

Complexity Theory

Part One

It may be that since one is customarily concerned with existence, [...] finiteness, and so forth, one is not inclined to take seriously the question of the existence of a *better-than-finite* algorithm.

- Jack Edmonds, "Paths, Trees, and Flowers"

A Decidable Problem

- Consider the following problem:

Given two regular expressions R_1 and R_2 , determine whether R_1 and R_2 have the same language.

- This problem is indeed decidable.
 - We autograded your regular expressions in Problem Set Seven. The algorithm we used is 100% accurate.
- **Theorem:** There is no algorithm for solving this problem whose runtime is $O(2^{m+n})$, where m and n are the lengths of the input regular expressions.

The Limits of Decidability

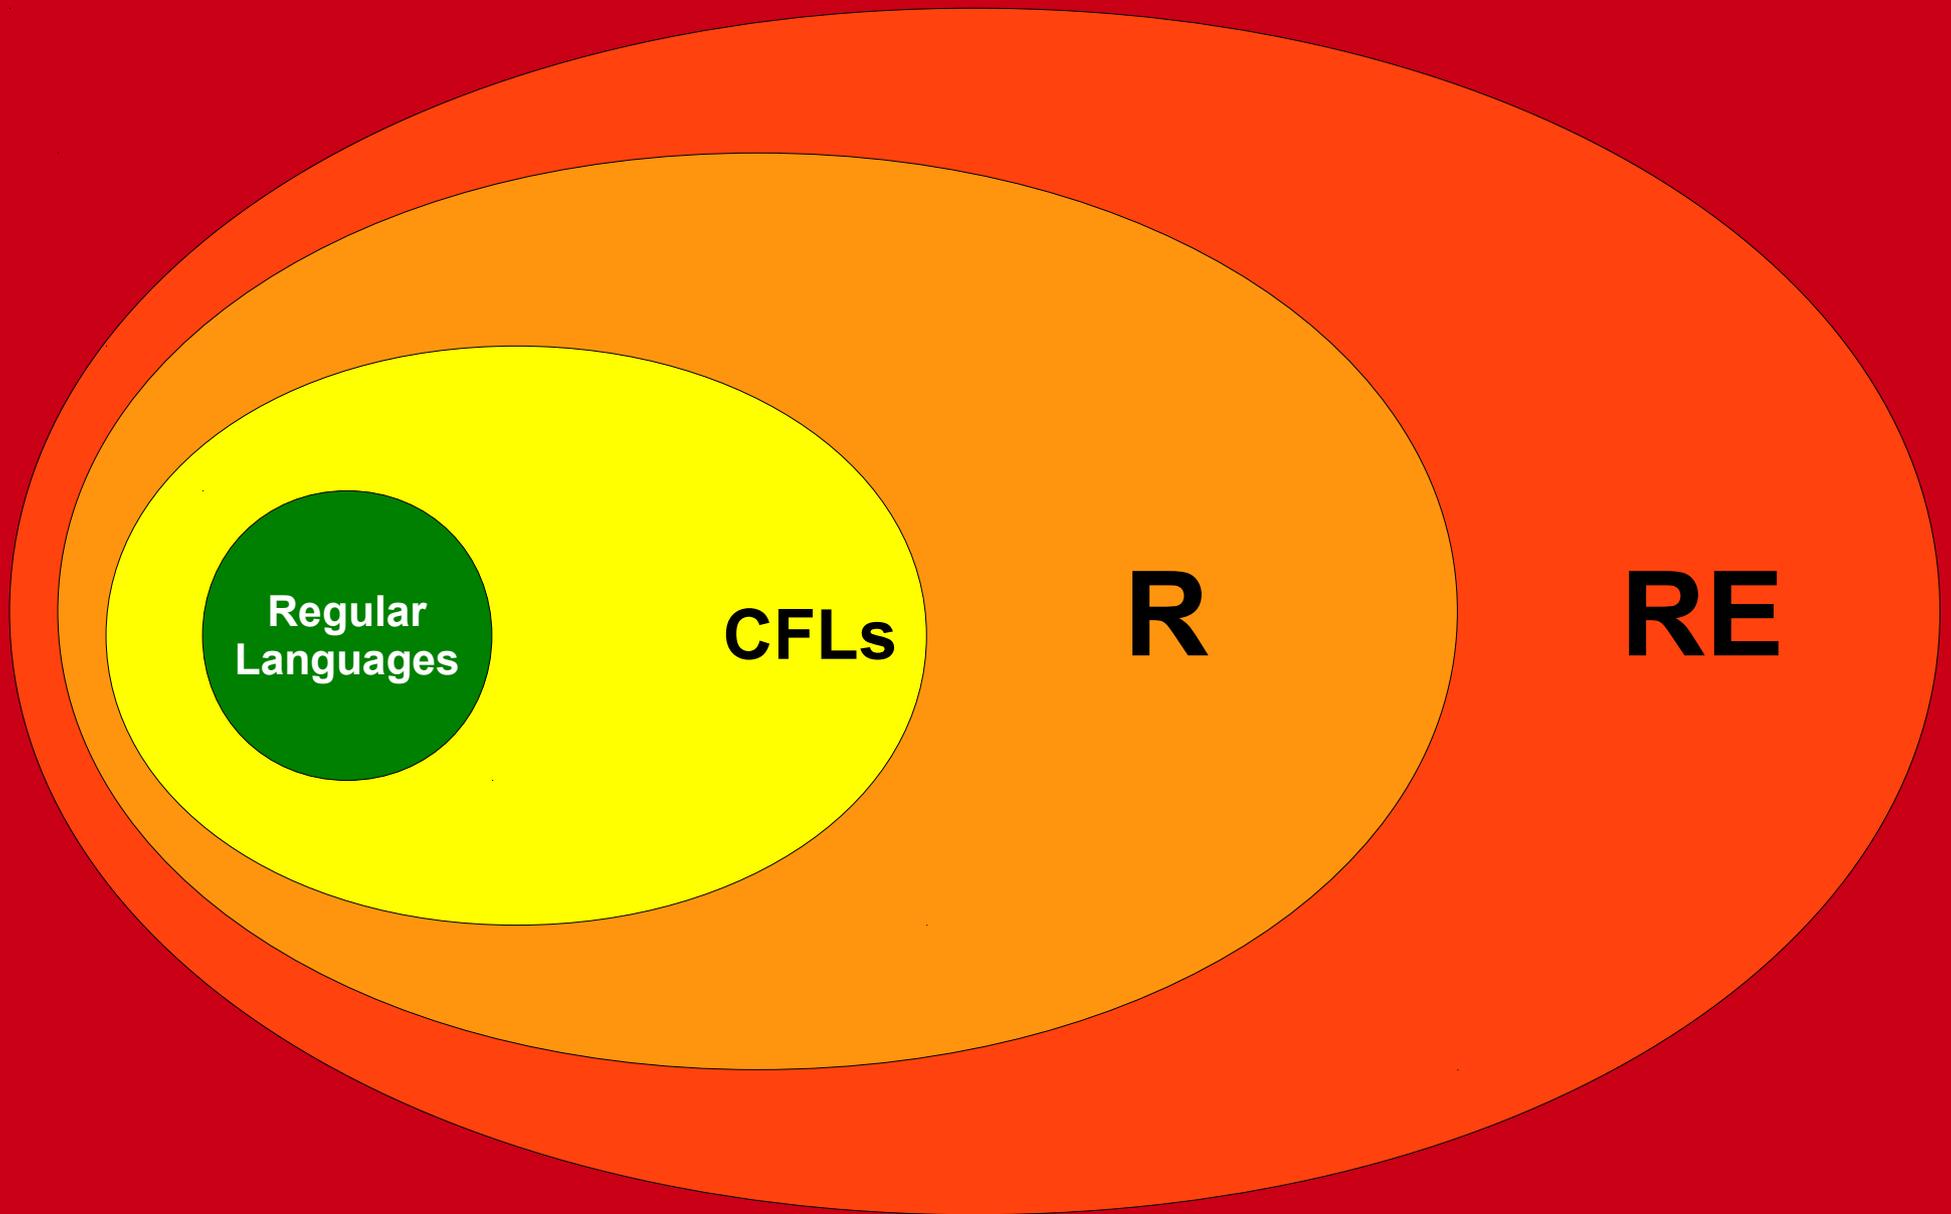
- The fact that a problem is decidable does not mean that it is *feasibly* decidable.
- In ***computability theory***, we ask the question
What problems can be solved by a computer?
- In ***complexity theory***, we ask the question
What problems can be solved ***efficiently*** by a computer?
- In the remainder of this course, we will explore this question in more detail.

Where We've Been

- The class **R** represents problems that can be solved by a computer.
- The class **RE** represents problems where “yes” answers can be verified by a computer.

Where We're Going

- The class **P** represents problems that can be solved *efficiently* by a computer.
- The class **NP** represents problems where “yes” answers can be verified *efficiently* by a computer.



