<div align="center">

CS/ECE 374 A ✧ Fall 2023

♫ Homework 1 ♫

Due Tuesday, August 29, 2023 at 9pm Central Time

</div>

- **Submit your written solutions electronically to Gradescope as PDF files.** Submit a separate PDF file for each numbered problem. If you plan to typeset your solutions, you are welcome to use the LaTeX solution template on the course web site. If you must submit scanned handwritten solutions, please use a black pen on blank white paper and a high-quality scanner app (or an actual scanner).

- Groups of up to three people can submit joint solutions on Gradescope. *Exactly* one student in each group should upload the solution and indicate their other group members.

- **You may use any source at your disposal**—paper, electronic,[1] or human—but you *must* cite *every* source that you use,[2] you *must* write everything yourself in your own words, and you are responsible for any errors in the sources you use.[3] See the academic integrity policies on the course web site for more details.

- Written homework is normally due every Tuesday at 9pm. In addition, guided problem sets on PrairieLearn are normally due every **Monday** at 9pm; each student must do these individually. In particular, Guided Problem Set 1 is due Monday, August 28!

- Both guided problem sets and homework may be submitted up to 24 hours late for 50% partial credit, or for full credit with an approved extension. See the grading policies on the course web site for more details.

- Each homework will include at least one fully solved problem, similar to that week's assigned problems, together with the rubric we would use to grade this problem if it appeared in an actual homework or exam. These model solutions show our recommendations for structure, presentation, and level of detail in your homework solutions. (Obviously, the actual *content* of your solutions won't match the model solutions, because your problems are different!) Homeworks may also include additional practice problems.

- Standard grading rubrics for many problem types can be found on the course web page. For example, the problems in this week's homework will be graded using the standard induction rubric. (Weak induction makes the baby Jesus cry.)

<div align="center">

See the course web site for more information.

If you have any questions about these policies,
please don't hesitate to ask in class, in office hours, or on Piazza.

</div>

---

[1]Yes, including ChatGPT.
[2]Yes, including ChatGPT.
[3]Yes, including ChatGPT.

1. Consider the following recursively defined function:

$$
stutter(w) := \begin{cases} \varepsilon & \text{if } w = \varepsilon \\ aa \bullet stutter(x) & \text{if } w = ax \end{cases}
$$

For example, $stutter(\texttt{MISSISSIPPI}) = \texttt{MMIISSSSIISSSSIIPPPPII}$.

   (a) Prove that $|stutter(w)| = 2|w|$ for every string w.

   (b) Prove that $stutter(x \bullet y) = stutter(x) \bullet stutter(y)$ for all strings x and y.

   (c) *Practice only. Do not submit solutions.*
   The **reversal w^R** of a string w is defined recursively as follows:

$$
w^R := \begin{cases} \varepsilon & \text{if } w = \varepsilon \\ x^R \bullet a & \text{if } w = ax \end{cases}
$$

   For example, $\texttt{MISSIPPIPPI}^R = \texttt{IPPIPPISSIM}$.

   Prove that $stutter(w)^R = stutter(w^R)$ for every string w.

You may freely use any result proved in lecture, in lab, or in the lecture notes. Otherwise your proofs must be formal and self-contained. In particular, your proofs *must* invoke the formal recursive definitions of string length and concatenation (and for part (c), reversal).

2. For each positive integer n, we define two strings p_n and v_n, respectively called the nth
   Piṅgala string and the nth *Virahāṅka string*. Piṅgala strings are defined by the following
   recurrence:

$$p_n = \begin{cases} 1 & \text{if } n = 1 \\ 0 & \text{if } n = 2 \\ p_{n-2} \bullet p_{n-1} & \text{otherwise} \end{cases}$$

For example:

$$p_7 = \overbrace{\underbrace{10010}_{p_5}\underbrace{010}_{p_4}\underbrace{10010}_{p_5}}^{p_6}.$$

Virahāṅka strings are defined more indirectly as

$$v_n = \begin{cases} 1 & \text{if } n = 1 \\ \text{grow}(v_{n-1}) & \text{otherwise} \end{cases}$$

where the string function grow is defined as follows:

$$\text{grow}(w) = \begin{cases} \varepsilon & \text{if } w = \varepsilon \\ 0 \cdot \text{grow}(x) & \text{if } w = 1x \\ 10 \bullet \text{grow}(x) & \text{if } w = 0x \end{cases}$$

For example:

$$\text{grow}(01010010) = 10 \bullet 0 \bullet 10 \bullet 0 \bullet 10 \bullet 10 \bullet 0 \bullet 10 = 1001001010010$$

Finally, recall that the Fibonacci numbers are defined recursively as follows:

$$F_n = \begin{cases} 0 & \text{if } n = 0 \\ 1 & \text{if } n = 1 \\ F_{n-1} + F_{n-2} & \text{otherwise} \end{cases}$$

   (a) Prove that $|p_n| = F_n$ for all $n \geq 1$.

   (b) Prove that $\text{grow}(w \bullet z) = \text{grow}(w) \bullet \text{grow}(z)$ for all strings w and z.

   (c) Prove that $p_n = v_n$ for all $n \geq 1$. *[Hint: Careful!]*

   (d) ***Practice only. Do not submit solutions.***
   Prove that $|v_n| = F_n$ for all $n \geq 1$.

As in problem 1, you may freely use any result that proved in lecture, in lab, or in the
lecture notes. Otherwise your proofs must be formal and self-contained. In particular,
your proofs *must* invoke the formal recursive definitions of the strings p_n and v_n, the grow
function, and the Fibonacci numbers F_n.

*3. *Practice only. Do not submit solutions.*

For each non-negative integer n, we recursively define two binary trees P_n and V_n, called the nth *Piṅgala tree* and the nth *Virahāṅka tree,* respectively.

- P_0 and V_0 are empty trees, with no nodes.
- P_1 and V_1 each consist of a single node.
- For any integer $n \geq 2$, the tree P_n consists of a root with two subtrees; the left subtree is a copy of P_{n-1}, and the right subtree is a copy of P_{n-2}.
- For any integer $n \geq 2$, the tree L_n is obtained from L_{n-1} by attaching a new right child to every leaf and attaching a new left child to every node that has only a right child.

The following figure shows the recursive construction of these two trees when $n = 7$.

Recall that the Fibonacci numbers are defined recursively as follows:

$$F_n = \begin{cases} 0 & \text{if } n = 0 \\ 1 & \text{if } n = 1 \\ F_{n-1} + F_{n-2} & \text{otherwise} \end{cases}$$

(a) Prove that the tree P_n has exactly F_n leaves.

(b) Prove that the tree V_n has exactly F_n leaves.
   [Hint: You need to prove a stronger result.]

(c) Prove that the trees P_n and V_n are identical, for all $n \geq 0$.
   [Hint: The hardest part of this proof is developing the right language/notation.]

As in problem 1, you may freely use any result that proved in lecture, in lab, or in the lecture notes. Otherwise your proofs must be formal and self-contained. In particular, your proofs *must* invoke the formal recursive definitions of the trees P_n and V_n and the Fibonacci numbers F_n.

Solved Problems

3. For any string $w \in \{0, 1\}^*$, let *swap*(w) denote the string obtained from w by swapping the first and second symbols, the third and fourth symbols, and so on. For example:

$$swap(10\,11\,00\,01\,10\,1) = 01\,11\,00\,10\,01\,1.$$

The *swap* function can be formally defined as follows:

$$swap(w) := \begin{cases} \varepsilon & \text{if } w = \varepsilon \\ w & \text{if } w = 0 \text{ or } w = 1 \\ ba \bullet swap(x) & \text{if } w = abx \text{ for some } a, b \in \{0, 1\} \text{ and } x \in \{0, 1\}^* \end{cases}$$

(a) Prove that $|swap(w)| = |w|$ for every string w.

> **Solution:** Let w be an arbitrary string.
>
> Assume $|swap(x)| = |x|$ for every string x that is shorter than w.
>
> There are three cases to consider (mirroring the definition of *swap*):
>
> - If $w = \varepsilon$, then
>
> $$\begin{aligned} |swap(w)| &= |swap(\varepsilon)| & \text{because } w = \varepsilon \\ &= |\varepsilon| & \text{by definition of } swap \\ &= |w| & \text{because } w = \varepsilon \end{aligned}$$
>
> - If $w = 0$ or $w = 1$, then
>
> $$|swap(w)| = |w| \qquad \text{by definition of } swap$$
>
> - Finally, if $w = abx$ for some $a, b \in \{0, 1\}$ and $x \in \{0, 1\}^*$, then
>
> $$\begin{aligned} |swap(w)| &= |swap(abx)| & \text{because } w = abx \\ &= |ba \bullet swap(x)| & \text{by definition of } swap \\ &= |ba| + |swap(x)| & \text{because } |y \bullet z| = |y| + |z| \\ &= |ba| + |x| & \text{by the induction hypothesis} \\ &= 2 + |x| & \text{by definition of } |\cdot| \\ &= |ab| + |x| & \text{by definition of } |\cdot| \\ &= |ab \bullet x| & \text{because } |y \bullet z| = |y| + |z| \\ &= |abx| & \text{by definition of } \bullet \\ &= |w| & \text{because } w = abx \end{aligned}$$
>
> In all cases, we conclude that $|swap(w)| = |w|$. ∎

Rubric: 5 points: Standard induction rubric (scaled). This is more detail than necessary for full credit.

