

P=NP?

Clay Math Inst.
\$1M

P - yes/no questions that can be answered
in time polynomial in input size.

Is this array sorted? — $O(n)$ time
↑
array size

Can all 42 Physical Kids
get from E to F → $O(\cancel{N}VE)$ time
in 24 hours?
n = const ↑
Adj. list


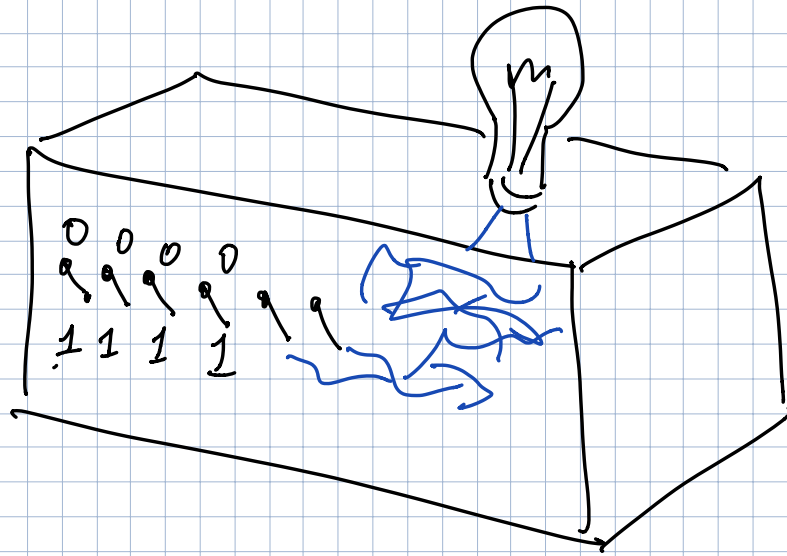
NP - Nondeterministic Poly time

Yes/no problems where YES instance
can be verified in poly time

Does this graph have a
Hamiltonian cycle?

YES: Show me the
cycle

NO: Um...



Circuit Satisfiability:

Given a Boolean circuit, are there inputs that produce output 1?

