

Reductions and Completeness

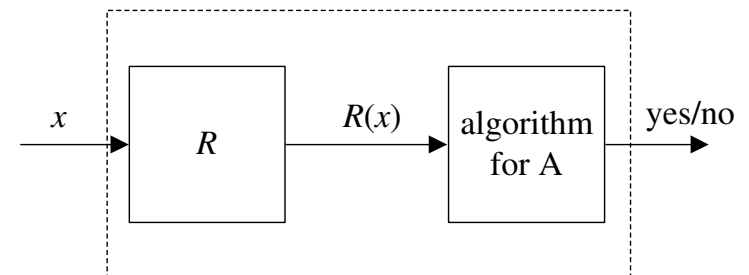
Degrees of Difficulty (concluded)

- Problem A is at least as hard as problem B if B reduces to A.
- This makes intuitive sense: If A is able to solve your problem B, then A must be at least as powerful.

Degrees of Difficulty

- When is a problem more difficult than another?
- **B reduces to A** if there is a transformation R which for every input x of B yields an equivalent input $R(x)$ of A.
 - The answer to x for B is the same as the answer to $R(x)$ for A.
 - There must be restrictions on the complexity of computing R .
 - Otherwise, $R(x)$ might as well solve B.

Reduction



Solving problem B by calling the algorithm for problem *once* and *without* further processing its answer.

Comments^a

- Suppose B reduces to A via a transformation R .
- The input x is an instance of B .
- The output $R(x)$ is an instance of A .
- $R(x)$ may not span all possible instances of A .
- So some instances of A may never appear in the reduction.

^aContributed by Mr. Ming-Feng Tsai (D92922003) on October 29, 2003.

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 between Languages

- Language L_1 is **reducible to** L_2 if there is a function R computable by a deterministic TM in polynomial time.
- Furthermore, for all inputs x , $x \in L_1$ if and only if $R(x) \in L_2$.
- R is said to be a **reduction** from L_1 to L_2 .
- If R is a reduction from L_1 to L_2 , then $R(x) \in L_2$ is a legitimate algorithm for $x \in L_1$.

Reduction of HAMILTONIAN PATH to SAT

- Given a graph G , we shall construct a CNF $R(G)$ such that $R(G)$ is satisfiable if and only if 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 .



$$x_{12} = x_{21} = x_{34} = x_{45} = x_{53} = x_{69} = x_{76} = x_{88} = x_{97} = 1.$$

The Proof

- $R(G)$ contains $O(n^3)$ clauses.
- $R(G)$ can be computed efficiently (simple exercise).
- Suppose $T \models R(G)$.
- From Clauses of 1 and 2, for each node j there is a unique position i such that $T \models x_{ij}$.
- From Clauses of 3 and 4, 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 Clauses of $R(G)$ and Their Intended Meanings

1. Each node j must appear in the path.
 - $x_{1j} \vee x_{2j} \vee \dots \vee x_{nj}$ 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.
 - $x_{i1} \vee x_{i2} \vee \dots \vee x_{in}$ 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 (concluded)

- Clauses of 5 furthermore guarantees 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

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

satisfies all clauses of $R(G)$.

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 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**.

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 $(i, j) \in E$.

Reduction of CIRCUIT SAT to SAT

- Given a circuit C , we shall construct a boolean expression $R(C)$ such that $R(C)$ is satisfiable if and only if C is satisfiable.
 - $R(C)$ will turn out to be a CNF.
- The variables of $R(C)$ are those of C plus g for each gate g of C .
- Each gate of C will be turned into equivalent clauses of $R(C)$.
- Recall that clauses are \wedge -ed together.

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: $g \Leftrightarrow \neg h$.

Composition of Reductions

Proposition 10 *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 .*

- Clearly $x \in L_1$ if and only if $R_{23}(R_{12}(x)) \in L_3$.
- It is also clear that $R_{12} \circ R_{23}$ can be computed in polynomial time.

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.

Completeness^a


- As reducibility is transitive, problems can be ordered with respect to their difficulty.
- Is there a *maximal* element?
- It is not altogether obvious that there should be a maximal element.
- Many infinite structures (such as integers and reals) 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).

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.

Illustration of Completeness and Hardness



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.

Closedness under Reduction

- 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, and EXP are all closed under reductions.

Complete Problems and Complexity Classes

Proposition 11 *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 a complete problem for \mathcal{C} . 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 $\mathcal{C} = \mathcal{C}'$.

Two Immediate Corollaries

Proposition 11 implies that

- $P = NP$ if and only if an NP-complete problem is in P.
- $L = P$ if and only if a P-complete problem is in L.

