

Recall that $TQBF$ is the set of all true, fully quantified Boolean formulas. We have already demonstrated that $TQBF \in PSPACE$. It remains to prove that any language $L \in PSPACE$ is polynomial-time-reducible to $TQBF$. My proof here is identical to the proof in the Sipser textbook, and to many proofs all over the web. My version is just a bit more explicit and verbose, especially about which variables are bound and free.

1. INFRASTRUCTURE

Let M be a deterministic Turing machine that decides L in polynomial space. There exists some constant k such that, for all n and all strings w of length n , M uses at most n^k space to decide whether $w \in L$. Let Q be the state set of M and Γ the tape alphabet of M . A configuration of M is a string $a_0 a_1 \cdots a_{n^k}$ over $Q \cup \Gamma$ in the usual way. When M is given input $w = w_1 \cdots w_n$, its starting configuration is

$$c_{\text{start}} = q_{\text{start}} w_1 \cdots w_n \sqcup \cdots \sqcup.$$

Without loss of generality, we assume that M , before halting, empties its tape and parks its tape head at the far left, so that its unique accepting configuration is

$$c_{\text{accept}} = q_{\text{accept}} \sqcup \cdots \sqcup.$$

Notice that $T = |Q \cup \Gamma|^{n^k+1}$ is an upper bound on the number of configurations of M , and hence on the time that M requires to accept or reject w .

We need to establish a preliminary concept. For any symbol c , we can form a set

$$\{c_{i,s} : 0 \leq i \leq n^k, s \in Q \cup \Gamma\}$$

of $|Q \cup \Gamma| \cdot (n^k + 1)$ variables. Any given configuration of M corresponds to a unique assignment of truth values to these variables: Namely, $c_{i,s}$ is true in the assignment if and only if cell i of the configuration contains symbol s . (On the other hand, there are some assignments of truth values that do not correspond to any configuration of M — for example, an assignment in which $c_{3,a}$ and $c_{3,b}$ are both true, or an assignment in which $c_{4,q}$ and $c_{6,q}$ are both true, for $q \in Q$.)

For all $t \geq 0$, our reduction will recursively construct a “nearly” fully quantified Boolean formula $\phi_{c,d,t}$, involving c -variables, d -variables, and perhaps many other variables. This formula will be free on the c - and d -variables but quantified on all of its other variables. The crucial property of $\phi_{c,d,t}$ will be that, if we assign truth values to the c -variables based on one configuration of M , and truth values to the d -variables based on another configuration of M , then $\phi_{c,d,t}$ is true if and only if M can go from the first configuration to the second in t or fewer steps. After recursively constructing this formula, the reduction will plug in the truth value assignments for $c = c_{\text{start}}$ and $d = c_{\text{accept}}$, and $t = T$. The resulting Boolean formula will be fully quantified, and will be true if and only if M accepts w .

2. BASE CASE: $t = 0$

Given two symbols c and d , we have two variable sets $\{c_{i,s}\}$ and $\{d_{i,s}\}$, and we can write the Boolean formula

$$\phi_{c,d,0} = \bigwedge_{i=0}^{n^k} \bigwedge_{s \in Q \cup \Gamma} (c_{i,s} \wedge d_{i,s}) \vee (\overline{c_{i,s}} \wedge \overline{d_{i,s}}).$$

Because the formula is unquantified, its truth or falsity depends on an assignment of truth values to its variables; under some assignments it is true, and under other assignments it is false. What we can say is this: If we assign truth values to the c -variables based on a certain configuration of M , and we assign truth values to the d -variables based on another configuration of M , then the formula is true if and only if the two configurations are identical — that is, iff M can go from the first configuration to the second in zero steps.

3. BASE CASE: $t = 1$

We now wish to write a formula that is true iff M can go from one configuration to another in one step. This formula will incorporate information about M 's transition function $\delta : Q \times \Gamma \rightarrow Q \times \Gamma \times \{\mathcal{L}, \mathcal{R}\}$. Namely, if $\delta(q, a) = (r, b, \mathcal{L})$, then consider the formula

$$\begin{aligned} \phi_{c,d,1}^{j,(q,a,r,b,\mathcal{L})} &= c_{j,q} \wedge c_{j+1,a} \wedge d_{j-1,r} \wedge d_{j+1,b} \wedge \left(\bigvee_{e \in \Gamma} c_{j-1,e} \wedge d_{j,e} \right) \\ &\quad \wedge \left(\bigwedge_{i \in \{0, \dots, j-2, j+2, \dots, n^k\}} \bigwedge_{s \in Q \cup \Gamma} (c_{i,s} \wedge d_{i,s}) \vee (\overline{c_{i,s}} \wedge \overline{d_{i,s}}) \right). \end{aligned}$$

If we assign truth values to the c - and d -variables based on two configurations of M , then $\phi_{c,d,1}^{j,(q,a,r,b,\mathcal{L})}$ is true if and only if the first configuration has its state marker in cell j and M moves from the first configuration to the second using the transition $\delta(q, a) = (r, b, \mathcal{L})$. The first line of the formula expresses the change of state, the movement of the tape head, and the modification of the tape around the tape head; the second line of the formula expresses the fact that the rest of the tape remains unchanged. One can invent a similar formula to express transitions of the form $\delta(q, a) = (r, b, \mathcal{R})$. Then let

$$\phi_{c,d,1} = \phi_{c,d,0} \vee \bigvee_j \bigvee_{(q,a,r,b,\mathcal{D}) \in \delta} \phi_{c,d,1}^{j,(q,a,r,b,\mathcal{D})}.$$

The variables in $\phi_{c,d,1}$ are all free (unquantified). If we assign truth values to the c - and d -variables based on two configurations of M , then $\phi_{c,d,1}$ is true if and only if M can transition from the first configuration to the second in one or fewer steps. This is the real base case of our recursion.

