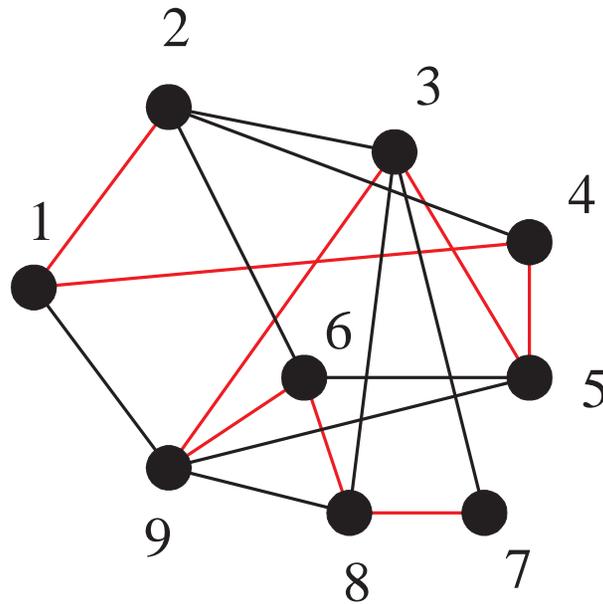


HAMILTONIAN PATH

- A **Hamiltonian path** of a graph is a path that visits every node of the graph exactly once.
- Suppose graph G has n nodes: $1, 2, \dots, n$.
- A Hamiltonian path can be expressed as a permutation π of $\{1, 2, \dots, n\}$ such that
 - $\pi(i) = j$ means the i th position is occupied by node j .
 - $(\pi(i), \pi(i + 1)) \in G$ for $i = 1, 2, \dots, n - 1$.
- HAMILTONIAN PATH asks if a graph has a Hamiltonian path.

Reduction of HAMILTONIAN PATH to SAT

- Given a graph G , we shall construct a CNF $R(G)$ such that $R(G)$ is satisfiable iff G has a Hamiltonian path.
- $R(G)$ has n^2 boolean variables x_{ij} , $1 \leq i, j \leq n$.
- x_{ij} means
“the i th position in the Hamiltonian path is occupied by node j .”



$$\begin{aligned}
 &x_{12} = x_{21} = x_{34} = x_{45} = x_{53} = x_{69} = x_{76} = x_{88} = x_{97} = 1; \\
 &\pi(1) = 2, \pi(2) = 1, \pi(3) = 4, \pi(4) = 5, \pi(5) = 3, \pi(6) = \\
 &9, \pi(7) = 6, \pi(8) = 8, \pi(9) = 7.
 \end{aligned}$$

The Clauses of $R(G)$ and Their Intended Meanings

1. Each node j must appear in the path.
 - $x_{1j} \vee x_{2j} \vee \cdots \vee x_{nj}$ for each j .
2. No node j appears twice in the path.
 - $\neg x_{ij} \vee \neg x_{kj}$ for all i, j, k with $i \neq k$.
3. Every position i on the path must be occupied.
 - $x_{i1} \vee x_{i2} \vee \cdots \vee x_{in}$ for each i .
4. No two nodes j and k occupy the same position in the path.
 - $\neg x_{ij} \vee \neg x_{ik}$ for all i, j, k with $j \neq k$.
5. Nonadjacent nodes i and j cannot be adjacent in the path.
 - $\neg x_{ki} \vee \neg x_{k+1,j}$ for all $(i, j) \notin G$ and $k = 1, 2, \dots, n - 1$.

The Proof

- $R(G)$ contains $O(n^3)$ clauses.
- $R(G)$ can be computed efficiently (simple exercise).
- Suppose $T \models R(G)$.
- From the 1st and 2nd types of clauses, for each node j there is a unique position i such that $T \models x_{ij}$.
- From the 3rd and 4th types of clauses, for each position i there is a unique node j such that $T \models x_{ij}$.
- So there is a permutation π of the nodes such that $\pi(i) = j$ if and only if $T \models x_{ij}$.

The Proof (concluded)

- The 5th type of clauses furthermore guarantee that $(\pi(1), \pi(2), \dots, \pi(n))$ is a Hamiltonian path.
- Conversely, suppose G has a Hamiltonian path

$$(\pi(1), \pi(2), \dots, \pi(n)),$$

where π is a permutation.

- Clearly, the truth assignment

$$T(x_{ij}) = \mathbf{true} \text{ if and only if } \pi(i) = j$$

satisfies all clauses of $R(G)$.

A Comment^a

- An answer to “Is $R(G)$ satisfiable?” does answer “Is G Hamiltonian?”
- But a positive answer does not give a Hamiltonian path for G .
 - Providing a witness is not a requirement of reduction.
- A positive answer to “Is $R(G)$ satisfiable?” plus a satisfying truth assignment does provide us with a Hamiltonian path for G .

^aContributed by Ms. Amy Liu (J94922016) on May 29, 2006.

Reduction of REACHABILITY to CIRCUIT VALUE

- Note that both problems are in P.
- Given a graph $G = (V, E)$, we shall construct a *variable-free* circuit $R(G)$.
- The output of $R(G)$ is true if and only if there is a path from node 1 to node n in G .
- Idea: the Floyd-Warshall algorithm.