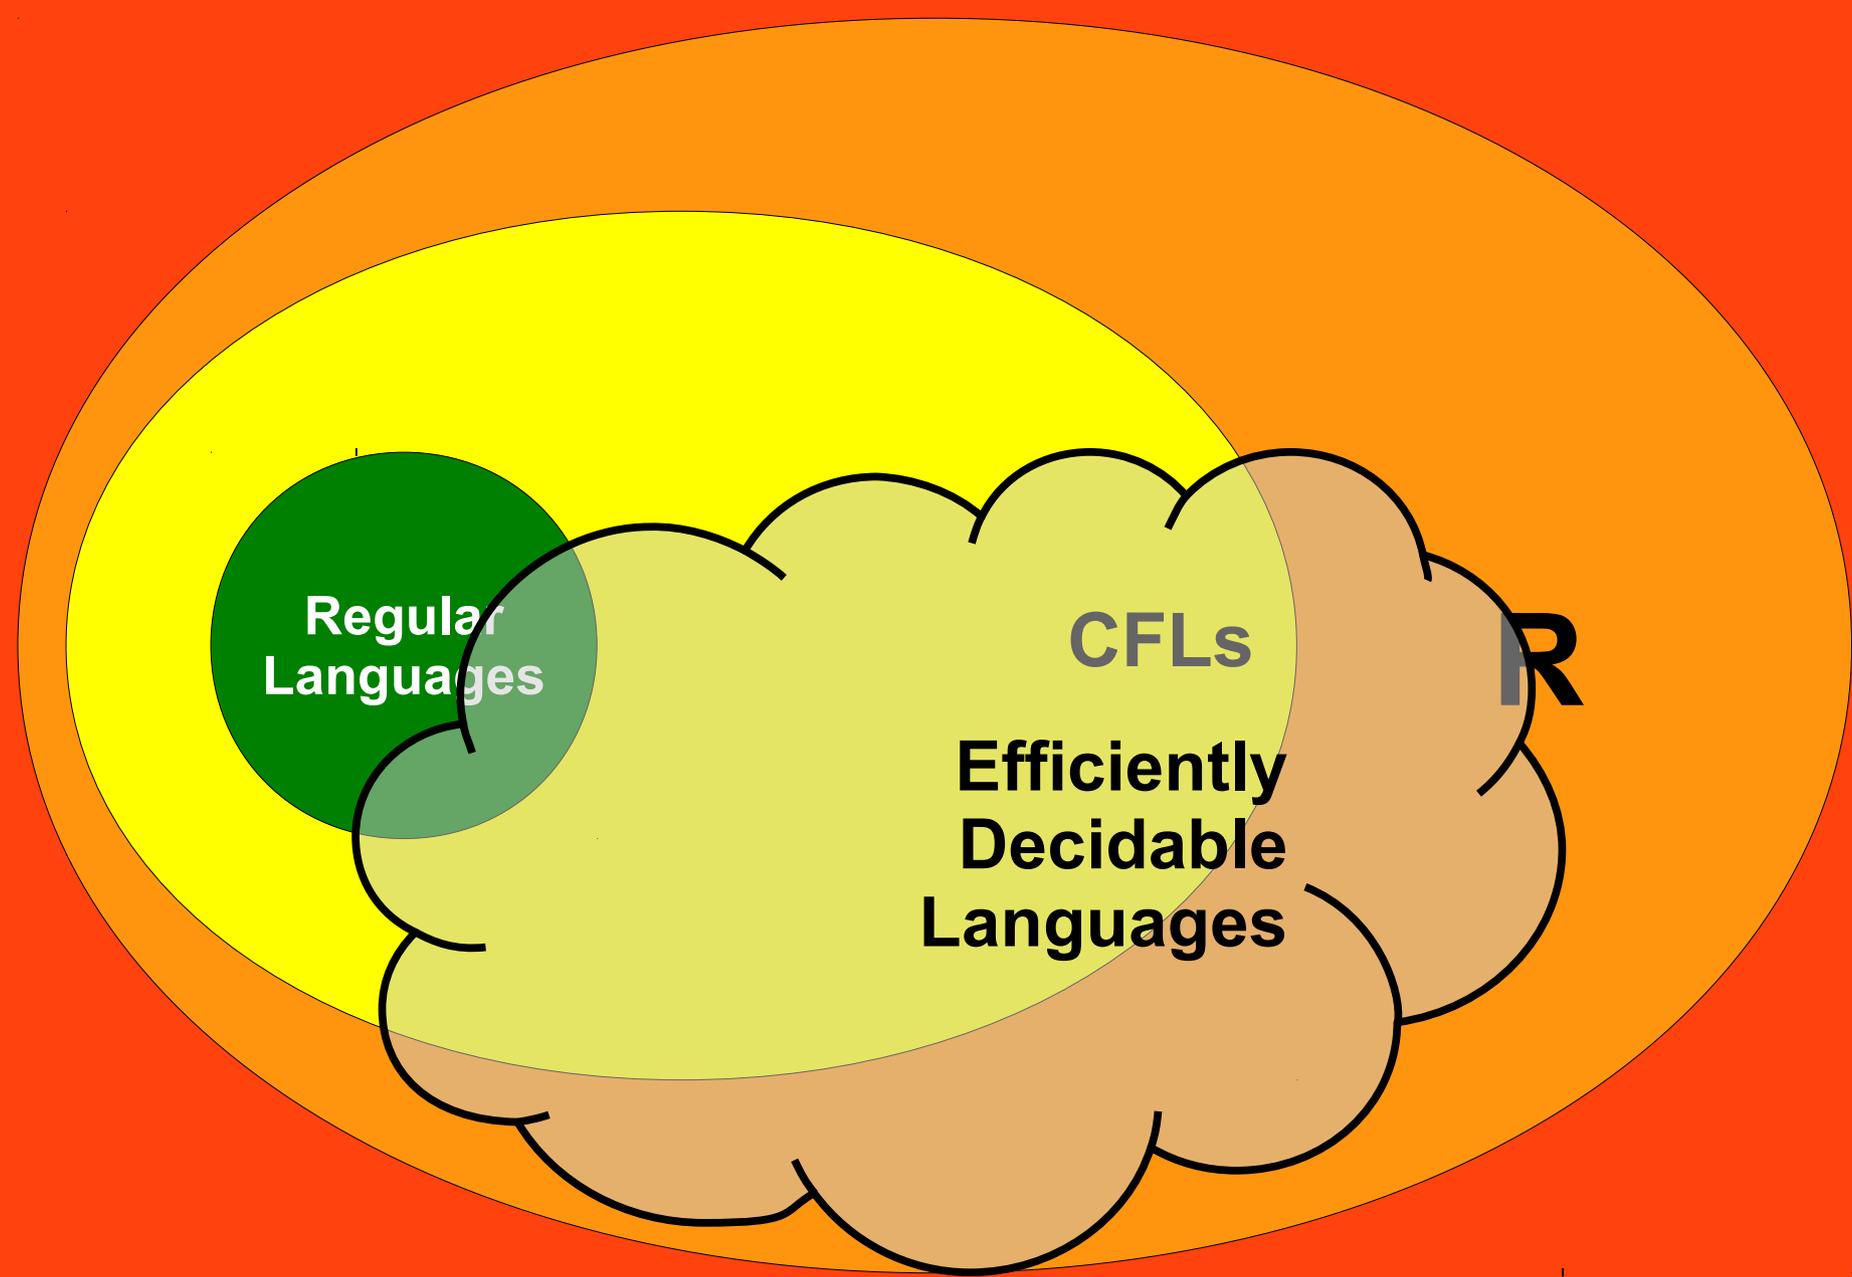
Regular
Languages

CFLs

R

RE

All Languages



Regular Languages

CFLs

Efficiently Decidable Languages

Undecidable Languages

The Setup

- In order to study computability, we needed to answer these questions:
 - What is “computation?”
 - What is a “problem?”
 - What does it mean to “solve” a problem?
- To study complexity, we need to answer these questions:
 - What does “complexity” even mean?
 - What is an “efficient” solution to a problem?

Measuring Complexity

- We have a program written in your Favorite Programming Language that's a decider for some problem.
- The program is correct in the sense that it always produces the right output for any given input.
- How might we measure the “complexity” of that solution?
 - The number of lines of code in the program.
 - How deeply-nested the loops or recursion in the program are.
 - How much time it takes for the program to solve the problem.
 - How much memory it takes for the program to solve the problem.
 - How much power it takes for the program to solve the problem.
 - How much network communication it takes for the program to solve the problem.
 - ...

Measuring Complexity

We have a program written in your Favorite Programming Language that's a decider for some problem.

The program is correct in the sense that it always produces the right output for any given input.

How might we measure the "complexity" of that solution?

The number of lines of code in the program.

How deeply-nested the loops or recursion in the program are.

- **How much time it takes for the program to solve the problem.**

How much memory it takes for the program to solve the problem.

How much power it takes for

How much network communication it takes to solve the problem.

...