The Proof (concluded)

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

Complete Problems and Complexity Classes

Proposition 12 *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 (recall that it is an upper 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 (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

- M halts after at most $|x|^k - 2$ steps.
 - The string length hence never exceeds $|x|^k$.
- Assume a large enough k to make it true for $|x| \geq 2$.
- Pad the table with \square 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 (concluded)

- If 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 j .

Comments

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

$$\triangleright \overbrace{0_s 10001 \square \square \dots \square}^{|x|^k}$$

- A typical row may be

$$\triangleright \overbrace{10100_q 01110100 \square \square \dots \square}^{|x|^k}$$

- The last rows must look like $\triangleright \dots \overbrace{\text{“yes”} \dots \square}^{|x|^k}$

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 \square , a \triangleright , and a \sqcup , respectively.
 - Three out of four of T 's borders are known.

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

A P-Complete Problem

Theorem 13 (Ladner (1975)) CIRCUIT VALUE is *P-complete*.

- It is easy to see that CIRCUIT VALUE $\in P$.
- For any $L \in P$, we will construct a reduction R from L to CIRCUIT 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)

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

$$\begin{array}{|c|c|c|} \hline T_{i-1,j-1} & T_{i-1,j} & T_{i-1,j+1} \\ \hline & T_{ij} & \\ \hline \end{array}$$

- 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

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

(state assignment in circuit design).

The Proof (continued)

- Let 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 $S_{ij\ell}$, where

$$\begin{aligned} 0 \leq i \leq n^k - 1, \\ 0 \leq j \leq n^k - 1, \\ 1 \leq \ell \leq m. \end{aligned}$$

The Proof (continued)

- These F_i 's depend on only M 's specification, not on x .
- Their sizes are fixed.
- These boolean functions can be turned into boolean circuits.
- Compose these m circuits in parallel to obtain circuit C with $3m$ -bit inputs and m -bit outputs.
 - Schematically, $C(T_{i-1,j-1}, T_{i-1,j}, T_{i-1,j+1}) = T_{ij}$.
 - C is like an ASIC (application-specific IC) chip.

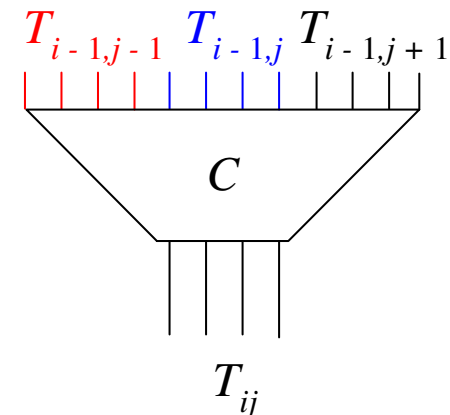
The Proof (continued)

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

$T_{i-1,j-1}$:	$S_{i-1,j-1,1}$	$S_{i-1,j-1,2}$	\cdots	$S_{i-1,j-1,m}$
$T_{i-1,j}$:	$S_{i-1,j,1}$	$S_{i-1,j,2}$	\cdots	$S_{i-1,j,m}$
$T_{i-1,j+1}$:	$S_{i-1,j+1,1}$	$S_{i-1,j+1,2}$	\cdots	$S_{i-1,j+1,m}$
- So there are m boolean functions F_1, F_2, \dots, F_m with $3m$ inputs each such that for all $i, j > 0$,

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

Circuit C



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” (coded as 1/0) appear at position $(|x|^k - 1, 1)$.

A Corollary

The construction in the above proof shows the following.

Corollary 14 *If $L \in TIME(T(n))$, then a circuit with $O(T^2(n))$ gates can decide if $x \in L$ for $|x| = n$.*

The Computation Tableau and $R(x)$



Cook's Theorem: the First NP-Complete Problem

Theorem 15 (Cook (1971)) *SAT is NP-complete.*


- $SAT \in NP$ (p. 49).
- CIRCUIT SAT reduces to SAT (p. 121).
- Now we only need to show that all languages in NP can be reduced to CIRCUIT SAT.

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 fixed, the computation is *deterministic*.


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



The Proof (continued)

- Each choice of B results in a deterministic polynomial-time computation, hence a table like the one on p. 147.
- 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 Proof (concluded)

- The overall circuit $R(x)$ (on p. 152) 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., $x \in L$.

NP-Complete Problems

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 $k \geq 1$.

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

3SAT

- k -SAT, where $k \in \mathbb{Z}^+$, is the special case of SAT.
- The formula is in CNF and all clauses have *exactly* k literals (repetition of literals is allowed).
- For example,

$$(x_1 \vee x_2 \vee \neg x_3) \wedge (x_1 \vee x_1 \vee \neg x_2) \wedge (x_1 \vee \neg x_2 \vee \neg x_3).$$

3SAT Is NP-Complete

- Recall Cook's Theorem (p. 149) and the reduction of CIRCUIT SAT to SAT (p. 121).
- The resulting CNF has at most 3 literals for each clause.
 - This shows that 3SAT where each clause has at most 3 literals is NP-complete.
- Finally, duplicate one literal once or twice to make it a 3SAT formula.