(b) Prove that $swap(swap(w)) = w$ for every string w.

Solution: Let w be an arbitrary string.

Assume $swap(swap(x)) = x$ for every string x that is shorter than w.

There are three cases to consider (mirroring the definition of $swap$):

- If $w = \varepsilon$, then

$$
\begin{aligned}
swap(swap(w)) &= swap(swap(\varepsilon)) && \text{because } w = \varepsilon \\
&= swap(\varepsilon) && \text{by definition of } swap \\
&= \varepsilon && \text{by definition of } swap \\
&= w && \text{because } w = \varepsilon
\end{aligned}
$$

- If $w = 0$ or $w = 1$, then

$$
\begin{aligned}
swap(swap(w)) &= swap(w) && \text{by definition of } swap \\
&= w && \text{by definition of } swap
\end{aligned}
$$

- Finally, if $w = abx$ for some $a, b \in \{0, 1\}$ and $x \in \{0, 1\}^*$, then

$$
\begin{aligned}
swap(swap(w)) &= swap(swap(abx)) && \text{because } w = abx \\
&= swap(ba \bullet swap(x)) && \text{by definition of } swap \\
&= swap(ba \bullet z) && \text{where } z = swap(x) \\
&= swap(baz) && \text{by definition of } \bullet \\
&= ab \bullet swap(z) && \text{by definition of } swap \\
&= ab \bullet swap(swap(x)) && \text{because } z = swap(x) \\
&= ab \bullet x && \text{by the induction hypothesis} \\
&= abx && \text{by definition of } \bullet \\
&= w && \text{because } w = abx
\end{aligned}
$$

In all cases, we conclude that $swap(swap(w)) = w$. ∎

Rubric: 5 points: Standard induction rubric (scaled). This is more detail than necessary for full credit.

4. The *reversal w^R* of a string w is defined recursively as follows:

$$w^R := \begin{cases} \varepsilon & \text{if } w = \varepsilon \\ x^R \bullet a & \text{if } w = a \cdot x \end{cases}$$

A *palindrome* is any string that is equal to its reversal, like AMANAPLANACANALPANAMA, RACECAR, POOP, I, and the empty string.

(a) Give a recursive definition of a palindrome over the alphabet Σ.

> **Solution:** A string $w \in \Sigma^*$ is a palindrome if and only if either
>
> - $w = \varepsilon$, or
> - $w = a$ for some symbol $a \in \Sigma$, or
> - $w = axa$ for some symbol $a \in \Sigma$ and some *palindrome $x \in \Sigma^*$*.
>
> ∎

> **Rubric:** 2 points = ½ for each base case + 1 for the recursive case. No credit for the rest of the problem unless this part is correct.

(b) Prove $w = w^R$ for every palindrome w (according to your recursive definition).
You may assume the following facts about all strings x, y, and z:

- Reversal reversal: $(x^R)^R = x$
- Concatenation reversal: $(x \bullet y)^R = y^R \bullet x^R$
- Right cancellation: If $x \bullet z = y \bullet z$, then $x = y$.

> **Solution:** Let w be an arbitrary palindrome.
> Assume that $x = x^R$ for every palindrome x such that $|x| < |w|$.
> There are three cases to consider (mirroring the definition of "palindrome"):
>
> - If $w = \varepsilon$, then $w^R = \varepsilon$ by definition, so $w = w^R$.
> - If $w = a$ for some symbol $a \in \Sigma$, then $w^R = a$ by definition, so $w = w^R$.
> - Finally, if $w = axa$ for some symbol $a \in \Sigma$ and some palindrome $x \in P$, then
>
> $$\begin{aligned} w^R &= (a \cdot x \bullet a)^R & \text{because } w = axa \\ &= (x \bullet a)^R \bullet a & \text{by definition of reversal} \\ &= a^R \bullet x^R \bullet a & \text{by concatenation reversal} \\ &= a \bullet x^R \bullet a & \text{by definition of reversal} \\ &= a \bullet x \bullet a & \text{by the inductive hypothesis} \\ &= w & \text{because } w = axa \end{aligned}$$
>
> In all three cases, we conclude that $w = w^R$. ∎

> **Rubric:** 4 points: standard induction rubric (scaled)

(c) Prove that every string w such that $w = w^R$ is a palindrome (according to your recursive definition).

Again, you may assume the following facts about all strings x, y, and z:

- Reversal reversal: $(x^R)^R = x$
- Concatenation reversal: $(x \bullet y)^R = y^R \bullet x^R$
- Right cancellation: If $x \bullet z = y \bullet z$, then $x = y$.

Solution: Let w be an arbitrary string such that $w = w^R$.

Assume that every string x such that $|x| < |w|$ and $x = x^R$ is a palindrome.

There are three cases to consider (mirroring the definition of "palindrome"):

- If $w = \varepsilon$, then w is a palindrome by definition.

- If $w = a$ for some symbol $a \in \Sigma$, then w is a palindrome by definition.

- Otherwise, we have $w = ax$ for some symbol a and some *non-empty* string x.
  The definition of reversal implies that $w^R = (ax)^R = x^R a$.
  Because x is non-empty, its reversal x^R is also non-empty.
  Thus, $x^R = by$ for some symbol b and some string y.
  It follows that $w^R = bya$, and therefore $w = (w^R)^R = (bya)^R = ay^Rb$.

  ⟨⟨At this point, we need to prove that $a = b$ and that y is a palindrome.⟩⟩

  Our assumption that $w = w^R$ implies that $bya = ay^Rb$.
  The recursive definition of string equality immediately implies $a = b$.

  Because $a = b$, we have $w = ay^Ra$ and $w^R = aya$.
  The recursive definition of string equality implies $y^Ra = ya$.
  Right cancellation implies $y^R = y$.
  The inductive hypothesis now implies that y is a palindrome.

  We conclude that w is a palindrome by definition.

In all three cases, we conclude that w is a palindrome. ∎

Rubric: 4 points: standard induction rubric (scaled).

5. Let $L \subseteq \{0, 1\}^*$ be the language defined recursively as follows:

   - The empty string ε is in L.
   - For any string $x \in L$, the strings $0101x$ and $1010x$ are also in L.
   - For all strings x and y such that $xy \in L$, the strings $x00y$ and $x11y$ are also in L.
     (In other words, inserting two consecutive 0s or two consecutive 1s anywhere in a string in L yields another string in L.)
   - These are the only strings in L.

Let EE denote the set of all strings $w \in \{0, 1\}^*$ such that $\#(0, w)$ and $\#(1, w)$ are both even.

In the following proofs, you may freely use any result proved in lecture, in lab, in the lecture notes, or earlier in your homework. Otherwise your proofs must be formal and self-contained; in particular, they must invoke the formal recursive definitions of $\#$ and L.

(a) Prove that $L \subseteq EE$.

> **Solution:** Let w be an arbitrary string in L. We need to prove that $\#(0, w)$ and $\#(1, w)$ are both even. Here I will prove only that $\#(0, w)$ is even; the proof that $\#(1, w)$ is even is symmetric.
>
> Assume for every string $x \in L$ such that $|x| < |w|$ that $\#(0, x)$ is even.
>
> There are several cases to consider, mirroring the definition of L.
>
> - Suppose $w = \varepsilon$. Then $\#(0, w) = 0$, and 0 is even.
> - Suppose $w = 0101x$ or $w = 1010x$ for some string $x \in L$. The definition of $\#$ (applied four times) implies $\#(0, w) = \#(0, x) + 2$. The inductive hypothesis implies $\#(0, x)$ is even. We conclude that $\#(0, w)$ is even.
> - Suppose $w = x00y$ for some strings x and y such that $xy \in L$. Then
>
> $$\begin{aligned} \#(0, w) &= \#(0, x00y) \\ &= \#(0, x) + \#(0, 00) + \#(0, y) \\ &= \#(0, x) + \#(0, y) + \#(0, 00) \\ &= \#(0, xy) + 2 \end{aligned}$$
>
> The induction hypothesis implies $\#(0, xy)$ is even. We conclude that $\#(0, w) = \#(0, xy) + 2$ is also even.
> - Finally, suppose $w = x11y$ for some strings x and y such that $xy \in L$. Then
>
> $$\begin{aligned} \#(0, w) &= \#(0, x11y) \\ &= \#(0, x) + \#(0, 11) + \#(0, y) \\ &= \#(0, x) + \#(0, y) \\ &= \#(0, xy) \end{aligned}$$
>
> The induction hypothesis implies $\#(0, w) = \#(0, xy)$ is even.

In all cases, we have shown that $\#(0, w)$ is even. Symmetric arguments imply that $\#(1, w)$ is even. We conclude that $w \in EE$.   ∎

Rubric: 5 points: standard induction rubric (scaled). Yes, this is enough detail for $\#(1, w)$. If explicit proofs are given for both $\#(0, w)$ and $\#(1, w)$, grade them independently, each for 2½ points.

(b) Prove that $EE \subseteq L$.

Solution: Let w be an arbitrary string in EE. We need to prove that $w \in L$.

Assume that for every string $x \in EE$ such that $|x| < |w|$, we have $x \in L$.

There are four (overlapping) cases to consider, depending on the first four symbols in w.