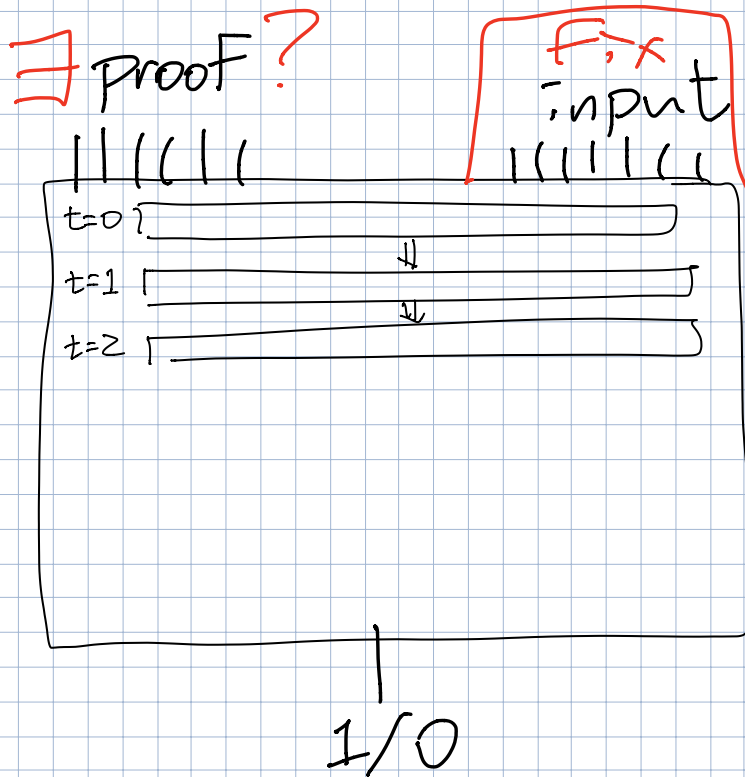

$O(2^n \cdot n)$ - brute force

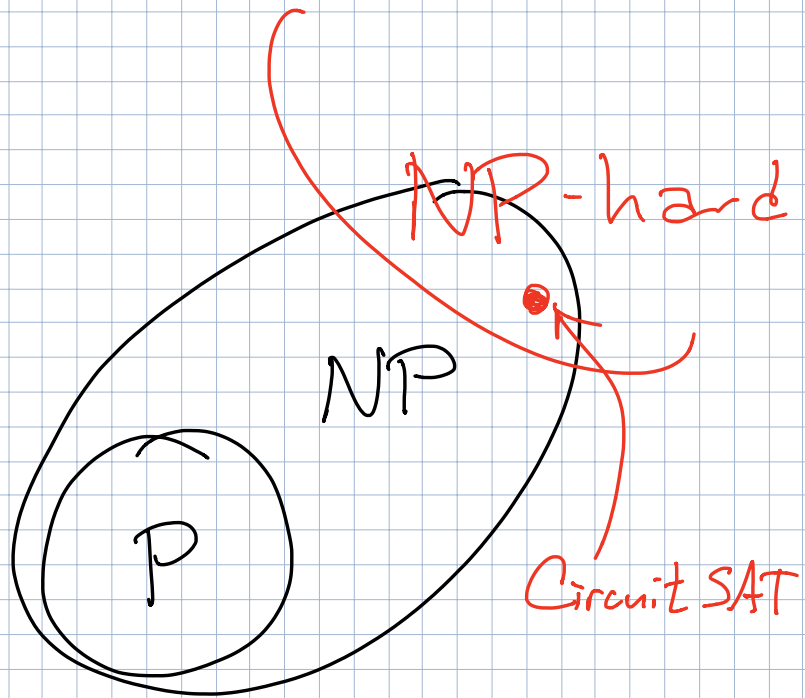
$O(2^n/n)$ - best known

Cook-Levin Theorem:

Circuit SAT is NP-hard.

If Circuit SAT \in P then $P=NP$.





If we accept

$P \neq NP$ as axiom / physical law

NP-hard \Rightarrow No poly time algo

[I]n his short and broken treatise he provides an eternal example—not of laws, or even of method, for there is no method except to be very intelligent, but of intelligence itself swiftly operating the analysis of sensation to the point of principle and definition.

— T. S. Eliot on Aristotle, “The Perfect Critic”, *The Sacred Wood* (1921)

The nice thing about standards is that you have so many to choose from; furthermore, if you do not like any of them, you can just wait for next year’s model.

— Andrew S. Tannenbaum, *Computer Networks* (1981)

Also attributed to Grace Murray Hopper and others


If a problem has no solution, it may not be a problem, but a fact — not to be solved, but to be coped with over time.

— Shimon Peres, as quoted by David Rumsfeld, *Rumsfeld’s Rules* (2001)


30 NP-Hard Problems

30.1 A Game You Can’t Win

A salesman in a red suit who looks suspiciously like Tom Waits presents you with a black steel box with n binary switches on the front and a light bulb on the top. The salesman tells you that the state of the light bulb is controlled by a complex *boolean circuit*—a collection of AND, OR, and NOT gates connected by wires, with one input wire for each switch and a single output wire for the light bulb. He then asks you the following question: Is there a way to set the switches so that the light bulb turns on? If you can answer this question correctly, he will give you the box and a ~~million billion~~ trillion dollars; if you answer incorrectly, or if you die without answering at all, he will take your soul.



An AND gate, an OR gate, and a NOT gate.



A boolean circuit. inputs enter from the left, and the output leaves to the right.

As far as you can tell, the Adversary hasn’t connected the switches to the light bulb at all, so no matter how you set the switches, the light bulb will stay off. If you declare that it is possible to turn on the light, the Adversary will open the box and reveal that there is no circuit at all. But if you declare that it is *not* possible to turn on the light, before testing all 2^n settings, the Adversary will magically create a circuit inside the box that turns on the light *if and only if* the

© Copyright 2016 Jeff Erickson.
 This work is licensed under a Creative Commons License (<http://creativecommons.org/licenses/by-nc-sa/4.0/>).
 Free distribution is strongly encouraged; commercial distribution is expressly forbidden.
 See <http://www.cs.uiuc.edu/~jeffe/teaching/algorithms/> for the most recent revision.

switches are in one of the settings you haven't tested, and then flip the switches to that setting, turning on the light. (You can't detect the Adversary's cheating, because you can't see inside the box until the end.) The only way to *provably* answer the Adversary's question correctly is to try all 2^n possible settings. You quickly realize that this will take *far* longer than you expect to live, so you gracefully decline the Adversary's offer.