The Proof (concluded)

- Add $(\neg x_1 \vee x_2) \wedge (\neg x_2 \vee x_3) \wedge \cdots \wedge (\neg x_k \vee x_1)$ to the expression.
 - This is logically equivalent to $x_1 \Rightarrow x_2 \Rightarrow \cdots \Rightarrow x_k \Rightarrow x_1$.
 - Note that each clause above has fewer than 3 literals.
- The resulting equivalent expression satisfies the condition for x .

Another Variant of 3SAT

Proposition 16 *3SAT is NP-complete for expressions in which each variable is restricted to appear at most three times, and each literal at most twice. (3SAT here requires only that each clause has at most 3 literals.)*

- Consider a general 3SAT expression in which x appears k times.
- Replace the first occurrence of x by x_1 , the second by x_2 , and so on, where x_1, x_2, \dots, x_k are k new variables.

NAESAT

- The NAESAT (for “not-all-equal” SAT) is like 3SAT.
- But we require additionally that there be a satisfying truth assignment under which no clauses have the three literals equal in truth value.
 - Each clause must have one literal assigned true and one literal assigned false.

NAESAT Is NP-Complete^a

- Recall the reduction of CIRCUIT SAT to SAT on p. 121.
- It produced a CNF ϕ in which each clause has at most 3 literals.
- Add the same variable z to all clauses with fewer than 3 literals to make it a 3SAT formula.
- Goal: The new formula $\phi(z)$ is NAE-satisfiable if and only if the original circuit is satisfiable.

^aKarp (1972).

The Proof (concluded)

- Suppose there is a truth assignment that satisfies the circuit.
 - Then there is a truth assignment T that satisfies every clause of ϕ .
 - Extend T by adding $T(z) = \text{false}$ to obtain T' .
 - T' satisfies $\phi(z)$.
 - So in no clauses are all three literals false under T' .
 - Under T' , in no clauses are all three literals true.
 - * Review the detailed construction on p. 122 and p. 123.

The Proof (continued)

- Suppose T NAE-satisfies $\phi(z)$.
 - \bar{T} also NAE-satisfies $\phi(z)$.
 - Under T or \bar{T} , variable z takes the value false.
 - This truth assignment must still satisfy all clauses of ϕ .
 - So it satisfies the original circuit.

Undirected Graphs

- An **undirected graph** $G = (V, E)$ has a finite set of nodes, V , and a set of *undirected* edges, E .
- It is like a directed graph except that the edges have no directions and there are no self-loops.
- We use $[i, j]$ to denote the fact that there is an edge between node i and node j .

Independent Sets

- Let $G = (V, E)$ be an undirected graph.
- $I \subseteq V$.
- I is **independent** if whenever $i, j \in I$, there is no edge between i and j .
- The INDEPENDENT SET problem: Given an undirected graph and a goal K , is there an independent set of size K ?
 - Many applications.


Reduction from 3SAT to INDEPENDENT SET

- Let ϕ be an instance of 3SAT with m clauses.
- We will construct graph G (with constraints as said) with $K = m$ such that ϕ is satisfiable if and only if G has an independent set of size K .
- There is a triangle for each clause with the literals as the nodes.
- Add additional edges between x and $\neg x$ for every variable x .

INDEPENDENT SET Is NP-Complete

- This problem is in NP: Guess a set of nodes and verify that it is independent and meets the count.
- If a graph contains a triangle, any independent set can contain at most one node of the triangle.
- We consider graphs whose nodes can be partitioned in m disjoint triangles.
 - If the special case is hard, the original problem must be at least as hard.

A Sample Construction



$$(x_1 \vee x_2 \vee x_3) \wedge (\neg x_1 \vee \neg x_2 \vee \neg x_3) \wedge (\neg x_1 \vee x_2 \vee x_3).$$

The Proof (continued)

- Suppose G has an independent set I of size $K = m$.
 - An independent set can contain at most m nodes, one from each triangle.
 - An independent set of size m exists if and only if it contains exactly one node from each triangle.
 - Truth assignment T assigns true to those literals in I .
 - T is consistent because contradictory literals are connected by an edge, hence not both in I .
 - T satisfies ϕ because it has a node from every triangle, thus satisfying every clause.

The Proof (concluded)

- Suppose a satisfying truth assignment T exists for ϕ .
 - Collect one node from each triangle whose literal is true under T .
 - The choice is arbitrary if there is more than one true literal.
 - This set of m nodes must be independent by construction.
 - * Literals x and $\neg x$ cannot be both assigned true.