1. Suppose you are given a magic black box that somehow answers the following decision problem *in polynomial time*:

   - INPUT: A directed graph G and a positive integer L. (The edges of G are not weighted, and G is not necessarily a dag.)
   - OUTPUT: TRUE if G contains a (simple) path of length at least L, and FALSE otherwise.

   (a) Using this black box as a subroutine, describe an algorithm that solves the following related optimization problem *in polynomial time*:

   - INPUT: A directed graph G.
   - OUTPUT: The length of the longest path in G.

Solution: Suppose LONGPATH(V, E, L) returns TRUE if the graph (V, E) contains a path of length at least L, and FALSE otherwise. Then the following algorithm returns the length of the longest path in G:

```
MAXPATHLENGTH(V, E):
  for L ← 1 to V
      if LONGPATH(V, E, L) = FALSE
          return L − 1
```

A graph with n vertices cannot contain a path with length greater than $n-1$, so this algorithm must return a value. If G contains a path of length at least L, then it also contains paths of every length less than L, so the algorithm is correct.

MAXPATHLENGTH clearly runs in polynomial time. Specifically, if the running time of LONGPATH(V, E, k) is $O(V^c)$, then MAXPATHLENGTH(V, E) runs in $O(V^{c+1})$ time.

Yes, using binary search instead of linear search would yield a faster algorithm that invokes LONGPATH only $O(\log V)$ times instead of $O(V)$. But we already have polynomial time, so why bother? ■

(b) Using this black box as a subroutine, describe an algorithm that solves the following related search problem *in polynomial time*:

- INPUT: A directed graph G.
- OUTPUT: The longest path in G

Solution (delete edges): I'll use the algorithm MAXPATHLENGTH as a black box instead. Let $G - e$ denote the graph obtained from G by deleting edge e, and let $G - v$ denote the graph obtained from G by deleting vertex v

```
LONGESTPATH(G):
  L ← MAXPATHLENGTH(G)
  for each edge e ∈ E
      if MAXPATHLENGTH(G − e) = L
          G ← G − e
  ⟨⟨At this point G is just a path plus isolated vertices.⟩⟩
  for each vertex v ∈ V
      if v has no incoming or outgoing edges in E
          G ← G − v
  ⟨⟨At this point G is just a path!⟩⟩
  return G
```

Correctness of this algorithm follows inductively from the following claims:

Claim 1. *MAXPATHLENGTH$(G - e) = L$ if and only if some longest path G excludes e.*

Proof: Every path in $G - e$ is also a path in G, so MAXPATHLENGTH$(G - e) \le L$.

Suppose G has a longest path π that excludes e. Then π is also a path of length L in $G - e$. It follows that MAXPATHLENGTH$(G - e) \ge L$, and therefore MAXPATHLENGTH$(G - e) = L$.

On the other hand, suppose MAXPATHLENGTH$(G - e) = L$. Then $G - e$ contains a path π with length L. It follows that G contains a path π of length L (that is, a longest path) that excludes e. □

The algorithm clearly runs in polynomial time. ■

2. An ***independent set*** in a graph G is a subset S of the vertices of G, such that no two vertices in S are connected by an edge in G. Suppose you are given a magic black box that somehow answers the following decision problem *in polynomial time*:

   - INPUT: An undirected graph G and an integer k.
   - OUTPUT: TRUE if G has an independent set of size k, and FALSE otherwise.

   (a) Using this black box as a subroutine, describe algorithms that solves the following related optimization problem *in polynomial time*:

   - INPUT: An undirected graph G.
   - OUTPUT: The size of the largest independent set in G.

Solution: Suppose INDSET(V, E, k) returns TRUE if the graph (V, E) has an independent set of size k, and FALSE otherwise. Then the following algorithm returns the size of the largest independent set in G:

```
MAXINDSETSIZE(V, E):
  for k ← 1 to V
      if INDSET(V, E, k + 1) = FALSE
          return k
```

A graph with n vertices cannot have an independent set of size larger than n, so this algorithm must return a value. If G has an independent set of size k, then it also has an independent set of size $k - 1$, so the algorithm is correct.

The algorithm clearly runs in polynomial time. Specifically, if INDSET(V, E, k) runs in $O(V^c)$ time, then MAXINDSETSIZE(V, E) runs in $O(V^{c+1})$ time.

Yes, we could have used binary search instead of linear search. Whatever. ■

(b) Using this black box as a subroutine, describe algorithms that solves the following related search problem *in polynomial time*:

- INPUT: An undirected graph G.
- OUTPUT: An independent set in G of maximum size.

