P and NP

CSE 332 25Sp
Lecture 26,27

Announcements

	Monday	Tuesday	Wed	Thursday	Friday
This Week	Ex 14 (P/NP, GS) out **TODAY**		Ex 13 (MST,prog) due		Ex 14 due
Next Week				Final! (12:30-2:20)	

Last call for conflict exam requests (will process them tomorrow)

↳ Extra video with more concurrency examples (bad interleavings) on Panopto tonight. (Optional)

Kruskal's Algorithm: Running Time

```
KruskalMST(Graph G)
   initialize each vertex to be a connected component
   sort the edges by weight
   foreach(edge (u, v) in sorted order){
     if(u and v are in different components){
        add (u,v) to the MST
        Update u and v to be in the same component
     }
   }
```

Kruskal's Algorithm: Running Time

How do we find connected components? Well BFS is our existing tool to do that, but...

Running a new BFS in the partial MST, at every step seems inefficient. The answer changes little by little, so we'll recompute work frequently.

Do we have an ADT that will work here?

Union-Find Crash Course

aka Disjoint Sets

Represents...well...disjoint sets.

Union-Find ADT

state

Set of Sets
- **Disjoint:** No element appears in multiple sets
- No required order
- Each set has representative

behavior

makeSet(x) – creates a new set where the only member (and the representative) is x.

findSet(x) – looks up the set containing element x, returns name of that set

union(x, y) – combines sets containing x and y. Picks new name for combined set.

Union-Find Running Time

What's important for us?

Amortized running times! We care about the total time across the entire set of unions and finds, not the running time of just one.

Operation	Worst-case Amortized	Worst-case Non-amortized
MakeSet()	$\Theta(1)$	$\Theta(1)$
Union()	$O(\log^* n)$	$O(\log n)$
Find()	$O(\log^* n)$	$O(\log n)$

Uses "forest of up-trees" implementation.

$\log^* n$

$\log^* n$

the number of times you need to apply $\log()$ to get a number at most 1.

E.g., $\log^*(16) = 3$

$\log(16) = 4$ $\log(4) = 2$ $\log(2) = 1.$

$\log^* n$ grows **ridiculously** slowly.

$\log^*(10^{80}) = 5.$

For all practical purposes these operations are constant time.
They're not constants (don't delete them from big-O notation), but you
will never worry about these in figuring out how many seconds a piece
of code takes.

Using Union-Find

Have each disjoint set represent a connected component
- A connected component is a "piece" of a (disconnected) undirected graph
- i.e. a vertex, and everything you can reach from that vertex.

When you add an edge, you **union** those connected components.

Try it Out

Operation	Worst-case Amortized
MakeSet()	$\Theta(1)$
Union()	$O(\log^* n)$
Find()	$O(\log^* n)$

```
KruskalMST(Graph G)
  initialize each vertex to be a connected component
  sort the edges by weight
  foreach(edge (u, v) in sorted order){
    if(find(u) != find(v)){
      add (u,v) to the MST
      union(find(u),find(v))
    }
  }
}
```

Running Time?

Operation	Worst-case Amortized
MakeSet()	$\Theta(1)$
Union()	$O(\log^* n)$
Find()	$O(\log^* n)$

```
KruskalMST(Graph G)
    initialize each vertex to be a connected component
    sort the edges by weight
    foreach(edge (u, v) in sorted order){
        if(find(u) != find(v)){
            add (u,v) to the MST
            union(find(u),find(v))
        }
    }
}
```

$E \log E$

$E \log^* V$

V

$V \log^* V$

$E \log E$ is the dominant term, but you'll usually see this written

$$\Theta(E \log V)$$

Since $E \le V^2$, those are equivalent.