The Gates

- The gates are
 - g_{ijk} with $1 \leq i, j \leq n$ and $0 \leq k \leq n$.
 - h_{ijk} with $1 \leq i, j, k \leq n$.
- g_{ijk} : There is a path from node i to node j without passing through a node bigger than k .
- h_{ijk} : There is a path from node i to node j passing through k but not any node bigger than k .
- Input gate $g_{ij0} = \text{true}$ if and only if $i = j$ or $(i, j) \in E$.

The Construction

- h_{ijk} is an AND gate with predecessors $g_{i,k,k-1}$ and $g_{k,j,k-1}$, where $k = 1, 2, \dots, n$.
- g_{ijk} is an OR gate with predecessors $g_{i,j,k-1}$ and $h_{i,j,k}$, where $k = 1, 2, \dots, n$.
- g_{1nn} is the output gate.
- Interestingly, $R(G)$ uses no \neg gates.
 - It is a **monotone circuit**.

Reduction of CIRCUIT SAT to SAT

- Given a circuit C , we will construct a boolean expression $R(C)$ such that $R(C)$ is satisfiable iff C is.
 - $R(C)$ will turn out to be a CNF.
 - $R(C)$ is a depth-2 circuit; furthermore, each gate has out-degree 1.
- The variables of $R(C)$ are those of C plus g for each gate g of C .
 - The g 's propagate the truth values for the CNF.
- Each gate of C will be turned into equivalent clauses.
- Recall that clauses are \wedge ed together by definition.

The Clauses of $R(C)$

g is a **variable gate** x : Add clauses $(\neg g \vee x)$ and $(g \vee \neg x)$.

- Meaning: $g \Leftrightarrow x$.

g is a **true gate**: Add clause (g) .

- Meaning: g must be true to make $R(C)$ true.

g is a **false gate**: Add clause $(\neg g)$.

- Meaning: g must be false to make $R(C)$ true.

g is a **\neg gate with predecessor gate h** : Add clauses $(\neg g \vee \neg h)$ and $(g \vee h)$.

- Meaning: $g \Leftrightarrow \neg h$.

The Clauses of $R(C)$ (concluded)

g is a \vee gate with predecessor gates h and h' : Add clauses $(\neg h \vee g)$, $(\neg h' \vee g)$, and $(h \vee h' \vee \neg g)$.

- Meaning: $g \Leftrightarrow (h \vee h')$.

g is a \wedge gate with predecessor gates h and h' : Add clauses $(\neg g \vee h)$, $(\neg g \vee h')$, and $(\neg h \vee \neg h' \vee g)$.

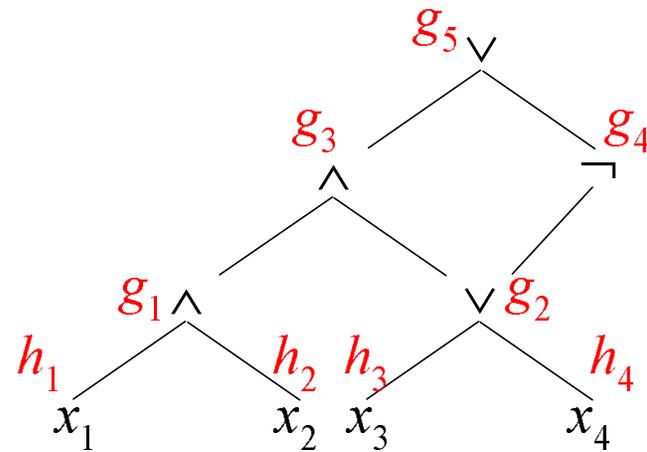
- Meaning: $g \Leftrightarrow (h \wedge h')$.

g is the output gate: Add clause (g) .

- Meaning: g must be true to make $R(C)$ true.

Note: If gate g feeds gates h_1, h_2, \dots , then variable g appears in the clauses for h_1, h_2, \dots in $R(C)$.

An Example



$$\begin{aligned}
 & (h_1 \Leftrightarrow x_1) \wedge (h_2 \Leftrightarrow x_2) \wedge (h_3 \Leftrightarrow x_3) \wedge (h_4 \Leftrightarrow x_4) \\
 \wedge & [g_1 \Leftrightarrow (h_1 \wedge h_2)] \wedge [g_2 \Leftrightarrow (h_3 \vee h_4)] \\
 \wedge & [g_3 \Leftrightarrow (g_1 \wedge g_2)] \wedge (g_4 \Leftrightarrow \neg g_2) \\
 \wedge & [g_5 \Leftrightarrow (g_3 \vee g_4)] \wedge g_5.
 \end{aligned}$$

An Example (concluded)

- In general, the result is a CNF.
- The CNF has size proportional to the circuit's number of gates.
- The CNF adds new variables to the circuit's original input variables.

Composition of Reductions

Proposition 25 *If R_{12} is a reduction from L_1 to L_2 and R_{23} is a reduction from L_2 to L_3 , then the composition $R_{12} \circ R_{23}$ is a reduction from L_1 to L_3 .*

- So reducibility is transitive.

