriage**: A stable marriage which minimizes the sum of w-ranks.
- **Egalitarian marriage**: A stable marriage which minimizes the sum of m-ranks and w-ranks.
- **Minimum regret marriage**: A stable marriage which minimizes the regret.

For the three stable marriages in Marx's problem, the various measures are summarized in Table 3.3. We see that M_1 is the men-optimal marriage and M_3 is the women-optimal marriage. Interestingly, all three marriages are Egalitarian, while M_2 and M_3 are both minimum-regret.

Marriage	Sum of m-ranks	Sum of w-ranks	Regret
M_1	5	12	4
M_2	7	10	3
M_3	9	8	3

There is an additional notion of fairness called the "median" marriage. We defer the discussion of this concept until the next section.

4 Stable marriage problem

In the previous section we formulated Marx's dilemma as the problem of finding a 0-1 vector satisfying certain linear constraints. We can do the same for a general stable marriage problem. In this section we will derive the formulation, discuss solution methods, and prove a variety of interesting properties.

4.1 Formulation

For clarity, we call the men $m_1, m_2, \ldots, m_n$ and the women $w_1, w_2, \ldots, w_n$. If man m_i prefers woman w_j to woman w_k, we write $w_j >_{m_i} w_k$. Let

$$x_{i,j} = \begin{cases} 1 & \text{if } m_i \text{ is matched to } w_j, \\ 0 & \text{otherwise.} \end{cases}$$

Consider the following system of linear constraints:

$$SM: \quad \begin{array}{rcl} \sum_j x_{i,j} = 1 & \forall\, i, & (33) \\ \sum_i x_{i,j} = 1 & \forall\, j, & (34) \\ x_{i,j} + \sum_{k: w_k <_{m_i} w_j} x_{i,k} + \sum_{k: m_k <_{w_j} m_i} x_{k,j} \leq 1 & \forall\, i, j, & (35) \\ x_{i,j} \in \{0, 1\} & \forall\, i, j. & (36) \end{array}$$

We now argue that SM is a correct formulation of the stable marriage problem. Equations (33) and (34) are, respectively, the matching constraints for the men and the women. Inequality (35) is clearly valid, otherwise, we must have $\sum_{k: m_k <_{w_j} m_i} x_{k,j} = 1$ and $\sum_{k: w_k <_{m_i} w_j} x_{i,k} = 1$. Thus both m_i and w_j are matched to less favorable mates in the marriage and can simultaneously improve by marrying each other, which makes (m_i, w_j) a blocking pair.

4.2 Solution

Now that we have a formulation of the stable marriage problem, is there a systematic way to solve it? We can think of an approach similar to the one we used in solving Marx's dilemma, and indeed, it can be done. However, we have overlooked the question of efficiency. Can we solve, say, problems with a million men and million women?

The formulation presented here for the stable marriage problem is an instance of an integer programming problem, a problem at the heart of discrete optimization. Integer programming offers a powerful modeling framework for a wide variety of problems arising in various branches of physical and social sciences, engineering etc. If we know of an efficient algorithm to solve general integer programming problems, we can appeal to that theory, and claim an efficient solution to the stable marriage problem. Unfortunately, no such algorithm is known as of now, and it is widely believed that such an algorithm is unlikely to exist! We will now take a brief detour and discuss concepts from the theory of linear and integer programming. Armed with this (newfound) knowledge, we will discover that stable marriage problems are very *special*; we will argue that even if we cannot solve general integer programming problems, the hardness results for integer programming do not affect problems in this special class. In short, we will see that *linear programming* will help us find a stable marriage (and *marital bliss*).

The problem of finding an integer vector satisfying linear constraints is well studied in discrete optimization. Typically, we also have an objective function, and our goal is to find a vector which optimizes (maximizes or minimizes, depending on the problem) the objective function, subject to linear constraints. This is called an Integer programming Problem. Integer programming is a powerful modeling framework that provides great flexibility to express a wide variety of problems. On the other hand, this flexibility comes with a price. To date, there is no known "efficient" algorithm to solve general integer programming problems. In fact, it is widely believed that such algorithms are unlikely to exist.

A linear programming problem is syntactically similar to an integer programming problem, except that the variables are not restricted to integer values only. Despite its similarity to integer programming, linear programming is an "easy" problem to solve and has many efficient algorithms.

Can we make use of our expertise in linear programs to solve integer programs efficiently? Sometimes, we can!

Consider the following integer programming problem

$$\begin{array}{rl} \max & \sum_{i=1}^{n} c_i x_i \\ \text{subject to} & Ax = b, \\ & x \geq 0, \\ & x \text{ integer.} \end{array}$$

With this integer programming problem (IP), we associate a linear programming problem (LP), called its *linear programming relaxation*. The associated LP is obtained by *relaxing* the integer constraints. Suppose we use an efficient LP algorithm to solve the relaxation. If we are incredibly lucky and find a solution x which is integral, we have solved the IP. This follows because every solution feasible to the IP is also feasible to the LP.

Under what conditions can we solve the integer program as a linear program? To answer this question, we need to discuss some elementary concepts from linear programming theory, which are briefly summarized in what follows.

Let v_1, v_2, $\ldots$, v_m be vectors in R^d. Vector $v = \lambda_1 v_1 + \lambda_2 v_2 + \ldots \lambda_m v_m$ is called a *convex* combination of $v_1, v_2, \ldots, v_n$ if $\sum_{i=1}^{i=m} \lambda_i = 1$, and $\lambda_i \geq 0$ for each i. The set of all convex combinations of a given set of vectors is called its convex hull. The convex hull of a finite number of vectors is called a *polytope*. Not all of these vectors may be essential in defining the polytope. In fact, for every polytope, there is a *unique minimal* set of defining vectors, the set of *vertices*. A vertex can also be characterized as a vector which is not contained in a segment joining two other vectors. A polytope is called *integral* if all of its vertices are integer-valued vectors. To put this in the context of the stable marriage problem, consider the set of vectors which represent stable marriages. Here, the vectors belong to R^{n^2}, and we do not know how many there are. The stable marriage polytope is the set of all convex combinations of stable marriages.