G 50 B 6
2 3 E
4 C 5
A
2 7 9
7
D 8 F

Try it Out

Edge	Include?	Reason

Try it Out

Edge	Include?	Reason
(A,B)	Yes	
(C,D)	Yes	
(B,F)	Yes	
(A,C)	Yes	
(C,E)	Yes	
(B,E)	No	Cycle A,C,D,B,A
(A,D)	No	Cycle A,D,C
(D,E)	No	Cycle C,D,E
(D,F)	No	Cycle A,B,F,D,C,A
(E,F)	No	Cycle E,F,B,A,C,E
(B,G)	Yes	

Some Extra Comments

Prim was the employee at Bell Labs in the 1950's

The mathematician in the 1920's was Boruvka
- He had a different *also greedy* algorithm for MSTs.
- Boruvka's algorithm is trickier to implement, but is useful in some cases.
- In particular it's the basis for fast **parallel** MST algorithms.

If all the edge weights are distinct, then the MST is unique.

If some edge weights are equal, there may be multiple spanning trees. Prim's/Kruskal's are only guaranteed to find you one of them.

Aside: A Graph of Trees

A tree is an undirected, connected, and acyclic graph.

How would we describe the graph Kruskal's builds.

It's not a tree until the end.

It's a forest!

A forest is any undirected and acyclic graph

EVERY TREE IS A FOREST.

P vs. NP

Taking a Big Step Back

What has this quarter been about?

We've taken problems you probably knew how to solve slowly,

And we figured out how to solve them faster.

You could have written a linked-list-dictionary. Would have been $O(n)$ not $O(\log n)$ for most operations, but it would've worked (eventually).

In some sense, that's the job of a computer scientist.

Figure out how to take our problems

And make the computer do the hard work for us.

Taking a Big Step Back

Let's take a big step back, and try to break problems into two types:

Those for which a computer might be able to help.

And those which would take so long to solve **even on a computer** we wouldn't expect to solve them.

This is not the same as asking for undecideable problems (like the Halting Problem you saw in 311).

There are problems we could solve in finite time…but we'll all be long dead before our computer tells us the answer.

Running Times

	n	$n \log_2 n$	n^2	n^3	1.5^n	2^n	$n!$
$n = 10$	< 1 sec	< 1 sec	< 1 sec	< 1 sec	< 1 sec	< 1 sec	4 sec
$n = 30$	< 1 sec	< 1 sec	< 1 sec	< 1 sec	< 1 sec	18 min	10^{25} years
$n = 50$	< 1 sec	< 1 sec	< 1 sec	< 1 sec	11 min	36 years	very long
$n = 100$	< 1 sec	< 1 sec	< 1 sec	1 sec	12,892 years	10^{17} years	very long
$n = 1,000$	< 1 sec	< 1 sec	1 sec	18 min	very long	very long	very long
$n = 10,000$	< 1 sec	< 1 sec	2 min	12 days	very long	very long	very long
$n = 100,000$	< 1 sec	2 sec	3 hours	32 years	very long	very long	very long
$n = 1,000,000$	1 sec	20 sec	12 days	31,710 years	very long	very long	very long

(somewhat old) table from Rosen. How big of a problem can we solve for an algorithm with the given running times.
"very long" means more than 10^{25} years.

Efficient

We'll consider a problem "efficiently solvable" if it has a polynomial time algorithm.

I.e. an algorithm that runs in time $O(n^k)$ where k is a constant.

Are these algorithms always actually efficient?

Well.........no

Your n^{10000} algorithm or even your $2^{2^{2^{2^2}}} \cdot n^3$ algorithm probably aren't going to finish anytime soon.

But these edge cases are rare, and polynomial time is good as a low bar
- If we can't even find an n^{10000} algorithm, we're probably not getting one that is efficient in practice anyway.

Definition Dump

Some definitions

A problem is a set of inputs and the correct outputs.

"Find a Minimum Spanning Tree" is a problem.
-Input is a graph, output is the MST.

-"Tell whether a list is sorted" is a problem.
-Input is an array, output is "yes" or "no"

-"Sort this array" is a problem.
-Input is an array, output is the same numbers, now in sorted order.