Completeness^a

- As reducibility is transitive, problems can be ordered with respect to their difficulty.
- Is there a *maximal* element?
- It is not obvious that there should be a maximal element.
 - Many infinite structures (such as integers and real numbers) do not have maximal elements.
- Hence it may surprise you that most of the complexity classes that we have seen so far have maximal elements.

^aCook (1971) and Levin (1973).

Completeness (concluded)

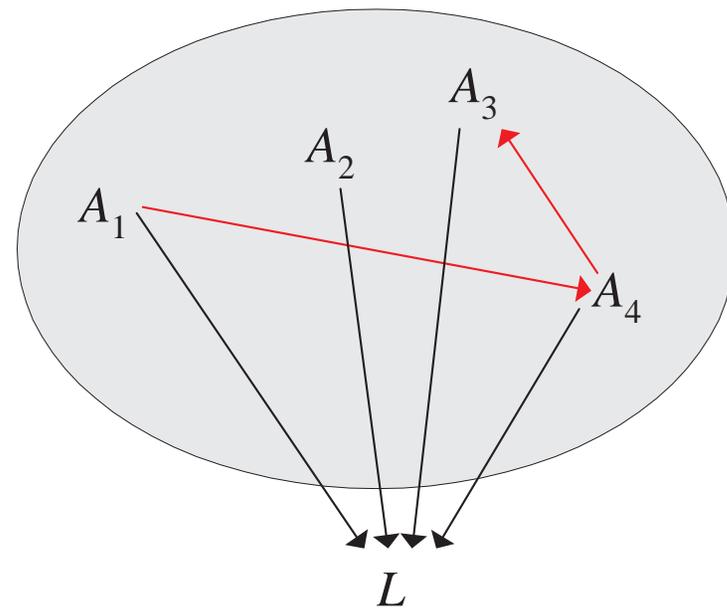
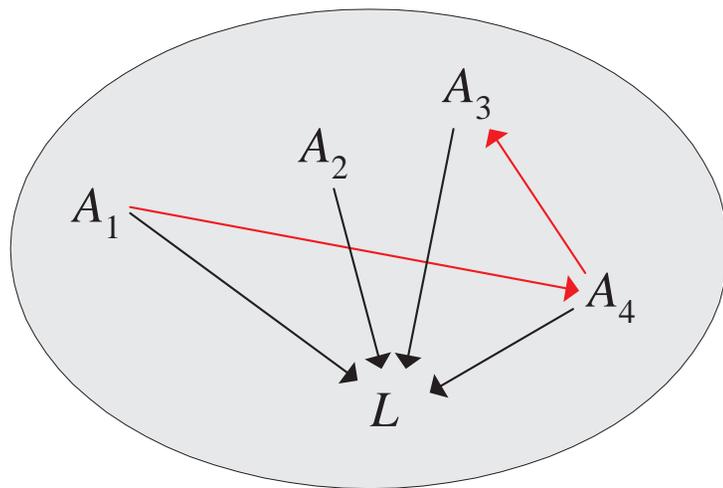
- Let \mathcal{C} be a complexity class and $L \in \mathcal{C}$.
- L is **\mathcal{C} -complete** if every $L' \in \mathcal{C}$ can be reduced to L .
 - Most complexity classes we have seen so far have complete problems!
- Complete problems capture the difficulty of a class because they are the hardest problems in the class.

Hardness

- Let \mathcal{C} be a complexity class.
- L is **\mathcal{C} -hard** if every $L' \in \mathcal{C}$ can be reduced to L .
- It is not required that $L \in \mathcal{C}$.
- If L is \mathcal{C} -hard, then by definition, every \mathcal{C} -complete problem can be reduced to L .^a

^aContributed by Mr. Ming-Feng Tsai (D92922003) on October 15, 2003.

Illustration of Completeness and Hardness



Closedness under Reductions

- A class \mathcal{C} is **closed under reductions** if whenever L is reducible to L' and $L' \in \mathcal{C}$, then $L \in \mathcal{C}$.
- P, NP, coNP, L, NL, PSPACE, and EXP are all closed under reductions.

Complete Problems and Complexity Classes

Proposition 26 *Let \mathcal{C}' and \mathcal{C} be two complexity classes such that $\mathcal{C}' \subseteq \mathcal{C}$. Assume \mathcal{C}' is closed under reductions and L is \mathcal{C} -complete. Then $\mathcal{C} = \mathcal{C}'$ if and only if $L \in \mathcal{C}'$.*

- Suppose $L \in \mathcal{C}'$ first.
- Every language $A \in \mathcal{C}$ reduces to $L \in \mathcal{C}'$.
- Because \mathcal{C}' is closed under reductions, $A \in \mathcal{C}'$.
- Hence $\mathcal{C} \subseteq \mathcal{C}'$.
- As $\mathcal{C}' \subseteq \mathcal{C}$, we conclude that $\mathcal{C} = \mathcal{C}'$.

The Proof (concluded)

- On the other hand, suppose $\mathcal{C} = \mathcal{C}'$.
- As L is \mathcal{C} -complete, $L \in \mathcal{C}$.
- Thus, trivially, $L \in \mathcal{C}'$.

Two Important Corollaries

Proposition 26 implies the following.

