

Tractability

15-110 – Wednesday 03/11

Quizlet5

Learning Goals

- Identify **brute force algorithms** to solve common problems that run in $O(n!)$ or $O(2^n)$, including solutions to **Travelling Salesperson, puzzle-solving, subset sum, Boolean satisfiability**, and **exam scheduling**
- Define the complexity classes **P** and **NP** and explain why these classes are important
- Identify whether an algorithm is **tractable** or **intractable**, and whether it is in **P**, **NP**, or **neither** complexity class
- Use **heuristics** to find good-enough solutions to NP problems in polynomial time

Big Idea: What *is* Efficient?

As we wrap up the unit on data structures and efficiency, we still need to answer a big question: **can all algorithms be made efficient?** And, importantly, **what does it mean to be efficient?**

To answer these questions, we'll consider a collection of important computational problems. While considering these problems, ask yourself: how efficient are these solutions? Could we make them better?

Computationally Difficult Problems


Example: Travelling Salesperson Problem

First, consider the **Travelling Salesperson problem**.

The program is given a **graph** that represents a map – nodes are cities, edges are distances between cities.

The goal is to find the **shortest possible route** that visits every city, then returns home.

Practical application: plan a route for a postal worker.




One Solution: Check All Paths

Intuitive algorithm: try **every possible route** from the starting city across all the others, then choose the shortest route of them all.

For example, starting from Pittsburgh in the graph to the right we have three possible first-stops. Each of those has two second-stop options, leading to six total possibilities.

When we compare the routes, the shortest route is PIT->DC->BALT->PHIL->PIT (or its reverse, PIT->PHIL->BALT->DC->PIT).



Brute Force Algorithms

This type of solution approach is called a **brute force algorithm**. Brute force algorithms are simple: you just generate every possible solution and check each of the generated solutions to see if any of them work based on the problem's constraints.

Brute force algorithms are easy to understand, implement, and test. They also apply to a wide range of problems, which makes them useful.

However, brute force algorithms have one major drawback: their **efficiency**.

Brute Force Efficiency

Consider the efficiency of our Travelling Salesperson algorithm. Let's say that generating a path of n stops counts as one action. How many possible paths are there in the worst case?

The worst case is a fully-connected graph (like the previous one). We have $n-1$ possible first stops on the route. For each of those routes, there are $n-2$ possible second stops, or $(n-1)*(n-2)$ routes so far. Then there are $n-3$ third stops per route, etc... until there is only one city left for the last stop.

This means that the number of possible routes is $(n-1) * (n-2) * (n-3) * \dots * 1$. **It's $O(n!)$. That's really inefficient!**


There are a lot of problems in computer science that share this property: we can solve them, but the intuitive algorithm takes a long time. Let's go through some examples.

Example: Puzzle Solving

Say we want to solve a basic puzzle by putting together square pieces (like the ones shown to the right) so that any two pieces that are touching each other make a figure with a head and feet of the same color.

To make this even simpler, let's make a rule that pieces cannot be rotated and the final result must be a $m \times m$ square.


Here's our question: given a set of pieces, is it possible to make a solution that follows these rules?



Brute Force on Puzzle Solving

We can again use brute force to solve the puzzle problem, just like we did with Travelling Salesperson. We can do this by trying all possible pieces for each location.

In the example to the right there are 9 options for the first position, 8 for the second, 7 for the third, etc.... it's **$O(n!)$** time again.



9 choices	8 choices	7 choices
6 choices	5 choices	4 choices
3 choices	2 choices	1 choice

$O(n!)$ is Really Bad

It turns out that $O(n!)$ is a *really bad* runtime. For example, let's assume that it takes 1 millisecond ($1/1000^{\text{th}}$ of a second) to set up a specific ordering of pieces of a puzzle and check if it's correct.

If we have 9 pieces (like in our example before), it will take **6.048 minutes** to solve the puzzle.

If we increase the size to a 4x4 puzzle (16 pieces), it will take **663.46 years!**

$O(n!)$ is awful. Let's see if we can find problems that do a bit better.

Example: Subset Sum

In the problem Subset Sum we are given a list of numbers and a target number, x . We want to determine if there's a subset of the list that sums to x .

Brute force algorithm: generate all possible subsets, see if any of them sum to x .

How do we generate all subsets? Use recursion! If we have all four subsets of the list $[2, 3]$ we can use them to create all 8 subsets of $[1, 2, 3]$. For each subset, make one version that includes 1 and one version that doesn't.


We double the number of subsets with each new number that is added, so this is $O(2^n)$.

