

Stable Marriage Problem

1 Introduction

Consider a world with two types of people. Let's call them male and female. If we start with n males and n females, a marriage is a 1-to-1 pairing of all the males and females.

Each person is assumed to have a preference order for members of the opposite sex. A male A might have preference order dcb , meaning that he prefers d over c , c over b etc. We could write that using a preference function Pref :

$$\text{Pref}_A(d) > \text{Pref}_A(c) > \text{Pref}_A(b) > \text{Pref}_A(a)$$

A marriage (i.e. a pairing of all the people) is stable if there is no dissatisfied pair. A dissatisfied pair is a male-female combination who like each other more than their current spouses, i.e. if $x - X$ is a married pair and $y - Y$ is another married pair, then $X - y$ would be a dissatisfied pair if

$$\text{Pref}_X(y) > \text{Pref}_X(x) \text{ and } \text{Pref}_y(X) > \text{Pref}_y(Y)$$

More than one matching possible: This is easily seen using the following example.

Man #:	Preference list	Woman #	Preference list
1	1 2	1	2 1
2	2 1	2	1 2

Both the matchings $(1, 1), (2, 2)$ and $(1, 2), (2, 1)$ are stable. The first is obtained when men make the proposals and the second is obtained when women make the proposals.

The Proposal Algorithm

Surprisingly enough, a stable marriage always exists for any group of equal numbers of males and females. And there is even a straightforward and fairly traditional way of finding it: the proposal algorithm.

Males At each step, some male (it can be any unmarried male) proposes to his most-preferred female who has not already turned him down.

Females Each female will accept a proposal unless she is married and prefers her current spouse more. That means that a single female always accepts a proposal.

```

Initialize all m and w to free
while there is a free man m who still has a woman w to propose to {
  w = m's highest ranked such woman to whom he has not yet proposed
  if w is free
    (m, w) become engaged
  else some pair (m', w) already exists
    if w prefers m to m'
      (m, w) become engaged
      m' becomes free
    else
      (m', w) remain engaged
}

```

Figure 1: Gale-Shapley men-optimal algorithm. The women-optimal version can be derived by reversing the roles of men and women.

A - a
 B - b
 C - a → d
 D - c

Example Find the stable marriage for 8 people:

- | | |
|--------------------|---------|
| A: abcd | a: BCDA |
| B: bcda | b: CADB |
| C: abc | c: ADBC |
| D: cabd | d: DBAC |

a - A
 b - B
 d - C
 c - D

2 Analysis

First of all, every male proposes at most once to each female. So the maximum number of proposals is just the number of male-female pairs, which is n^2 . If the preferences of the males are all the same (homework problem) there will actually be $\Theta(n^2)$ proposals.

Correctness Its actually fairly easy to show that the proposal algorithm works. The proof is by contradiction.

Suppose there is a dissatisfied pair $X - y$. Then we must have:

1. $\text{Pref}_X(y) > \text{Pref}_X(x)$ and
2. $\text{Pref}_y(X) > \text{Pref}_y(Y)$

Statement 1. means that male X likes y more than his current spouse. He must have proposed to her before, but he isn't married to her now, so either:

- She rejected X , which means she liked her then-husband (call him Z) more than X . She is married to Y (who might be the same person as Z), so she must have liked him even more. i.e.

$$\text{Pref}_y(X) < \text{Pref}_y(Z) \leq \text{Pref}_y(Y)$$

But that creates a contradiction.

- She accepted X at first, but left him later for someone else (call him W) later. Eventually, she was married to Y (who was perhaps W all along), so that means:

$$\text{Pref}_y(X) < \text{Pref}_y(W) \leq \text{Pref}_y(Y)$$

and again we get a contradiction.

So we always get a contradiction, and our assumption that a dissatisfied pair exists must be false. Thus the marriage is stable.

Average-case analysis

For some orderings of preferences, the proposal algorithm takes $\Theta(n^2)$ steps. But that was with a rather special ordering of preferences (the same for all males). What about a more "typical" ordering? By typical ordering, we choose each person's preference as a random, uniformly distributed permutation of $1, \dots, n$.