Corollary 27 *$P = NP$ if and only if an NP-complete problem is in P .*

Corollary 28 *$L = P$ if and only if a P-complete problem is in L .*

Complete Problems and Complexity Classes

Proposition 29 *Let \mathcal{C}' and \mathcal{C} be two complexity classes closed under reductions. If L is complete for both \mathcal{C} and \mathcal{C}' , then $\mathcal{C} = \mathcal{C}'$.*

- All languages $\mathcal{L} \in \mathcal{C}$ reduce to $L \in \mathcal{C}'$.
- Since \mathcal{C}' is closed under reductions, $\mathcal{L} \in \mathcal{C}'$.
- Hence $\mathcal{C} \subseteq \mathcal{C}'$.
- The proof for $\mathcal{C}' \subseteq \mathcal{C}$ is symmetric.

Table of Computation

- Let $M = (K, \Sigma, \delta, s)$ be a single-string polynomial-time deterministic TM deciding L .
- Its computation on input x can be thought of as a $|x|^k \times |x|^k$ table, where $|x|^k$ is the time bound.
 - It is a sequence of configurations.
- Rows correspond to time steps 0 to $|x|^k - 1$.
- Columns are positions in the string of M .
- The (i, j) th table entry represents the contents of position j of the string *after* i steps of computation.

Some Conventions To Simplify the Table

- M halts after at most $|x|^k - 2$ steps.
- Assume a large enough k to make it true for $|x| \geq 2$.
- Pad the table with \sqcup s so that each row has length $|x|^k$.
 - The computation will never reach the right end of the table for lack of time.
- If the cursor scans the j th position at time i when M is at state q and the symbol is σ , then the (i, j) th entry is a *new* symbol σ_q .

Some Conventions To Simplify the Table (continued)

- If q is “yes” or “no,” simply use “yes” or “no” instead of σ_q .
- Modify M so that the cursor starts not at \triangleright but at the first symbol of the input.
- The cursor never visits the leftmost \triangleright by telescoping two moves of M each time the cursor is about to move to the leftmost \triangleright .
- So the first symbol in every row is a \triangleright and not a \triangleright_q .

Some Conventions To Simplify the Table (concluded)

- Suppose M has halted before its time bound of $|x|^k$, so that “yes” or “no” appears at a row before the last.
- Then all subsequent rows will be identical to that row.
- M accepts x if and only if the $(|x|^k - 1, j)$ th entry is “yes” for some position j .

Comments

- Each row is essentially a configuration.
- If the input $x = 010001$, then the first row is

$$\begin{array}{c} |x|^k \\ \hline \triangleright 0_s 10001 \square \square \cdots \square \end{array}$$

- A typical row may look like

$$\begin{array}{c} |x|^k \\ \hline \triangleright 10100_q 01110100 \square \square \cdots \square \end{array}$$

Comments (concluded)

- The last rows must look like

$$\overbrace{\triangleright \dots \text{“yes”} \dots \square}^{|x|^k} \quad \text{or} \quad \overbrace{\triangleright \dots \text{“no”} \dots \square}^{|x|^k}$$

- Three out of the table's 4 borders are known:

$$\begin{array}{cccccc} \triangleright & \mathbf{a} & \mathbf{b} & \mathbf{c} & \mathbf{d} & \mathbf{e} & \mathbf{f} & \square \\ \triangleright & & & & & & & \square \\ \triangleright & & & & & & & \square \\ \triangleright & & & & & & & \square \\ \triangleright & & & & & & & \square \\ & & & & \vdots & & & \square \end{array}$$

A P-Complete Problem

Theorem 30 (Ladner (1975)) CIRCUI T VALUE *is P-complete.*

- It is easy to see that CIRCUI T VALUE \in P.
- For *any* $L \in$ P, we will construct a reduction R from L to CIRCUI T VALUE.
- Given any input x , $R(x)$ is a variable-free circuit such that $x \in L$ if and only if $R(x)$ evaluates to true.
- Let M decide L in time n^k .
- Let T be the computation table of M on x .

The Proof (continued)

- When $i = 0$, or $j = 0$, or $j = |x|^k - 1$, then the value of T_{ij} is known.
 - The j th symbol of x or \sqcup , a \triangleright , and a \sqcup , respectively.
 - Recall that three out of T 's 4 borders are known.

The Proof (continued)

- Consider *other* entries T_{ij} .
- T_{ij} depends on only $T_{i-1,j-1}$, $T_{i-1,j}$, and $T_{i-1,j+1}$:

$T_{i-1,j-1}$	$T_{i-1,j}$	$T_{i-1,j+1}$
	T_{ij}	

- Let Γ denote the set of all symbols that can appear on the table: $\Gamma = \Sigma \cup \{\sigma_q : \sigma \in \Sigma, q \in K\}$.
- Encode each symbol of Γ as an m -bit number, where^a

$$m = \lceil \log_2 |\Gamma| \rceil.$$

^aCalled **state assignment** in circuit design.