- Suppose $|w| < 4$. Then w must be one of the strings ε, 00, or 11; brute force inspection implies that every other string of length at most 3 (0, 1, 01, 10, 000, 001, 010, 011, 100, 101, 110, 111) has an odd number of 0s or an odd number of 1s (or both). All three strings ε, 00, and 11 are in L. In all other cases, we can assume that $|w| \geq 4$, so the "first four symbols of w" are well-defined.

- Suppose the first four symbols of w are 0000 or 0001 or 0010 or 0011 or 0100 or 1000 or 1001 or 1100. Then $w = x00y$ for some (possibly empty) strings x and y. Arguments in part (a) imply that $\#(0, xy) = \#(0, w) - 2$ and $\#(1, xy) = \#(1, w)$ are both even. Thus $xy \in EE$ by definition of EE. So the induction hypothesis implies $xy \in L$. We conclude that $w = x00y \in L$ by definition of L.

- Suppose the first four symbols of w are 0011 or 0110 or 0111 or 1011 or 1100 or 1101 or 1110 or 1111.) After swapping 0s and 1s, the argument in the previous case implies that $w \in L$.

- Finally, suppose the first four symbols of w are 0101 or 1010; in other words, suppose $w = 0101x$ or $w = 1010x$ for some (possibly empty) string x. Then $\#(0, x) = \#(0, w) - 2$ and $\#(1, x) = \#(1, w) - 2$ are both even, so $x \in EE$ by definition. The induction hypothesis implies $x \in L$. We conclude that $w \in L$ by definition of L.

Each of the 16 possible choices for the first four symbols of w is considered in at least one of the last three cases.

In all cases, we conclude that $w \in L$.   ∎

Rubric: 5 points: standard induction rubric (scaled). This is not the only correct proof. This is not the only correct way to express this particular case analysis.

CS/ECE 374 A ✦ Fall 2023
⌒ Homework 2 ⌒

Due Wednesday, September 6, 2023 at 9pm Central Time

(One day later than usual because of Labor Day)

- Starting with this homework, please start your solution to each *lettered subproblem* ((a), (b), (c), etc.) on a new page. Yes, even if the previous subproblem is only one line long. Please also remember to tell Gradescope which page(s) are relevant for which subproblems.

---

1. For each of the following languages over the alphabet $\{0, 1\}^*$, describe an equivalent regular expression, and briefly explain why your regular expression is correct. There are infinitely many correct answers for each language.

   (a) All strings in 1^*01^* whose length is a multiple of 3.

   (b) All strings that begin with the prefix 001, end with the suffix 100, and contain an odd number of 1s.

   (c) All strings that contain both 0011 and 1100 as substrings.

   (d) All strings that contain the substring 01 an odd number of times.

   (e) $\{0^a 1^b 0^c \mid a \geq 0 \text{ and } b \geq 0 \text{ and } c \geq 0 \text{ and } a \equiv b + c \pmod 2\}$.

2. For each of the following languages over the alphabet $\Sigma = \{0, 1\}$, describe a DFA that accepts the language, and briefly describe the purpose of each state. You can describe your DFA using a drawing, or using formal mathematical notation, or using a product construction; see the standard DFA rubric.

   (a) All strings in 1^*01^* whose length is a multiple of 3.

   (b) All strings that represent a multiple of 5 in base 3. For example, this language contains the string 10100, because $10100_3 = 90_{10}$ is a multiple of 5. (Yes, base 3 allows the digits 0, 1, and 2, but your input string will never contain a 2.)

   (c) All strings containing the substring 01010010. (The required substring is $p_6 = v_6$ from Homework 1.)

   (d) All strings whose ninth-to-last symbol is 0, or equivalently, the set

   $$\left\{ x0z \mid x \in \Sigma^* \text{ and } z \in \Sigma^8 \right\}.$$

   (e) All strings w such that $(\#(0, w) \bmod 3) + (\#(1, w) \bmod 7) = (|w| \bmod 4)$.

   [Hint: Don't try to draw the last two.]

3. *Practice only. Do not submit solutions.*

This question asks about strings over the set of *pairs* of bits, which we will write vertically. Let Σ_2 denote the set of all bit-pairs:

$$\Sigma_2 = \left\{ \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right\}$$

We can interpret any string w of bit-pairs as a $2 \times |w|$ matrix of bits; each row of this matrix is the binary representation of some non-negative integer, possibly with leading 0s. Let $hi(w)$ and $lo(w)$ respectively denote the *numerical values* of the top and bottom row of this matrix. For example, $hi(\varepsilon) = lo(\varepsilon) = 0$, and if

$$w = \begin{bmatrix} 0 \\ 0 \end{bmatrix}\begin{bmatrix} 0 \\ 1 \end{bmatrix}\begin{bmatrix} 1 \\ 0 \end{bmatrix}\begin{bmatrix} 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 0011 \\ 0101 \end{bmatrix}$$

then $hi(w) = 3$ and $lo(w) = 5$.

(a) Describe a DFA that accepts the language $L_{+1} = \{w \in \Sigma_2^* \mid hi(w) = lo(w) + 1\}$.
   For example, $w = \begin{bmatrix} 1 \\ 1 \end{bmatrix}\begin{bmatrix} 1 \\ 0 \end{bmatrix}\begin{bmatrix} 0 \\ 1 \end{bmatrix}\begin{bmatrix} 0 \\ 1 \end{bmatrix} = \begin{bmatrix} 1100 \\ 1011 \end{bmatrix} \in L_{+1}$, because $hi(w) = 12$ and $lo(w) = 11$.

(b) Describe a regular expression for L_{+1}.

(c) Describe a DFA that accepts the language $L_{\times 3} = \{w \in \Sigma_2^* \mid hi(w) = 3 \cdot lo(w)\}$.
   For example, $w = \begin{bmatrix} 1 \\ 0 \end{bmatrix}\begin{bmatrix} 0 \\ 0 \end{bmatrix}\begin{bmatrix} 0 \\ 1 \end{bmatrix}\begin{bmatrix} 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 1001 \\ 0011 \end{bmatrix} \in L_3$, because $hi(w) = 9$ and $lo(w) = 3$.

(d) Describe a regular expression for $L_{\times 3}$.

★(e) Describe a DFA that accepts the language $L_{\times 3/2} = \{w \in \Sigma_2^* \mid 2 \cdot hi(w) = 3 \cdot lo(w)\}$.
   For example, $w = \begin{bmatrix} 1 \\ 0 \end{bmatrix}\begin{bmatrix} 0 \\ 0 \end{bmatrix}\begin{bmatrix} 0 \\ 1 \end{bmatrix}\begin{bmatrix} 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 1001 \\ 0110 \end{bmatrix} \in L_{\times 3/2}$, because $hi(w) = 9$ and $lo(w) = 6$.

(Don't bother with the regular expression for this one.)

Solved problem

4. **C comments** are the set of strings over alphabet $\Sigma = \{*, /, A, \diamond, \downarrow\}$ that form a proper comment in the C program language and its descendants, like C++ and Java. Here $\downarrow$ represents the newline character, $\diamond$ represents any other whitespace character (like the space and tab characters), and A represents any non-whitespace character other than $*$ or $/$.[1] There are two types of C comments:

- Line comments: Strings of the form $//\cdots\downarrow$
- Block comments: Strings of the form $/*\cdots*/$

Following the C99 standard, we explicitly disallow **nesting** comments of the same type. A line comment starts with $//$ and ends at the first $\downarrow$ after the opening $//$. A block comment starts with $/*$ and ends at the the first $*/$ completely after the opening $/*$; in particular, every block comment has at least two $*$s. For example, each of the following strings is a valid C comment:

$$/***/ \qquad //\diamond//\diamond\downarrow \qquad /*///\diamond*\diamond\downarrow**/ \qquad /*\diamond//\diamond\downarrow\diamond*/$$

On the other hand, *none* of the following strings is a valid C comment:

$$/*/ \qquad //\diamond//\diamond\downarrow\diamond\downarrow \qquad /*\diamond/*\diamond*/\diamond*/$$

(Questions about C comments start on the next page.)

---

[1]The actual C commenting syntax is considerably more complex than described here, because of character and string literals.
- The opening $/*$ or $//$ of a comment must not be inside a string literal ($"\cdots"$) or a (multi-)character literal ($'\cdots'$).
- The opening double-quote of a string literal must not be inside a character literal ($'"'$) or a comment.
- The closing double-quote of a string literal must not be escaped ($\backslash"$)
- The opening single-quote of a character literal must not be inside a string literal ($"\cdots'\cdots"$) or a comment.
- The closing single-quote of a character literal must not be escaped ($\backslash'$)
- A backslash escapes the next symbol if and only if it is not itself escaped ($\backslash\backslash$) or inside a comment.

For example, the string $"/*\backslash\backslash\backslash"*/"/*"/*\backslash"/*"*/$ is a valid string literal (representing the 5-character string $/*\backslash"\backslash*/$, which is itself a valid block comment!) followed immediately by a valid block comment. **For this homework question, just pretend that the characters ', ", and $\backslash$ don't exist.**

Commenting in C++ is even more complicated, thanks to the addition of *raw* string literals. Don't ask.

Some C and C++ compilers do support nested block comments, in violation of the language specification. A few other languages, like OCaml, explicitly allow nesting block comments.

(a) Describe a regular expression for the set of all C comments.