Now it's rather hard to analyze this. Instead, rather than picking all those random preference permutations at the beginning, we use a random version of the algorithm to simulate them. This random algorithm, at each step, allows a male to propose to a random female. The name of the female gets added to the end of his preference list as though she had been there on his preference list from the beginning. So it's not really the random algorithm that we're interested in. It's just that the random algorithm behaves just the same as running the normal algorithm with a set of preferences that are uniformly-distributed random permutations. The preference lists get built up during the algorithm. You should be able to see that you still end up with random, uniformly distributed permutations as the preferences anyway. Also, the random choice isn't different from choosing the highest preferred female that hasn't been proposed to yet. The random choice with this scheme defines the highest preferred female who hasn't been proposed to yet.

The only odd case is that since a male proposes to a random female, he might come up with the same female in his preference list twice. If so, the second choice is just removed from his list (as though she had not been picked the second time) and he tries again. The random algorithm is now pretty easy to analyze. We make a few observations:

- A female, once married, stays married forever, although not necessarily to the same male.
- A female is married if and only if she has been proposed to some time.
- Every female is married, and the algorithm stops, if and only if every female has been proposed to.

So to figure out how long the algorithm takes, it's enough to figure out how many random choices for proposals to n females it takes before all of them has been proposed to. This is equivalent to the following occupancy problem: If m balls are placed into n bins, how long does it take before every bin is non-empty with high probability?

The Coupon Collectors Problem

That occupancy problem is usually stated as the coupon collectors problem. If there are n coupons at a supermarket and you get a random one each time you visit, how many (m) do you need to have high probability of getting all the coupons?

Let E_i be the property that bin i is empty in the final arrangement. Then $\Pr[E_i] = (1 - 1/n)^m$. The probability that *some bin* (i.e. any of the bins) is empty is:

$$\Pr[E_1 \vee E_2 \vee \dots \vee E_n] \leq \Pr[E_1] + \Pr[E_2] + \dots + \Pr[E_n]$$

Which is less than or equal to

$$\sum_{i=1}^n \left(1 - \frac{1}{n}\right)^m = n \left(1 - \frac{1}{n}\right)^m = n \left(\left(1 - \frac{1}{n}\right)^n\right)^{m/n} \approx n(e^{-1})^{m/n} = n \exp(-m/n)$$

If we set that probability to say 0.01, then we have an inequality

$$0.01 > n \exp(-m/n)$$

taking logs to base e,

$$-5 > \ln(n) - m/n$$

and rearranging:

$$m > n \ln(n) + 5n$$

So if m is somewhat bigger than $n \ln(n)$, we have a very good chance of hitting all the bins. And if you make about $n \ln n$ visits to the supermarket, you have a good chance to get all n coupons.

Finally, the result that we were interested in: After about $n \ln n$ proposals under the random algorithm, every female will have been proposed to, and a stable marriage will result. That means if you started with random permutations as preferences, the proposal algorithm would run for about $n \ln n$ steps.

3 Expected Rank of Spouse

Lets return to the proposal algorithm when run on males and females with random preferences. We have shown that the total number of proposals is less than about $n \ln n$ with high probability. In the next lecture we will show that the expected value of the number of proposals is also $n \ln n$. So if X_i is the number of proposals made by male i , then $E[X_i]$ is the same for each male and is

$$E[X_i] = E[X]/n \approx \ln n$$

Now each male proposes to females exactly in the order 1, 2, 3, ... on his preference list. So if he makes m proposals, his last proposal, and the rank of his final wife, is m . Therefore the expected rank of a male's final wife is $\approx \ln n$.

The situation for females is more complex. Each female still receives an expected number $\approx \ln n$ of proposals. But each proposal has a random rank on her list. Let Y_i be random variable which is the rank of the i^{th} proposal that some female receives. Then the Y_i are independent random variables drawn from $\{1, \dots, n\}$ with the uniform distribution. If a female receives m proposals, she will end up married to the lowest ranked one (low numerical rank means high preference). That is, if Y is the rank of the female's final husband, then

$$Y = \min(Y_1, \dots, Y_m)$$

To compute $E[Y]$, we use this form of the expected value formula, which applies to integer-valued r.v.'s:

$$E[Y] = \sum_{k=1}^n \Pr\{Y \geq k\}$$

and notice that $Y \geq k$ means $\min(Y_1, \dots, Y_m) \geq k$, which is equivalent to $Y_1, \dots, Y_m \geq k$. Now $\Pr\{Y_i \geq k\} = (n - k + 1)/n$, and since the Y_i are independent:

$$\Pr\{Y \geq k\} = \left(\frac{n - k + 1}{n}\right)^m$$

and so

$$E[Y] = \sum_{k=1}^n \left(\frac{n - k + 1}{n}\right)^m$$

To simplify this, we substitute $j = n - k + 1$, giving:

$$E[Y] = \frac{1}{n^m} \sum_{j=1}^n j^m$$

Now the last sum is a standard sum which you can look up in a text, or solve using generating functions, combinatorial identities, or mathematical induction. It evaluates to:

$$\sum_{j=1}^n j^m = \frac{n^{m+1}}{m+1} + O(n^m)$$

and substituting, we find that:

$$E[Y] \approx \frac{1}{n^m} \frac{n^{m+1}}{m+1} = \frac{n}{m+1}$$

So if a female receives exactly $\ln n$ proposals, the expected rank of her final husband is $n/(\ln n + 1)$. Now the number of proposals that a female receives is a random variable with mean $\ln n$, but as n grows, the distribution of this number converges to its mean in ratio. So the error from using the mean is small.

So the proposal algorithm with random preference lists is *much* worse for females than for males. Males end up with spouses near the beginning of their lists ($\ln n$), while females end up with spouse way down on their lists ($\approx n/\ln n$). Notice that the longer the proposal algorithm runs, the better things get for females, and the worse things get for males. This is true for *any* preference order (not just random), because each new proposal is one notch further along some males preference list. Conversely, the rank of a female's husband can only improve over time. In the general case, if the proposal algorithm runs for nm steps, the expected rank of a male's spouse is m . For females, it is harder to figure out, unless females have random preferences, in which case the expected rank of a female's spouse is n/m .

4 Generalization of the Stable Marriage Problem

The original Stable Marriage Problem is very restrictive for practical applications so it is natural to consider about variations that are not so restrictive.

Incomplete lists

The first generalization of *SMP* is that some partners may find other partners unacceptable. In this case we enable that preference lists may be incomplete.

In the Stable Marriage problem with Incomplete lists (*SMI*), the numbers of men and women need not be the same, and each person's preference list consists of a subset of the members of the opposite sex in strict order.

A pair (m, w) is acceptable if each member of the pair appears on the preference list of the other.

A matching M is now a set of acceptable pairs such that each person belongs to at most one pair. As in the classical case, there is always at least one stable matching for an instance of *SMI*, and it is straightforward to extend the Gale / Shapley algorithm to give an algorithm for this case.

Given an instance of *SMI*. A marriage M is stable if there is no unmatched pair (m, w) such that:

1. m and w is unmatched in M , and finds each other acceptable
2. or each prefers the other to his/her partner in M .

Ties

We can also consider a situation when some partner does not distinguish between two or more partners. In this case we allow that partners do not have to rank all partners of opposite sex in strict order, the preference list may contain ties. Let *SMT* denote such a generalization of *SM*.

In this context, (m, w) is a blocking pair for a matching M if

- (a) (m, w) is an acceptable pair,
- (b) m is either unmatched or prefers w to $M(m)$, and
- (c) likewise w is either unmatched or prefers m to $M(w)$.

A marriage M is:

1. *weakly stable* if there is no blocking pair for M .
2. *strongly stable* a pair which does not appear in M , blocks M if one member of the pair prefers the other to their partner in M while the other is at least indifferent. If there is no such blocking pair for M , M is said to be strongly stable
3. *super-stable*, a pair which does not appear in M blocks M if each member of the pair prefers the other to their partner in M , or is indifferent between them. If there is no such blocking pair for M , M is said to be super-stable

Incomplete lists and Ties

Both of these variants (*SMI* and *SMT*) are polynomial to solve. In real world situations, both ties and incomplete preference lists may be needed. Unfortunately, when we allow both, the problem becomes NP-hard. In a *SMTI* (Stable Marriage with Ties and Incomplete lists) problem, there may be several stable marriages of different sizes, and solving the problem means finding a stable marriage of maximum size.

References

- [1] Wikipedia, Stable marriage problem http://en.wikipedia.org/wiki/Stable_marriage_problem
- [2] The Stable Marriage Problem, Lane Department of Computer Science and Electrical Engineering, West Virginia University, William Hunt's Lecture notes www.csee.wvu.edu/~ksmani/courses/fa01/random/lecnotes/lecture5.pdf
- [3] <http://www.dcs.gla.ac.uk/publications/PAPERS/8482/thesis.pdf>
- [4] <http://www.ceid.upatras.gr/courses/pithmeth/slides/lecture5.pdf>