The Proof (continued)

- Let the m -bit binary string $S_{ij1}S_{ij2} \cdots S_{ijm}$ encode T_{ij} .
- We may treat them interchangeably without ambiguity.
- The computation table is now a table of binary entries S_{ijl} , where

$$0 \leq i \leq n^k - 1,$$

$$0 \leq j \leq n^k - 1,$$

$$1 \leq l \leq m.$$

The Proof (continued)

- Each bit $S_{ij\ell}$ depends on only $3m$ other bits:

$$T_{i-1,j-1}: \quad S_{i-1,j-1,1} \quad S_{i-1,j-1,2} \quad \cdots \quad S_{i-1,j-1,m}$$

$$T_{i-1,j}: \quad S_{i-1,j,1} \quad S_{i-1,j,2} \quad \cdots \quad S_{i-1,j,m}$$

$$T_{i-1,j+1}: \quad S_{i-1,j+1,1} \quad S_{i-1,j+1,2} \quad \cdots \quad S_{i-1,j+1,m}$$

The Proof (continued)

- There is a boolean function F_ℓ with $3m$ inputs such that

$$\begin{aligned}
 & S_{ij\ell} \\
 = & F_\ell \left(\overbrace{S_{i-1,j-1,1}, S_{i-1,j-1,2}, \dots, S_{i-1,j-1,m}}^{T_{i-1,j-1}}, \right. \\
 & \quad \left. \overbrace{S_{i-1,j,1}, S_{i-1,j,2}, \dots, S_{i-1,j,m}}^{T_{i-1,j}}, \right. \\
 & \quad \left. \overbrace{S_{i-1,j+1,1}, S_{i-1,j+1,2}, \dots, S_{i-1,j+1,m}}^{T_{i-1,j+1}} \right),
 \end{aligned}$$

where for all $i, j > 0$ and $1 \leq \ell \leq m$.

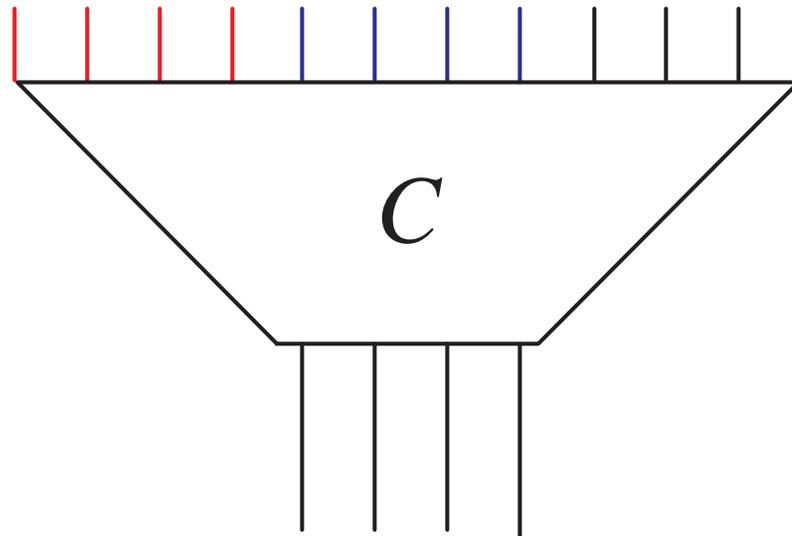
The Proof (continued)

- These F_i 's depend only on M 's specification, not on x .
- Their sizes are constant.
- These boolean functions can be turned into boolean circuits (see p. 166).
- Compose these m circuits in parallel to obtain circuit C with $3m$ -bit inputs and m -bit outputs.
 - Schematically, $C(T_{i-1,j-1}, T_{i-1,j}, T_{i-1,j+1}) = T_{ij}$.^a

^a C is like an ASIC (application-specific IC) chip.