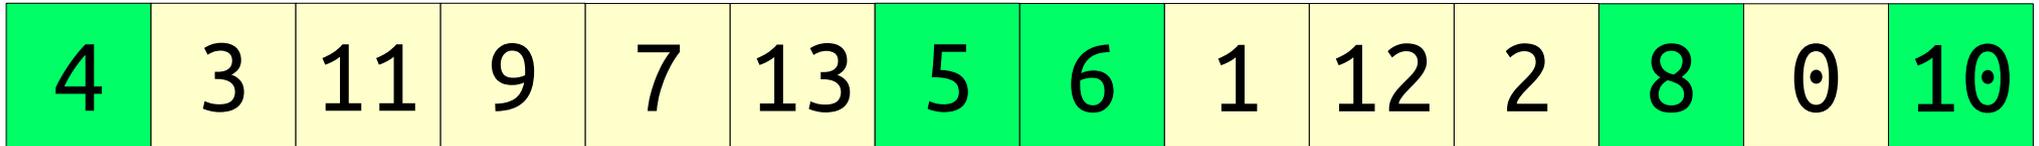
We're going to focus on this measure of "complexity," but that doesn't mean these other ones aren't interesting! There's tons of research on them.

What is an efficient algorithm?

Searching Finite Spaces

- Many decidable problems can be solved by searching over a large but finite space of possible options.
- Searching this space might take a staggeringly long time, but only finite time.
- From a decidability perspective, this is totally fine.
- From a complexity perspective, this may be totally unacceptable.

A Sample Problem



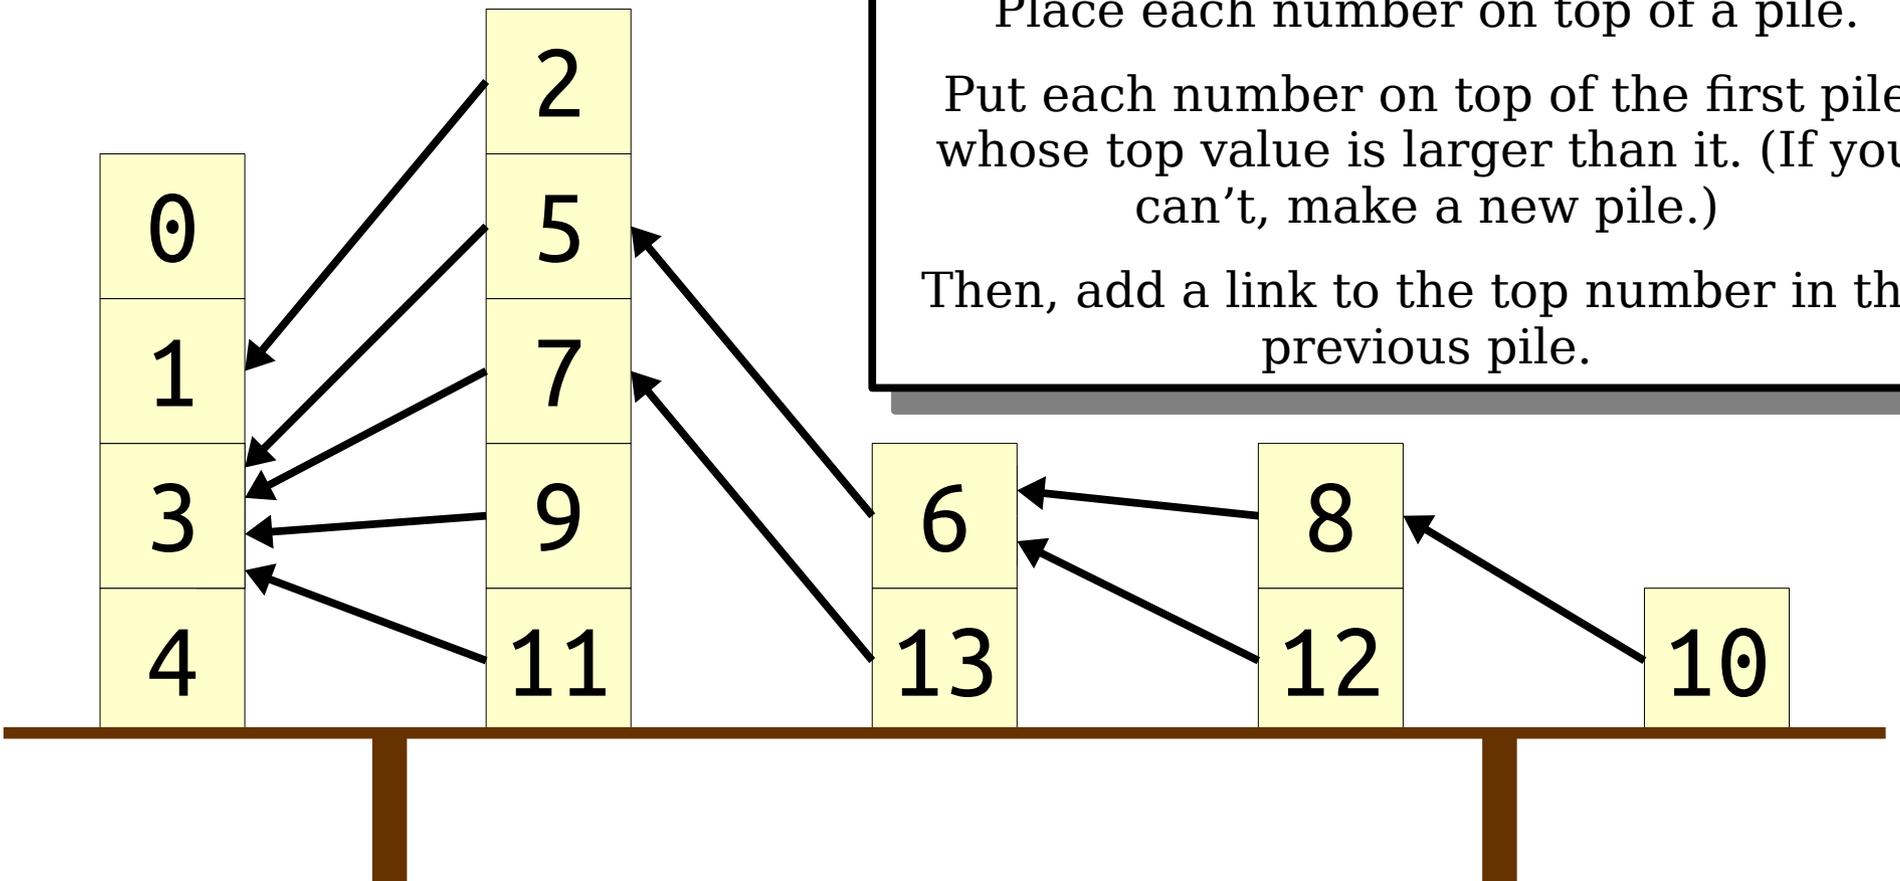
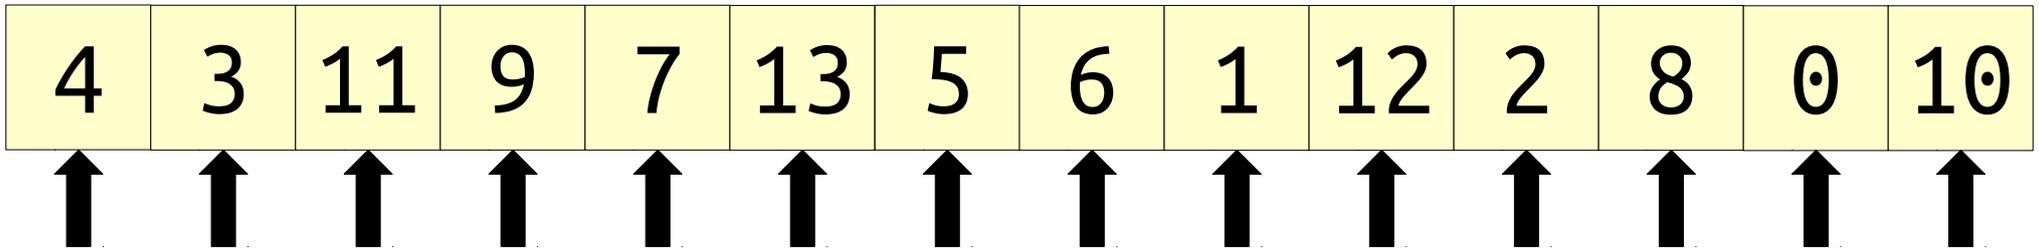
Goal: Find the length of the longest increasing subsequence of this sequence.

Longest Increasing Subsequences

- ***One possible algorithm:*** try all subsequences, find the longest one that's increasing, and return that.
- There are 2^n subsequences of an array of length n .
 - (Each subset of the elements gives back a subsequence.)
- Checking all of them to find the longest increasing subsequence will take time $O(n \cdot 2^n)$.
- ***Fact:*** the age of the universe is about 4.3×10^{26} nanoseconds. That's about 2^{85} nanoseconds.
- Practically speaking, this algorithm doesn't terminate if you give it an input of size 100 or more.