> **Solution:**
>
> $$//(/ + * + A + \diamond)^* \hookleftarrow \quad + \quad /*\left(/ + A + \diamond + \hookleftarrow + **^*(A + \diamond + \hookleftarrow)\right)^* **^*/$$
>
> The first subexpression matches all line comments, and the second subexpression matches all block comments. Within a block comment, we can freely use any symbol other than $*$, but any run of $*$s must be followed by a character in $(A + \diamond + \hookleftarrow)$ or by the closing slash of the comment.   ■

> **Rubric:** Standard regular expression rubric. This is not the only correct solution.

(b) Describe a regular expression for the set of all strings composed entirely of blanks ($\diamond$), newlines ($\hookleftarrow$), and C comments.

> **Solution:**
>
> $$\left(\diamond + \hookleftarrow + //(/ + * + A + \diamond)^* \hookleftarrow + /*(/ + A + \diamond + \hookleftarrow + **^*(A + \diamond + \hookleftarrow))^* **^*/\right)^*$$
>
> This regular expression has the form $(\langle\text{whitespace}\rangle + \langle\text{comment}\rangle)^*$, where $\langle\text{whitespace}\rangle$ is the regular expression $\diamond + \hookleftarrow$ and $\langle\text{comment}\rangle$ is the regular expression from part (a).   ■

> **Rubric:** Standard regular expression rubric. This is not the only correct solution.

(c) Describe a DFA that accepts the set of all C comments.

Solution: The following eight-state DFA recognizes the language of C comments. All missing transitions lead to a hidden reject state.

The states are labeled mnemonically as follows:

- s — We have not read anything.
- / — We just read the initial /.
- // — We are reading a line comment.
- L — We have just read a complete line comment.
- /* — We are reading a block comment, and we did not just read a * after the opening /*.
- /** — We are reading a block comment, and we just read a * after the opening /*.
- B — We have just read a complete block comment.

∎

Rubric: Standard DFA design rubric. This is not the only correct solution, or even the simplest correct solution. (We don't need two distinct accepting states.)

(d) Describe a DFA that accepts the set of all strings composed entirely of blanks (◇), newlines (↵), and C comments.

Solution: By merging the accepting states of the previous DFA with the start state and adding white-space transitions at the start state, we obtain the following six-state DFA. Again, all missing transitions lead to a hidden reject state.

The states are labeled mnemonically as follows:

- *s* — We are between comments.
- / — We just read the initial / of a comment.
- // — We are reading a line comment.
- /* — We are reading a block comment, and we did not just read a * after the opening /*.
- /** — We are reading a block comment, and we just read a * after the opening /*.

∎

Rubric: Standard DFA design rubric. This is not the only correct solution, but it is the simplest correct solution.

⋆5. Recall that the reversal w^R of a string w is defined recursively as follows:

$$w^R := \begin{cases} \varepsilon & \text{if } w = \varepsilon \\ x^R \bullet a & \text{if } w = a \cdot x \end{cases}$$

The reversal L^R of any *language* L is the set of reversals of all strings in L:

$$L^R := \left\{ w^R \mid w \in L \right\}.$$

Prove that the reversal of every regular language is regular.

Solution: Let r be an arbitrary regular expression. We want to derive a regular expression r' such that $L(r') = L(r)^R$.

Assume for every regular expression s smaller than r that there is a regular expression s' such that $L(s') = L(s)^R$.

There are five cases to consider (mirroring the definition of regular expressions).

(a) If $r = \varnothing$, then we set $r' = \varnothing$, so that

$$\begin{aligned} L(r)^R &= L(\varnothing)^R & \text{because } r = \varnothing \\ &= \varnothing^R & \text{because } L(\varnothing) = \varnothing \\ &= \varnothing & \text{because } \varnothing^R = \varnothing \\ &= L(\varnothing) & \text{because } L(\varnothing) = \varnothing \\ &= L(r') & \text{because } r = \varnothing \end{aligned}$$

(b) If $r = w$ for some string $w \in \Sigma^*$, then we set $r' := w^R$, so that

$$\begin{aligned} L(r)^R &= L(w)^R & \text{because } r = w \\ &= \{w\}^R & \text{because } L(\langle\text{string}\rangle) = \{\langle\text{string}\rangle\} \\ &= \{w^R\} & \text{by definition of } L^R \\ &= L(w^R) & \text{because } L(\langle\text{string}\rangle) = \{\langle\text{string}\rangle\} \\ &= L(r') & \text{because } r = w^R \end{aligned}$$

(c) Suppose $r = s^*$ for some regular expression s. The inductive hypothesis implies a regular expressions s' such that $L(s') = L(s)^R$. Let $\mathbf{r' = (s')^*}$; then we have

$$\begin{aligned} L(r)^R &= L(s^*)^R & \text{because } r = s^* \\ &= (L(s)^*)^R & \text{by definition of } {}^* \\ &= (L(s)^R)^* & \text{because } (L^R)^* = (L^*)^R \\ &= (L(s'))^* & \text{by definition of } s' \\ &= L((s')^*) & \text{by definition of } {}^* \\ &= L(r') & \text{by definition of } r' \end{aligned}$$

(d) Suppose $r = s + t$ for some regular expressions s and t. The inductive hypothesis implies regular expressions s' and t' such that $L(s') = L(s)^R$ and $L(t') = L(t)^R$.

Set $r' := s' + t'$; then we have

$$
\begin{aligned}
L(r)^R &= L(s+t)^R && \text{because } r = s+t \\
&= (L(s) \cup L(t))^R && \text{by definition of } + \\
&= \{w^R \mid w \in (L(s) \cup L(t))\} && \text{by definition of } L^R \\
&= \{w^R \mid w \in L(s) \text{ or } w \cup L(t)\} && \text{by definition of } \cup \\
&= \{w^R \mid w \in L(s)\} \cup \{w^R \mid w \cup L(t)\} && \text{by definition of } \cup \\
&= L(s)^R \cup L(t)^R && \text{by definition of } L^R \\
&= L(s') \cup L(t') && \text{by definition of } s' \text{ and } t' \\
&= L(s' + t') && \text{by definition of } + \\
&= L(r') && \text{by definition of } r'
\end{aligned}
$$

(e) Suppose $r = s \bullet t$ for some regular expressions s and t. The inductive hypothesis implies regular expressions s' and t' such that $L(s') = L(s)^R$ and $L(t') = L(t)^R$. Set $r' = t' \bullet s'$; then we have

$$
\begin{aligned}
L(r)^R &= L(st)^R && \text{because } r = s+t \\
&= (L(s) \bullet L(t))^R && \text{by definition of } \bullet \\
&= \{w^R \mid w \in (L(s) \bullet L(t))\} && \text{by definition of } L^R \\
&= \{(x \bullet y)^R \mid x \in L(s) \text{ and } y \in L(t)\} && \text{by definition of } \bullet \\
&= \{y^R \bullet x^R \mid x \in L(s) \text{ and } y \in L(t)\} && \text{concatenation reversal} \\
&= \{y' \bullet x' \mid x' \in L(s)^R \text{ and } y' \in L(t)^R\} && \text{by definition of } L^R \\
&= \{y' \bullet x' \mid x' \in L(s') \text{ and } y' \in L(t')\} && \text{by definition of } s' \text{ and } t' \\
&= L(t') \bullet L(s') && \text{by definition of } \bullet \\
&= L(t' \bullet s') && \text{by definition of } \bullet \\
&= L(r') && \text{by definition of } r'
\end{aligned}
$$

In all five cases, we have found a regular expression r' such that $L(r') = L(r)^R$. It follows that $L(r)^R$ is regular. ∎

Rubric: Standard induction rubric!!

CS/ECE 374 A ✧ Fall 2023

♫ Homework 3 ∿

Due Tuesday, September 12, 2023 at 9pm Central Time

- Please start your solution to each *lettered subproblem* ((a), (b), (c), etc.) on a new page. Please also remember to tell Gradescope which page(s) are relevant for which subproblems.

---

1. Prove that the following languages over the alphabet $\Sigma = \{0, 1\}$ are *not* regular.

   (a) $\left\{0^a 1 0^b 1 0^c \mid 2b = a + c\right\}$.

   (b) The set of all palindromes in Σ^* whose lengths are divisible by 7.

   (c) $\left\{1^m 0^n \mid m + n > 0 \text{ and } \gcd(m, n) = 1\right\}$

     Here $\gcd(m, n)$ denotes the *greatest common divisor* of m and n: the largest integer d such that both m/d and n/d are integers. In particular, $\gcd(1, n) = 1$ and $\gcd(0, n) = n$ for every positive integer n.

2. For each of the following languages over the alphabet $\Sigma = \{0, 1\}$, either prove that the language is regular (by constructing an appropriate DFA, NFA, or regular expression) or prove that the language is not regular (by constructing an infinite fooling set). Recall that Σ^+ denotes the set of all *nonempty* strings over Σ.

   (a) Strings in which the substrings 01 and 10 appear the same number of times. For example, 1100011 $\in L$ because both substrings appear once, but 01000011 $\notin L$.

   (b) Strings in which the substrings 00 and 11 appear the same number of times. For example, 1100011 $\in L$ because both substrings appear twice, but 01000011 $\notin L$.

   (c) $\left\{xyyx \mid x, y \in \Sigma^+\right\}$

   (d) $\left\{xyyz \mid x, y, z \in \Sigma^+\right\}$

   [Hint: Exactly two of these languages are regular.]