Circuit C

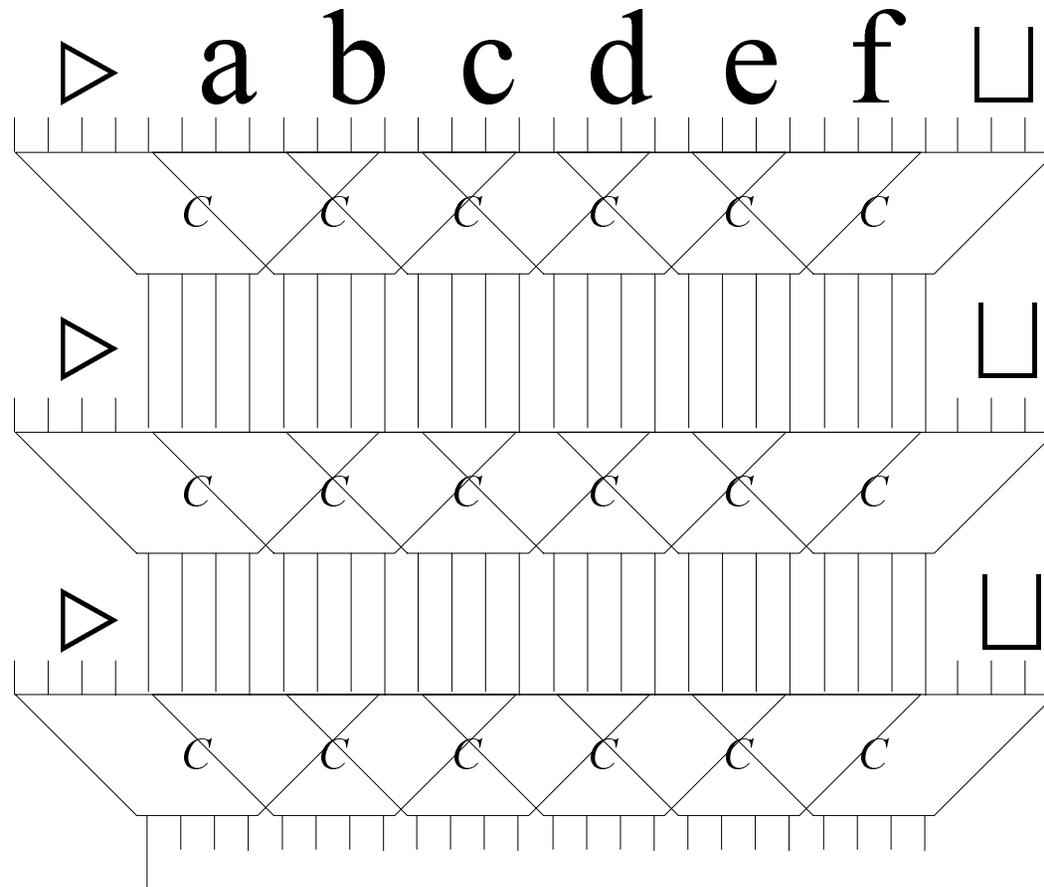
$T_{i-1,j-1}$ $T_{i-1,j}$ $T_{i-1,j+1}$



The Proof (concluded)

- A copy of circuit C is placed at each entry of the table.
 - Exceptions are the top row and the two extreme columns.
- $R(x)$ consists of $(|x|^k - 1)(|x|^k - 2)$ copies of circuit C .
- Without loss of generality, assume the output “yes” / “no” appear at position $(|x|^k - 1, 1)$.
- Encode “yes” as 1 and “no” as 0.

The Computation Tableau and $R(x)$



A Corollary

The construction in the above proof yields the following, more general result.

Corollary 31 *If $L \in TIME(T(n))$, then a circuit with $O(T^2(n))$ gates can decide if $x \in L$ for $|x| = n$.*

MONOTONE CIRCUIT VALUE

- A **monotone** boolean circuit's output cannot change from true to false when one input changes from false to true.
- Monotone boolean circuits are hence less expressive than general circuits.
 - They can compute only *monotone* boolean functions.
- Monotone circuits do not contain \neg gates (prove it).
- MONOTONE CIRCUIT VALUE is CIRCUIT VALUE applied to monotone circuits.

MONOTONE CIRCUIT VALUE Is P-Complete

Despite their limitations, MONOTONE CIRCUIT VALUE is as hard as CIRCUIT VALUE.

Corollary 32 MONOTONE CIRCUIT VALUE *is P-complete.*

- Given any general circuit, we can “move the \neg 's downwards” using de Morgan's laws. (Why?)

Cook's Theorem: the First NP-Complete Problem

Theorem 33 (Cook (1971)) *SAT is NP-complete.*

- $\text{SAT} \in \text{NP}$ (p. 87).
- CIRCUIT SAT reduces to SAT (p. 229).
- Now we only need to show that all languages in NP can be reduced to CIRCUIT SAT .^a

^aThis also shows CIRCUIT SAT is NP-complete.

The Proof (continued)

- Let single-string NTM M decide $L \in \text{NP}$ in time n^k .
- Assume M has exactly *two* nondeterministic choices at each step: choices 0 and 1.
- For each input x , we construct circuit $R(x)$ such that $x \in L$ if and only if $R(x)$ is satisfiable.
- A sequence of nondeterministic choices is a bit string

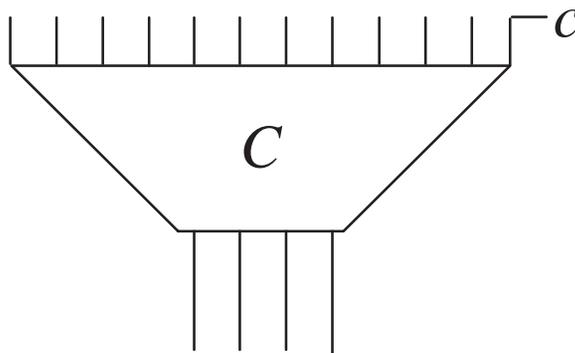
$$B = (c_1, c_2, \dots, c_{|x|^k-1}) \in \{0, 1\}^{|x|^k-1}.$$

- Once B is given, the computation is *deterministic*.