Subsets of $[2, 3]$:

- $[]$
- $[2]$
- $[3]$
- $[2, 3]$

Subsets of $[1, 2, 3]$:

- $[]$
- $[1]$
- $[2]$
- $[1, 2]$
- $[3]$
- $[1, 3]$
- $[2, 3]$
- $[1, 2, 3]$



Example: Boolean Satisfiability


A similar problem commonly encountered in computer science, called **Boolean Satisfiability**, asks: for a given circuit with n inputs (X_1 to X_n), is there a set of assignments of X_i to 0 or 1 that makes the whole circuit output 1?

Instead of generating all possible subsets, we generate all possible combinations of input values (like generating a truth table!).

This also doubles every time we add a new input as we must try all possible combinations with the input set to 0, then set to 1. It's still $O(2^n)$.

Inputs for 2 elements

Inputs for 3 elements



Real-life Example: Exam Scheduling

Here's one final example: scheduling final exams. Given a list of classes, a dictionary mapping students to their classes, and a list of timeslots over the period of a week, generate a schedule that fits within the period and results in no student having two exams in the same slot.

We can generate all possible schedules using a similar approach to subset sum. Then we just need to look for one schedule that has no conflicts by checking every student. However, every time we add a new class we need to try adding it to **every** possible schedule in **every** possible timeslot.

If we say there are k timeslots (where k is some constant number) and n classes, we turn one schedule into k different schedules for every new class added. This is $O(k^n)$!

Semester & Mini-2 Final Exams: December 9, 10, 12, 13, 15 & 16(Make-Up Day)					
Course	Section	Title	Date	Time (USA EST)	Classroom(s)
Architecture					
48116	A	BUILDING PHYSICS	Sunday, December 15, 2019	01:00 pm - 04:00 pm	To Be Announced (TBA)
48315	1	ENVIR I: CLIM & ENG	Thursday, December 12, 2019	08:30 am - 11:30 am	To Be Announced (TBA)
48432	A	ENV II	Thursday, December 12, 2019	08:30 am - 11:30 am	To Be Announced (TBA)
48531	A	FABRICATNG CUSTOMZTN	Monday, December 9, 2019	01:00 pm - 04:00 pm	To Be Announced (TBA)
48558	A	RLT COMP	Thursday, December 12, 2019	08:30 am - 11:30 am	To Be Announced (TBA)
48568	A	ADV CAD BIM 3D VISLZ	Tuesday, December 10, 2019	08:30 am - 11:30 am	To Be Announced (TBA)
48635	1	ENVIRO I MARCH	Thursday, December 12, 2019	08:30 am - 11:30 am	To Be Announced (TBA)
48655	A	ENV II GRAD	Thursday, December 12, 2019	08:30 am - 11:30 am	To Be Announced (TBA)
48714	A	DATA ANL URBN DSN	Friday, December 13, 2019	01:00 pm - 04:00 pm	To Be Announced (TBA)
48729	A	PROD HLTH QUAL BLDGS	Thursday, December 12, 2019	01:00 pm - 04:00 pm	To Be Announced (TBA)
48734	A	RCTV SP MD ARC	Friday, December 13, 2019	05:30 pm - 08:30 pm	To Be Announced (TBA)
48743	A	INTRO ECO DES	Friday, December 13, 2019	01:00 pm - 04:00 pm	To Be Announced (TBA)
48749	A	CD SPECIAL TOPICS	Tuesday, December 10, 2019	01:00 pm - 04:00 pm	To Be Announced (TBA)
48785	A	MAAD RES PROJ	Sunday, December 15, 2019	05:30 pm - 08:30 pm	To Be Announced (TBA)
48798	A	HVAC & PS LOW CARB B	Monday, December 9, 2019	05:30 pm - 08:30 pm	To Be Announced (TBA)
Art					
60157	A	DRAWING NON-MAJORS	Tuesday, December 10, 2019	05:30 pm - 08:30 pm	CFA TBD
60218	A	REAL-TIME ANIMATION	Monday, December 9, 2019	08:30 am - 11:30 am	To Be Announced (TBA)
60220	A	TECH CHARACTER ANIM	Thursday, December 12, 2019	05:30 pm - 08:30 pm	To Be Announced (TBA)
60220	B	TECH CHARACTER ANIM	Thursday, December 12, 2019	05:30 pm - 08:30 pm	To Be Announced (TBA)
60333	A	CHARACTER RIGGING	Sunday, December 15, 2019	08:30 am - 11:30 am	BH 140F

$O(2^n)$ and $O(k^n)$ are Still Really Slow

$O(2^n)$ is a bit better than $O(n!)$, but not *that* much better. Let's say we want to solve the subset sum problem and it again takes us 1 millisecond to generate a specific subset and see if it is equal to the target.