Some definitions

An <u>instance</u> is a single input to a problem.

A single, particular graph is an instance of the MST problem

-A single, particular graph with vertices s and t is an instance of the Shortest Path problem.

-A single, particular array is an instance of the "is the array sorted?" problem.

13	24	31	30	35	39	51

Not Sorted

Decision Problems

Let's go back to dividing problems into solvable/not solvable.
For today, we're going to talk about **decision problems**.

Problems that have a "yes" or "no" answer.

Why?

Theory reasons (ask me later).

But also most problems can be rephrased as very similar decision problems.

E.g. instead of "find the shortest path from s to t" ask
Is there a path from s to t of length at most k?

P

P (stands for "Polynomial")

The set of all decision problems that have an algorithm that runs in time $O(n^k)$ for some constant k.

The decision version of all problems we've solved in this class are in P.

P is an example of a "complexity class"
A set of problems that can be solved under some limitations (e.g. with some amount of memory or in some amount of time).

Remember the decision part! It's important

Be careful looking through old finals, prior quarters didn't include the decision requirement in the definition!

I'll know it when I see it.

Another class of problems we want to talk about.

"I'll know it when I see it" Problems.

Decision Problems such that:

If the answer is YES, you can prove the answer is yes by
- Being given a "proof" or a "certificate"
- Verifying that certificate in polynomial time.

What certificate would be convenient for short paths?
- The path itself. Easy to check the path is really in the graph and really short.

I'll know it when I see it.

More formally,

NP (stands for "nondeterministic polynomial")

The set of all decision problems such that if the answer is YES, there is a proof of that which can be verified in polynomial time.

Intuitively: you can show me why the answer is yes, and I can verify it.

3-COLOR is an NP problem

"Is there a Spanning Tree of cost at most 25?" is an NP problem

I'll know it when I see it.

More formally,

NP (stands for "nondeterministic polynomial")

The set of all decision problems such that if the answer is YES, there is a proof of that which can be verified in polynomial time.

Intuitively: you can show me why the answer is yes, and I can verify it.

3-COLOR is an NP problem

Give me the coloring (u is red, v is blue,...) and check each edge.

"Is there a spanning tree of cost at most 25?" is an NP problem

Give me the tree, I'll see if it's a spanning tree (run BFS, every vertex visited, no cycles, etc.); and see if the weight is small enough.

I'll know it when I see it.

More formally,

NP (stands for "nondeterministic polynomial")

The set of all decision problems such that if the answer is YES, there is a proof of that which can be verified in polynomial time.

It's a common misconception that NP stands for "not polynomial" Please never ever ever ever say that.

Please.

Every time you do a theoretical computer scientist sheds a single tear.

(That theoretical computer scientist is me)

NP

We can **verify** YES instances of NP problems efficiently, but can we **decide** whether the answer is YES or NO efficiently?

That is, can we do it without the hint?

I.e. can you bootstrap the ability to check a certificate into the ability to find a certificate efficiently?

We don't know.

This is the P vs. NP problem.

P vs. NP

P (stands for "Polynomial")

The set of all decision problems that have an algorithm that runs in time $O(n^k)$ for some constant k.

NP (stands for "nondeterministic polynomial")

The set of all decision problems such that if the answer is YES, there is a proof of that which can be verified in polynomial time.

Claim: P $\subseteq$ NP (do you see why?)

P vs. NP

Some problems in NP we know how to **solve** in polynomial time (solve from scratch, not just verify)

"Is there a spanning tree of cost at most 25?" can be solved with Prim's.
- It's in P.

Other problems we don't know how to solve in polynomial time.

We don't know whether 3-COLOR is in P (most people don't think it is).

But maybe it is, and we just don't know the algorithm.

P vs. NP asks this question in general: does knowing you can verify a solution guarantee that you can find a solution?