Solution (delete vertices): I'll use the algorithm MAXINDSETSIZE(V, E) from part (a) as a black box instead. Let $G - v$ denote the graph obtained from G by deleting vertex v, and let $G - N(v)$ denote the graph obtained from G by deleting v and all neighbors of v.

```
MaxIndSet(G):
  S ← ∅
  k ← MaxIndSetSize(G)
  for all vertices v of G
      if MaxIndSetSize(G − v) = k
          G ← G − v
      else
          G ← G − N(v)
          add v to S
          k ← k − 1
  return S
```

Correctness of this algorithm follows inductively from the following claims:

Claim 2. *MAXINDSETSIZE$(G - v) = k$ if and only if G has an independent set of size k that excludes v.*

Proof: Every independent set in $G - v$ is also an independent set in G; it follows that MAXINDSETSIZE$(G - v) \leq k$.

Suppose G has an independent set S of size k that does excludes v. Then S is also an independent set of size k in $G - v$, so MAXINDSETSIZE$(G - v)$ is at least k, and therefore equal to k.

On the other hand, suppose $G - v$ has an independent set S of size k. Then S is also a maximum independent set of G (because $|S| = k$) that excludes v. □

The algorithm clearly runs in polynomial time. ■

Solution (add edges): I'll use the algorithm MaxIndSetSize(V, E) from part (a) as a black box instead. Let $G + uv$ denote the graph obtained from G by adding edge uv.

```
MaxIndSet(G):
  k ← MaxIndSetSize(G)
  if k = 1
      return any vertex
  for all vertices u
      for all vertices v
          if u ≠ v and uv is not an edge
              if MaxIndSetSize(G + uv) = k
                  G ← G + uv
  S ← ∅
  for all vertices v
      if deg(v) < V − 1
          add v to S
  return S
```

The algorithm adds every edge it can without changing the maximum independent set size. Let G' denote the final graph. Any independent set in G' is also an independent set in the original input graph G. Moreover, the *largest* independent set in G' is also a largest independent set in G. Thus, to prove the algorithm correct, we need to prove the following claims about the final graph G':

Claim 3. *The maximum independent set in G' is unique.*

Proof: Suppose the final graph G' has more than two maximum independent sets A and B. Pick any vertex $u \in A \setminus B$ and any other vertex $v \in A$. The set B is still an independent set in the graph $G' + uv$. Thus, when the algorithm considered edge uv, it would have added uv to the graph, contradicting the assumption that A is an independent set. □

Claim 4. *Suppose $k > 1$. The unique maximum independent set of G' contains vertex v if and only if* $\deg(v) < V - 1$.

Proof: Let S be the unique maximum independent set of G', and let v be any vertex of G. If $v \in S$, then v has degree at most $V - k < V - 1$, because v is disconnected from every other vertex in S.

On the other hand, suppose $\deg(v) < V - 1$ but $v \notin S$. Then there must be at least vertex u such that uv is not an edge in G'. Because $v \notin S$, the set S is still an independent set in $G' + uv$. Thus, when the algorithm considered edge uv, it would have added uv to the graph, and we have a contradiction. □

The algorithm clearly runs in polynomial time. ■

To think about later:

3. Suppose you are given a magic black box that somehow answers the following decision problem *in polynomial time*:

   - INPUT: An undirected graph G and an integer k.
   - OUTPUT: TRUE if G has a proper coloring with k colors, and FALSE otherwise.

   Using this black box as a subroutine, describe an algorithm that solves the following related search problem *in polynomial time*:

   - INPUT: An undirected graph G.
   - OUTPUT: A valid coloring of G using the minimum possible number of colors.

   *[Hint: You can use the magic box more than once. The input to the magic box is a graph and **only** a graph, meaning **only** vertices and edges.]*

Solution: First we describe an algorithm to compute the minimum number of colors in any valid coloring.

```
ChromaticNumber(G):
  for k ← V down to 1
      if Colorable(G, k − 1) = False
          return k
```

Given a graph $G = (V, E)$ with n vertices $v_1, v_2, \ldots, v_n$, the following algorithm computes an array $color[1..n]$ describing a valid coloring of G with the minimum number of colors.

```
Coloring(G):
  k ← ChromaticNumber(G)
  ⟨⟨—— add a disjoint clique of size k ——⟩⟩
  H ← G
  for c ← 1 to k
      add vertex z_c to G
      for i ← 1 to c − 1
          add edge z_i z_c to H
  ⟨⟨—— for each vertex, try each color ——⟩⟩
  for i ← 1 to n
      for c ← 1 to k
          add edge v_i z_c to H
      for c ← 1 to k
          remove edge v_i z_c from H
          if Colorable(H, k) = True
              color[i] ← c
              break inner loop
          add edge v_i z_c to H
  return color[1..n]
```