For our purposes, linear programming is the problem of of minimizing a linear function over a polytope. One of the fundamental theorems of linear programming says that the optimal value of a linear programming problem is always attained at a vertex. So, one algorithm to solve linear programs is to enumerate all the vertices, evaluate the objective at each vertex, and pick the best vertex. (This is very similar to the approach we used in solving Marx's problem.) However, the number of vertices can be extremely high, which makes this scheme impractical. A standard algorithm for linear programming is called the Simplex method due to Dantzig. This is a simple, and efficient algorithm which always produces an optimal vertex (assuming the polytope is non-empty).

This discussion suggests an attractive scheme to solve integer programs. Suppose we can show that the LP relaxation of an IP is an integral polytope. If we solve the LP relaxation, we will get an optimal vertex solution, which is integral, and so is feasible to the IP; since the solution we found is optimal for the LP, it has to be optimal for the IP as well. The hardness of integer programming suggests that the LP relaxation of an IP may not always be integral. Fortunately for us, the stable marriage polytope is integral and so we can use this approach. In what follows, we will provide an elegant proof of this fact.

Let P_{SM} denote the linear programming relaxation of SM, the integer programming formulation of the stable marriage problem.

Theorem 1 *The polytope* (P_{SM}) *is the convex hull of the stable marriage solutions.*

To give a simple proof of this result, we need the following property of the stable marriage polytope:

Lemma 1 *Let* x *be a feasible solution in the stable marriage polytope. Then*

$$x_{i,j} > 0 \textit{ implies } x_{i,j} + \sum_{k:w_k <_{m_i} w_j} x_{i,k} + \sum_{k:m_k <_{w_j} m_i} x_{k,j} = 1.$$

Proof: Consider the (primal) problem

$$\begin{array}{lcl}
\min & \sum_{i,j} x_{i,j} & \\
\text{subject to} & \sum_j x_{i,j} = 1 & \forall\, i, \\
 & \sum_i x_{i,j} = 1 & \forall\, j, \\
 & x_{i,j} + \sum_{k:w_k <_{m_i} w_j} x_{i,k} + \sum_{k:m_k <_{w_j} m_i} x_{k,j} \leq 1 & \forall\, i,j, \\
 & 0 \leq x_{i,j} \leq 1 & \forall\, i,j.
\end{array} \tag{37}$$

Using the duality theory of linear programming, we can formulate another problem called the dual linear program, as follows. For each constraint in the primal, associate a dual variable, and for each variable in the primal, associate a dual constraint. Depending on whether the primal constraint is an equality or an inequality, the corresponding dual variable is either free or sign-constrained. Also, depending on whether a primal variable is free or sign-constrained, we have an equality or an inequality constraint in the dual.

$$\begin{array}{lcl}
\max & \sum_i \alpha_i + \sum_j \beta_j - \sum_{i,j} \gamma_{i,j} & \\
\text{subject to} & \alpha_i + \beta_j - \sum_{k:w_k \geq_{m_i} w_j} \gamma_{i,k} - \sum_{k:m_k >_{w_j} m_i} \gamma_{k,j} & \leq 1,\ \forall\, i,j, \qquad (38)\\
 & \gamma_{i,j} & \geq 0,\ \forall\, i,j. \qquad (39)
\end{array}$$

Not surprisingly, there is a close connection between the primal problem and its dual. The value of the dual is never more than the value of the primal—this is called the *weak duality* theorem. The strong duality theorem of linear programming asserts that at optimality, the values of the primal and dual problems are the same. The property that is most relevant here is *complementary slackness.* Complementary slackness conditions relate the primal and dual optimal solutions in a nice way—they assert that if an inequality constraint in one problem is not tight (i.e. not satisfied as an equality), its associated variable in the other problem is zero. The proofs of all of these statements are standard, and are omitted here.

Set $\alpha_i = \sum_j x_{i,j}$, $\beta_j = \sum_i x_{i,j}$, $\gamma_{i,j} = x_{i,j}$. Then the constraints in the dual reduce to

$$x_{i,j} + \sum_{k:w_k <_{m_i} w_j} x_{i,k} + \sum_{k:m_k <_{w_j} m_i} x_{k,j} \leq 1\,, \quad \forall\, i,j,$$

which is clearly true if x is feasible in the primal. Furthermore, the above is a dual solution with objective function value $\sum_{i,j} x_{i,j}$. Thus, we have found primal and dual optimal solutions. The lemma follows directly from the complementary slackness property of optimal primal and dual solutions. ■

We show next that this lemma gives rise to a particularly interesting geometry for the stable marriage problem. Note that since $\sum_j x_{i,j} = 1$, then for each man m_i, we can construct n intervals of the type $(a, b]$ (left-open, right-closed) of length $x_{i,j}$ — one for each woman w_j (some of these intervals can be empty). We can arrange these n intervals (with no overlaps) to cover the line segment $(0, 1]$ in *any order.* Likewise we can do the same for the women. By slight abuse of terminology, we will also use $x_{i,j}$ to refer to the subinterval corresponding to man m_i and woman w_j in the arrangement.

Proof of Theorem 1 : Let x be any feasible solution in P_{SM}. We show that x can be written as a convex combination of stable marriages. Consider the following table of $2n$ rows: for each man m_i, $x_{i,j}$, $j = 1, ..., n$ are arranged in *decreasing* preference of m_i to cover the interval (0,1]. For each woman w_j, $x_{i,j}$, $i = 1, ..., n$ are arranged in *increasing* preference of w_j to cover the interval (0,1]. Note that the lemma suggests a nice property of this arrangement, i.e.,

> if $x_{i,j} > 0$, then the subinterval spanned by $x_{i,j}$ in the arrangement corresponding to m_i and w_j coincides in (0,1] (See Figure 1).