If $n = 10$, we solve the problem in **1.024 seconds**. Much better!

But if $n = 20$, we solve the problem in **17.48 minutes**...

And if $n = 30$, it will take us **12.43 days**. By the time $n = 40$, it takes **35 years**.

$O(2^n)$ is not as bad as $O(n!)$, but it's still really bad.

Tractability

This leads us to a new concept: **tractability**. A problem is said to be **tractable** if it has a reasonably efficient runtime so that we can use it for practical input sizes.


We say that a runtime is reasonable if it can be expressed as a **polynomial** equation. This means an equation of the form:

$$c_k x^k + c_{k-1} x^{k-1} + \dots + c_1 x + c_0$$

where x is a variable and c_i & k are constants.

$O(1)$, $O(\log n)$, $O(n)$, $O(n^2)$, and $O(n^k)$ are all tractable. $O(2^n)$, $O(k^n)$, and $O(n!)$ are not - they're **intractable**.

We can see the difference in growth quickly using the graph to the right.



Caveat: logarithms are tractable even though they aren't polynomial, because they're faster than $O(n)$.

Activity: Identify the Solution Runtime

If you consider how a brute-force solution generates solutions, and how that algorithm would be affected by increasing the input size, you can often determine whether the solution will be tractable or intractable without digging deeply into the exact runtime.

You do:

- solve a $n \times n$ Sudoku puzzle by trying every possible combination of numbers. Is that tractable or intractable?
- check every pair of elements in a n -element list to see if there are any duplicates. Is that tractable or intractable?

Complexity Classes

Goal: Find Tractable Solutions

Now we know just how bad the brute-force solutions to this set of problems are when it comes to efficiency. Maybe we can design a different algorithm that doesn't require us to generate every possible answer.

That will be our goal for the rest of the lecture: to see if we can find a **tractable** solution to these hard problems.

Designing New Solutions is Hard!

Discuss: when we talked about the brute-force solutions to the previous problems, did you think of a way to solve the problems a lot faster?

You might have come up with ways to shave some time off the algorithms, but most likely your new solutions are still **intractable**. Coming up with fast solutions to these problems is hard!


Until now, we've only discussed how long it takes to **find** the solution to a problem. Let's take a different approach.

Magical Schedule-Making Box

Suppose a magical black box descends from the sky onto campus one day.

Someone discovers that if you feed the box a list of all the classes in a semester, all the final exam timeslots, and every student's schedule, the box will spit out a final exam schedule for CMU.

If CMU has n classes, how long would it take us to check if this schedule has any conflicts in it?



Semester & Mini-2 Final Exams: December 9, 10, 12, 13, 15 & 16(Make-Up Day)					
Course	Section	Title	Date	Time (USA EST)	Classroom(s)
Architecture					
48116	A	BUILDING PHYSICS	Sunday, December 15, 2019	01:00 pm - 04:00 pm	To Be Announced (TBA)
48315	1	ENVIR I: CLIM & ENG	Thursday, December 12, 2019	08:30 am - 11:30 am	To Be Announced (TBA)
48432	A	ENV II	Thursday, December 12, 2019	08:30 am - 11:30 am	To Be Announced (TBA)
48531	A	FABRICATING CUSTOMZTN	Monday, December 9, 2019	01:00 pm - 04:00 pm	To Be Announced (TBA)
48558	A	RLT COMP	Thursday, December 12, 2019	08:30 am - 11:30 am	To Be Announced (TBA)
48568	A	ADV CAD BIM 3D VISLZ	Tuesday, December 10, 2019	08:30 am - 11:30 am	To Be Announced (TBA)
48635	1	ENVIRO I MARCH	Thursday, December 12, 2019	08:30 am - 11:30 am	To Be Announced (TBA)
48655	A	ENV II GRAD	Thursday, December 12, 2019	08:30 am - 11:30 am	To Be Announced (TBA)
48714	A	DATA ANL URBN DSNG	Friday, December 13, 2019	01:00 pm - 04:00 pm	To Be Announced (TBA)
48729	A	PROD HLTH QUAL BLDGS	Thursday, December 12, 2019	01:00 pm - 04:00 pm	To Be Announced (TBA)
48734	A	RCTV SP MD ARC	Friday, December 13, 2019	05:30 pm - 08:30 pm	To Be Announced (TBA)
48743	A	INTRO ECO DES	Friday, December 13, 2019	01:00 pm - 04:00 pm	To Be Announced (TBA)
48749	A	CD SPECIAL TOPICS	Tuesday, December 10, 2019	01:00 pm - 04:00 pm	To Be Announced (TBA)
48785	A	MAAD RES PROJ	Sunday, December 15, 2019	05:30 pm - 08:30 pm	To Be Announced (TBA)
48798	A	HVAC & PS LOW CARB B	Monday, December 9, 2019	05:30 pm - 08:30 pm	To Be Announced (TBA)
Art					
60157	A	DRAWING NON-MAJORS	Tuesday, December 10, 2019	05:30 pm - 08:30 pm	CFA TBD
60218	A	REAL-TIME ANIMATION	Monday, December 9, 2019	08:30 am - 11:30 am	To Be Announced (TBA)
60220	A	TECH CHARACTER ANIM	Thursday, December 12, 2019	05:30 pm - 08:30 pm	To Be Announced (TBA)
60220	B	TECH CHARACTER ANIM	Thursday, December 12, 2019	05:30 pm - 08:30 pm	To Be Announced (TBA)
60333	A	CHARACTER RIGGING	Sunday, December 15, 2019	08:30 am - 11:30 am	BH 140F