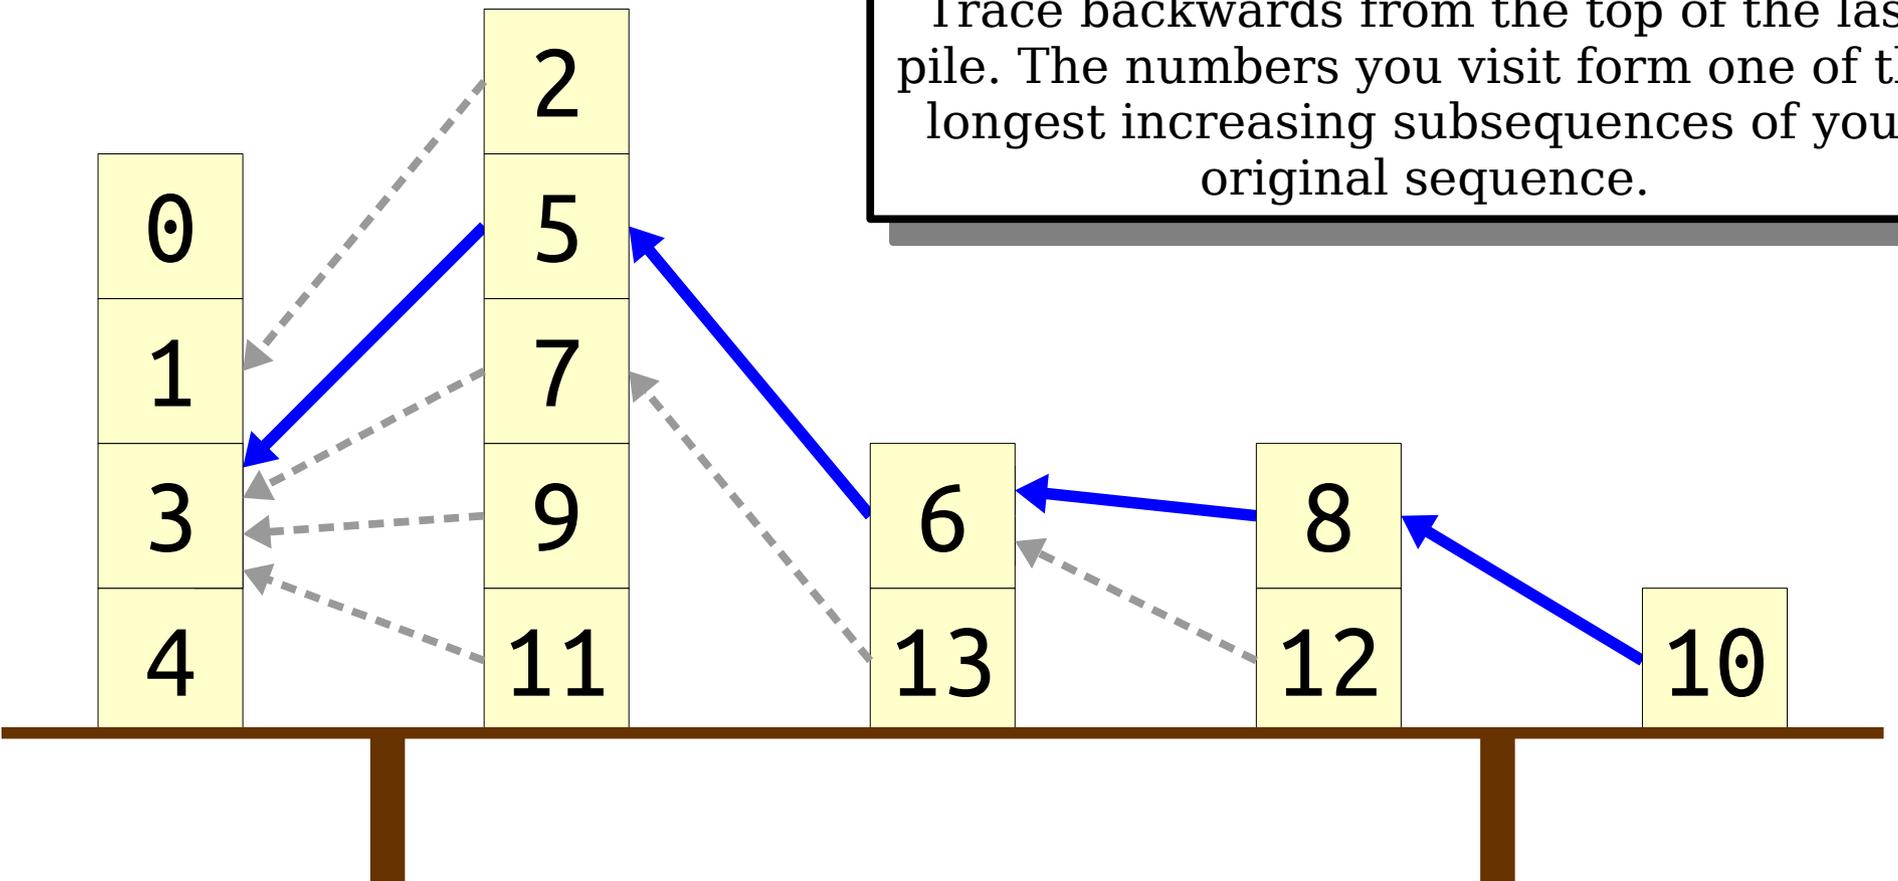
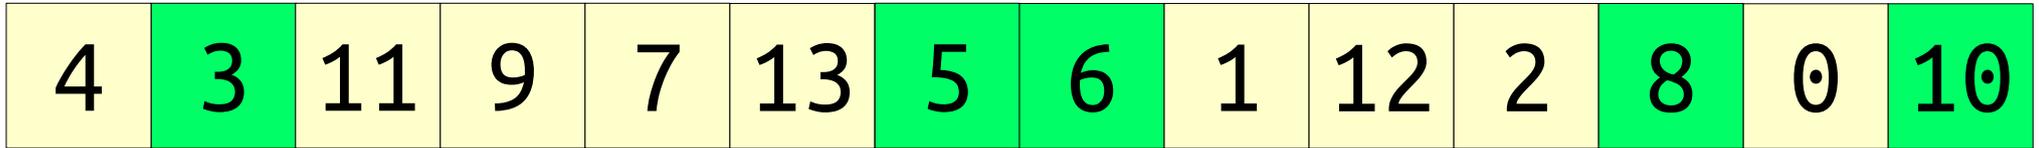
A Different Approach

Patience Sorting



Place each number on top of a pile.
Put each number on top of the first pile whose top value is larger than it. (If you can't, make a new pile.)
Then, add a link to the top number in the previous pile.