EXP

There is an algorithm to solve 3-COLOR, it's just slow

Think of a correct (just inefficient) algorithm to solve 3-COLOR

Generate all 3^n possible colorings, if one of them works, great! Return true.

If none of them work, return false.

This algorithm takes exponential time

EXP

EXP (stands for "Exponential")
The set of all decision problems that have an algorithm that runs in time $O(2^{n^{\wedge}k})$ for some constant k.

3-COLOR is in EXP (we just saw why on the last slide)

So is

Claim: NP $\subseteq$ EXP (do you see why?)

Reductions

Let's say we want to prove that some problem in NP needs exponential time (i.e. that P is not equal to NP).

Ideally we'd start with a really hard problem in NP.

What does it mean for one problem to be harder than another?

Polynomial Time Reducible

We say A reduces to B in polynomial time, if there is an algorithm that, using a (hypothetical) polynomial-time algorithm for B, solves problem A in polynomial-time.

I could solve problem A efficiently, if you give me a library that solves problem B efficiently

Reductions

Polynomial Time Reducible

We say A reduces to B in polynomial time, if there is an algorithm that, using a (hypothetical) polynomial-time algorithm for B, solves problem A in polynomial-time.

If A reduces to B then A should be "easier" than B. (for us as algorithm designers)
- If we can solve B, we can definitely solve A.

Usually denoted $A \leq_P B$.

The Direction Matters!

Direction matters, and is often confusing:
$A \leq B$ "A reduces to B"

I wrote an algorithm to solve problem A using a library designed to solve problem B

"A is no harder than B" (solving B guarantees you can solve A, but maybe there's a different way to solve A)

How do you remember the direction?
Robbie recommends you repeat "Reduction from A to B means writing an algorithm for problem A using an algorithm designed for problem B" to yourself 50 times until it's stuck in your brain.

We reduced shortest paths on (integer-weighted) graphs to shortest paths on unweighted graphs

NP-complete

Let's say we want to prove that some problem in NP needs exponential time (i.e. that P is not equal to NP).

Ideally we'd start with a really hard problem in NP.

What is the hardest problem in NP?

NP-complete

A problem B is NP-complete if B is in NP and
for all problems A in NP, A reduces to B in polynomial time.

NP-complete

An NP-complete problem is a hardest problem in NP.

Seems like the right place to start for proving P≠NP.

It's also the right place to start for proving P=NP.

A polynomial time algorithm for one NP-complete problem, gives you a polynomial time algorithm for **every** problem in NP.

Reductions Redux

When we reduced A to B before, it was because we had an algorithm for B already, and wanted to solve A.

We knew how to handle unweighted graphs, now we want to see if we can handle weighted.

In complexity theory (where we're trying to show algorithms don't exist) we reduce well-studied A to new problem B.

Goal is a proof by contradiction.

1. Suppose (for sake of contradiction) new problem B has a nice algorithm.
2. But then we can use that for an algorithm well-studied problem A.
3. But, uh, no one knows an algorithm for well-studied problem A.
4. "contradiction"

Reductions Redux

To show problem B is NP-hard

Reduce from A, a known NP-hard problem, to B.

From the known-hard problem to your new problem—must be that direction!

How do you remember the direction?
Robbie recommends you memorize "reduce from known problem to new problem" by repeating it to yourself 50 times.

Alternatively reconstruct that proof by contradiction from the last slide to see which direction is needed.

Examples

There are literally thousands of NP-complete problems.
And some of them look weirdly similar to problems we do know efficient algorithms for.

In P

Short Path
Given a directed graph, report if there is a path from s to t of length at most k.

NP-Complete

Long Path
Given a directed graph, report if there is a path from s to t of length at least k.

Examples

In P

NP-Complete

Light Spanning Tree
Given a weighted graph, is there a spanning tree (a set of edges that connect all vertices) of weight at most k.

Traveling Salesperson
Given a weighted graph, is there a tour (a walk that visits every vertex and returns to its start) of weight at most k.