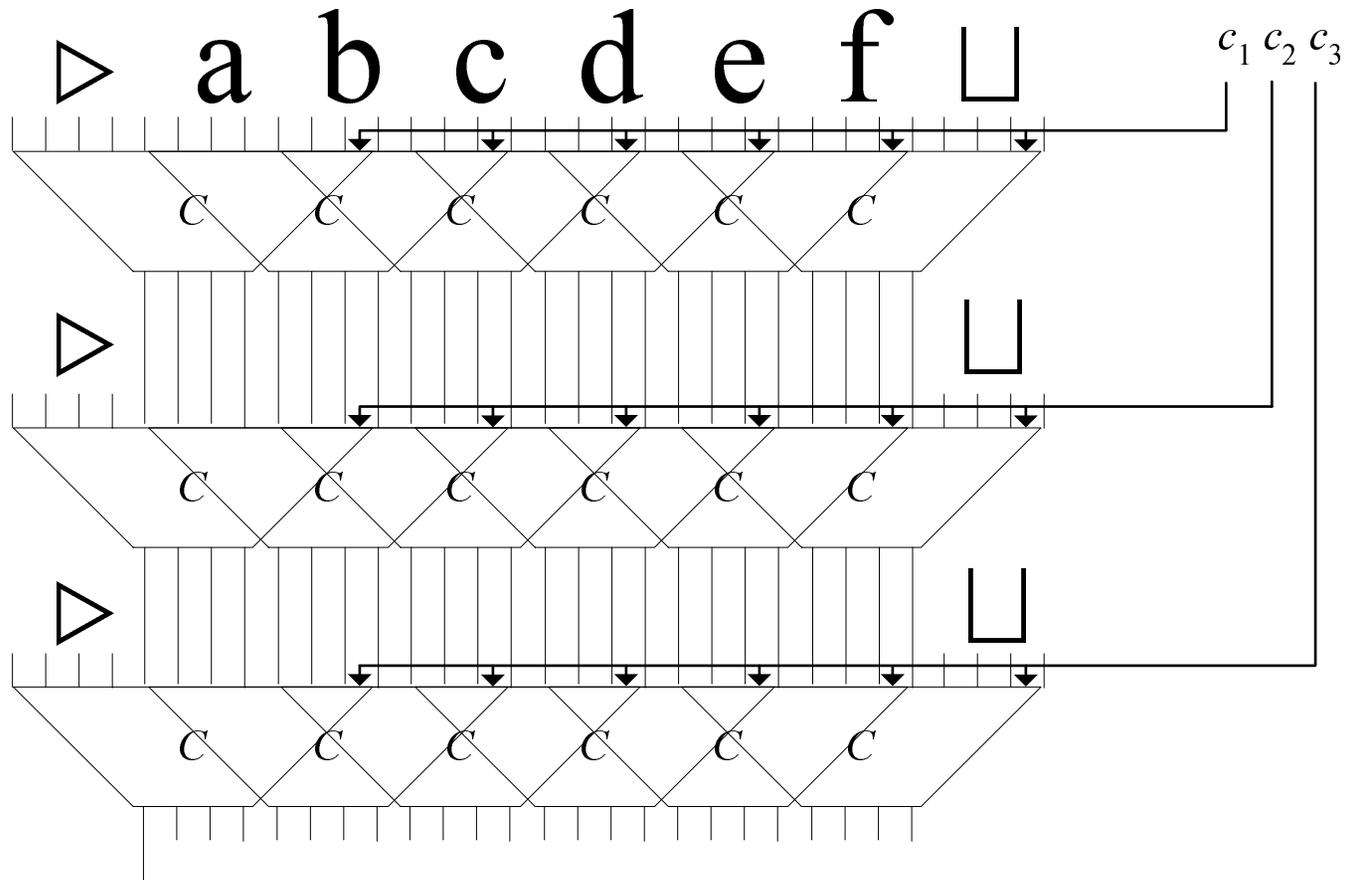
The Proof (continued)

- Each choice of B results in a deterministic polynomial-time computation.
- Each circuit C at time i has an extra binary input c corresponding to the nondeterministic choice:

$$C(T_{i-1,j-1}, T_{i-1,j}, T_{i-1,j+1}, c) = T_{ij}.$$



The Computation Tableau for NTMs and $R(x)$

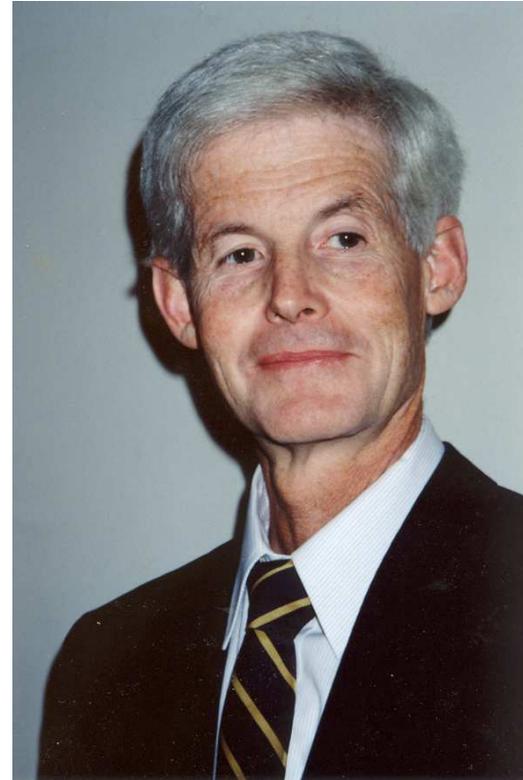


The Proof (concluded)

- The overall circuit $R(x)$ (on p. 266) is satisfiable if there is a truth assignment B such that the computation table accepts.
- This happens if and only if M accepts x , i.e., $x \in L$.