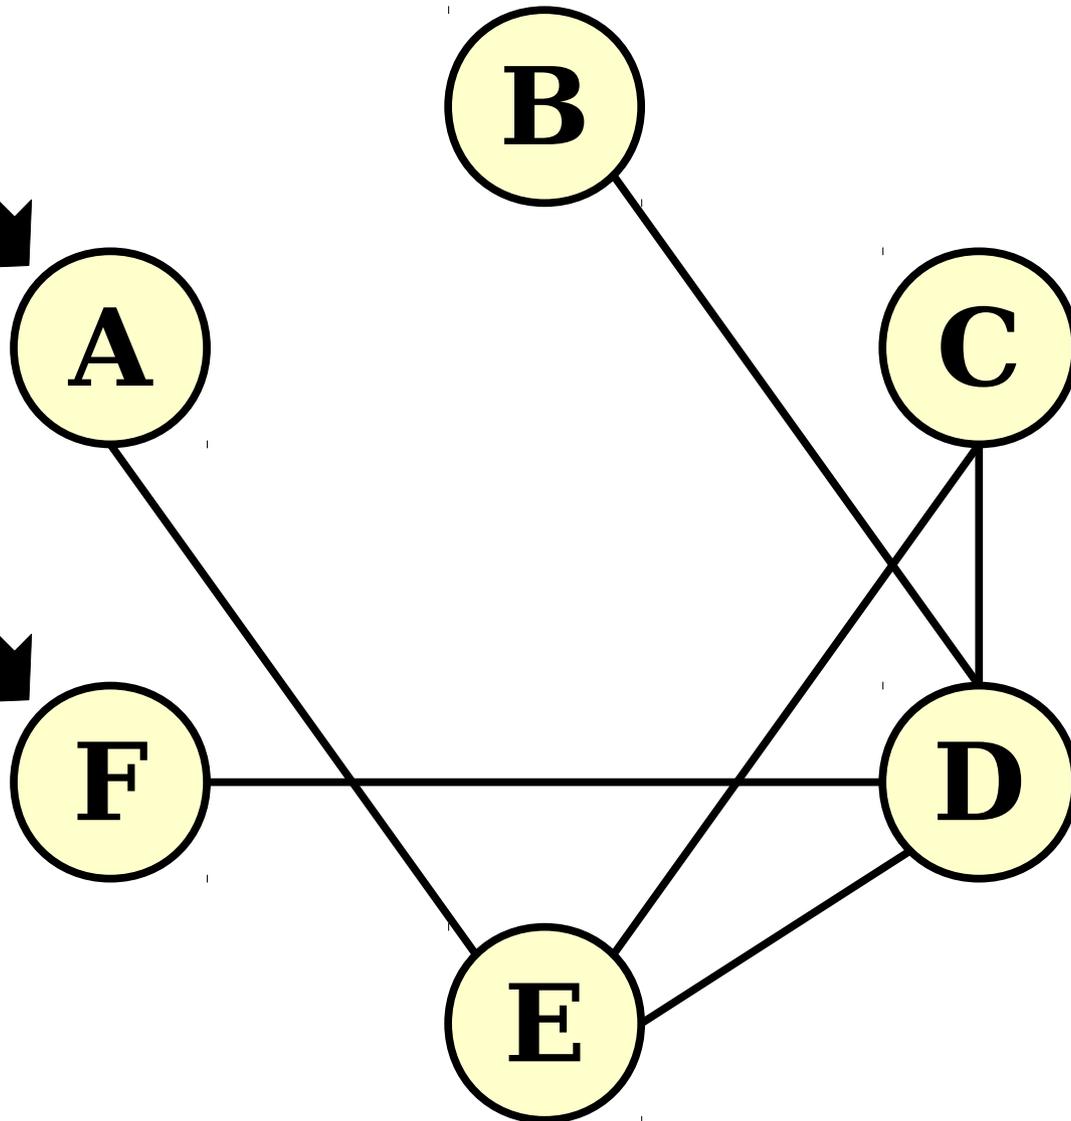
Patience Sorting



Longest Increasing Subsequences

- **Theorem:** There is an algorithm that can find the longest increasing subsequence of an array in time $O(n^2)$.
 - It's the previous **patience sorting** algorithm, with some clever implementation tricks.
- This algorithm works by exploiting particular aspects of how longest increasing subsequences are constructed. It's not immediately obvious that it works correctly.
- **CS161-Style Exercise 1:** Prove that this procedure always works!
- **CS161-Style Exercise 2:** Show that you can implement this algorithm in time $O(n \log n)$.

Another Problem



Goal: Determine the length of the shortest path from **F** to **A** in this graph.

Shortest Paths

- It is possible to find the shortest path in a graph by listing off all sequences of nodes in the graph in ascending order of length and finding the first that's a path.
- This takes time $O(n \cdot n!)$ in an n -node graph.
- For reference: $29!$ nanoseconds is longer than the lifetime of the universe.

Shortest Paths

- ***Theorem:*** It's possible to find the shortest path between two nodes in an n -node, m -edge graph in time $O(m + n)$.
- ***Proof idea:*** Use breadth-first search!
- The algorithm is a bit nuanced. It uses some specific properties of shortest paths and the proof of correctness is nontrivial.

For Comparison

- ***Longest increasing subsequence:***
 - Naive: $O(n \cdot 2^n)$
 - Fast: $O(n^2)$
- ***Shortest path problem:***
 - Naive: $O(n \cdot n!)$
 - Fast: $O(n + m)$.

Defining Efficiency

- When dealing with problems that search for the “best” object of some sort, there are often at least exponentially many possible options.
- Brute-force solutions tend to take at least exponential time to complete.
- Clever algorithms often run in time $O(n)$, or $O(n^2)$, or $O(n^3)$, etc.

Polynomials and Exponentials

- An algorithm runs in ***polynomial time*** if its runtime is some polynomial in n .
 - That is, time $O(n^k)$ for some constant k .
- Polynomial functions “scale well.”
 - Small changes to the size of the input do not typically induce enormous changes to the overall runtime.
- Exponential functions scale terribly.
 - Small changes to the size of the input induce huge changes in the overall runtime.

The Cobham-Edmonds Thesis

A language L can be ***decided efficiently*** if there is a TM that decides it in polynomial time.

Equivalently, L can be decided efficiently if it can be decided in time $O(n^k)$ for some $k \in \mathbb{N}$.

Like the Church-Turing thesis, this is ***not*** a theorem!

It's an assumption about the nature of efficient computation, and it is somewhat controversial.

The Cobham-Edmonds Thesis

- Efficient runtimes:
 - $4n + 13$
 - $n^3 - 2n^2 + 4n$
 - $n \log \log n$
- “Efficient” runtimes:
 - $n^{1,000,000,000,000}$
 - 10^{500}
- Inefficient runtimes:
 - 2^n
 - $n!$
 - n^n
- “Inefficient” runtimes:
 - $n^{0.0001 \log n}$
 - 1.0000000001^n

Why Polynomials?

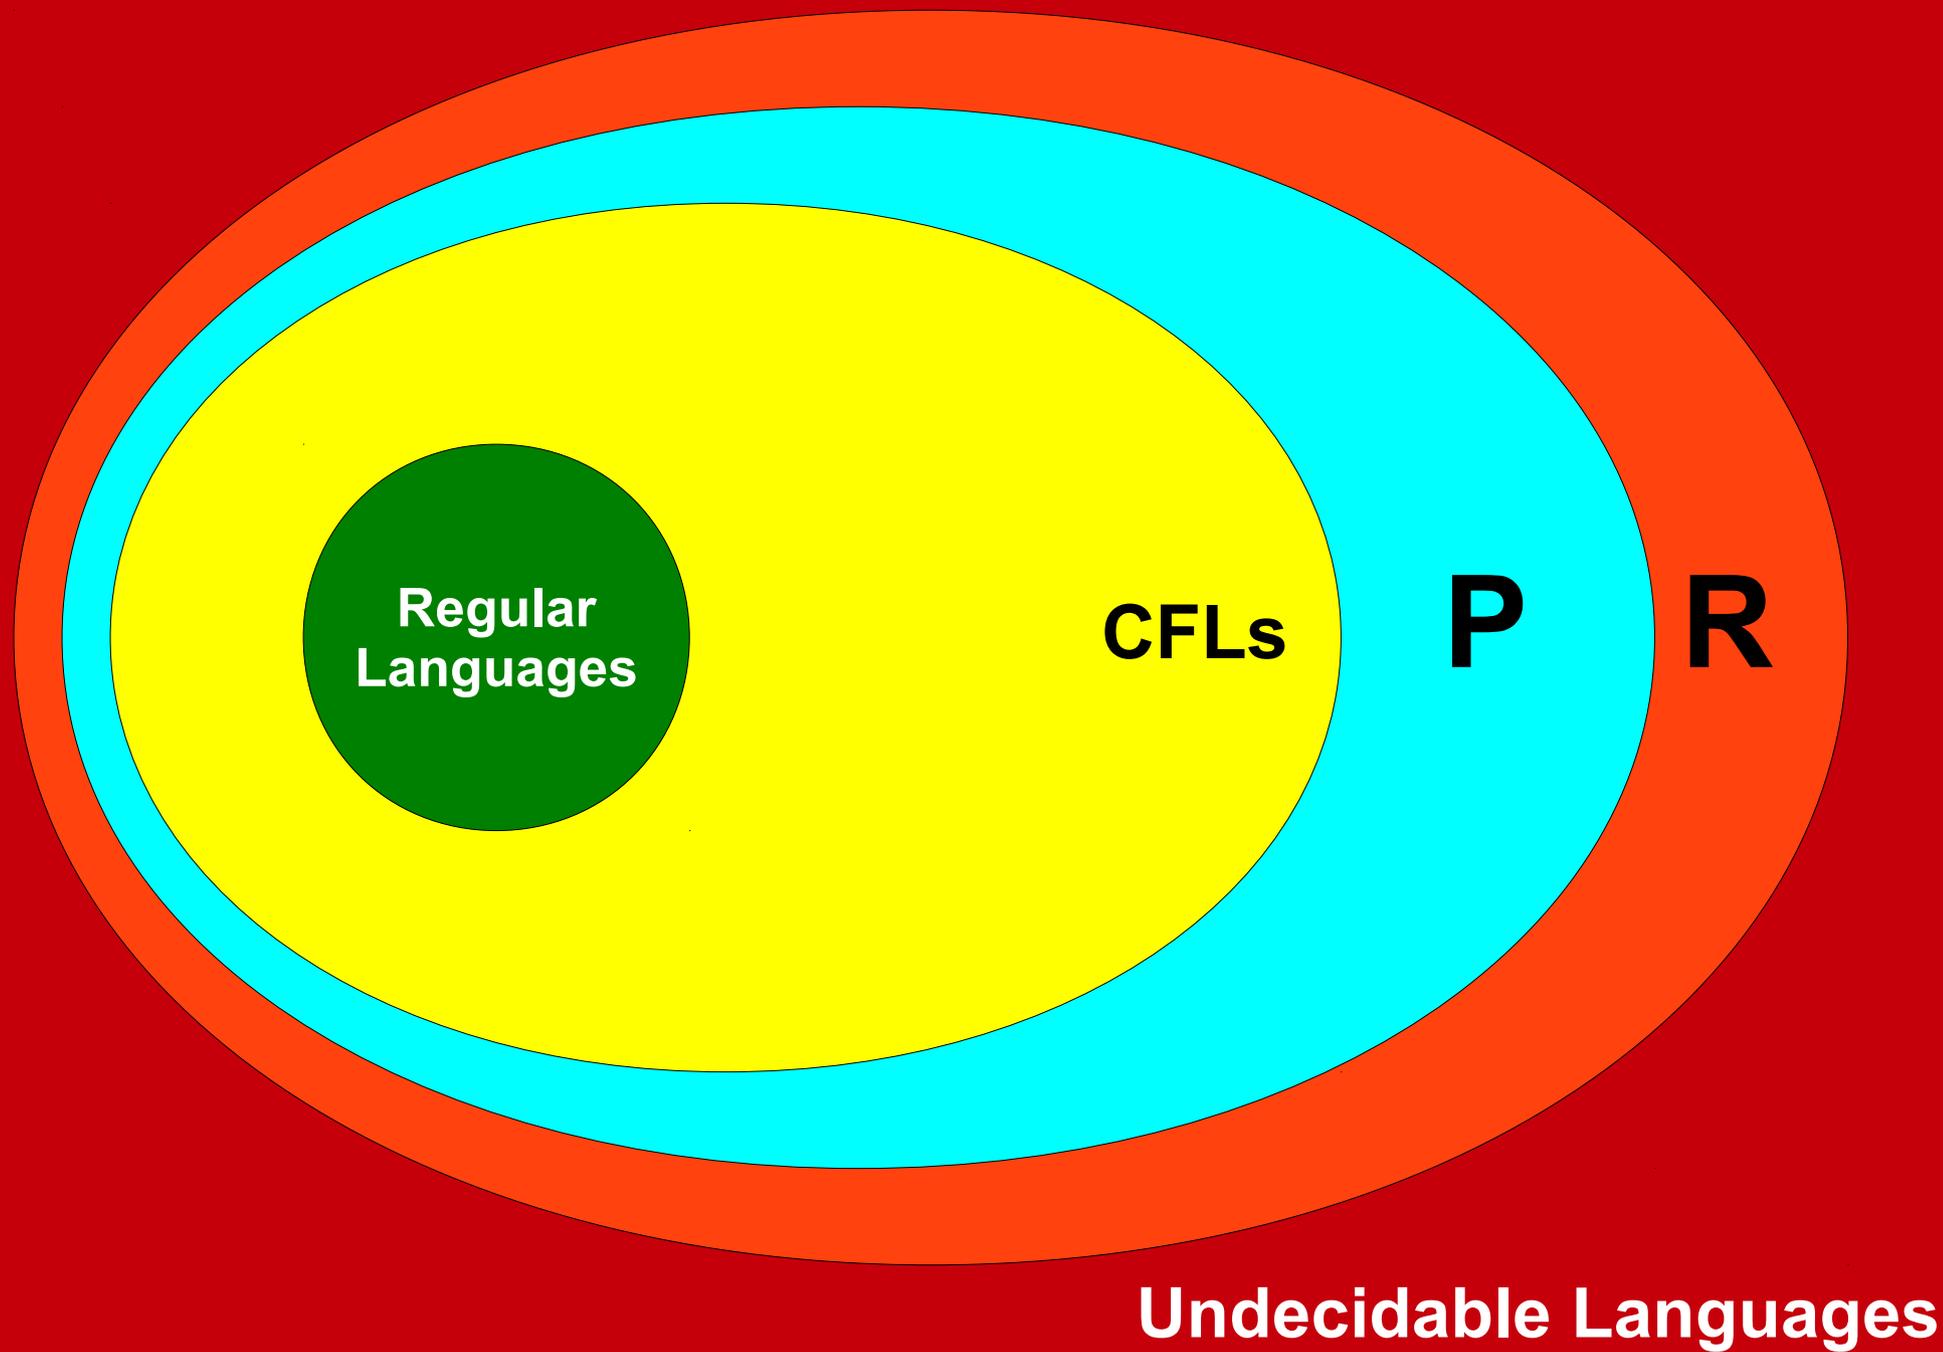
- Polynomial time *somewhat* captures efficient computation, but has a few edge cases.
- However, polynomials have very nice mathematical properties:
 - The sum of two polynomials is a polynomial. (Running one efficient algorithm, then another, gives an efficient algorithm.)
 - The product of two polynomials is a polynomial. (Running one efficient algorithm a “reasonable” number of times gives an efficient algorithm.)
 - The *composition* of two polynomials is a polynomial. (Using the output of one efficient algorithm as the input to another efficient algorithm gives an efficient algorithm.)