In any k-coloring of H, the new vertices $z_1, \ldots, z_k$ must have k distinct colors, because every pair of those vertices is connected. We assign $color[i] \leftarrow c$ to indicate that there is a k-coloring of H in which v_i has the same color as z_c. When the algorithm terminates, $color[1..n]$ describes a valid k-coloring of G.

To prove that the algorithm is correct, we must prove that for all i, when the ith iteration of the outer loop ends, G has a valid k-coloring that is consistent with the partial coloring $color[1..i]$. Fix an integer i. The inductive hypothesis implies that when the ith iteration of the outer loop *begins*, G has a k-coloring consistent with the first $i-1$ assigned colors. (The base case $i = 0$ is trivial.) If we connect v_i to every new vertices except z_c, then v_i must have the same color as z_c in any valid k-coloring. Thus, the call to COLORABLE inside the inner loop returns TRUE if and only if H has a k-coloring in which v_i has the same color as z_c (and the previous $i-1$ vertices are also colored). So COLORABLE must return TRUE during the second inner loop, which completes the inductive proof.

This algorithm makes $O(kn) = O(n^2)$ calls to COLORABLE, and therefore runs in polynomial time. ■

4. Suppose you are given a magic black box that somehow answers the following decision problem in *polynomial time*:

   - INPUT: A boolean circuit K with n inputs and one output.
   - OUTPUT: TRUE if there are input values $x_1, x_2, \ldots, x_n \in \{\text{TRUE}, \text{FALSE}\}$ that make K output TRUE, and FALSE otherwise.

   Using this black box as a subroutine, describe an algorithm that solves the following related search problem *in polynomial time*:

   - INPUT: A boolean circuit K with n inputs and one output.
   - OUTPUT: Input values $x_1, x_2, \ldots, x_n \in \{\text{TRUE}, \text{FALSE}\}$ that make K output TRUE, or NONE if there are no such inputs.

   [Hint: You can use the magic box more than once.]

Solution: For any boolean circuit K with inputs $x_1, \ldots, x_n$, let $K \wedge x_i$ be the boolean circuit obtained from K by adding a new AND gate, with one input connected to the output of K and the other to the input x_i. Similarly, let $K \wedge \overline{x}_i$ be the boolean circuit obtained from K by adding a NOT gate, with input connected to x_i, and an AND gate, with one input connected to the output of K and the other to the NOT gate. For both of these circuits, the output of the new AND gate is the output of the circuit.

Suppose CIRCUITSAT(K) returns TRUE if K is satisfiable and FALSE otherwise. Then the following algorithm constructs a satisfying input assignment for K or correctly reports that no such assignment exists.

```
SATASSIGNMENT(K):
  if CIRCUITSAT(K) = FALSE
      return NONE
  for i ← 1 to n
      if CIRCUITSAT(K ∧ x_i)
          K ← K ∧ x_i
          A[i] ← TRUE
      else
          K ← K ∧ x̄_i
          A[i] ← FALSE
  return A[1..n]
```

The correctness of this algorithm follows by induction from the following observation:

Claim 5. *The circuit $K \wedge x_i$ is satisfiable if and only if K has a satisfying assignment where $x_i = \text{TRUE}$.*

Proof: First, if $K \wedge x_i$ has a satisfying assignment, then that input assignment must satisfy K and must have have $x_i = \text{TRUE}$, because otherwise the AND gate would output FALSE.

On the other hand, if K has a satisfying assignment where $x_i = \textsc{True}$, then that input assignment also satisfies $K \wedge x_i$, because that's how AND gates do. □

The algorithm runs in polynomial time. Specifically, suppose CIRCUITSAT(K) runs in $O(N^c)$ time, where N is the total number of vertices and edges in dag representing K. (The vertices consist of the inputs, the internal gates, and the output; the edges are the wires between those points.) Then SATASSIGNMENT(K) runs in time

$$O(N^c) + \sum_{i=1}^{n} O((N+5i)^c) \leq (N+1) \cdot O((6N)^c) = O(N^{c+1}),$$

which is polynomial in N. ■