The Adversary smiles and says, "Ah, yes, of course, you have no reason to trust me. But perhaps I can set your mind at ease." He hands you a large roll of parchment—which you hope was made from sheep skin—with a circuit diagram drawn (or perhaps tattooed) on it. "Here are the complete plans for the circuit inside the box. Feel free to poke around inside the box to make sure the plans are correct. Or build your own circuit from these plans. Or write a computer program to simulate the circuit. Whatever you like. If you discover that the plans don't match the actual circuit in the box, you win the trillion bucks." A few spot checks convince you that the plans have no obvious flaws; subtle cheating appears to be impossible.

But you should still decline the Adversary's generous offer. The problem that the Adversary is posing is called **circuit satisfiability** or **CIRCUITSAT**: Given a boolean circuit, is there is a set of inputs that makes the circuit output TRUE, or conversely, whether the circuit *always* outputs FALSE. For any particular input setting, we can calculate the output of the circuit in polynomial (actually, *linear*) time using depth-first-search. But nobody knows how to solve CIRCUITSAT faster than just trying all 2^n possible inputs to the circuit, but this requires exponential time. On the other hand, nobody has actually *proved* that this is the best we can do; maybe there's a clever algorithm that just hasn't been discovered yet!

30.2 P versus NP

A minimal requirement for an algorithm to be considered "efficient" is that its running time is polynomial: $O(n^c)$ for some constant c , where n is the size of the input.¹ Researchers recognized early on that not all problems can be solved this quickly, but had a hard time figuring out exactly which ones could and which ones couldn't. There are several so-called **NP-hard** problems, which most people believe *cannot* be solved in polynomial time, even though nobody can prove a super-polynomial lower bound.

A *decision problem* is a problem whose output is a single boolean value: YES or NO. Let me define three classes of decision problems:

- **P** is the set of decision problems that can be solved in polynomial time. Intuitively, P is the set of problems that can be solved quickly.
- **NP** is the set of decision problems with the following property: If the answer is YES, then there is a *proof* of this fact that can be checked in polynomial time. Intuitively, NP is the set of decision problems where we can verify a YES answer quickly if we have the solution in front of us.
- **co-NP** is essentially the opposite of NP. If the answer to a problem in co-NP is NO, then there is a proof of this fact that can be checked in polynomial time.

¹This notion of efficiency was independently formalized by Alan Cobham (The intrinsic computational difficulty of functions. *Logic, Methodology, and Philosophy of Science (Proc. Int. Congress)*, 24–30, 1965), Jack Edmonds (Paths, trees, and flowers. *Canadian Journal of Mathematics* 17:449–467, 1965), and Michael Rabin (Mathematical theory of automata. *Proceedings of the 19th ACM Symposium in Applied Mathematics*, 153–175, 1966), although similar notions were considered more than a decade earlier by Kurt Gödel and John von Neumann.


For example, the circuit satisfiability problem is in NP. If a given boolean circuit is satisfiable, then any set of m input values that produces TRUE output is a proof that the circuit is satisfiable; we can check the proof by evaluating the circuit in polynomial time. It is widely believed that circuit satisfiability is *not* in P or in co-NP, but nobody actually knows.

Every decision problem in P is also in NP. If a problem is in P, we can verify YES answers in polynomial time recomputing the answer from scratch! Similarly, every problem in P is also in co-NP.

Perhaps the single most important unanswered question in theoretical computer science—if not all of computer science—if not all of *science*—is whether the complexity classes P and NP are actually different. Intuitively, it seems obvious to most people that $P \neq NP$; the homeworks and exams in this class and others have (I hope) convinced you that problems can be incredibly hard to solve, even when the solutions are obvious in retrospect. It's completely obvious; *of course* solving problems from scratch is harder than just checking that a solution is correct. We can quite reasonably accept the statement " $P \neq NP$ " as a law of nature.

But nobody knows how to *prove* $P \neq NP$. In fact, there has been little or no real progress toward a proof for decades.² The Clay Mathematics Institute lists P versus NP as the first of its seven Millennium Prize Problems, offering a \$1,000,000 reward for its solution. And yes, in fact, several people *have* lost their souls attempting to solve this problem.