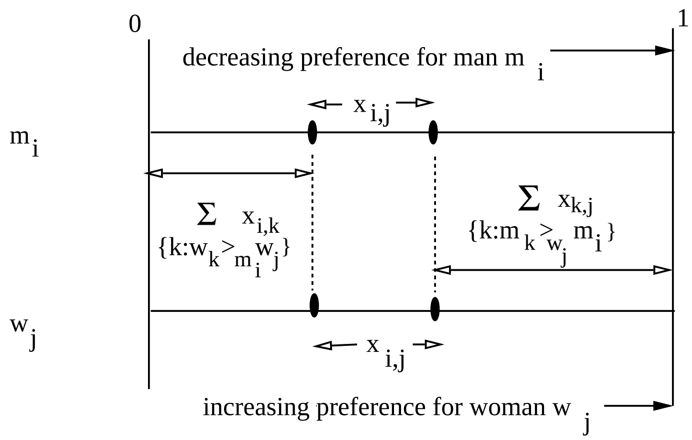

Figure 1: Geometry of the table of arrangement

We generate a random number U uniformly in (0,1] and construct a matching in the following way:

> If, in the row corresponding to m_i, U lies in the sub-interval spanned by $x_{i,k}$ in (0,1], then we match m_i with w_k. Similarly, match w_j with m_l if, in the row corresponding to w_j, U lies in the sub-interval spanned by $x_{l,j}$ in $(0, 1]$.

By Lemma 1, and by the way we construct the arrangement (see Figure 1), m_i is matched to w_j if and only if w_j is matched to m_i. Furthermore, no two men can be matched to the same woman, and similarly no two women can be matched to the same man. So the above assignment gives rise to a legitimate marriage. Any woman (say w_k) who is preferred by m_i to his mate w_j in this assignment (i.e., the subinterval $x_{i,k}$ is on the left of $x_{i,j}$ in the arrangement by m_i) is assigned a mate whom she strictly prefers to m_i, since in the row (in the table) corresponding to w_k, the random number U lies strictly to the right of the subinterval $x_{i,k}$. Hence the marriage is stable.

Let us again study Marx's problem. Suppose we solve the LP relaxation of Marx's problem and obtain the following solution.

$$\begin{array}{llll} x_{1,1} = 9/10 & x_{2,1} = 1/10 & x_{3,3} = 6/10 & x_{4,2} = 1 \\ x_{1,3} = 1/10 & x_{2,3} = 3/10 & x_{3,4} = 4/10 & \\ & x_{2,4} = 6/10 & & \end{array}$$

The intervals corresponding to this fractional are shown in Figure 2.

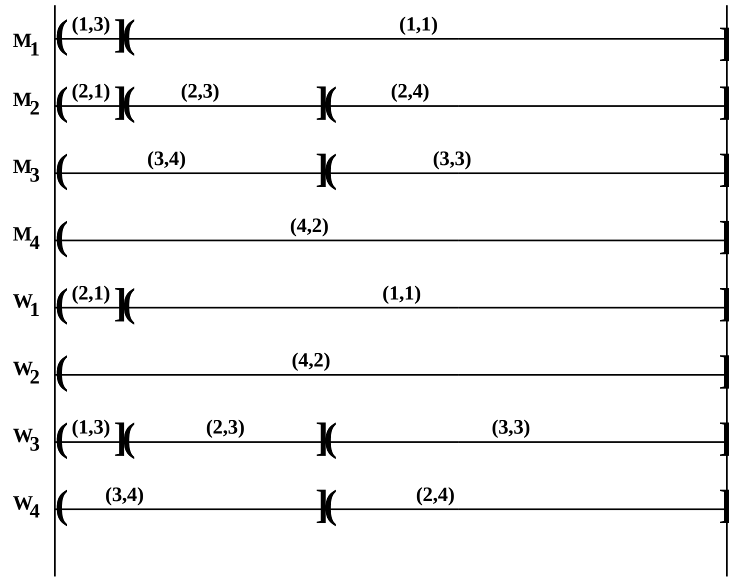

Figure 2: Marx's Dilemma: table of arrangement

The stable marriage corresponding to this $u = 0.3$ is $\{(1,1),(2,3),(3,4),(4,2)\}$. For $u = 0.05$, we get the men-optimal marriage M_1, and for $u = 0.95$, we get the women-optimal marriage M_3. Thus, we can generate all of the stable marriages in this picture by varying the u's, one of the hallmarks of the randomization approach.

Let X_U be the assignment obtained. $X_U(i,j) = 1$ if m_i and w_j are matched with each other under the above assignment. Then

$$E(X_U(i,j)) = P(U \text{ lies in the subinterval spanned by } x_{i,j}) = x_{i,j}.$$

Thus $x(i,j) = \int_0^1 X_u(i,j)du$, and we have shown that x can be written as convex combination of X_u as u varies over the interval $(0,1]$. We know that the cost of any integral solution is at least as much as the cost of the optimal LP solution. We have also shown that the expected value of the integral solution obtained by the rounding process is equal to the optimal value of the LP. However, the expected value is a "weighted" average, with the "weights" adding to one. The only possibility is if each of the integral solutions obtained by the rounding process has the same cost as that of the optimal LP solution. This shows that the integral solution obtained by the rounding process is optimal, and also shows that any fractional solution is a convex combination of the integral solutions that arise by choosing different values of u. Note that there are at most $O(n^2)$ distinct assignments arising from X_u as u varies. ■

The above argument gives rise to an efficient constructive procedure to express any fractional stable marriage solution as a convex combination of integral stable marriage solutions:

- Let x be a feasible fractional solution in the stable marriage polytope. Using the values of x, we construct the table of $2n$ rows as in the proof.
- We sort the co-ordinates of the endpoints of all the subintervals. Each row gives rise to at most $n+1$ distinct co-ordinates, so there are at most $O(n^2)$ distinct co-ordinates obtained in this way. Let $I_1, I_2, \ldots, I_p$ be the distinct co-ordinates for some integer p; without loss of generality, we assume $0 = I_1 < I_2 < \ldots < I_p = 1$, for some integer p.
- For each random number U in the interval $(I_k, I_{k+1}]$, the assignment produced by U in the randomized algorithm is identical. Let $X_{I_{k+1}}$ be the (integral) stable marriage solution obtained with $U = I_{k+1}$. Then
$$x = \sum_{j=1}^{p-1} (I_{j+1} - I_j) X_{I_{j+1}}.$$

Note: We still have not proved that we can always find a stable marriage. All we have proved is that if there is a stable marriage, we will find one by solving a linear programming problem. Is it possible for the stable marriage polytope to be empty? (Note that none of the theorems or properties established here will be contradicted if this were the case.) The answer, remarkably, is no! In fact, Gale and Shapley provided a simple algorithm in their first paper on stable marriage and established that the algorithm always terminates with a stable marriage. Hence, a stable marriage *always* exists, a rather surprising and neat result. However, a linear-algebraic proof of this fact is not known, and is left as an open problem for the curious reader.

4.3 Fairness

One approach to find "fair" stable marriage solutions is to incorporate an appropriate objective function into the problem. The power of the algebraic approach is best seen here. Finding fair stable matchings is a (at least conceptually) simple matter in this framework—all we need to do is to find a linear function which captures the "fairness" measure we have in mind. Combinatorial approaches to finding fair stable marriage solutions require a lot more ingenuity and can be very non-intuitive.

Let $M(x), W(x)$ denote respectively the worst ranking obtained by a man and a woman under the stable marriage solution x. In this section, we propose a linear programming model to handle this problem when the objective function is of the type

$$\sum_{i,j}(c(i,j)+c(j,i))x_{i,j}+f(M(x))+g(W(x))$$

where $(c(i,j)+c(j,i))$ is the weight on the edge $\{i,j\}$ and f,g are *non-decreasing* functions of $M(x)$ and $W(x)$ respectively. This class of objective functions is interesting because it includes as special cases several notions of fairness proposed in the literature. For instance, the minimum regret stable marriage problem can be reduced to the above with $f(k)=g(k)=k^{2n}$ for $k=1,\ldots,n$, and $c(i,j)=c(j,i)=0$ for all i,j.

Let x be a stable marriage solution. For man m_i, let $w_{i(1)}, w_{i(2)}, \ldots, w_{i(n)}$ be the ordering of the potential partners under his preferences. Similarly, we order the potential partners of w_j by $m_{j(1)}, \ldots, m_{j(n)}$. We model the worst ranking obtained by a man and a woman under x by

$$M(x)=\sum_{k=1}^{n} y_k,\ W(x)=\sum_{k=1}^{n} z_k,$$

where

$$y_k=\max_{m_i}(\sum_{l\geq k} x_{i,i(l)}),\ z_k=\max_{w_j}(\sum_{l\geq k} x_{j(l),j}),\quad k=1,2,\ldots,n.$$

Note that in this way, $y_k=1$ only when some man has been assigned to a partner whom he ranked k or worse in his preference list. Hence $\sum_{k=1}^{n} y_k$ and $\sum_{k=1}^{n} z_k$ model the worst ranking obtained by a man and a woman in x.

Let $f(0)=g(0)=0$. A natural LP relaxation for the problem is as follows:

$$(MM)\quad \min\quad \sum_{i,j}(c(i,j)+c(j,i))x_{i,j}+\sum_{k=1}^{n}\Big((f(k)-f(k-1))y_k+(g(k)-g(k-1))z_k\Big)$$

$$
\begin{aligned}
\text{subject to :}\quad & \textstyle\sum_j x_{i,j}=1 && \forall\, i, && (40)\\
& \textstyle\sum_i x_{i,j}=1 && \forall\, j, && (41)\\
& x_{i,j}+\textstyle\sum_{k:w_k<_{m_i}w_j} x_{i,k}+\sum_{k:m_k<_{w_j}m_i} x_{k,j}\leq 1 && \forall\, i,j, && (42)\\
& y_k\geq \textstyle\sum_{l\geq k} x_{i,i(l)} && \forall i, && (43)\\
& z_k\geq \textstyle\sum_{l\geq k} x_{j(l),j} && \forall j, && (44)\\
& x_{i,j}, y_k, z_k\geq 0, && \forall\, i,j,k. && (45)
\end{aligned}
$$

Theorem 2 $Z_{LP}(MM)=Z_{IP}(MM)$

Proof: Let (x^*,y^*,z^*) be an optimal solution to the LP relaxation of (MM). Note that at optimality, we may assume

$$y_k^*=\max_{m_i}\sum_{l\geq k} x^*_{i,i(l)},\ z_k^*=\max_{w_j}\sum_{l\geq k} x^*_{j(l),j}.$$

We round x^* to a stable marriage solution $\overline{x}$ using the rounding procedure in the proof of Theorem 1. Hence we have $E(\overline{x}_{i,j}) = x^*_{i,j}$. Let $y_k = \max_{m_i} \sum_{l \geq k} \overline{x}_{i,i(l)}$, and $z_k = \max_{w_j} \sum_{l \geq k} \overline{x}_{j(l),j}$. Let U be the random number generated in the rounding procedure.