⋆3. *Practice only. Do not submit solutions.*

　　A **Moore machine** is a variant of a finite-state automaton that produces output; Moore machines are sometimes called finite-state *transducers*. For purposes of this problem, a Moore machine formally consists of six components:

- A finite set Σ called the input alphabet
- A finite set Γ called the output alphabet
- A finite set Q whose elements are called states
- A start state $s \in Q$
- A transition function $\delta: Q \times \Sigma \to Q$
- An output function $\omega: Q \to \Gamma$

More intuitively, a Moore machine is a graph with a special start vertex, where every node (state) has one outgoing edge labeled with each symbol from the input alphabet, and each node (state) is additionally labeled with a symbol from the output alphabet.

　　The Moore machine reads an input string $w \in \Sigma^*$ one symbol at a time. For each symbol, the machine changes its state according to the transition function δ, and then outputs the symbol $\omega(q)$, where q is the new state. Formally, we recursively define a *transducer* function $\omega^*: Q \times \Sigma^* \to \Gamma^*$ as follows:

$$\omega^*(q,w) = \begin{cases} \varepsilon & \text{if } w = \varepsilon \\ \omega(\delta(q,a)) \cdot \omega^*(\delta(q,a), x) & \text{if } w = ax \end{cases}$$

Given input string $w \in \Sigma^*$, the machine outputs the string $\omega^*(w,s) \in \Gamma^*$. The **output language $L^\circ(M)$** of a Moore machine M is the set of all strings that the machine can output:

$$L^\circ(M) := \{\omega^*(s,w) \mid w \in \Sigma^*\}$$

(a) Let M be an arbitrary Moore machine. Prove that $L^\circ(M)$ is a regular language.

(b) Let M be an arbitrary Moore machine whose input alphabet Σ and output alphabet Γ are identical. Prove that the language

$$L^=(M) = \{w \in \Sigma^* \mid w = \omega^*(s,w)\}$$

is regular. $L^=(M)$ consists of all strings w such that M outputs w when given input w; these are also called *fixed points* for the transducer function ω^*.

[Hint: These problems are easier than they look!]

Solved problems

4. For each of the following languages, either prove that the language is regular (by constructing an appropriate DFA, NFA, or regular expression) or prove that the language is not regular (by constructing an infinite fooling set).

   Recall that a *palindrome* is a string that equals its own reversal: $w = w^R$. Every string of length 0 or 1 is a palindrome.

   (a) Strings in $(0 + 1)^*$ in which no prefix of length at least 2 is a palindrome.

   > **Solution: Regular:** $\varepsilon + 01^* + 10^*$. Call this language L_a.
   >
   > Let w be an arbitrary non-empty string in $(0 + 1)^*$. Without loss of generality, assume $w = 0x$ for some string x. There are two cases to consider.
   >
   > - If x contains a 0, then we can write $w = 01^n0y$ for some integer n and some string y. The prefix 01^n0 is a palindrome of length at least 2. Thus, $w \notin L_a$.
   > - Otherwise, $x \in 1^*$. Every non-empty prefix of w is equal to 01^n for some non-negative integer $n \le |x|$. Every palindrome that starts with 0 also ends with 0, so the only palindrome prefixes of w are ε and 0, both of which have length less than 2. Thus, $w \in L_a$.
   >
   > We conclude that $0x \in L_a$ if and only if $x \in 1^*$. A similar argument implies that $1x \in L_a$ if and only if $x \in 0^*$. Finally, trivially, $\varepsilon \in L_a$.   ∎

   > **Rubric:** 2½ points = ½ for "regular" + 1 for regular expression + 1 for justification. This is more detail than necessary for full credit.

   (b) Strings in $(0 + 1 + 2)^*$ in which no prefix of length at least 2 is a palindrome.

   > **Solution: Not regular.** Call this language L_b.
   >
   > Consider the set $F = (012)^+$.
   >
   > Let x and y be arbitrary distinct strings in F.
   >
   > Then $x = (012)^i$ and $y = (012)^j$ for some positive integers $i \ne j$.
   >
   > Without loss of generality, assume $i < j$.
   >
   > Let z be the suffix $(210)^i$.
   >
   > - $xz = (012)^i(210)^i$ is a palindrome of length $6i \ge 2$, so $xz \notin L_b$.
   > - $yz = (012)^j(210)^i$ has no palindrome prefixes except ε and 0, because $i < j$, so $yz \in L_b$.
   >
   > Thus, z is a distinguishing suffix for x and y.
   >
   > We conclude that F is a fooling set for L_b.
   >
   > Because F is infinite, L_b cannot be regular.   ∎

   > **Rubric:** 2½ points = ½ for "not regular" + 2 for fooling set proof (standard rubric, scaled).

(c) Strings in $(0+1)^*$ in which no prefix of length at least 3 is a palindrome.

> **Solution: Not regular.** Call this language L_c.
>
> Consider the set $F = (001101)^+$.
>
> Let x and y be arbitrary distinct strings in F.
>
> Then $x = (001101)^i$ and $y = (001101)^j$ for some positive integers $i \neq j$.
>
> Without loss of generality, assume $i < j$.
>
> Let z be the suffix $(101100)^i$.
>
> - $xz = (001101)^i(101100)^i$ is a palindrome of length $12i \geq 2$, so $xz \notin L_b$.
> - $yz = (001101)^j(101100)^i$ has no palindrome prefixes except ε and 0 and 00, because $i < j$, so $yz \in L_b$.
>
> Thus, z is a distinguishing suffix for x and y.
>
> We conclude that F is a fooling set for L_c.
>
> Because F is infinite, L_c cannot be regular. ∎

> **Rubric:** 2½ points = ½ for "not regular" + 2 for fooling set proof (standard rubric, scaled).

(d) Strings in $(0+1)^*$ in which no *substring* of length at least 3 is a palindrome.

> **Solution: Regular.** Call this language L_d.
>
> Every palindrome of length at least 3 contains a palindrome substring of length 3 or 4. Thus, the complement language $\overline{L_d}$ is described by the regular expression
>
> $$(0+1)^*(000 + 010 + 101 + 111 + 0110 + 1001)(0+1)^*$$
>
> Thus, $\overline{L_d}$ is regular, so its complement L_d is also regular. ∎

> **Solution: Regular.** Call this language L_d.
>
> In fact, L_d is *finite*! Appending either 0 or 1 to any of the underlined strings creates a palindrome suffix of length 3 or 4.
>
> $$\varepsilon + 0 + 1 + 00 + 01 + 10 + 11 + 001 + \underline{011} + \underline{100} + 110 + \underline{0011} + \underline{1100}$$
>
> ∎

> **Rubric:** 2½ points = ½ for "regular" + 2 for proof:
> - 1 for expression for $\overline{L_d}$ + 1 for applying closure
> - 1 for regular expression + 1 for justification

CS/ECE 374 A ✦ Fall 2023

ꕥ Homework 4 ꕥ

Due Tuesday, September 19, 2023 at 9pm Central Time

This is the last homework before Midterm 1.

---

1. Recall the following string functions from Homework 1:

$$stutter(w) := \begin{cases} \varepsilon & \text{if } w = \varepsilon \\ aa \bullet stutter(x) & \text{if } w = ax \end{cases} \qquad grow(w) := \begin{cases} \varepsilon & \text{if } w = \varepsilon \\ \texttt{0} \cdot grow(x) & \text{if } w = \texttt{1}x \\ \texttt{10} \bullet grow(x) & \text{if } w = \texttt{0}x \end{cases}$$

For example, $stutter(\texttt{1001}) = \texttt{11000011}$, and $grow(\texttt{1001}) = \texttt{0} \bullet \texttt{10} \bullet \texttt{10} \bullet \texttt{0} = \texttt{010100}$.

Let L be an arbitrary regular language over the alphabet $\Sigma = \{\texttt{0}, \texttt{1}\}$. Prove that the following languages are also regular.

   (a) $\text{STUTTER}(L) = \{stutter(w) \mid w \in L\}$

   (b) $\text{UNSTUTTER}(L) = \{w \mid stutter(w) \in L\}$

   (c) $\text{GROW}(L) = \{grow(w) \mid w \in L\}$

   (d) $\text{UNGROW}(L) = \{w \mid grow(w) \in L\}$

2. Give context-free grammars for the following languages, and clearly explain how they work and the set of strings generated by each nonterminal. Grammars with unclear or missing explanations may receive little or no credit. On the other hand, we do *not* want formal proofs of correctness.

   (a) $\left\{ \texttt{0}^a\texttt{10}^b\texttt{10}^c \mid b = 2a + 2c \right\}$.

   (b) $\left\{ \texttt{0}^a\texttt{10}^b\texttt{10}^c \mid 2b = a + c \right\}$.

   (c) The set of all palindromes in Σ^* whose lengths are divisible by 7.

   ⋆(d) *Practice only. Do not submit solutions.*

   Strings in which the substrings $\texttt{00}$ and $\texttt{11}$ appear the same number of times. For example, $\texttt{1100011} \in L$ because both substrings appear twice, but $\texttt{01000011} \notin L$.

   Yes, you've seen most of these languages before.

*3. *Practice only. Do not submit solutions.*

Let L_1 and L_2 be arbitrary regular languages over the alphabet $\Sigma = \{0,1\}$. Prove that the following languages are also regular.

