

We've seen graphs involving boys and girls and connections between them in the context of our sexual demographics calculation and study. A similar problem comes up in terms of what I call stable matching, which is again the issue of matching up boys and girls in a special way according to some constraints. It turns out to have a lot of applications which we'll discuss toward the end.

Let's just look at what the problem is. The setup is that there's some number of boys, in this case 5, 1 through 5, and an equal number of girls labeled A through E. And each of the boys has a ranking of the girls. Different rankings, because different boys have different preferences. And likewise, the girls have rankings of the boys. Different girls have different preferences. So here, girl A likes boy 3 best and boy 5 second best, and boy 1 likes girls C best and girl D least.

So the problem, basically, is that we want to get all the boys married to all the girls. We want to form 5 monogamous heterosexual marriages, a boy and a girl, and we'd like, in some vague way, to acknowledge these preferences and satisfy as many as we can. I'll be more specific about that in a minute. But let's just play with that idea of trying to accommodate people's preferences.

So one way to do it is let's just decide, well, we'll favor the boys this time. Let's try a greedy strategy for the boys. Let's just look at the boy preferences. And a greedy strategy means I'm going to try to give each boy the best possible choice that he can make.

So I'm going to start off by deciding that let's let boy 1 have his first choice, girl C. I'm going to marry them off. And once I've married them off, I'll just stop considering 1 and C. And now I have a reduced problem. I have four remaining boys and four remaining girls.

And proceeding in this way-- greedy way for boys-- I'm going to now give boy 2 his next choice that remains, namely girl A. And I'll marry them off. And again, now 2 and A can be eliminated from consideration. I continue in this way and I wind up with this set of five marriages ending with boy 5 married to girl E.

OK. Now if we look at this set of marriages, there's a problem which motivates the whole stable marriage problem that we're going to be examining. Namely, we've married off boy 1 to his first choice, girl C. He should be happy, but she may not be. And we've also married off boy 4 to girl B.

Now a difficulty here is that if you look at the preferences, girl C actually is more desirable to boy 4 than girl B. Girl C, or boy 4, likes somebody else's wife better than his own. And what makes it really bad is that girl C, the other person's wife, likes boy 4 better than her husband. Each of them would be better off if they ran off together. They

are, whether they do or not, they certainly are under tremendous pressure. It makes this set of marriages unstable.

So they're called a rogue couple. When you have in a set of marriages a boy and a girl who prefer each other over their current spouses, they are said to be a rogue couple and a source of instability. So the stable marriage problem is let's see if we can get everybody married off and have no rogue couples. It'll be a stable set of marriages.

Now people may not be happy, but it doesn't matter because they'll never find anybody else that is unhappy in the same way that would be willing to run off with them and make them happier. So it's stable. And it turns out that there always is a way to find a stable set of marriages. A couple of ways. But why don't we just play with the idea.

Here is display of those preferences again and you can stop the video and fiddle with a piece of paper and see if you can come up with a stable set of marriages between the boys and girls. We used to do this in class in real time. We would give five different boys a chart of preferences of girls, and we'd give the five different girls a chart of preferences of boys. They were not supposed to tell each other what their preferences were, but the girls were supposed to be interviewing the boys and the boys interviewing the girls simultaneously in a parallel, and try to agree to get married and see whether the set of marriages that they wound up with were stable.

Most of the time they actually did wind up with a stable set of marriages, but not always. Just an amusing exercise. And it does illustrate something about the fact that the procedures that we're going to be going through to find stable marriages work very nicely if you choose to do them in parallel.

Anyway, there are, it turns out, two sets of stable marriages that we can find in this particular set of preferences. The simplest one to understand is all the girls get their first choice. It so happens, if you look at that chart, all of the girls have different first choice boys. If we simply give them their first choice, no girl is going to be tempted to be part of a rogue couple because she's got her first choice. It's absolutely stable.

But of course, that's a very special circumstance. It would be unusual that all the girls' first choices were different, or likewise, it would be unusual if all the boys' first choices were different. But if they were, then you instantly get a stable set of marriages.