Stephen Arthur Cook^a (1939–)

Richard Karp, “It is to our everlasting shame that we were unable to persuade the math department [of UC-Berkeley] to give him tenure.”



^aTuring Award (1982).

NP-Complete Problems

Wir müssen wissen, wir werden wissen.
(We must know, we shall know.)
— David Hilbert (1900)

I predict that scientists will one day adopt a new principle: “NP-complete problems are hard.”
That is, solving those problems efficiently is impossible on any device that could be built in the real world, whatever the final laws of physics turn out to be.
— Scott Aaronson (2008)

Two Notions

- Let $R \subseteq \Sigma^* \times \Sigma^*$ be a binary relation on strings.
- R is called **polynomially decidable** if

$$\{x; y : (x, y) \in R\}$$

is in P.

- R is said to be **polynomially balanced** if $(x, y) \in R$ implies $|y| \leq |x|^k$ for some $k \geq 1$.

An Alternative Characterization of NP

Proposition 34 (Edmonds (1965)) *Let $L \subseteq \Sigma^*$ be a language. Then $L \in NP$ if and only if there is a polynomially decidable and polynomially balanced relation R such that*

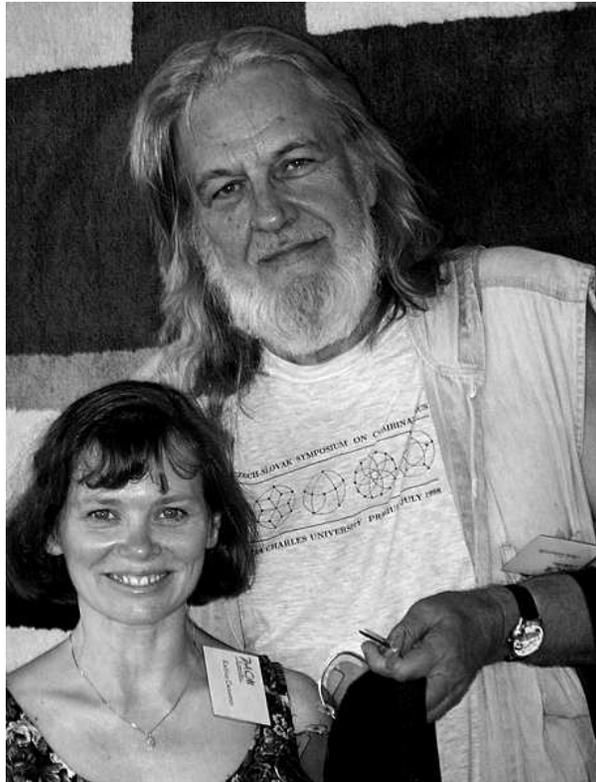
$$L = \{x : \exists y (x, y) \in R\}.$$

- Suppose such an R exists.
- L can be decided by this NTM:
 - On input x , the NTM guesses a y of length $\leq |x|^k$.
 - It then tests if $(x, y) \in R$ in polynomial time.
 - It returns “yes” if the test is positive.

The Proof (concluded)

- Now suppose $L \in \text{NP}$.
- NTM N decides L in time $|x|^k$.
- Define R as follows: $(x, y) \in R$ if and only if y is the encoding of an accepting computation of N on input x .
- R is polynomially balanced as N is polynomially bounded.
- R is polynomially decidable because it can be efficiently verified by consulting N 's transition function.
- Finally $L = \{x : (x, y) \in R \text{ for some } y\}$ because N decides L .

Jack Edmonds (1934–)



Comments

- Any “yes” instance x of an NP problem has at least one **succinct certificate** or **polynomial witness** y .
- “No” instances have none.
- Certificates are short and easy to verify.
 - An alleged satisfying truth assignment for SAT, an alleged Hamiltonian path for HAMILTONIAN PATH, etc.
- Certificates may be hard to generate,^a but verification must be easy.
- NP is the class of *easy-to-verify* (i.e., in P) problems.

^aUnless P equals NP.

Levin Reduction and Parsimonious Reductions

- The reduction R in Cook's theorem (p. 263) is such that
 - Each satisfying truth assignment for circuit $R(x)$ corresponds to an accepting computation path for $M(x)$.
- It actually yields an efficient way to transform a certificate for x to a satisfying assignment for $R(x)$, and vice versa.
- A reduction with this property is called a **Levin reduction**.^a

^aLevin is the co-inventor of NP-completeness, in 1973.

Leonid Levin (1948–)

Leonid Levin (1998), “Mathematicians often think that historical evidence is that NP is exponential. Historical evidence is quite strongly in the other direction.”



Levin Reduction and Parsimonious Reductions (concluded)

- Furthermore, the proof gives a one-to-one and onto mapping between the set of certificates for x and the set of satisfying assignments for $R(x)$.
- So the number of satisfying truth assignments for $R(x)$ equals that of $M(x)$'s accepting computation paths.
- This kind of reduction is called **parsimonious**.
- We will loosen the requirement for parsimonious reduction: It runs in deterministic polynomial time.