A more subtle but still open question is whether the complexity classes NP and co-NP are different. Even if we can verify every YES answer quickly, there's no reason to believe we can also verify No answers quickly. For example, as far as we know, there is no short proof that a boolean circuit is *not* satisfiable. It is generally believed that $NP \neq co-NP$, but again, nobody knows how to prove it.



What we *think* the world looks like.

30.3 NP-hard, NP-easy, and NP-complete

A problem Π is **NP-hard** if a polynomial-time algorithm for Π would imply a polynomial-time algorithm for *every problem in NP*. In other words:


Π is NP-hard \iff If Π can be solved in polynomial time, then $P=NP$

Intuitively, if we could solve one particular NP-hard problem quickly, then we could quickly solve *any* problem whose solution is easy to understand, using the solution to that one special problem as a subroutine. NP-hard problems are at least as hard as every problem in NP.

Finally, a problem is **NP-complete** if it is both NP-hard and an element of NP (or “NP-easy”). Informally, NP-complete problems are the hardest problems in NP. A polynomial-time algorithm for even one NP-complete problem would immediately imply a polynomial-time algorithm for *every* NP-complete problem. Literally *thousands* of problems have been shown to be NP-complete, so a polynomial-time algorithm for one (and therefore all) of them seems incredibly unlikely.

²Perhaps the most significant progress has taken the form of *barrier* results, which imply that entire avenues of attack are doomed to fail. In a very real sense, these results actually *prove* that we have no idea how to prove $P \neq NP$!

Calling a problem NP-hard is like saying “If I own a dog, then it can speak fluent English.” You probably don’t know whether or not I own a dog, but I bet you’re pretty sure that I don’t own a *talking* dog. Nobody has a mathematical *proof* that dogs can’t speak English—the fact that no one has ever heard a dog speak English is evidence, as are the hundreds of examinations of dogs that lacked the proper mouth shape and brainpower, but mere evidence is not a mathematical proof. Nevertheless, no sane person would believe me if I said I owned a dog that spoke fluent English. So the statement “If I own a dog, then it can speak fluent English” has a natural corollary: No one in their right mind should believe that I own a dog! Likewise, if a problem is NP-hard, no one in their right mind should believe it can be solved in polynomial time.



More of what we *think* the world looks like.

It is not immediately clear that *any* problems are NP-hard. The following remarkable theorem was first published by Steve Cook in 1971 and independently by Leonid Levin in 1973.³ I won’t even sketch the proof here, since I’ve been (deliberately) vague about the definitions; interested readers find a proof in my lecture notes on nondeterministic Turing machines.

The Cook-Levin Theorem. *Circuit satisfiability is NP-hard.*

*30.4 Formal Definitions (*HC SVNT DRACONES*)

Formally, the complexity classes P, NP, and co-NP are defined in terms of *languages* and *Turing machines*. A language is just a set of strings over some finite alphabet Σ ; without loss of generality, we can assume that $\Sigma = \{0, 1\}$. P is the set of languages that can be decided in **P**olynomial time by a deterministic single-tape Turing machine. Similarly, NP is the set of all languages that can be decided in polynomial time by a nondeterministic Turing machine; NP is an abbreviation for **N**ondeterministic **P**olynomial-time.

Polynomial time is a sufficient crude requirement that the precise form of Turing machine (number of heads, number of tracks, and so on) is unimportant. In fact, careful application and analysis of the techniques described in the Turing machine notes imply that any algorithm that runs on a random-access machine⁴ in $T(n)$ time can be simulated by a single-tape, single-track, single-head Turing machine that runs in $O(T(n)^3)$ time. This simulation result allows us to

³Levin first reported his results at seminars in Moscow in 1971, while still a PhD student. News of Cook’s result did not reach the Soviet Union until at least 1973, after Levin’s announcement of his results had been published; in accordance with Stigler’s Law, this result is often called “Cook’s Theorem”. Levin was denied his PhD at Moscow University for political reasons; he emigrated to the US in 1978 and earned a PhD at MIT a year later. Cook was denied tenure at Berkeley in 1970, just one year before publishing his seminal paper; he (but not Levin) later won the Turing award for his proof.