The electric company just needs a greedy algorithm to lay its wires. Amazon doesn't know a way to optimally route its delivery trucks.

Examples

In P

2-Coloring

Given a graph, decide if you can color the vertices red and blue so every edge has different colored endpoints

NP-Complete

3-Coloring

Given a graph, decide if you can color the vertices red, blue, and green so every edge has different colored endpoints.

2-Coloring can be done with a modification of BFS (or DFS). Color the start vertex red, its neighbors must be blue, their neighbors red, etc. No one knows how to tell if a graph is 3-colorable efficiently.

NP-hard

One more class:

NP-hard

Problem B is NP-hard if
for all problems A in NP, A reduces to B in polynomial time.

An NP-hard problem need not be in NP.
Examples?
Find the "best possible" certificate for an NP-hard problem.
 Instead of a path of length at least k, find the longest path.
 Instead of a tour of weight at most k, find the shortest tour.

NP-hard

Other Examples:
The halting problem is NP-hard (but not NP-complete).
So is n x n chess.

What The World Looks Like (We Think)

What The World Looks Like (If P=NP)

Still hard:
nxn chess

P
Short Paths, Light
Spanning Tree, 2-SAT
TSP, 3-SAT, Long Paths

Still impossible:
Halting Problem
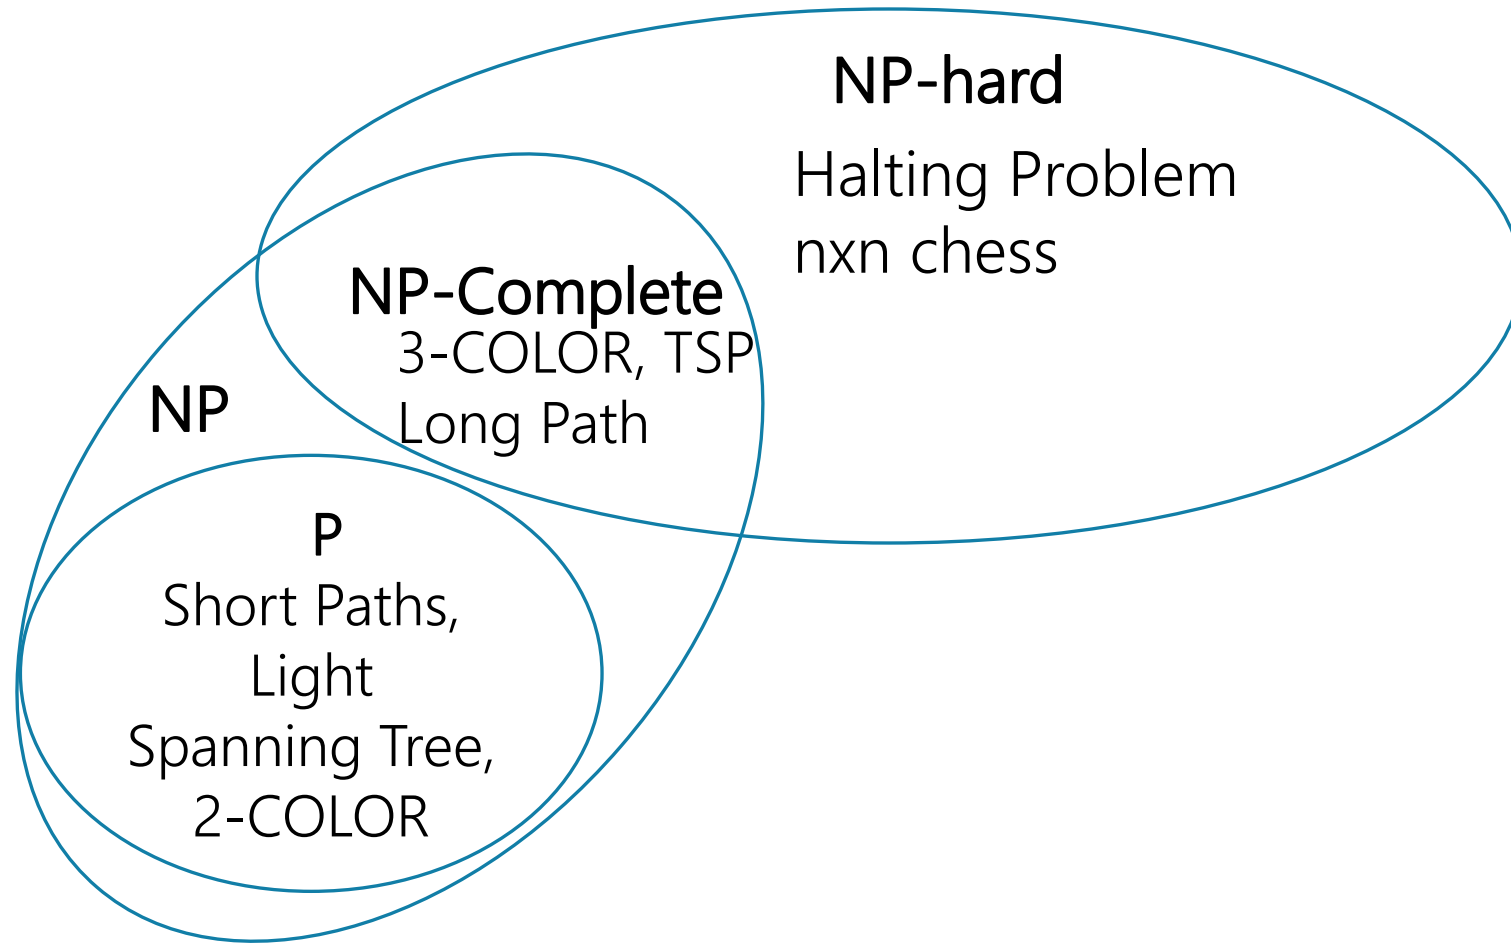
Why P vs. NP matters