The Complexity Class **P**

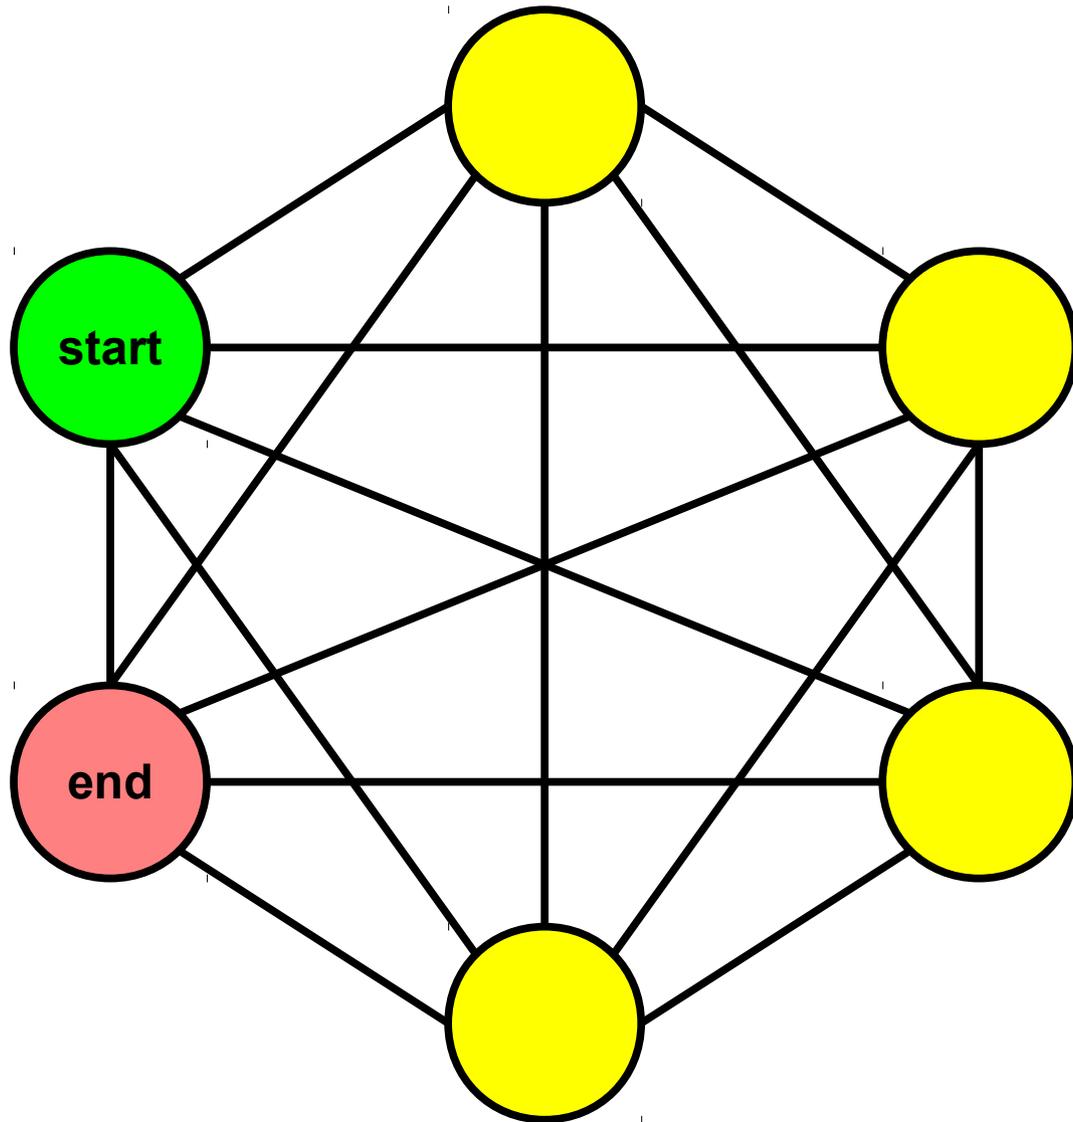
- The ***complexity class P*** (for *p*olynomial time) contains all problems that can be solved in polynomial time.
- Formally:
$$\mathbf{P} = \{ L \mid \text{There is a polynomial-time decider for } L \}$$
- Assuming the Cobham-Edmonds thesis, a language is in **P** if it can be decided efficiently.

Examples of Problems in **P**

- All regular languages are in **P**.
 - All have linear-time TMs.
- All CFLs are in **P**.
 - Requires a more nuanced argument (the *CYK algorithm* or *Earley's algorithm*.)
- And a *ton* of other problems are in **P** as well.
 - Curious? Take CS161!



What *can't* you do in polynomial time?



How many simple paths are there from the start node to the end node?



How many
subsets of this
set are there?

An Interesting Observation

- There are (at least) exponentially many objects of each of the preceding types.
- However, each of those objects is not very large.
 - Each simple path has length no longer than the number of nodes in the graph.
 - Each subset of a set has no more elements than the original set.
- This brings us to our next topic...

What if you need to search a large space for a single object?