⁴Random-access machines are a model of computation that more faithfully models physical computers. A random-access machine has unbounded random-access memory, modeled as an array $M[0.. \infty]$ where each address $M[i]$ holds a single w -bit integer, for some fixed integer w , and can read to or write from any memory addresses in constant time. RAM algorithms are formally written in assembly-like language, using instructions like **ADD** i, j, k (meaning “ $M[i] \leftarrow M[j] + M[k]$ ”), **INDIR** i, j (meaning “ $M[i] \leftarrow M[M[j]]$ ”), and **IFZGOTO** i, ℓ (meaning “if $M[i] = 0$, go to line ℓ ”). In practice, RAM algorithms can be faithfully described using higher-level pseudocode, as long as we’re careful about arithmetic precision.

argue formally about computational complexity in terms of standard high-level programming constructs like for-loops and recursion, instead of describing everything directly in terms of Turing machines.

A problem Π is formally NP-hard if and only if, for every language $\Pi' \in \text{NP}$, there is a polynomial-time **Turing reduction** from Π' to Π . A Turing reduction just means a reduction that can be executed on a Turing machine; that is, a Turing machine M that can solve Π' using another Turing machine M' for Π as a black-box subroutine. Turing reductions are also called *oracle reductions*; polynomial-time Turing reductions are also called *Cook reductions*.

Researchers in complexity theory prefer to define NP-hardness in terms of polynomial-time **many-one reductions**, which are also called *Karp reductions*. A *many-one* reduction from one language $L' \subseteq \Sigma^*$ to another language $L \subseteq \Sigma^*$ is an function $f : \Sigma^* \rightarrow \Sigma^*$ such that $x \in L'$ if and only if $f(x) \in L$. Then we can define a *language* L to be NP-hard if and only if, for any language $L' \in \text{NP}$, there is a many-one reduction from L' to L that can be computed in polynomial time.

Every Karp reduction “is” a Cook reduction, but not vice versa. Specifically, any Karp reduction from one decision problem Π to another decision Π' is equivalent to transforming the input to Π into the input for Π' , invoking an oracle (that is, a subroutine) for Π' , and then returning the answer verbatim. However, as far as we know, not every Cook reduction can be simulated by a Karp reduction.

Complexity theorists prefer Karp reductions primarily because NP is closed under Karp reductions, but is *not* closed under Cook reductions (unless $\text{NP}=\text{co-NP}$, which is considered unlikely). There are natural problems that are (1) NP-hard with respect to Cook reductions, but (2) NP-hard with respect to Karp reductions only if $\text{P}=\text{NP}$. One trivial example is of such a problem is UNSAT: Given a boolean formula, is it *always false*? On the other hand, many-one reductions apply *only* to decision problems (or more formally, to languages); formally, no optimization or construction problem is Karp-NP-hard.

To make things even more confusing, both Cook and Karp originally defined NP-hardness in terms of **logarithmic-space** reductions. Every logarithmic-space reduction is a polynomial-time reduction, but (as far as we know) not vice versa. It is an open question whether relaxing the set of allowed (Cook or Karp) reductions from logarithmic-space to polynomial-time changes the set of NP-hard problems.

Fortunately, none of these subtleties raise their ugly heads in practice—in particular, every algorithmic reduction described in these notes can be formalized as a logarithmic-space many-one reduction—so you can wake up now.

30.5 Reductions and SAT

To prove that any problem other than Circuit satisfiability is NP-hard, we use a *reduction argument*. Reducing problem A to another problem B means describing an algorithm to solve problem A under the assumption that an algorithm for problem B already exists. You’re already used to doing reductions, only you probably call it something else, like writing subroutines or utility functions, or modular programming. To prove something is NP-hard, we describe a similar transformation between problems, but not in the direction that most people expect.

You should tattoo the following rule of onto the back of your hand, right next to your mom’s birthday and the *actual* rules of Monopoly.⁵

⁵If a player lands on an available property and declines (or is unable) to buy it, that property is immediately auctioned off to the highest bidder; the player who originally declined the property may bid, and bids may be arbitrarily higher or lower than the list price. Players in Jail can still buy and sell property, buy and sell houses and hotels, and collect rent. The game has 32 houses and 12 hotels; once they’re gone, they’re gone. In particular, if all

To prove that problem A is NP-hard, reduce a known NP-hard problem to A.