There's another stable set that's not quite so obvious. And you can check that all of these pairs have no instability. There's no rogue couples in here when I marry 5 to A and 1 to E. This is a so-called "boy optimal" set of marriages. It turns out that in this set of marriages, every boy gets the best possible spouse that he could possibly get in any set of stable marriages. There's no set of stable marriages in which boy 5 gets a more desirable girl than A. There's no set of stable marriages in which boy 1 gets a girl that's more desirable to him than girl E.

The sad news is that it's simultaneously pessimal for the girls. That is, each girl is getting their worst possible spouse among all sets of stable marriages. We'll examine that further in a minute.

But let me just point out that this is more than a puzzle. I mean, it's fun, and it's a nice puzzle, but it's more than a puzzle. Because the original case where it was studied or published first was in a paper by Gale and Shapley in 1962. You may remember the name David Gale from the subset game that we played early in the term when we were practicing with set relations.

And what they were dealing with was the problem of college admissions where students have rankings of colleges that they've applied to and their preferences and colleges have rankings of students that have applied to them. And we're trying to get matching up between college offers and student preferences.

And in a circumstance where a college made an offer and a student sort of accepted, but then later the student got another offer from a place they preferred more, and they were changing their mind, and withdrawing acceptances and so on, it was making everybody crazy, the administrators and the students themselves. And the desire was, let's get some stable set of offers on the table. And Gale and Shapley proposed a protocol to get stable marriages that would apply to college admissions.

It turns out, interestingly enough, that although Gale and Shapley are credited with the stable marriage solution that we're going to discuss, they did that because they were the first to publish it. But, in fact, it had been discovered and put into practice at least 20 years earlier by a national board whose job was to match interns in hospitals, that is graduating medical students who were about to start their further clinical training as interns, now called residents in modern language, had to be matched up with hospitals. And the residents had preferred hospitals that they'd like to go to, and the hospitals had rankings of residents that met their criteria.

And again, the issue was, how do you assign residents to hospitals in a stable way? Before they had discovered this stability algorithm, it was a mess. Again, there's a wonderful story in the book by Gusfield and Irving that is an entire book about the stable marriage problem published by *MIT Press*, I think in '89. And as a matter of fact, I was the editor of the series in which it appears. This stable marriage problem turns out to have a lot of structure.

And they describe a wonderful anecdote in their preface about the problems that were happening between the hospitals and the residents and the measures that were taken to try to achieve stability before they discovered this algorithm.

Another genuine computer science application is one that was described by Tom Leighton who was a co-author of the text and then now the CEO of Akamai, an internet infrastructure plumbing company that has some large number of servers, I think 65,000 servers in 2010 around the world. And it basically is providing cached web

pages so that they can respond more quickly to local calls.

They have a problem that they get something like 62,000, 65,000 servers and they get I think 200 billion web requests a day. So the web requests are kind of acting like the boys and the servers are kind of acting like the girls, or the hospitals and the residents. And the web requests have preferences based on proximity and speed of server, and the servers have preference based on where they're located and the magnitude of the web request that they're coming to.

And the question is, how do you assign web requests to servers so that things get done expeditiously? And it turns out that the stable marriage method gave a satisfactory way to accomplish that kind of matching. And in particular, because there's such large numbers involved that the stable marriage ritual, which we'll describe shortly, is very amenable to being run in parallel.

Another application, it turns out, to come up was in matching dance partners when I was teaching this course ten years ago with a co-instructor who was a member of the Indian dance team.

She said, we could use this, because it turns out that, again, there are boy and girl partners in the dance, and it was constantly the case that one boy would like another boy's partner better, and vice versa. And they would start pairing up and leaving the other people hanging and there were bad feelings, and it was a source of disruption in the society.

There's a picture of that Indian dance group. My co-instructor's not actually there, but it gives you some sense of the reality of the problem here at MIT. And she told me that it was actually being used by that group.