

PROFESSOR: The example of scheduling courses in terms is really a special case of a general problem that you can probably see of scheduling a bunch of tasks or jobs under constraints of which ones have to be done before other ones, which is a topic that comes up actually in lots of applications. But you can see applications in computer science where you might have a complex calculation, pieces of which could be done in parallel and other parts had to be done in order because later results depended on the results of an earlier computation.

It leads us to the general discussion of parallel scheduling.

And we've already worked out some theory of that really just from the example. Namely, if we look at the minimum number of terms to graduate, this corresponds to the minimum amount of number of stages or the minimum amount of time that it takes to process a bunch of tasks, assuming that you can do tasks in parallel and as many in parallel as you need to-- that there's no limit on the amount of parallelism allowed.

In that case, what we can say is that the minimum parallel time for a bunch of constrained tasks is simply the maximum chain size in the constraint graph.

We saw that example with the course prerequisites where we had five. And in general, this is the theorem. Minimum parallel time is exactly equal to maximum chain size for chains in the graph that constrains the order in which tasks can be completed.

Now what about the maximum term load? Well, that corresponds to the number of processors you need to be doing tasks in parallel.

So for the course scheduling example, it means how many subjects can you take in one term? But if you were, say, doing computations, how many separate CPUs would you need in order to be able to fully utilize the parallelism to as much in parallel as you possibly could and abound on the number of processors that are needed for minimum time is simply the maximum antichain size, which in the example from the previous segment on course scheduling, it turns out there were five courses you could take in one term, the second term. And that was, in fact, the maximum antichain size.

So that's an upper bound on the number of processors that you need to achieve minimum time. But in fact, it's a course upper bound because although the number of processors

needed to achieve minimum parallel time is at most the maximum antichain size. In fact, in the previous example, it turns out you could get away with three processors. It was possible to schedule the subjects so you only took three courses a term and still finished in minimum time.

So can you do better than three subjects? Well, there's a trivial argument that says, no, you can't. Because in that previous example, we had 13 subjects to schedule. The maximum chain size was 5. So it was going to take at least five terms.

So that means you have to distribute these 13 subjects among five terms. There has to be some term that has at least the average number of subjects, namely 13 divided by 5.

So that means there has to be a term in which you're taking 13 divided by 5 subjects. Of course, you round up because it has to be an integer.

So the minimum number of terms to finish and graduate-- finishing these 13 subjects in five terms-- is 3 because 13 divided by 5 rounded up is 3.

And this is a general phenomenon that applies. And what we can say is that if you have a DAG with n vertices and the maximum chain size is c -- so that's how deep it can be at most-- and the maximum antichain size is a -- that's the largest number of things that you could ever possibly do in parallel-- then clearly, the total number of vertices is c times a , at most.

So the total number of tasks that you can do where you are going to finish in c steps using at most a processors is bounded by c times a .

So what that tells you is that you can't both have the antichain size and the chain size be too small because their product has to be at least n .

That can be rephrased as a lemma that is credited to a guy named Dilworth. Dilworth is actually famous for Dilworth's theorem of which this Dilworth's lemma is a special case, but we don't need the general theorem.

Dilworth's lemma says that if you have an n -vertex DAG, then for any number t , it either has a chain of size bigger than t , or it has an antichain of size greater than or equal to n over t . And we proved this on the previous slide.

The product of these two things has to be at least n , and the general case is t times n over t is greater than or equal to n . And this holds for all t between 1 and n .

Well, let's think of a simple application of that. If I choose the t that balances antichain size from chain size, then I choose t to be the square root of n .

So over here, I have square root of n , and here I have n divided by the square root of n , which is also square root of n .

And what we can conclude is that every end vertex DAG has either a chain of size at least the square root of n or an anti chain of size at least square root of n .

This turns out to actually have a few applications, but we're just going to look at a fun application of this remark that you have to have a chain or an antichain of size at least square root of n .

You might have only one of these. You might have both. But one or the other has to be at least as big as square root of n .

Let's think of a new DAG that I'm going to construct as follows. I'm going to draw an edge between students in the class, and I'm going to think of one student as having a direct edge to another student if the first student is both shorter and younger-- actually meaning no taller than and no older than the other. Let's just say shorter-- meaning shorter or possibly the same height-- younger or possibly the same age.

And so the rule is if I think of a student as being represented by their shortness s and their age a , then a student with a height s_1 and age a_1 has a direct arrow to another student with height s_2 and age a_2 , providing that the first pair is less than or equal to the second pair in both coordinates. S_1 is less than or equal to s_2 . And A_1 is less than or equal to A_2 .

Now, we don't want ties here because that would break the DAG property if I have two students with exactly the same age and height. So let's assume that we're measuring age in microseconds and the height in micrometers. And with that kind of a fineness, the likelihood of a tie is pretty low. So then it becomes a DAG again.

So this is the definition of taking a DAG built out of pairs-- there's a pure DAG for height, and there's a pure DAG for age. And I combine them into pairs, and I get a new DAG by looking at how the coordinates behave together.

This is called the product graph. It's a general construction that comes up, and we will talk a little bit more about when we reexamined DAGs in the context of the language of relations and

partial orders.

Anyway, this is the product graph.

According to Dilworth's lemma, in a class like ours of 141 students, it means that we're going to have a chain or an antichain in this product DAG of size square root of 141 rounded up, or 12.

According to Dilworth's lemma, in this particular age-height graph, what does it mean for this to be an antichain? Suppose I take a bunch of students and I line them up in order of size with the tallest on the left and the shortest on the right.

If this is going to be an antichain, it means that they have to be getting older as they get shorter. Because if I ever had a case where somebody to the right was both shorter and younger than somebody to the left, it wouldn't be an antichain because they'd be comparable.

So an antichain-- according to Dilworth's lemma-- if you'd sort the students by height, they have to be getting older and as they get shorter. If it was a chain, they would be getting younger as they got shorter. But the more interesting one is the antichain in this height-birthday example.

So we should be looking at-- we'll either have a chain or an antichain in this class according to this product DAG.

As a matter of fact, we really had an antichain. Here's a quick list of a dozen students. And indeed, if you look at the birthdays, there's somebody who's 6'1 was born in August '94 and then somebody who was born in April '94 and is 6'0 all the way down to somebody who was born in 1991 and who's five feet tall.

So we lucked out. We could have only had to chain, but we actually had the antichain in this case.