4. INDUCTIVE CASE: FIRST TRY

Now suppose that $t > 1$. Consider the formula

$$\phi_{c,d,t} = \exists m (\phi_{c,m,t/2} \wedge \phi_{m,d,t/2}),$$

where m is a symbol distinct from c and d , and “ $\exists m$ ” is shorthand for the concatenation of the $|Q \cup \Gamma| \cdot (n^k + 1)$ quantifiers $\exists m_{i,s}$, for $i = 0, \dots, n^k$ and $s \in Q \cup \Gamma$. This formula is free on the c - and d -variables but quantified on all of its other variables. If we assign truth values to the c -variables based on one configuration of M , and truth values to the d -variables based on another configuration of M , then we obtain a fully quantified Boolean formula that is true if and only if M can go from the first configuration to the second in t or fewer steps. If we use c_{start} and c_{accept} as our two configurations, and $t = T$ for our time bound, then $\phi_{c,d,t}$ is true if and only if M accepts w . In other words, this algorithm is a reduction from L to TQBF:

- (1) For some symbols c and d , recursively compute $\phi_{c,d,T}$.

- (2) Replace the d -variables with their truth values determined from c_{accept} .
- (3) For the given w , replace the c -variables with their truth values determined from the starting configuration c_{start} of M on input w .

The problem with this reduction is that the resulting formula is too large. On each step of the recursion, we halve t , but we double the number of ϕ -formulas involved. In the end, the number of subformulas of the form $\phi_{c,d,1}$ will be proportional to T , and hence exponential in n . So the space required will be exponential in n , and the time required must be (at least) exponential in n . We wanted a polynomial-time reduction.

5. INDUCTIVE CASE: SECOND TRY

To improve our answer, we will take one step back and then two steps forward. Let c' and d' be new symbols. First we rewrite the $\phi_{c,d,t}$ above as

$$\exists m \forall c' \forall d' \left(((c' = c \wedge d' = m) \rightarrow \phi_{c',d',t/2}) \wedge ((c' = m \wedge d' = d) \rightarrow \phi_{c',d',t/2}) \right).$$

The notion of “=” being used here can be implemented using the $t = 0$ base case:

$$\exists m \forall c' \forall d' \left(((\phi_{c',c,0} \wedge \phi_{d',m,0}) \rightarrow \phi_{c',d',t/2}) \wedge ((\phi_{c',m,0} \wedge \phi_{d',d,0}) \rightarrow \phi_{c',d',t/2}) \right).$$

Thus far, our new approach is even worse than the old approach. It still has the double recursion, and it's adding on even more quantifiers and other verbiage. But notice that, after the quantifiers, the expression has the form

$$\begin{aligned} (A \rightarrow C) \wedge (B \rightarrow C) &\equiv (\neg A \vee C) \wedge (\neg B \vee C) \\ &\equiv (\neg A \wedge \neg B) \vee C \\ &\equiv \neg(A \vee B) \vee C \\ &\equiv (A \vee B) \rightarrow C. \end{aligned}$$

So, in the end, we can write $\phi_{c,d,t}$ as

$$\exists m \forall c' \forall d' \left((\phi_{c',c,0} \wedge \phi_{d',m,0}) \vee (\phi_{c',m,0} \wedge \phi_{d',d,0}) \right) \rightarrow \phi_{c',d',t/2}.$$

Keep in mind that each quantifier here is actually shorthand for $|Q \cup \Gamma| \cdot (n^k + 1)$ quantifiers. By increasing the number of quantifiers, but not too badly, we have reduced double recursion to single recursion. That's quite clever.

How long is this formula? Well, before the \rightarrow symbol, there are three sets of $\mathcal{O}(n^k)$ quantifiers each, and there are four formulas of the form $\phi_{c,d,0}$, each of which is $\mathcal{O}(n^k)$. So each level of recursion adds $\mathcal{O}(n^k)$ symbols to the formula. There are

$$\log_2 T = \log_2 |Q \cup \Gamma|^{n^k+1} = (n^k + 1) \log_2 |Q \cup \Gamma| \in \mathcal{O}(n^k)$$

levels of recursion. So the total length of the formula is $\mathcal{O}(n^{2k})$. In addition, computing this polynomial-length formula requires only polynomial time, because at each level of recursion the algorithm simply selects three new symbols m, c', d' , writes some quantifiers and $\phi_{c,d,0}$ -terms based on those symbols, and divides t by two. This completes the proof.

6. SOME DETAILS

The divisions by two are a bit messy, unless T is a power of two. So make T be the smallest power of two that is at least $|Q \cup \Gamma|^{n^k+1}$. This tactic comes up in many divide-and-conquer algorithms.

Writing a symbol with subscripts requires more than one tape cell. The size of the subscripts is logarithmic in n^k , so this detail can't ruin the fact that the space and time are polynomial in n .