$$\begin{aligned} E(y_k) &= P(\cup_{m_i}\{U \text{ lies in the subinterval spanned by } x^*_{i,i(l)}, \text{ for some } l \geq k\}) \\ &= \max_{m_i} \sum_{l \geq k} x^*_{i,i(l)} = y^*_k. \end{aligned}$$

Similarly, $E(z_k) = z^*_k$, and so $Z_{LP}(MM) = Z_{IP}(MM)$. ■

To see that the above models the minimum-regret version of the stable marriage problem, we consider the case with n men and n women. Let $m = 2n$. If we set $f(k) = g(k) = k^m$ for $k = 1, \ldots, n$, and $c_{i,j} = 0$ for all i, j, the problem reduces to finding a stable marriage solution x^*, which is optimal for the objective function $\min(M(x)^m + W(x)^m)$. We claim that the solution $\overline{x}$ is also optimal for the objective function $\min\{\max(M(\overline{x}), W(\overline{x}))\}$. Suppose otherwise, then let x' be a stable marriage solution that is optimal for the latter objective function and

$$\{\max(M(x'), W(x'))\} + 1 \leq \{\max(M(x^*), W(x^*))\}.$$

Since

$$\max(M(x^*).W(x^*))^m \leq (M(x^*)^m + W(x^*)^m) \leq (M(x')^m + W(x')^m) \leq 2\max(M(x'), W(x'))^m,$$

we have, by combining the two conditions that

$$(1 + \frac{1}{\max(M(x'), W(x'))})^m \leq 2.$$

Since $\max(M(\overline{x}), W(\overline{x})) \leq n$, we obtain $(1 + \frac{1}{n})^{2n} \leq 2$, which is a contradiction. That the objective function we consider can model the men-optimal, women-optimal and egalitarian versions of the stable marriage problem is left as an (easy) exercise to the reader.

We now define a new notion of "fairness" called the median stable marriage. Suppose $X_1, X_2, \ldots, X_l$ are all the distinct solutions to an instance of the stable marriage problem. Suppose further that l is odd. Each person has a list of l (possibly non-distinct) partners in the matching. For each person, rank-order all of his/her partners in these l marriages consistent with the person's preference list, and pick the person in position $\frac{l+1}{2}$ as the partner. It is not even clear if this process leads to a legitimate marriage (maybe man 5 is assigned to woman 3, but woman 3 is assigned to someone other than man 5). Remarkably, not only does this process define a legitimate marriage, but the resulting marriage is stable! This is an attractive result—it says that there is a marriage which is "fair" to both sexes—each person gets his "median" partner. In fact, the rounding approach can also be used to derive the following generalization.

Theorem 3 *Let $X_1, X_2, \ldots, X_l$ be l distinct stable marriage solutions. Each man m_i has l possible partners in these marriages; so does each woman w_j. For each participant order the l partners in increasing order of desirability. Match m_i with the woman who appears in position k among his l (possibly non-distinct) partners. Similarly, match woman w_j with the man who appears in position $l+1-k$ among her l partners in these solutions. This assignment gives rise to a stable marriage.*

Proof: Let $x(i,j) = \frac{1}{l}\sum_{p=1}^{l} X_p(i,j)$. Then x is a fractional point in the stable marriage polytope. Consider the case $U = \frac{k}{l} - \epsilon$ under the randomized rounding scheme, where $\epsilon > 0$ is chosen to be very small. Every man is then assigned to the woman whose rank is k among the l matchings. Similarly, every woman is assigned to the man whose rank is $l+1-k$ among the l matchings. Furthermore, it follows from the proof of Theorem 1 that the assignment is a stable marriage solution. ■

When $l = 3$ and $k = 2$, the above reduces to the well-known median property. (This says that the set of all solutions to a stable marriage problem forms a distributive lattice.) When l corresponds to the total number of distinct stable marriage solutions in the problem, and $k = 1$, the above gives rise to the men-optimal assignment considered by Gale and Shapley. When the total number of solutions, l, is odd, the case $k = \frac{l+1}{2}$ if of special interest — there is a stable marriage solution in which each person is assigned to the "median" among all possible mates!

Example: Consider the following instance of the stable marriage problem:

man	decreasing preferences	woman	decreasing preferences
1	3 1 5 7 4 2 8 6	1	4 3 8 1 2 5 7 6
2	6 1 3 4 8 7 5 2	2	3 7 5 8 6 4 1 2
3	7 4 3 6 5 1 2 8	3	7 5 8 3 6 2 1 4
4	5 3 8 2 6 1 4 7	4	6 4 2 7 3 1 5 8
5	4 1 2 8 7 3 6 5	5	8 7 1 5 6 4 3 2
6	6 2 5 7 8 4 3 1	6	5 4 7 6 2 8 3 1
7	7 8 1 6 2 3 4 5	7	1 4 5 6 2 8 3 7
8	2 6 7 1 8 3 4 5	8	2 5 4 3 7 8 1 6

We list below all the 23 stable marriages given in terms of permutations. (For instance, in marriage M_2, man 1 marries woman 1, man 2 marries woman 3, man 3 marries woman 7, man 7 marries woman 8 etc.)