In other words, to prove that your problem is hard, you need to describe an algorithm to solve a *different* problem, which you already know is hard, using a magical mystery algorithm for *your* problem as a subroutine. The essential logic is a proof by contradiction. The reduction shows implies that if your problem were easy, then the other problem would be easy, too. Equivalently, since you know the other problem is hard, the reduction implies that your problem must also be hard.


For example, consider the *formula satisfiability* problem, usually just called **SAT**. The input to SAT is a boolean *formula* like

$$(a \vee b \vee c \vee \bar{d}) \Leftrightarrow ((b \wedge \bar{c}) \vee (\bar{a} \Rightarrow \bar{d}) \vee (c \neq a \wedge b)),$$

and the question is whether it is possible to assign boolean values to the variables a, b, c, \dots so that the entire formula evaluates to TRUE.

To prove that SAT is NP-hard, we need to give a reduction from a known NP-hard problem. The only problem we know is NP-hard so far is CIRCUITSAT, so let's start there.

Let K be an arbitrary boolean circuit. We can transform K into a boolean formula Φ by creating new output variables for each gate, and then just writing down the list of gates separated by ANDs. For example, our example circuit would be transformed into a formula as follows:



$$(y_1 = x_1 \wedge x_2) \wedge (y_2 = \bar{x}_4) \wedge (y_3 = x_3 \wedge y_2) \wedge (y_4 = y_1 \vee x_2) \wedge (y_5 = \bar{x}_2) \wedge (y_6 = \bar{x}_5) \wedge (y_7 = y_3 \vee y_5) \wedge (z = y_4 \wedge y_7) \wedge z$$

Now we claim that the original circuit K is satisfiable if and only if the resulting formula Φ is satisfiable. We prove this claim in two steps:


- Given a set of inputs that satisfy the circuit K , we can obtain a satisfying assignment for the formula Φ by computing the output of every gate in K .
- Given a satisfying assignment for the formula Φ , we can obtain a satisfying input the the circuit by simply ignoring the internal gate variables y_i and the output variable z .

The entire transformation from circuit to formula can be carried out in linear time. Moreover, the size of the resulting formula is at most a constant factor larger than any reasonable representation of the circuit.

Now suppose, for the sake of argument, there is a magical mystery algorithm that can determine in polynomial time whether a given boolean formula is satisfiable. Then given any

houses are already on the board, you cannot downgrade a hotel to four houses; you must sell all three hotels in the group. Players can sell/exchange undeveloped properties, but not buildings or cash. A player landing on Free Parking does not win anything. A player landing on Go gets \$200, no more. Railroads are not magic transporters. Finally, Jeff always gets the car.

boolean circuit K , we can decide whether K is satisfiable by first transforming K into a boolean formula Φ as described above, and then asking our magical mystery SAT algorithm whether Φ is satisfiable, as suggested by the following cartoon. Each box represents an algorithm. The red box on the left is the transformation subroutine; the box on the right the magical SAT algorithm. It *must* be magic, because it has a rainbow on it!⁶



If you prefer pseudocode to rainbows:

```

CIRCUIVSAT(K):
  transcribe K into a boolean formula Φ
  return SAT(Φ)      <<Magic!>>
    
```

Transcribing K into Φ requires only polynomial time (in fact, only *linear* time, but whatever), so the entire CIRCUIVSAT algorithm also runs in polynomial time.

$$T_{\text{CIRCUIVSAT}}(n) \leq O(n) + T_{\text{SAT}}(O(n))$$

We conclude that any polynomial-time algorithm for SAT would give us a polynomial-time algorithm for CIRCUIVSAT, which in turn would imply $P=NP$. So SAT is NP-hard!

30.6 3SAT (from SAT)

A special case of SAT that is particularly useful in proving NP-hardness results is called 3SAT.

A boolean formula is in *conjunctive normal form* (CNF) if it is a conjunction (AND) of several *clauses*, each of which is the disjunction (OR) of several *literals*, each of which is either a variable or its negation. For example:

$$(a \vee b \vee c \vee d) \wedge (b \vee \bar{c} \vee \bar{d}) \wedge (\bar{a} \vee c \vee d) \wedge (a \vee \bar{b})$$


A 3CNF formula is a CNF formula with *exactly three literals per clause*; the previous example is not a 3CNF formula, since its first clause has four literals and its last clause has only two. 3SAT is just SAT restricted to 3CNF formulas: Given a 3CNF formula, is there an assignment to the variables that makes the formula evaluate to TRUE?