(a) $\textsc{Faro}(L_1, L_2) := \{faro(x,z) \mid x \in L_1 \text{ and } z \in L_2 \text{ with } |x| = |z|\}$, where

$$faro(x,z) := \begin{cases} z & \text{if } x = \varepsilon \\ a \cdot faro(z,y) & \text{if } x = ay \end{cases}$$

For example, $faro(0011, 0101) = 00011011$ and $\textsc{Faro}(0^*, 1^*) = (01)^*$.

(b) $\textsc{Shuffles}(L_1, L_2) := \bigcup_{w \in L_1, y \in L_2} shuffles(w,y)$, where $shuffles(w,y)$ is the set of all strings obtained by shuffling w and y, or equivalently, all strings in which w and y are complementary subsequences. Formally:

$$shuffles(w,y) = \begin{cases} \{y\} & \text{if } w = \varepsilon \\ \{w\} & \text{if } y = \varepsilon \\ \{a\} \bullet shuffles(x,y) \cup \{b\} \bullet shuffles(w,z) & \text{if } w = ax \text{ and } y = bz \end{cases}$$

For example, $shuffles(001, 1) = \{0011, 0101, 1001\}$ and $shuffles(00, 11) = \{0011, 0101, 0110, 1001, 1010, 1100\}$. Finally, $\textsc{Shuffles}(0^*, 1^*) = (0 + 1)^*$.

Both of these names are taken from methods of mixing a deck of playing cards. A *shuffle* divides the deck into two smaller stacks, and then interleaves those two stacks arbitrarily. A *Faro shuffle* or *perfect shuffle* divides the pack of cards exactly in half, and then interleaves them perfectly; the final deck alternates between cards from one half and cards from the other half. Faro shuffles are the basis of several card tricks.

Solved problems

4. (a) Fix an arbitrary regular language L. Prove that the language $half(L) := \{w \mid ww \in L\}$ is also regular.

> **Solution:** Let $M = (\Sigma, Q, s, A, \delta)$ be an arbitrary DFA that accepts L. We define a new NFA $M' = (\Sigma, Q', s', A', \delta')$ with ε-transitions that accepts $half(L)$, as follows:
>
> $$Q' = (Q \times Q \times Q) \cup \{s'\}$$
> $$s' \text{ is an explicit state in } Q'$$
> $$A' = \{(h, h, q) \mid h \in Q \text{ and } q \in A\}$$
>
> $$\delta'(s', \varepsilon) = \{(s, h, h) \mid h \in Q\}$$
> $$\delta'(s', a) = \varnothing$$
>
> $$\delta'((p, h, q), \varepsilon) = \varnothing$$
> $$\delta'((p, h, q), a) = \{(\delta(p, a), h, \delta(q, a))\}$$
>
> M' reads its input string w and simulates M reading the input string ww. Specifically, M' simultaneously simulates two copies of M, one reading the left half of ww starting at the usual start state s, and the other reading the right half of ww starting at some intermediate state h.
>
> - The new start state s' non-deterministically guesses the "halfway" state $h = \delta^*(s, w)$ without reading any input; this is the only non-determinism in M'.
> - State (p, h, q) means the following:
>   - The left copy of M (which started at state s) is now in state p.
>   - The initial guess for the halfway state is h.
>   - The right copy of M (which started at state h) is now in state q.
> - M' accepts if and only if the left copy of M ends at state h (so the initial non-deterministic guess $h = \delta^*(s, w)$ was correct) and the right copy of M ends in an accepting state.
>
> ■

> **Solution (smartass):** A complete solution is given in the lecture notes.   ■

> **Rubric:** 5 points: standard langage transformation rubric (scaled). Yes, the smartass solution would be worth full credit.

(b) Describe a regular language L such that the language $double(L) := \{ww \mid w \in L\}$ is *not* regular. Prove your answer is correct.

> **Solution:** Consider the regular language $L = 0^*1$.
>
> Expanding the regular expression lets us rewrite $L = \{0^n1 \mid n \geq 0\}$. It follows that $double(L) = \{0^n10^n1 \mid n \geq 0\}$. I claim that this language is not regular.
>
> Let x and y be arbitrary distinct strings in L.
>
> Then $x = 0^i1$ and $y = 0^j1$ for some integers $i \neq j$.
>
> Then x is a distinguishing suffix of these two strings, because
>
> - $xx \in double(L)$ by definition, but
> - $yx = 0^i10^j1 \notin double(L)$ because $i \neq j$.
>
> We conclude that L is a fooling set for $double(L)$.
>
> Because L is infinite, $double(L)$ cannot be regular. ∎

> **Solution:** Consider the regular language $L = \Sigma^* = (0 + 1)^*$.
>
> I claim that the language $double(\Sigma^*) = \{ww \mid w \in \Sigma^*\}$ is not regular.
>
> Let F be the infinite language 01^*0.
>
> Let x and y be arbitrary distinct strings in F.
>
> Then $x = 01^i0$ and $y = 01^j0$ for some integers $i \neq j$.
>
> The string $z = 1^i$ is a distinguishing suffix of these two strings, because
>
> - $xz = 01^i01^i = ww$ where $w = 01^i$, so $xz \in double(\Sigma^*)$, but
> - $yx = 01^j01^i \notin double(\Sigma^*)$ because $i \neq j$.
>
> We conclude that F is a fooling set for $double(\Sigma^*)$.
>
> Because F is infinite, $double(\Sigma^*)$ cannot be regular. ∎

> **Rubric:** 5 points:
>
> - 2 points for describing a regular language L such that $double(L)$ is not regular.
> - 3 point for the fooling set proof (standard fooling set rubric, scaled and rounded)
>
> These are not the only correct solutions. These are not the only fooling sets for these languages.

5. Give context-free grammars for the following languages over the alphabet $\Sigma = \{0, 1\}$. Clearly explain how they work and the set of strings generated by each nonterminal. Grammars with unclear or missing explanations may receive little or no credit; on the other hand, we do *not* want formal proofs of correctness.

   (a) In any string, a **run** is a maximal non-empty substring of identical symbols. For example, the string $\texttt{0111000011001} = \texttt{0}^1\texttt{1}^3\texttt{0}^4\texttt{1}^2\texttt{0}^2\texttt{1}^1$ consists of six runs.

   Let L_a be the set of all strings in Σ^* that contain two runs of $\texttt{0}$s of equal length. For example, L_a contains the strings $\texttt{01101111}$ and $\texttt{01001011100010}$ (because each of those strings contains more than one run of $\texttt{0}$s of length 1) but L_a does not contain the strings $\texttt{000110011011}$ and $\texttt{00000000111}$.

   > **Solution:**
   >
   > $$\begin{aligned} S &\to ACB & &\text{strings with two blocks of } \texttt{0}\text{s of same length} \\ A &\to \varepsilon \mid X1 & &\text{empty or ends with } \texttt{1} \\ B &\to \varepsilon \mid 1X & &\text{empty or starts with } \texttt{1} \\ C &\to \texttt{0}C\texttt{0} \mid \texttt{0}D\texttt{0} & &\texttt{0}^n y \texttt{0}^n, \text{ where } y \text{ starts and ends with } \texttt{1} \\ D &\to 1 \mid 1X1 & &\text{starts and ends with } \texttt{1} \\ X &\to \varepsilon \mid 1X \mid \texttt{0}X & &\text{all strings: } (\texttt{0}+\texttt{1})^* \end{aligned}$$
   >
   > Every string in L has the form $x\texttt{0}^n y \texttt{0}^n z$, where x is either empty or ends with $\texttt{1}$, y starts and ends with $\texttt{1}$, and z is either empty or begins with $\texttt{1}$. Nonterminal A generates the prefix x; non-terminal B generates the suffix z; nonterminal C generates the matching runs of $\texttt{0}$s, and nonterminal D generates the interior string y.
   >
   > The same decomposition can be expressed more compactly as follows:
   >
   > $$\begin{aligned} S &\to B \mid B1A \mid A1B \mid A1B1A & &\text{strings with two blocks of } \texttt{0}\text{s of same length} \\ A &\to 1A \mid \texttt{0}A \mid \varepsilon & &\text{all strings: } (\texttt{0}+\texttt{1})^* \\ B &\to \texttt{0}B\texttt{0} \mid \texttt{010} \mid \texttt{01}A\texttt{10} & &\texttt{0}^n y \texttt{0}^n, \text{ where } y \text{ starts and ends with } \texttt{1} \end{aligned}$$
   >
   > ∎

   > **Rubric:** 5 points = 3 for clearly correct grammar + 2 for clear explanation. These are not the only correct solutions.

(b) $L_b = \{w \in \Sigma^* \mid w \text{ is } not \text{ a palindrome}\}$.