M_1	=	3 1 7 5 4 6 8 2	M_2	=	1 3 7 5 4 6 8 2
M_3	=	3 1 4 5 2 6 8 7	M_4	=	3 1 7 8 4 5 6 2
M_5	=	1 3 4 5 2 6 8 7	M_6	=	1 3 7 8 4 5 6 2
M_7	=	3 1 4 8 2 5 6 7	M_8	=	1 3 4 8 2 5 6 7
M_9	=	5 3 4 8 2 7 6 1	M_{10}	=	1 4 3 5 2 6 8 7
M_{11}	=	3 1 4 6 8 5 2 7	M_{12}	=	5 4 3 8 2 7 6 1
M_{13}	=	5 3 4 6 8 7 2 1	M_{14}	=	1 4 3 6 8 5 2 7
M_{15}	=	5 4 3 6 8 7 2 1	M_{16}	=	5 4 1 8 2 7 6 3
M_{17}	=	5 8 3 6 7 4 2 1	M_{18}	=	5 8 1 6 7 4 2 3
M_{19}	=	5 4 1 6 8 7 6 3	M_{20}	=	7 8 1 6 3 4 2 5
M_{21}	=	7 8 2 1 6 4 3 5	M_{22}	=	1 4 3 8 2 5 6 7
M_{23}	=	1 3 4 6 8 5 2 7			

Now, we can arrange all 23 partners for each man according to his order of preference (with repetitions). That list is as follows:

man		
1	:	3 3 3 3 3 1 1 1 1 1 1 1 1 5 5 5 5 5 5 5 5 7 7
2	:	1 1 1 1 1 3 3 3 3 3 3 3 4 4 4 4 4 4 4 8 8 8 8
3	:	7 7 7 7 4 4 4 4 4 4 4 4 3 3 3 3 3 3 1 1 1 1 2
4	:	5 5 5 5 5 8 8 8 8 8 8 8 8 6 6 6 6 6 6 6 6 6 1
5	:	4 4 4 4 2 2 2 2 2 2 2 2 2 8 8 8 8 8 8 7 7 3 6
6	:	6 6 6 6 6 5 5 5 5 5 5 5 5 7 7 7 7 7 7 4 4 4 4
7	:	8 8 8 8 8 6 6 6 6 6 6 6 6 6 2 2 2 2 2 2 2 2 3
8	:	2 2 2 2 7 7 7 7 7 7 7 7 7 1 1 1 1 1 3 3 3 5 5

In this arrangement, any column will yield a stable marriage; the median stable marriage is just the 12^{th} column.

It is an interesting open problem to determine whether the median solution can be computed efficiently.

5 Extensions, generalizations and open problems

The worst thing you can do to a problem is to solve it completely.
—D. J. Kleitman

In this section, we consider some simple extensions of the stable marriage problem. We discuss the cases where the number of men is different from the number of women, preference lists may be incomplete, or preferences may not be strict. We briefly discuss deceit—is it possible for someone to gain an advantage by deliberately lying about his/her true preferences ? We then discuss the stable roommates problem, which is a generalization of the stable marriage problem. Essentially, we now have a single set of people, who have to be paired up as roommates. (Why is this a generalization?) This gives rise to some interesting issues which do not arise in the stable marriage problem. Along the way, we point out several (simple?) open questions for future research.

Stable marriage with sets of unequal size: Suppose we have m men and n women, and suppose $m < n$. How can we define stability in this setting? A natural definition of a blocking pair in a marriage M is a pair (m, w) such that (a) m and w are unmatched in M, (b) m is either unmatched in M or prefers w to his partner in M, and (c) w is either unmatched in M or prefers m to her current partner in M. A stable marriage is one with no blocking pairs. With this definition of stability, we can argue that there is always a stable marriage which matches all of the members of the smaller size. Furthermore, the set of women can be partitioned into two sets, the members of one subset being matched in all stable marriages, and the members of the other subset being matched in none of the stable marriages. (Try proving this statement!)

Stable marriage with unacceptable partners: In this case, we allow each person the freedom to rank only some of the members of the opposite sex. Clearly, we cannot hope to find a match for everyone in this case. A pair (m, w) blocks a marriage M if (a) m and w are unmatched in M, but each is acceptable to the other, (b) m is either unmatched in M or prefers w to his partner in M, and (c) w is either unmatched in M or prefers m to her current partner in M. A stable marriage, as always, is one with no blocking pairs. With these notions, we can prove that the men and the women can be partitioned into two subsets—those that have partners in all stable marriages and those that have partners in none. This again suggests that the notion of stability is robust.

Stable marriage with Indifference: Suppose we do not require that the preference lists be strict. This allows a person to be "indifferent" to some of the members of the opposite sex. It is not apriori clear how one extends the notion of stability. There is more than one way to define stability, but the most reasonable definition is as follows: a marriage is unstable if there is a man and a woman who are not partners, each of whom *strictly* prefers the other to his/her partners. With this definition of stability, it is possible to argue that ties can be broken arbitrarily; a stable marriage can always be found. (However, different ways of breaking ties may lead to different stable marriages.) (Think of other *reasonable* definitions of stability, and try proving results using your definitions.)

Deceit: Suppose women make all the decisions in society, subject, of course, to stability constraints. Clearly, among all stable marriages, the women-optimal marriage will be the one that will be picked; this, however, is the worst possible marriage for the men. Is it possible for the men to mislead the women by presenting false preference lists? If so, what is the best that men can do? More specifically, can the men force the women to pick the men-optimal marriage?

At least in some cases, the answer is yes! Consider the following example.

Alan	Brenda	Alice	Cindy
Bob	Alice	Brenda	Cindy
Carl	Alice	Brenda	Cindy

Men's preference lists

Alice	Carl	Alan	Bob
Brenda	Bob	Alan	Carl
Cindy	Alan	Bob	Carl

Women's preference lists

The women-optimal marriage is {(Alan, Alice), (Bob, Brenda), (Carl, Cindy)}, and the men-optimal marriage is {(Alan, Brenda), (Bob, Alice), (Carl, Cindy)}.

If Alan changes his preference list to Brenda, Cindy, and Alice, then the only stable matching for the new problem is {(Alan, Brenda), (Bob, Alice), (Carl, Cindy)}. which was the men-optimal stable marriage for the original problem.

Can this always be done? If we allow the men to declare some of the women as unacceptable, then this can be done trivially. Each man declares any woman appearing after his optimal partner as unacceptable. However, can this be done in the standard model, where each man has to rank all of the women? The answer is not known!