We could prove that 3SAT is NP-hard by a reduction from the more general SAT problem, but it's easier just to start over from scratch, by reducing directly from CIRCUIVSAT.

Let K be an arbitrary boolean circuit. We transform K into a 3CNF formula in several stages.

1. Make sure every AND and OR gate in K has exactly two inputs. If any gate has $k > 2$ inputs, replace it with a binary tree of $k - 1$ two-input gates. Call the resulting circuit K' .

⁶Katherine Z. Erickson. Personal communication, 2011.




To prove that X is NP-hard.

- ① Pick a known NP-hard problem Y
- ② Assume (for the sake of argument) a poly-time algo for X
- ③ Derive a poly-time algo for Y using X as a subroutine
- ④ HED ASPLUDE.

Circuit SAT
SAT

"Reduce Y to X "



Polynomial-time reduction from CIRCUITSAT to 3SAT

- Transcribe K' into a boolean formula Φ_1 with one clause per gate, exactly as in our previous reduction to SAT. $b \text{---} \text{AND} \text{---} c \Rightarrow (a = b \wedge c)$
- Replace each clause in Φ_1 with a CNF formula. There are only three types of clauses in Φ_1 , one for each type of gate in K' :

$$\begin{aligned}
 a = b \wedge c &\mapsto (a \vee \bar{b} \vee \bar{c}) \wedge (\bar{a} \vee b) \wedge (\bar{a} \vee c) \\
 a = b \vee c &\mapsto (\bar{a} \vee b \vee c) \wedge (a \vee \bar{b}) \wedge (a \vee \bar{c}) \\
 a = \bar{b} &\mapsto (a \vee b) \wedge (\bar{a} \vee \bar{b})
 \end{aligned}$$

Call the resulting CNF formula Φ_2 .

- Replace each clause in Φ_2 with a 3CNF formula. Every clause in Φ_2 has at most three literals. We can keep the three-literal clauses as-is. We expand each two-literal clause into two three-literal clauses by introducing a new variable. Finally, we expand any one-literal clause into four three-literal clauses by introducing two new variables.

$$\begin{aligned}
 a \vee b &\mapsto (a \vee b \vee x) \wedge (a \vee b \vee \bar{x}) \\
 a &\mapsto (a \vee x \vee y) \wedge (a \vee \bar{x} \vee y) \wedge (a \vee x \vee \bar{y}) \wedge (a \vee \bar{x} \vee \bar{y})
 \end{aligned}$$

Call the final 3CNF formula Φ_3 .

For example, if we start with the same example circuit we used earlier, we obtain the following 3CNF formula Φ_3 .