Verifying a Final Exam Schedule

For every student, we need to go through all pairs of their classes to see if any of their classes are in the same timeslot. Each student is likely enrolled in no more than 5 classes, so that's a constant number of checks – 10.

How many students are there? We can probably find a constant relation between the number of classes in a semester and the number of students enrolled. Let's say if there are n classes, there are $6*n$ students.

That means that overall we have to do $\text{students} * \text{conflict-checks} = (6*n) * 10$ work. That's $60n$, which is $O(n)$. **Verifying the solution is tractable!**

Complexity Classes

Now that we've talked about both solving and verifying problems, we can start putting problems into different groups.

We call these groups **complexity classes**. These are collections of problems that have similar efficiency. For this class, we'll say that every algorithm in a certain complexity class is **bounded by a certain runtime**.


For example, we could design a complexity class called 'Fast' that only includes algorithms which run in $O(n)$ time or faster. This would also include $O(\log n)$ and $O(1)$.

Complexity Class P

First we define the complexity class P to be **the set of problems that we know can be solved in polynomial time**. Recall that an algorithm is polynomial if it can be expressed as:

$$c_k x^k + c_{k-1} x^{k-1} + \dots + c_1 x + c_0$$

Our earlier examples (subset sum, puzzle solving, exam scheduling) don't fall into this category yet. But plenty of other algorithms do, such as linear search, summing a list, etc.




Complexity Class NP

Next we define the complexity class NP to be **the set of problems that can be verified in polynomial time.**

This includes all problems in P, because if you can solve something in polynomial time, you can check it as well.

It also includes most of the problems we discussed before! We already showed that we can check exam scheduling in linear time. We can also check subset sum, Boolean satisfiability, and puzzle solving this way.




Not all Problems are in P or NP

Some problems are so difficult we can't even verify them in polynomial time.

Travelling Salesperson is an example of this. If we're given a solution, we can't verify that it's the **best** path; it's just one possible path that exists. In general, trying to find the 'best' solution takes a long time to verify.

We can turn Travelling Salesperson into an NP problem by changing the prompt: instead of finding the best path, just try to find a path that is less than X total distance for some number X. This is easy to verify.



Quick P/NP Guide


	P	NP	All Problems
Solve in Polynomial Time?	Yes!	???	???
Verify in Polynomial Time?	Yes!	Yes!	???

P vs NP

Big Question: Does $P = NP$?

Here's our big idea for the day.
Wouldn't it be nice if the set of problems P was the same as the set of problems NP ?

If this was true, we could find an algorithm that would put together CMU's final exam schedule in a day instead of waiting half a semester to find out when exams will happen. We'd be able to solve a lot of hard problems really quickly without having to think hard about clever new approaches!



Does $P = NP$? We Don't Know.

Whether or not $P = NP$ is a core question in the field of computer science, but it's still unsolved.

The first person who proves whether or not $P = NP$ will win [a million dollars](#), but no one has proved it yet...

Proving $P \neq NP$

Let's assume that $P \neq NP$. How would we prove this?

You'd need to definitively prove that a problem in NP exists that **cannot** be solved in polynomial time. But how can we show that it's impossible to come up with a clever new algorithm?

This is tricky!

Proving $P = NP$

Let's assume $P = NP$. How would we prove this?

You need to show that **every** problem in NP can be solved in polynomial time. That's a lot of problems!

