Linear Programming Brings Marital Bliss

V. J. Sethuraman and C.P. 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 it to Marx.

Men's preference lists

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

Women's preference lists

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

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 arranged. 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." We conclude that the proposed marriage in this case is "unstable". Obviously, such a situation is undesirable for Marx, and so the least he should look for is a *stable* marriage, where no pair of man and woman will find it beneficial to divorce their respective spouses and marry each other. Can Marx always find a stable marriage? How many stable marriages are there? Are some stable marriages better than the others? Can people misrepresent their true preferences and thereby gain an advantage? We will answer some of these questions in the next section.

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 [62]. 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. It has been applied to, for example, the matching of graduating medical students to hopsitals and college admissions.

* *This article is edited version of a paper of the same title by the first author. The paper was submitted for the "Writing and Humanistic Studies Prize for Scientific and Engineering Writing" competition organised by Massachusetts Institute of Technology in 1997; it won a second prize in the competition.*

2. Helping Marx out

We shall translate Marx's problem to a problem of finding value of boolean variables satisfying linear constraints, and find all the stable marriage solutions to this problem.

Formulation

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. 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.

For each man-woman pair (m,w), we introduce a boolean variable $x_{m,w}$: thus, $x_{m,w}$ is either zero or one and $x_{m,w} = 1$ if m and w have been paired in M and $x_{m,w} = 0$ otherwise. For (mathematical) convenience, we "rename" the players involved. We use 1, 2, 3, 4 for Alan, Bob, Carl and Dan respectively; similarly, 1, 2, 3, 4 will also 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:**

These express the fact that each person has exactly one partner (of the opposite sex).

- **Stability Constraints:**

If man m and woman w have not been paired up, then at least one of them has been paired with a "better" partner. 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 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

(1) $$x_{1,1} + x_{1,2} + x_{1,3} + x_{1,4} = 1.$$

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

(2) $$x_{2,1} + x_{2,2} + x_{2,3} + x_{2,4} = 1$$

(3) $$x_{3,1} + x_{3,2} + x_{3,3} + x_{3,4} = 1$$

(4) $$x_{4,1} + x_{4,2} + x_{4,3} + x_{4,4} = 1$$

(5) $$x_{1,1} + x_{2,1} + x_{3,1} + x_{4,1} = 1$$

(6) $$x_{1,2} + x_{2,2} + x_{3,2} + x_{4,2} = 1$$

(7) $$x_{1,3} + x_{2,3} + x_{3,3} + x_{4,3} = 1$$

(8) $$x_{1,4} + x_{2,4} + x_{3,4} + x_{4,4} = 1.$$

To write down the stability constraints, we use a similar reasoning. Notice that we have sixteen stability constraints - one for each (man, woman) pair. We shall illustrate here how one obtains 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 gets a "better" partner. In our example, Bob prefers Alice and Cindy to Debbie, whereas Debbie prefers Dan to Bob. In order for a marriage to be stable, 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 Carl or Alan, is unstable. A mathematical translation of the last statement is

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

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

By applying similar argument, one can write down the remaining fifteen constraints.

The problem of finding a stable marriage now reduces to finding the values of $x_{m,w}$ satisfying the above eight equations and the sixteen stability constraints.

Solution

Let us now find all possible $x_{m,w}$ satisfying the above eight matching constraints and the sixteen stability constraints. It turns out that certain variables are forced to be zero, namely
$x_{4,3} = x_{4,4} = x_{3,1} = x_{3,2} = x_{1,4} = x_{1,2} = x_{2,2} = x_{3,2} = x_{4,1} = 0$;
while $x_{4,2} = 1$ And the remaining variables are subjected to the following seven constraints:

(10) $x_{1,1} + x_{1,3} = 1$

(11) $x_{1,1} + x_{2,1} = 1$

(12) $x_{1,3} + x_{2,3} + x_{2,4} \leq 1$

(13) $x_{1,3} + x_{2,3} + x_{3,3} = 1$

(14) $x_{2,1} + x_{2,3} + x_{2,4} = 1$

(15) $x_{2,4} + x_{3,4} = 1$

(16) $x_{3,3} + x_{3,4} = 1$

We finally arrive at the following three possible solutions:

- Suppose $x_{1,3} = 1$. Then (10) forces $x_{1,1}$ to be zero, which in turn forces $x_{2,1} = 1$ (by equation (11)). By (14), $x_{2,3} = x_{2,4} = 0$, and so (15) forces $x_{3,4} = 1$. This specifies a complete solution, which is

  $M_1 = \{(1,3), (2,1), (3,4), (4,2)\}$.

- Suppose $x_{2,3} = 1$. By similar consideration as above, we obtain a complete solution

  $M_2 = \{(1,1), (2,3), (3,4), (4,2)\}$.

- Suppose $x_{3,3} = 1$. 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 (13)), the three cases discussed above are exhaustive. Thus, we conclude that Marx's problem has exactly three stable marriage solutions.

3. Linear programming and the stable marriage problem

In the previous section, we formulated Marx's problem as a problem of finding value of boolean variables satisfying certain linear constraints. We can do the same for a general stable marriage problem. The formulation presented here is an example of an integer programming problem, a problem which has been well-studied in discrete optimization. Typically, the goal is to find an integer vector which optimizes (maximizes or minimizes, depending on the problem) a given objective function, subject to linear constraints. Integer programming is a powerful modeling framework that provides great flexibility to express a wide variety of probelms. 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 porblem 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 programming to solve integer programming problems efficiently? Sometimes we can!

With an 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.

Fortunately for us, we can solve the stable marriage problem using this approach. There are other variations of the stable marriage problem. Interested readers are urged to consult the references.

References

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

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

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

[4] **Teo, C. P., J. Sethuraman (1997).** *LP based approach to optimal stablematchings. Proc. of the 8th ACM - SIAM Symposium on Discrete Algorithms 710-719.*

Mr Jay Sethuraman is a graduate student in operations research at *Massachusetts Institute of Technology*. His research interests are in applied probability and optimization.

Dr Teo Chung Piaw is a lecturer in the *Department of Decision Sciences, Faculty of Business administration* at NUS. His main research interests are in optimization, logistics and operations research.