$$\begin{aligned}
 &(y_1 \vee \bar{x}_1 \vee \bar{x}_4) \wedge (\bar{y}_1 \vee x_1 \vee z_1) \wedge (\bar{y}_1 \vee x_1 \vee \bar{z}_1) \wedge (\bar{y}_1 \vee x_4 \vee z_2) \wedge (\bar{y}_1 \vee x_4 \vee \bar{z}_2) \\
 &\wedge (y_2 \vee x_4 \vee z_3) \wedge (y_2 \vee x_4 \vee \bar{z}_3) \wedge (\bar{y}_2 \vee \bar{x}_4 \vee z_4) \wedge (\bar{y}_2 \vee \bar{x}_4 \vee \bar{z}_4) \\
 &\wedge (y_3 \vee \bar{x}_3 \vee \bar{y}_2) \wedge (\bar{y}_3 \vee x_3 \vee z_5) \wedge (\bar{y}_3 \vee x_3 \vee \bar{z}_5) \wedge (\bar{y}_3 \vee y_2 \vee z_6) \wedge (\bar{y}_3 \vee y_2 \vee \bar{z}_6) \\
 &\wedge (\bar{y}_4 \vee y_1 \vee x_2) \wedge (y_4 \vee \bar{x}_2 \vee z_7) \wedge (y_4 \vee \bar{x}_2 \vee \bar{z}_7) \wedge (y_4 \vee \bar{y}_1 \vee z_8) \wedge (y_4 \vee \bar{y}_1 \vee \bar{z}_8) \\
 &\wedge (y_5 \vee x_2 \vee z_9) \wedge (y_5 \vee x_2 \vee \bar{z}_9) \wedge (\bar{y}_5 \vee \bar{x}_2 \vee z_{10}) \wedge (\bar{y}_5 \vee \bar{x}_2 \vee \bar{z}_{10}) \\
 &\wedge (y_6 \vee x_5 \vee z_{11}) \wedge (y_6 \vee x_5 \vee \bar{z}_{11}) \wedge (\bar{y}_6 \vee \bar{x}_5 \vee z_{12}) \wedge (\bar{y}_6 \vee \bar{x}_5 \vee \bar{z}_{12}) \\
 &\wedge (\bar{y}_7 \vee y_3 \vee y_5) \wedge (y_7 \vee \bar{y}_3 \vee z_{13}) \wedge (y_7 \vee \bar{y}_3 \vee \bar{z}_{13}) \wedge (y_7 \vee \bar{y}_5 \vee z_{14}) \wedge (y_7 \vee \bar{y}_5 \vee \bar{z}_{14}) \\
 &\wedge (y_8 \vee \bar{y}_4 \vee \bar{y}_7) \wedge (\bar{y}_8 \vee y_4 \vee z_{15}) \wedge (\bar{y}_8 \vee y_4 \vee \bar{z}_{15}) \wedge (\bar{y}_8 \vee y_7 \vee z_{16}) \wedge (\bar{y}_8 \vee y_7 \vee \bar{z}_{16}) \\
 &\wedge (y_9 \vee \bar{y}_8 \vee \bar{y}_6) \wedge (\bar{y}_9 \vee y_8 \vee z_{17}) \wedge (\bar{y}_9 \vee y_8 \vee \bar{z}_{17}) \wedge (\bar{y}_9 \vee y_6 \vee z_{18}) \wedge (\bar{y}_9 \vee y_6 \vee \bar{z}_{18}) \\
 &\wedge (y_9 \vee z_{19} \vee z_{20}) \wedge (y_9 \vee \bar{z}_{19} \vee z_{20}) \wedge (y_9 \vee z_{19} \vee \bar{z}_{20}) \wedge (y_9 \vee \bar{z}_{19} \vee \bar{z}_{20})
 \end{aligned}$$

Although this formula may look a lot more ugly and complicated than the original circuit at first glance, it's actually only a constant factor larger—every binary gate in the original circuit has


been transformed into at most five clauses. Even if the formula size were a large *polynomial* function (like n^{573}) of the circuit size, we would still have a valid reduction.

This process transforms the circuit into an equivalent 3CNF formula; the output formula is satisfiable if and only if the input circuit is satisfiable. As with the more general SAT problem, the output formula Φ_3 is only a constant factor larger than any reasonable description of the original circuit K , and the reduction can be carried out in polynomial time. Thus, if 3SAT can be solved in polynomial time, then CIRCUITSAT can be solved in polynomial time, which implies that $P = NP$. We conclude 3SAT is NP-hard.

30.7 Maximum Independent Set (from 3SAT)


For the next few problems we consider, the input is a simple, unweighted graph, and the problem asks for the size of the largest or smallest subgraph satisfying some structural property.

Let G be an arbitrary graph. An *independent set* in G is a subset of the vertices of G with no edges between them. The *maximum independent set* problem, or simply **MAXINDSET**, asks for the size of the largest independent set in a given graph. I will prove that MAXINDSET is NP-hard using a reduction from 3SAT, as suggested by the following figure.



Polynomial-time reduction from 3SAT to MAXINDSET

Given an arbitrary 3CNF formula Φ , we construct a graph G as follows. Let k denote the number of clauses in Φ . The graph G contains exactly $3k$ vertices, one for each literal in Φ . Two vertices in G are connected by an edge if and only if either (1) they correspond to literals in the same clause, or (2) they correspond to a variable and its inverse. For example, the formula $(a \vee b \vee c) \wedge (b \vee \bar{c} \vee \bar{d}) \wedge (\bar{a} \vee c \vee d) \wedge (a \vee \bar{b} \vee \bar{d})$ is transformed into the following graph.



A graph derived from a satisfiable 3CNF formula, and an independent set of size 4. Black edges join literals from the same clause; red (heavier) edges join contradictory literals.

Any independent set in G contains at most one vertex from each clause triangle, because any two vertices in each triangle are connected. Thus, the largest independent set in G has size at