P (stands for "Polynomial")
The set of all decision problems that have an algorithm that runs in time $O(n^k)$ for some constant k.

NP (stands for "nondeterministic polynomial")
The set of all decision problems such that if the answer is YES, there is a proof of that which can be verified in polynomial time.

NP-complete
Problem B is NP-complete if B is in NP and
for all problems A in NP, A reduces to B in polynomial time.

NP-hard
Problem B is NP-hard if
for all problems A in NP, A reduces to B in polynomial time.

Why is it called NP?

You've seen nondeterministic computation before.
Back in 311.

NFAs would "magically" decide among a set of valid transitions.
Always choosing one that would lead to an accept state (if such a transition exists).

An NFA and a DFA for the language "binary strings with a 1 in the 3rd position from the end."

From Kevin & Paul's 311 Lecture 23.

Nondeterminism

What would a nondeterministic **computer** look like?

It can run all the usual commands,

But it can also magically (i.e. nondeterministically) decide to set any bit of memory to 0 or 1.

Always choosing 0 or 1 to cause the computer to output YES,

(if such a choice exists).

If we had a nondeterministic computer...

Can you think of a polynomial time algorithm on a nondeterministic computer to:

Solve 2-COLOR?

Solve 3-COLOR?

If we had a nondeterministic computer...

Can you think of a polynomial time algorithm on a nondeterministic computer to:

Solve 2-COLOR?

Just run our regular deterministic polynomial time algorithm

Or nondeterministically guess colors, output if they work.

Solve 3-COLOR?

nondeterministically guess colors, output if they work.

Analogue of NFA/DFA equivalence

You showed in 311 that the set of languages decided by NFAs and DFAs were the same.

I.e. NFAs didn't let you solve more problems than DFAs.

But it did sometimes make the process a lot easier.

There are languages such that the best DFA is exponentially larger than the best NFA. (like the one from a few slides ago).