Verifiers - Again

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

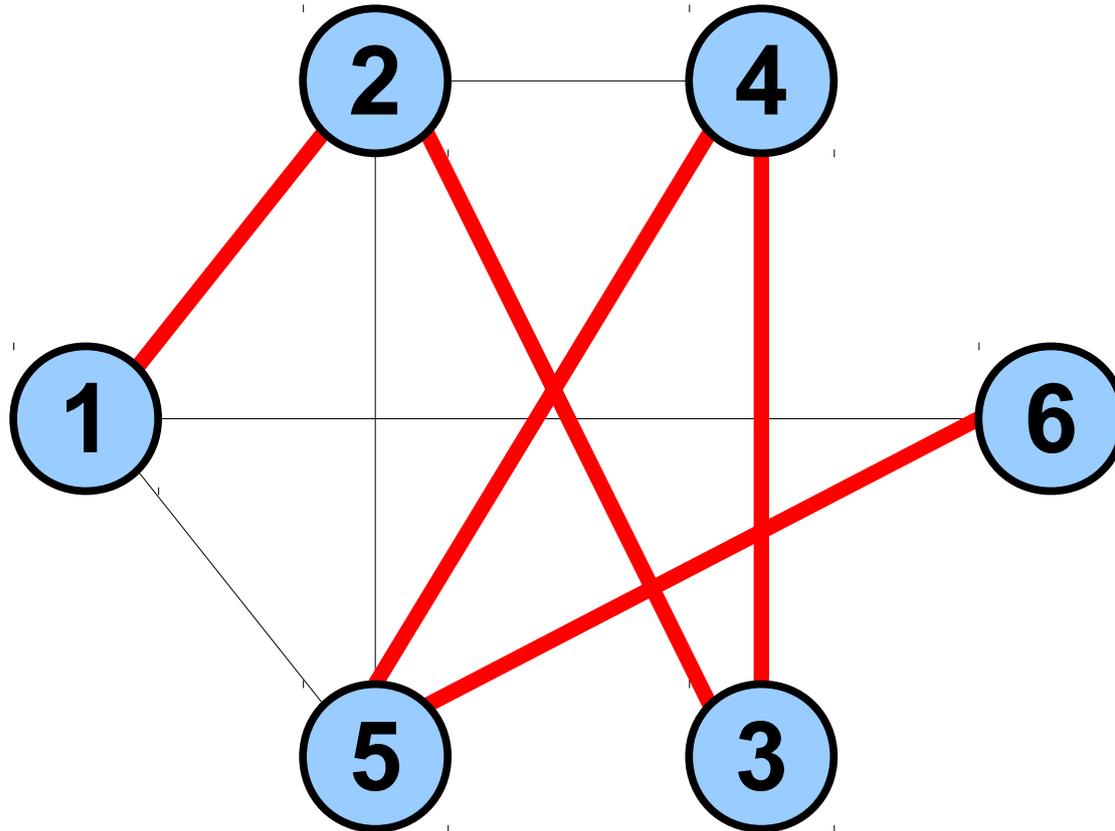
Does this Sudoku problem
have a solution?

Verifiers - Again



Is there an ascending subsequence of length at least 5?

Verifiers - Again



Is there a simple path that goes through every node exactly once?

Polynomial-Time Verifiers

- A ***polynomial-time verifier*** for L is a TM V such that
 - V halts on all inputs.
 - $w \in L \iff \exists c \in \Sigma^*. V$ accepts $\langle w, c \rangle$.
 - V runs “efficiently” (its runtime is $O(|w|^k)$ for some $k \in \mathbb{N}$).
 - All strings in L have “short” certificates (their lengths are $O(|w|^r)$ for some $r \in \mathbb{N}$).

The Complexity Class **NP**

- The complexity class **NP** (*nondeterministic polynomial time*) contains all problems that can be verified in polynomial time.
- Formally:

$$\mathbf{NP} = \{ L \mid \text{There is a polynomial-time verifier for } L \}$$

- The name **NP** comes from another way of characterizing **NP**. If you introduce *nondeterministic Turing machines* and appropriately define “polynomial time,” then **NP** is the set of problems that an NTM can solve in polynomial time.
- *Useful fact:* **NP** \subsetneq **R**. Come talk to me after class if you’re curious why!

P = { L | there is a polynomial-time decider for L }

NP = { L | there is a polynomial-time verifier for L }

R = { L | there is a ~~polynomial-time~~ decider for L }

RE = { L | there is a ~~polynomial-time~~ verifier for L }

We know that $\mathbf{R} \neq \mathbf{RE}$.

So does that mean $\mathbf{P} \neq \mathbf{NP}$?

And now...

The

Biggest Question

in

Theoretical Computer Science

P $\stackrel{?}{=}$ **NP**

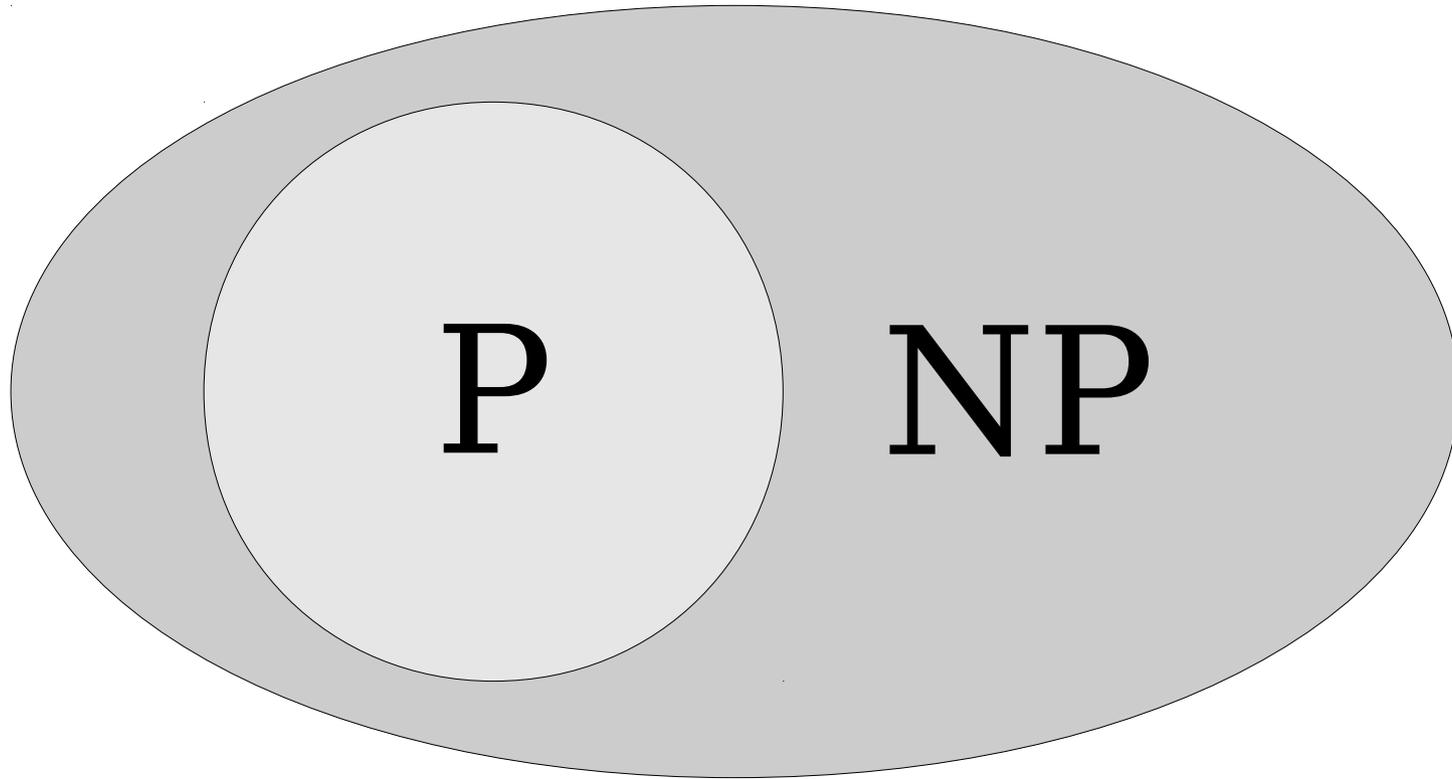
$\mathbf{P} = \{ L \mid \text{There is a polynomial-time decider for } L \}$