> **Solution:**
>
> $$S \rightarrow 0S0 \mid 0S1 \mid 1S0 \mid 1S1 \mid A \qquad \text{non-palindromes}$$
> $$A \rightarrow 0B1 \mid 1B0 \qquad\qquad\qquad \text{start and end with different symbols}$$
> $$B \rightarrow 0B \mid 1B \mid \varepsilon \qquad\qquad\quad \text{all strings}$$
>
> Every non-palindrome w can be decomposed as either $w = x0y1z$ or $w = x1y0z$, for some substrings x, y, z such that $|x| = |z|$. Non-terminal S generates the prefix x and matching-length suffix z; non-terminal A generates the distinct symbols, and non-terminal B generates the interior substring y. ∎

> **Solution:**
>
> $$S \rightarrow 0S0 \mid 1S1 \mid A \qquad \text{non-palindromes}$$
> $$A \rightarrow 0B1 \mid 1B0 \qquad\quad \text{start and end with different symbols}$$
> $$B \rightarrow 0B \mid 1B \mid \varepsilon \qquad \text{all strings}$$
>
> Every non-palindrome w must have a prefix x and a substring y such that either $w = x0y1x^R$ or $w = x1y0x^R$. Specifically, x is the longest common prefix of w and w^R. In the first case, the grammar generates w as follows:
>
> $$S \rightsquigarrow^* x A x^R \rightsquigarrow x 0B1 x^R \rightsquigarrow^* x0y1x^R = w$$
>
> The derivation for $w = x1y0x^R$ is similar. ∎

> **Rubric:** 5 points = 3 for clearly correct grammar + 2 for clear explanation. These are not the only correct solutions.

♪ Homework 5 ♫

Due Tuesday, October 3, 2023 at 9pm Central Time

---

1. In the lab on Wednesday, you'll see an algorithm that finds a local minimum in a one-dimensional array in $O(\log n)$ time. This question asks you to consider two higher-dimensional versions of this problem.

   (a) Suppose we are given a two-dimensional array $A[1..n, 1..n]$ of distinct integers. An array element $A[i, j]$ is called a **local minimum** if it is smaller than its four immediate neighbors:

   $$A[i, j] < \min \left\{ A[i-1, j],\ A[i+1, j],\ A[i, j-1],\ A[i, j+1] \right\}$$

   To avoid edge cases, we assume all cells in row 1, row n, column 1, and column n have value $+\infty$.

   Describe and analyze an algorithm to find a local minimum in A as quickly as possible. (Remember that faster algorithms are worth more points, but only if they are correct.)

   [Hint: Suppose $A[i, j]$ is the smallest element in row i. If $A[i, j]$ is smaller than both of its vertical neighbors $A[i-1, j]$ and $A[i+1, j]$, we are clearly done. But what if $A[i, j] > A[i+1, j]$?]

   *[Hint: This problem is more subtle than it appears at first glance; many published solutions for this problem on the internet are incorrect. The main issue is that a local minimum in a rectangular subarray is not necessarily a local minimum in the original array. Design a recursive algorithm for the following more general problem: Given a two-dimensional array that contains a local minimum **whose value is less than the value of every border cell**, find such a local minimum.]*

   (b) Now suppose we are given a three-dimensional array $A[1..n, 1..n, 1..n]$ of distinct integers. An array element $A[i, j, k]$ is called a **local minimum** if it is smaller than its six immediate neighbors:

   $$A[i, j] < \min \left\{ \begin{array}{l} A[i-1, j, k],\ A[i+1, j, k], \\ A[i, j-1, k],\ A[i, j+1, k], \\ A[i, j, k-1],\ A[i, j, k+1] \end{array} \right\}$$

   To avoid edge cases, we assume all cells on the boundary of the array have value $+\infty$.

   Describe and analyze an algorithm to find a local minimum in A as quickly as possible.

   (Remember that faster algorithms are worth more points, but only if they are correct.)

2. Suppose we have n points scattered inside a two-dimensional box. A *kd-tree* recursively subdivides the points as follows. First we split the box into two smaller boxes with a *vertical* line, then we split each of those boxes with *horizontal* lines, and so on, always alternating between horizontal and vertical splits. Each time we split a box, the splitting line partitions the rest of the interior points *as evenly as possible* by passing through a median point in the interior of the box (*not* on its boundary). If a box doesn't contain any points, we don't split it any more; these final empty boxes are called *cells*.

A kd-tree for 15 points. The dashed line crosses the four shaded cells.

(a) How many cells does the kd-tree have, as a function of n? Prove that your answer is correct.

(b) In the worst case, *exactly* how many cells can a horizontal line cross, as a function of n? Prove that your answer is correct. Assume that $n = 2^k - 1$ for some integer k. *[Hint: There is more than one function f such that $f(15) = 4$.]*

(c) Suppose we have n points stored in a kd-tree. Describe and analyze an algorithm that counts the number of points above a given horizontal line (such as the dashed line in the figure) as quickly as possible. *[Hint: Use part (b).]*

I should have specified that the following information is stored in each internal node v in the kd-tree:

- $v.x$ and $v.y$: The coordinates of the point defining the cut at v
- $v.dir \in \{vertical, horizontal\}$: The direction of the cut at v.
- $v.left$ and $v.right$: The children of v if $v.dir = vertical$
- $v.up$ and $v.down$: The children of v if $v.dir = horizontal$
- $v.size$: the number of points=cuts in the subtree rooted at v.

Instead I allowed arbitrary information to be computed in preprocessing; that freedom allows a much simpler and more efficient query algorithm!

(d) Describe and analyze an efficient algorithm that counts, given a kd-tree storing n points, the number of points that lie inside a given rectangle R with horizontal and vertical sides. *[Hint: Use part (c).]*

Assume that all x-coordinates and y-coordinates are distinct; that is, no two points lie on the same horizontal line or the same vertical line, no point lies on the query line in part (c), and no point lies on the boundary of the query rectangle in part (d).

*3. **_Practice only. Do not submit solutions._**

The following variant of the infamous StoogeSort algorithm[1] was discovered by the British actor Patrick Troughton during rehearsals for the 20th anniversary *Doctor Who* special "The Five Doctors".[2]

WHOSORT($A[1..n]$) :
  if $n < 13$
    sort A by brute force
  else
    $k = \lceil n/5 \rceil$
    WHOSORT($A[1..3k]$)   ⟨⟨*Hartnell*⟩⟩
    WHOSORT($A[2k+1..n]$)   ⟨⟨*Troughton*⟩⟩
    WHOSORT($A[1..3k]$)   ⟨⟨*Pertwee*⟩⟩
    WHOSORT($A[k+1..4k]$)   ⟨⟨*Davison*⟩⟩

(a) Prove by induction that WHOSORT correctly sorts its input. *[Hint: Where can the smallest k elements be?]*

(b) Would WHOSORT still sort correctly if we replaced "if $n < 13$" with "if $n < 4$"? Justify your answer.

(c) Would WHOSORT still sort correctly if we replaced "$k = \lceil n/5 \rceil$" with "$k = \lfloor n/5 \rfloor$"? Justify your answer.

(d) What is the running time of WHOSORT? (Set up a running-time recurrence and then solve it, ignoring the floors and ceilings.)

(e) Forty years later, 15th Doctor Ncuti Gatwa discovered the following optimization to WHOSORT, which uses the standard MERGE subroutine from mergesort, which merges two sorted arrays into one sorted array.

| |
|---|
| NUWHOSORT($A[1..n]$) :
  if $n < 13$
    sort A by brute force
  else
    $k = \lceil n/5 \rceil$
    NUWHOSORT($A[1..3k]$)   ⟨⟨*Grant*⟩⟩
    NUWHOSORT($A[2k+1..n]$)   ⟨⟨*Whittaker*⟩⟩
    MERGE($A[1..2k]$, $A[2k+1..4k]$)   ⟨⟨*Tennant*⟩⟩ |

What is the running time of NUWHOSORT?

---

[1] https://en.wikipedia.org/wiki/Stooge_sort

[2] Tom Baker, the fourth Doctor, declined to return for the reunion; hence, only four Doctors appeared in "The Five Doctors". (Well, okay, technically the BBC used excerpts of the unfinished episode "Shada" to include Baker, but he wasn't really *there*—to the extent that any fictional character in a television show about a time traveling wizard arguing with several other versions of himself about immortality can be said to be "really" "there".)

Solved problems

4. Suppose we are given two sets of n points, one set $\{p_1, p_2, \ldots, p_n\}$ on the line $y = 0$ and the other set $\{q_1, q_2, \ldots, q_n\}$ on the line $y = 1$. Consider the n line segments connecting each point p_i to the corresponding point q_i. Describe and analyze a divide-and-conquer algorithm to determine how many pairs of these line segments intersect, in $O(n \log n)$ time. See the example below.

Seven segments with endpoints on parallel lines, with 11 intersecting pairs.

Your input consists of two arrays $P[1..n]$ and $Q[1..n]$ of x-coordinates; you may assume that all $2n$ of these numbers are distinct. No proof of correctness is necessary, but you should justify the running time.

Solution: We begin by sorting the array $P[1..n]$ and permuting the array $Q[1..n]$ to maintain correspondence between endpoints, in $O(n \log n)$ time. Then for any indices $i < j$, segments i and j intersect if and only if $Q[i] > Q[j]$. Thus, our goal is to compute the number of pairs of indices $i < j$ such that $Q[i] > Q[j]$. Such a pair is called an **inversion**.

We count the number of inversions in Q using the following extension of mergesort; as a side effect, this algorithm also sorts Q. If $n < 100$, we use brute force in $O(1)$ time. Otherwise:

- Color the elements in the Left half $Q[1..\lfloor n/2 \rfloor]$ bLue.
- Color the elements in the Right half $Q[\lfloor n/2 \rfloor + 1..n]$ Red.
- Recursively count inversions in (and sort) the blue subarray $Q[1..\lfloor n/2 \rfloor]$.
- Recursively count inversions in (and sort) the red subarray $Q[\lfloor n/2 \rfloor + 1..n]$.
- Count red/blue inversions as follows:
  - MERGE the sorted subarrays $Q[1..n/2]$ and $Q[n/2+1..n]$, maintaining the element colors.
  - For each blue element $Q[i]$ of the now-sorted array $Q[1..n]$, count the number of smaller red elements $Q[j]$.