Stable roommates: In an instance of the stable roommates problem we have a set of $2n$ persons (say college students) who need to be paired up as roommates. Each person has a preference list ranking *all* the other persons in the set. We assume all preferences are *strict*. A matching is unstable when two people are not paired in the matching, but prefer each other to their assigned roommates. Such a pair is called a blocking pair. A matching is stable if and only if there are no blocking pairs.

Using the same ideas, we can formulate the stable roommates problem as an integer programming problem.

$$(FSM) \qquad \sum_j x_{i,j} = 1 \qquad \forall i, \tag{46}$$

$$x_{i,j} + \sum_{l <_j i} x_{l,j} + \sum_{l <_i j} x_{i,l} \leq 1 \qquad \forall i, j. \tag{47}$$

$$x_{i,j} \in 0, 1 \qquad \forall i, j. \tag{48}$$

Constraints 47 must be valid, since otherwise, j is matched to someone inferior to i whereas i is matched to someone inferior to j and so (i, j) is a blocking pair. We call constraints (47) the *paired* inequalities. Note that in the roommates case, the variable $x_{i,j}$ models the decision whether person i and person j are roommates.

In sharp contrast to the stable marriage problem, there are instances of the stable roommates problem which have no solution.

Consider the following stable roommates problem, involving four people. The preference lists of the individuals are:

P(1): 2, 3, 4
P(2): 3, 1, 4
P(3): 1, 2, 4
P(4): 1, 2, 3

We leave it to the reader to check that the above problem does not have a solution (no one can be matched to 4), and our formulation of this problem (FSM) does not have a solution either.

Can we identify instances of the stable roommates problem which have solutions? It turns out that the naive LP approach is not good enough.

There are infeasible instances of the stable roommates problem for which (FSM) is non-empty (see the 6 person example below), and so (FSM) cannot be used to address the feasibility question for the stable roommates problem. A natural way to improve the above formulation is to include the odd-set constraints from the matching polytope $\sum_{i,j\in S} x_{i,j} \leq \lfloor |S|/2 \rfloor$ for every set S of odd cardinality. This does not help either; we leave it to the reader to verify this observation on the following example.

Example: Consider the stable roommates problem with 6 people, with the following preference lists:

P(1): 2, 3, 5, 6, 4
P(2): 3, 4, 6, 1, 5
P(3): 4, 5, 1, 2, 6
P(4): 5, 6, 2, 3, 1
P(5): 6, 1, 3, 4, 2
P(6): 1, 2, 4, 5, 3

It can be verified easily that vector x with $x_{i,j} = \frac{1}{2}$ for $\{i, j\}$ in the set $\{\{1, 2\}, \{2, 3\}, \{3, 4\}, \{4, 5\}, \{5, 6\}, \{6, 1\}\}$ and $x_{i,j} = 0$ otherwise is a feasible solution in (FSM) but there does not exist any stable roommates solution.

Using the proof technique of the previous section, we can show that the polytope (FSM) is half integral (i.e. all its vertices have entries 0, 1 or 1/2).

The LP approach can be used to address the feasibility question. By adding more inequalities which are violated by fractional solutions, but satisfied by the integral solutions, FSM can be strenghtened to address the feasibility question. However, it has been shown that optimizing a linear function on the stable roommates polytope is a significantly harder problem. The key difference is that we do not have a (compact) description of the convex hull of solutions of the stable roommates problem. A number of open questions remain: can we solve the optimization problem approximately? If so, can we give a performance guarantee on the solution obtained (say, at most a factor of 2 from the optimal)? If we generate instances of this problem uniformly at random, what is the probability of finding a feasible instance?

Median solution: Perhaps the most interesting question is whether the median stable marriage solution can be found efficiently. This corresponds to finding an "average" of all the vertices of the stable marriage polytope. At this time, the only way we know of determining the median solution is to explicitly enumerate the vertices and average them. A more efficient scheme would have significance in economic contexts.

Finally, a linear algebraic proof of the existence of a stable marriage would be very welcome!

6 Acknowledgement

We are grateful to Dimitris Bertsimas, Rakesh Vohra and Viji Srinivasan for their comments on an earlier version of this essay. The title of this essay is borrowed from, and is a tribute to, Vande Vate's original paper, which started it all.

7 Notes and Sources

Gale and Shapley initiated the study of stable matching problems in their pioneering paper [6]. In their paper, they introduced the stable marriage problem, the stable roommates problem and the college admissions problem. For the marriage problem, they presented an elegant algorithm which always produces a (men-optimal) stable marriage. They showed that this algorithm could be adapted to solve the college admissions problem. They also showed that some instances of the stable roommates problem do not have stable solutions.

Knuth[11] published the first book on stable marriage. Knuth's primary aim was to use the stable marriage problem to illustrate a variety of techniques to analyze algorithms. Nevertheless, this book, especially the twelve open problems in its final section, spurred a lot of interest in the stable marriage problem and its variants.

McVitie and Wilson[12] first gave an algorithm ot generate all stable marriages for a given problem. Knuth[11] was the first to observe that the number of stable marriages could be exponential in the number of participants. Gusfield, Irving and their co-workers have studied the structure of all stable solutions to the marriage problem. In a series of papers, they reveal some fascinating connections to other well-known combinatorial problems and derive efficient algorithms for a variety of problems. In particular, they provide an efficient scheme to generate all stable marriages.

The lattice property of the set of all stable marriage solutions is attributed to John Conway.

Several papers in the literature deal with finding fair stable marriage solutions. Irving et al. [10] introduced the *egalitarian* stable marriage problem, and discovered an efficient algorithm to solve this problem. Knuth [11] described and solved (attributed to S. Selkow) the *minimum-regret* stable marriage problem.