$\mathbf{NP} = \{ L \mid \text{There is a polynomial-time verifier for } L \}$

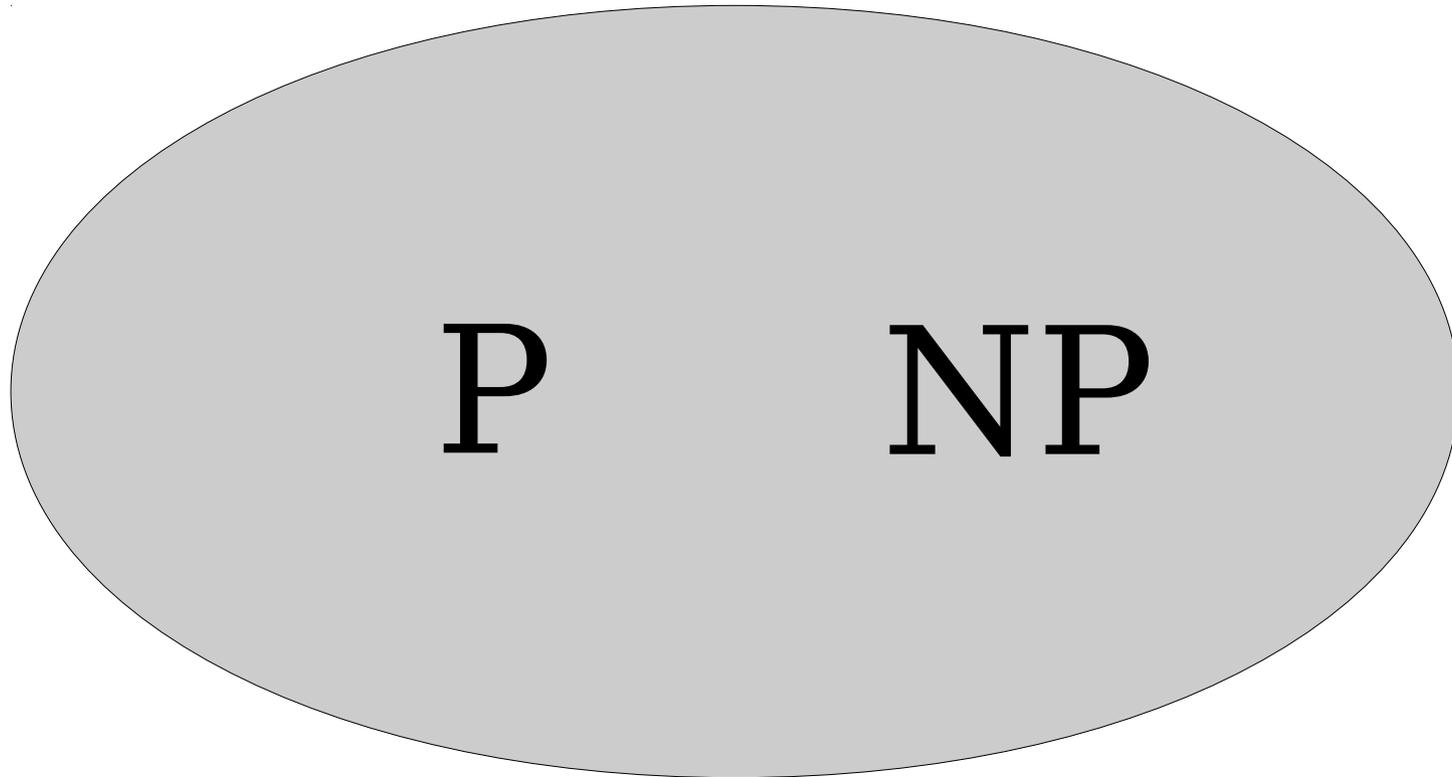


$\mathbf{P} \subseteq \mathbf{NP}$

Which Picture is Correct?



Which Picture is Correct?



$\mathbf{P} \stackrel{?}{=} \mathbf{NP}$

- The $\mathbf{P} \stackrel{?}{=} \mathbf{NP}$ question is the most important question in theoretical computer science.
- With the verifier definition of \mathbf{NP} , one way of phrasing this question is

*If a solution to a problem can be **checked** efficiently, can that problem be **solved** efficiently?*
- An answer either way will give fundamental insights into the nature of computation.

Why This Matters

- The following problems are known to be efficiently verifiable, but have no known efficient solutions:
 - Determining whether an electrical grid can be built to link up some number of houses for some price (Steiner tree problem).
 - Determining whether a simple DNA strand exists that multiple gene sequences could be a part of (shortest common supersequence).
 - Determining the best way to assign hardware resources in a compiler (optimal register allocation).
 - Determining the best way to distribute tasks to multiple workers to minimize completion time (job scheduling).
 - *And many more.*
- If $P = NP$, *all* of these problems have efficient solutions.
- If $P \neq NP$, *none* of these problems have efficient solutions.

Why This Matters

- If **P = NP**:
 - A huge number of seemingly difficult problems could be solved efficiently.
 - Our capacity to solve many problems will scale well with the size of the problems we want to solve.
- If **P \neq NP**:
 - Enormous computational power would be required to solve many seemingly easy tasks.
 - Our capacity to solve problems will fail to keep up with our curiosity.

What We Know

- Resolving $\mathbf{P} \stackrel{?}{=} \mathbf{NP}$ has proven *extremely difficult*.
- In the past 47 years:
 - Not a single correct proof either way has been found.
 - Many types of proofs have been shown to be insufficiently powerful to determine whether $\mathbf{P} \stackrel{?}{=} \mathbf{NP}$.
 - A majority of computer scientists believe $\mathbf{P} \neq \mathbf{NP}$, but this isn't a large majority.
- Interesting read: Interviews with leading thinkers about $\mathbf{P} \stackrel{?}{=} \mathbf{NP}$:
 - <http://web.eng.puc.cl/~jabaier/iic2212/poll-1.pdf>

The Million-Dollar Question

CHALLENGE ACCEPTED



The Clay Mathematics Institute has offered a ***\$1,000,000 prize*** to anyone who proves or disproves **$P = NP$** .

“My hunch is that [**P** $\stackrel{?}{=}$ **NP**] will be solved by a young researcher who is not encumbered by too much conventional wisdom about how to attack the problem.”

– Prof. Richard Karp

(The guy who first popularized the P $\stackrel{?}{=}$ NP problem.)

Time-Out for Announcements!

Please evaluate this course in Axess.
Your comments really make a difference.

Don Knuth Lecture

- Don Knuth, Professor Emeritus of the Art of Computer Programming, will be giving a public lecture on Thursday, December 5th, right here in Nvidia Auditorium, from 6:30PM – 7:30PM.

- Here's his description of the lecture:

The number π appears thousands of times in The Art of Computer Programming, in many different contexts. Dr. I. J.

Matrix has remarked that its digits, “when properly interpreted,” actually convey the entire history of the human race! [See page 41 of Volume 2.] This lecture will examine many interpretations of those digits, both proper and improper.

- Highly recommended!

Problem Set Nine

- Problem Set Nine is due this Friday at 2:30PM.
 - As a reminder, ***no late submissions will be accepted***. Please budget enough time to get your submission in!
 - ***Very smart idea***: submit at least three hours early.
- As always, feel free to ask questions in office hours or online via Piazza.

Midterm Regrades

- As a reminder, regrade requests for the second midterm are due on Friday.
- ***Please look over your feedback*** on the second midterm, and ***ask questions if you have them!*** The point of the exams is to give you a sense of where you are and what you need to improve on.

Final Exam Logistics

- Our final exam is Tuesday, December 10th from 8:30AM – 11:30AM in ***Dinkelspiel Auditorium***.
- The exam is cumulative. You're responsible for topics from PS1 – PS9 and all of the lectures.
- As with the midterms, the exam is closed-book, closed-computer, and limited-note. You can bring one double-sided sheet of 8.5" × 11" notes with you to the exam, decorated any way you'd like.
- Students with OAE accommodations: if we don't yet have your OAE letter, please send it to us ASAP.

Preparing for the Final

- On the course website you'll find
 - **eight** practice final exams, which are all real exams with minor modifications, with solutions, and
 - a giant set of practice problems (EPP3), with solutions.
- Our recommendation: Look back over the exams and problem sets and redo any problems that you didn't really get the first time around.
- Keep the TAs in the loop: stop by office hours to have them review your answers and offer feedback.

Preparing for the Final

- The amazing Amy Liu will be holding a final exam review session this Saturday from 2:00PM - 4:00PM in Gates 104.
- We'll aim to record that session and make it available shortly after this time window.
- You're also welcome to peruse Amy's CS103A materials on her course website, <https://cs103a.stanford.edu>.

Your Questions

“What are your thoughts on **P** vs **NP**? Have you ever tried solving it?”

I've accidentally tried solving it multiple times without realizing that's what I was doing. We'll talk about how that happened next time. 😊

I'm excited to see how this problem ends up getting resolved. I think I'll see it in my lifetime. And I think there's a good chance that it'll be a former CS103 student who cracks it, and I'll be able to point to a chair and say “this is known as X's theorem, and she used to sit right there.”

Back to CS103!

What do we know about $\mathbf{P} \stackrel{?}{=} \mathbf{NP}$?

Adapting our Techniques

A Problem

- The **R** and **RE** languages correspond to problems that can be decided and verified, *period*, without any time bounds.
- To reason about what's in **R** and what's in **RE**, we used two key techniques:
 - **Universality**: TMs can simulate other TMs.
 - **Self-Reference**: TMs can get their own source code.
- Why can't we just do that for **P** and **NP**?

Theorem (Baker-Gill-Solovay): Any proof that purely relies on universality and self-reference cannot resolve $\mathbf{P} \stackrel{?}{=} \mathbf{NP}$.

Proof: Take CS154!

So how *are* we going to
reason about **P** and **NP**?

Next Time

- ***Reducibility***
 - A technique for connecting problems to one another.
- ***NP-Completeness***
 - What are the hardest problems in **NP**?