P vs. NP is an analogous question. Does non-determinism let us use exponentially fewer resources to solve some problems?

History, and Why P vs. NP?

NP-Completeness

An NP-complete problem is a **universal language** for encoding "I'll know it when I see it" problems.

If you find an efficient algorithm for an NP-complete problem, you have an algorithm for **every** problem in NP

Cook-Levin Theorem (1971)

SAT is NP-complete

Theorem 1: If a set S of strings is accepted by some nondeterministic Turing machine within polynomial time, then S is P-reducible to {DNF tautologies}.

NP-Complete Problems

But Wait! There's more!

Karp's Theorem (1972)

A lot of problems people care about are NP-complete

NP-Complete Problems

But Wait! There's more!

By 1979, at least 300 problems had been proven NP-complete.

Garey and Johnson put a list of all the NP-complete problems they could find in this textbook.

Took them almost 100 pages to just list them all.

No one has made a comprehensive list since.

COMPUTERS AND INTRACTABILITY
A Guide to the Theory of NP-Completeness

Michael R. Garey / David S. Johnson

NP-Complete Problems

But Wait! There's more!

In the last month, mathematicians and computer scientists have put papers on the arXiv claiming to show (at least) 10 more problems are NP-complete.

If you spend enough time trying to use computers to solve your problems, you will run into an NP-complete problem sooner or later.

What do you do?

Dealing with NP-Completeness

Option 1: Maybe it's a special case we understand

Maybe you don't need to solve the general problem, just a special case
- 2-COLOR vs. 3-COLOR

Option 2: Maybe it's a special case we *don't* understand (yet)

There are algorithms that are known to run quickly on "nice" instances. Maybe your problem has one of those.

One approach: Turn your problem into a SAT instance, find a solver and cross your fingers.

Dealing with NP-Completeness

Option 3: Approximation Algorithms

You might not be able to get an exact answer, but you might be able to get close.

> **Optimization version of Traveling Salesperson**
>
> Given a weighted graph, find a tour (a walk that visits every vertex and returns to its start) of minimum weight.

Algorithm:

Find a minimum spanning tree.

Have the tour follow the visitation order of a DFS of the spanning tree.

Theorem: This tour is at most twice as long as the best one.

Why should you care about P vs. NP

Most computer scientists are convinced that P≠NP.

Why should you care about this problem?

It's your chance for:

$1,000,000. The Clay Mathematics Institute will give $1,000,000 to whoever solves P vs. NP (or any of the 5 remaining problems they listed)

To get a Turing Award

Why should you care about P vs. NP

Most computer scientists are convinced that P≠NP.

Why should you care about this problem?

It's your chance for:

$1,000,000. The Clay Mathematics Institute will give $1,000,000 to whoever solves P vs. NP (or any of the 5 remaining problems they listed)

To get ~~a Turing Award~~ the Turing Award renamed after you.

Why Should You Care if P=NP?

Suppose P=NP.

Specifically that we found a genuinely in-practice efficient algorithm for an NP-complete problem. What would you do?
- $1,000,000 from the Clay Math Institute obviously, but what's next?

Why Should You Care if P=NP?

We found a genuinely in-practice efficient algorithm for an NP-complete problem. What would you do?

- Another $5,000,000 from the Clay Math Institute

- Put mathematicians out of work.

- Decrypt (essentially) all current internet communication.

- A world where P=NP is a very very different place from the world we live in now.

Why Should You Care if P≠NP?

We already expect P≠NP. Why should you care when we finally prove it?

P≠NP says something fundamental about the universe.

For some questions there is not a clever way to find the right answer
- Even though you'll know it when you see it.

There is actually a way to obscure information.

Why Should You Care if P≠NP?

To prove P≠NP we need to better understand the differences between problems.

- Why do some problems allow easy solutions and others don't?
- What is the structure of these problems?

We don't care about P vs NP just because it has a huge effect about what the world looks like.

We will learn a lot about computation along the way.