Interest within the Economics community was aroused by Roth's discovery [16] (in 1984) that a version of the stable marriage algorithm was used in the National Resident Matching Program since 1952. Not surprisingly this has given rise to a huge body of research on strategic and game-theoretic aspects of the stable marriage problem. In this approach, the basic stable marriage model has been generalized to allow for the introduction of money and prices. Of particular interest to Economists are issues such as fairness and deceit. A comprehensive introduction to these aspects of the problem is in Roth and Sotomayor [18].

Vande Vate [22] initiated the study of the marriage problem using a linear programming approach. He obtained a complete characterization of the convex hull of the stable marriage solutions. This polyhedral description was extended by Rothblum [19] to the more general market model which permits incomplete preference lists. Since the original proof by Vande Vate is rather complicated, several simpler proofs have been proposed in the literature (Rothblum [19], Roth et al. [17]). The simple proof based on randomization is due to Teo and Sethuraman and appears in [21].

Knuth [11] asked if we can efficiently determine whether an instance of the stable roommates problem has a solution. This question has been answered in the affirmative by Irving [9], and also more recently by Subramanian [20]. On a related note, Feder [4] showed that finding an optimal solution to the stable roommates problem is NP-hard.

The polyhedral structure of the stable roommates problem was first studied by Abeledo and Rothblum [1]. Generalizing Vande Vate's formulation for the marriage problem, they showed that many structural properties of the stable roommates problem can be derived via linear algebraic arguments, mainly via linear programming duality. However, there are infeasible instances of the stable roommates problem for which the linear programming polytope associated with their formulation is non-empty. Hence, their formulation is not strong enough to be used to decide the existence of a stable matching. This formulation was strengthned by Teo and Sethuraman[21].

An interesting connection between the stable matching problem and network stability was discovered by [20]. Thus approach has been extended in many ways by Feder [4]. A comprehensive treatment of this approach appears in the recent monograph of Feder[5].

Generalizations of the stable marriage problem (including the ones we briefly discussed) have been studied extensively. Most of the results on the stable marriage problem and its variants are treated in detail in the delightful book by Gusfield and Irving[8].

There are many excellent books on linear programming—the books by Chvatal[3] and Bertsimas & Tsitsiklis[2] are especially recommended. Standard references for integer programming and combinatorial optimization include the books by Papadimitriou & Steiglitz[14], and Nemhauser & Wolsey [13]. The book by Grötschel, Lovász, & Schrijver[7], although more advanced, is a useful reference.

The references provided in this paper are by no means exhaustive. We urge the reader to consult the books by Gusfield & Irving[8], Roth & Sotomayor[18] and Feder[5].

References

[1] Abeledo, H. G., U. G. Rothblum (1994). Stable matching and linear inequalities. *Discrete Applied Mathematics* **54** 1-27.

[2] Bertsimas, D. J., J. N. Tsitsiklis (1997). *Introduction to Linear Programming*, Athena Scientific, Massachusetts.

[3] Chvatal, V. (1983). *Linear Programming*, W. H. Freeman, San Francisco.

[4] Feder, T. (1992). A new fixed point approach for stable networks and stable marriages. *Journal of Computer and System Sciences* **45** 233-284.

[5] Feder, T. (1995). *Stable Networks and Product Graphs*, Memoirs of the American Mathematical Society, **555**.

[6] Gale, D., L. S. Shapley (1962). College admissions and the stability of marriage. *American Mathematical Monthly* **69** 9-15.

[7] Grötschel, M., L. Lovász, A. Schrijver (1988). *Geometric algorithms and combinatorial optimization*, Springer-Verlag, Berlin.

[8] Gusfield, D., R. W. Irving (1989). *The Stable Marriage Problem: Structure and Algorithms*, MIT Press, Massachusetts.

[9] Irving, R. W. (1985). An efficient algorithm for the stable roommates problem. *Journal of Algorithms* **6** 577-595.

[10] Irving, R. W., P. Leather, D. Gusfield (1987). An efficient algorithm for the optimal stable marriage problem. *Journal of the Association of Computing Machinery* **34** 532-543.

[11] Knuth, D. E. (1976). *Marriages Stables*, Les Presses de l'Université de Montreal, Montreal.

[12] McVitie, D., L. B. Wilson (1971). The stable marriage problem. *Communications of the Association of Computing Machinery* **14** 486-490.

[13] Nemhauser, G. L, L. A. Wolsey (1988). *Integer Programming and Combinatorial Optimization*, Wiley, New York.

[14] Papadimitriou, C. H., K. Steiglitz (1982). *Combinatorial Optimization : algorithms and complexity*, Prentice Hall, Englewood Cliffs, N.J.

[15] Raghavan, P., C. Thompson (1987). Randomized rounding: a technique for provably good algorithms and algorithmic proofs. *Combinatorica* **7** 365-374.

[16] Roth, A. E. (1984). The evolution of the labor market for medical interns and residents: a case study in game theory. *Journal of Political Economy* **92** 991-1016.

[17] Roth, A. E., U. G. Rothblum, J. H. Vande Vate (1993). Stable matching, optimal assignments and linear programming. *Mathematics of Operations Research* **18** 808-828.

[18] Roth, A. E., M. Sotomayor (1991). *Two-sided matching: a study in game-theoretic modeling and analysis*, Cambridge University Press, Cambridge.

[19] Rothblum, U. G. (1992). Characterization of stable matchings as extreme points of a polytope. *Mathematical Programming* **54** 57-67.

[20] Subramanian, A. (1994). A new approach to stable matching problems. *SIAM Journal of Computing* **23** 671-700.

[21] Teo, C. P., J. Sethuraman (1997). LP based approach to optimal stable matchings. *Proc. Of the* 8^{th} *ACM-SIAM Symposium on Discrete Algorithms* 710-719.

[22] Vande Vate, J. H. (1989). Linear programming brings marital bliss. *Operations Research Letters* **8** 147-153.