To make this easier, computer scientists try to find problems in NP that are related to each other.

Transforming Problems


Consider subset sum and Boolean satisfiability. We can **transform** subset sum into satisfiability. We just need to make a circuit that uses each value in the list as an input (0 if it isn't included, 1 if it is) and make the circuit output 1 if the included values sum to the target.

In fact, this mapping can be done in **polynomial time**. This means that if we can find a tractable solution to Boolean satisfiability, we can also use it to make a tractable solution to subset sum.

Find a subset of [4, 2, 7, 13] that sums to 8



Set the inputs so that the circuit outputs 1



Useful NP Problems

Computer scientists have identified a set of problems that have this problem-transformation capacity for **all** NP problems. If we can find a tractable solution to one of them, **we can make all problems in NP tractable**. That will mean that $P = NP$!

In fact, if you use the limited version of the Travelling Salesperson problem, **all the problems we discussed today are in this set of problems**.

If you can find a tractable solution to any of these problems, you'll prove $P = NP$ and will become rich and famous.

Possible Outcomes

What happens if we prove $P = NP$?

We'll be able to solve a lot of hard problems very quickly. NP problems show up everywhere, so nearly everything in the world will get radically faster!

On the other hand, this might also wreck how modern security and encryption is implemented (as it will get easier to break cryptography).

What happens if we prove $P \neq NP$?

Not much; we'll still be in our current situation. But a lot of computer scientists can turn their focus to other problems.

Most people think $P \neq NP$, but we don't know how to prove it.

Heuristics

Speeding Up Slow Algorithms

There are lots of useful problems that fall into the NP class. How can we solve these problems practically when they're intractable?

Instead of trying to find the **ideal** solution to a problem, we can change our standards to say we only need to find a **good-enough** solution. For example:

- In exam scheduling, maybe it's okay if there's a small number of conflicts that affect $< 1\%$ of the student body
- In subset sum, maybe it's okay if we find a subset that is *almost* equal to the target, instead of exactly equal

When we're willing to compromise on optimality or accuracy, or put other restrictions on the data, we can use **heuristics** to speed up the process a great deal.

Heuristics Provide Approximate Answers


A **heuristic** is a search technique used by an algorithm to find a **good-enough solution** to a problem. Heuristics may not find the best answer to an NP problem, but they often achieve good results.

A heuristic can generate **scores** to rank potential next steps that the algorithm can take at each decision point. By choosing the highest-scored next step, the algorithm is more likely to find a working solution quickly.

Heuristics Example: Travelling Salesperson

For the Travelling Salesperson problem, we could generate a heuristic that ranks next-possible paths based on their length. The algorithm can always choose the next city to visit by trying the shorter paths first.

With this approach, we can generate a pretty decent path in **polynomial time**. This path might not be the *best* path, but it's likely better than a random path.



Example: Applying a Heuristic

Let's try applying a heuristic to **subset sum**. Order all the values in the list from largest to smallest. Always try adding the largest available value to the subset first.

We'll simplify the problem further by assuming all values in the list are positive, so as soon as the subset is larger than the target, we can backtrack and try something else.

We'll also sacrifice some optimality by accepting any answer that comes within 2 of our desired target value.

How many subsets do we need to try to determine if there's a subset of [13, 14, 7, 10, 7, 16, 2, 8, 3, 5] that sums to ~25?

Sort the list: [16, 14, 13, 10, 8, 7, 7, 5, 3, 2]

[16] – too small

[16, 14] – too big, backtrack!

[16, 13] – still too big...

[16, 10] – this is 26, it works!

We missed the optimal solution – [16, 7, 2] would have been perfect. But we found [16, 10] much faster.

Sidebar: Additional Watching

Want to learn more about these topics? Check out the following videos recommended by prior students!

P vs. NP and the Computational Complexity Zoo:

<https://www.youtube.com/watch?v=YX40hbAHx3s>

P vs. NP - The Biggest Unsolved Problem in Computer Science:

<https://www.youtube.com/watch?v=EHp4FPyajKQ>

Learning Goals

- Identify **brute force algorithms** to solve common problems that run in $O(n!)$ or $O(2^n)$, including solutions to **Travelling Salesperson**, **puzzle-solving**, **subset sum**, **Boolean satisfiability**, and **exam scheduling**
- Define the complexity classes **P** and **NP** and explain why these classes are important
- Identify whether an algorithm is **tractable** or **intractable**, and whether it is in **P**, **NP**, or **neither** complexity class
- Use **heuristics** to find good-enough solutions to NP problems in polynomial time