The last substep can be performed in $O(n)$ time using a simple for-loop:

```
COUNTREDBLUE(A[1..n]):
    count ← 0
    total ← 0
    for i ← 1 to n
        if A[i] is red
            count ← count + 1
        else
            total ← total + count
    return total
```

MERGE and COUNTREDBLUE each run in $O(n)$ time. Thus, the running time of our inversion-counting algorithm obeys the mergesort recurrence $T(n) = 2T(n/2) + O(n)$. (We can safely ignore the floors and ceilings in the recursive arguments.) We conclude that the overall running time of our algorithm is $O(n \log n)$, as required.

Rubric: This is enough for full credit.

In fact, we can execute the third merge-and-count step directly by modifying the MERGE algorithm, without any need for "colors". Here changes to the standard MERGE algorithm are indicated in red.

```
MERGEANDCOUNT(A[1..n], m):
    i ← 1; j ← m + 1; count ← 0; total ← 0
    for k ← 1 to n
        if j > n
            B[k] ← A[i]; i ← i + 1; total ← total + count
        else if i > m
            B[k] ← A[j]; j ← j + 1; count ← count + 1
        else if A[i] < A[j]
            B[k] ← A[i]; i ← i + 1; total ← total + count
        else
            B[k] ← A[j]; j ← j + 1; count ← count + 1
    for k ← 1 to n
        A[k] ← B[k]
    return total
```

We can further optimize MERGEANDCOUNT by observing that *count* is always equal to $j - m - 1$, so we don't need an additional variable. (Proof: Initially, $j = m + 1$ and *count* = 0, and we always increment j and *count* together.)

```
MERGEANDCOUNT2(A[1..n], m):
    i ← 1;  j ← m + 1;  total ← 0
    for k ← 1 to n
        if j > n
            B[k] ← A[i];  i ← i + 1;  total ← total + j − m − 1
        else if i > m
            B[k] ← A[j];  j ← j + 1
        else if A[i] < A[j]
            B[k] ← A[i];  i ← i + 1;  total ← total + j − m − 1
        else
            B[k] ← A[j];  j ← j + 1
    for k ← 1 to n
        A[k] ← B[k]
    return total
```

MERGEANDCOUNT2 still runs in $O(n)$ time, so the overall running time is still $O(n \log n)$, as required.   ∎

Rubric: 10 points = 2 for base case + 2 for divide (split and recurse) + 4 for conquer (merge and count) + 2 for time analysis. This is neither the only way to correctly describe this algorithm nor the only correct $O(n \log n)$-time algorithm. No proof of correctness is required.

Max 3 points for a correct $O(n^2)$-time algorithm.

Notice that each boxed algorithm is preceded by a clear English description of the task that algorithm performs—not how the algorithm works, but the relationship between its input and its output. **Each English description is worth 25% of the credit for that algorithm** (rounding to the nearest half-point). For example, the COUNTREDBLUE algorithm is worth 4 points ("conquer"); the English description alone ("For each blue element $Q[i]$ of the now-sorted array $Q[1..n]$, count the number of smaller red elements $Q[j]$.") is worth 1 point.

๑ Homework 6 ๖

Due Tuesday, October 10, 2023 at 9pm Central Time

---

1. Satya is in charge of establishing a new testing center for the Standardized Awesomeness Test (SAT), and found an old conference hall that is perfect. The conference hall has n rooms of various sizes along a single long hallway, numbered in order from 1 through n. Satya knows exactly how many students fit into each room, and he wants to use a subset of the rooms to host as many students as possible for testing.

   Unfortunately, there have been several incidents of students cheating at other testing centers by tapping secret codes through walls. To prevent this type of cheating, Satya can use two adjacent rooms only if he demolishes the wall between them. The city's chief architect has determined that demolishing the walls on both sides of the same room would threaten the building's structural integrity. For this reason, Satya can never host students in three consecutive rooms.

   Describe an efficient algorithm that computes the largest number of students that Satya can host for testing without using three consecutive rooms. The input to your algorithm is an array $S[1..n]$, where each $S[i]$ is the (non-negative integer) number of students that can fit in room i.

2. As a typical overworked college student, you occasionally pull all-nighters to get more work done. Painful experience has taught you that the longer you stay awake, the less productive you are.

   Suppose there are n days left in the semester. For each of the next n days, you can either stay awake and work, or you can sleep. You have an array $Score[1..n]$, where $Score[i]$ is the (always positive) number of points you will earn on day i *if you are awake and well-rested*.

   However, staying awake for several days in a row has a price: Each consecutive day you stay awake cuts the quality of your work in half. Thus, if you are awake on day i, and you most recently slept on day $i - k$, then you will actually earn $Score[i]/2^{k-1}$ points on day i. (You've already decided to sleep on day 0.)

   For example, suppose $n = 6$ and $Score = [3, 7, 4, 3, 9, 1]$.

   - If you work on all six days, you will earn $3 + \frac{7}{2} + \frac{4}{4} + \frac{3}{8} + \frac{9}{16} + \frac{1}{32} = 8.46875$ points.
   - If you work only on days 1, 3, and 5, you will earn $3 + 4 + 9 = 16$ points.
   - If you work only on days 2, 3, 5, and 6, you will earn $7 + \frac{4}{2} + 9 + \frac{1}{2} = 18.5$ points.

   Design and analyze an algorithm that computes the maximum number of points you can earn, given the array $Score[1..n]$ as input. For example, given the input array $[3, 7, 4, 3, 9, 1]$, your algorithm should return the number 18.5.

   VERY IMPORTANT: Do not actually do this in real life!

3. *Practice only. Do not submit solutions.*

   (a) Any string can be decomposed into a sequence of palindromes. For example, the
       string BUBBASEESABANANA ("Bubba sees a banana.") can be broken into palindromes
       in the following ways (and 65 others):

       <div align="center">

       BUB • BASEESAB • ANANA

       B • U • BB • ASEESA • B • ANANA

       BUB • B • A • SEES • ABA • N • ANA

       B • U • BB • A • S • EE • S • A • B • A • NAN • A

       B • U • B • B • A • S • E • E • S • A • B • A • N • A • N • A

       </div>

       Describe and analyze an efficient algorithm to find the smallest number of palin-
       dromes that make up a given input string. For example, given the input string
       BUBBASEESABANANA, your algorithm should return 3.

   (b) A *metapalindrome* is a decomposition of a string into a sequence of palindromes,
       such that the sequence of palindrome lengths is itself a palindrome. For example,
       the string BOBSMAMASEESAUKULELE ("Bob's mama sees a ukulele") has the following
       metapalindromes (among others):

       <div align="center">

       BOB • S • MAM • ASEESA • UKU • L • ELE

       B • O • B • S • M • A • M • A • S • E • E • S • A • U • K • U • L • E • L • E

       </div>

       The length sequences of these metapalindromes are $(3, 1, 3, 6, 3, 1, 3)$ and $(1, 1, 1, 1, 1,$
       $1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1)$; notice that both of these sequences are themselves
       palindromes.

       Describe and analyze an efficient algorithm to find the smallest number of
       palindromes in any metapalindrome for a given string. For example, given the input
       string BOBSMAMASEESAUKULELE, your algorithm should return 7.

Solved Problems

3. A *shuffle* of two strings X and Y is formed by interspersing the characters into a new string, keeping the characters of X and Y in the same order. For example, the string BANANAANANAS is a shuffle of the strings **BANANA** and ANANAS in several different ways.

$$\text{BANANA}_{\text{ANANAS}} \qquad \text{BAN}_{\text{ANA}}\text{ANA}_{\text{NAS}} \qquad \text{B}_{\text{AN}}\text{AN}_{\text{A}}\text{ANA}_{\text{NA}}\text{S}$$

Similarly, the strings PRODGYRNAMAMMIINCG and DYPRONGARMAMMICING are both shuffles of the strings **DYNAMIC** and PROGRAMMING:

$$\text{PRO}\text{DGYRNAM}_{\text{AMMI}}\text{INC}\text{G} \qquad \text{DY}\text{PRONGARM}_{\text{AMM}}\text{ICING}$$

(a) Given three strings $A[1..m]$, $B[1..n]$, and $C[1..m+n]$, describe and analyze an algorithm to determine whether C is a shuffle of A and B.

> **Solution:** We define a boolean function $Shuf(i,j)$, which is TRUE if and only if the prefix $C[1..i+j]$ is a shuffle of the prefixes $A[1..i]$ and $B[1..j]$. We need to compute $Shuf(m,n)$. The function $Shuf$ satisfies the following recurrence:
>
> $$Shuf(i,j) = \begin{cases} \text{TRUE} & \text{if } i = j = 0 \\ Shuf(0,j-1) \wedge (B[j] = C[j]) & \text{if } i = 0 \text{ and } j > 0 \\ Shuf(i-1,0) \wedge (A[i] = C[i]) & \text{if } i > 0 \text{ and } j = 0 \\ \big(Shuf(i-1,j) \wedge (A[i] = C[i+j])\big) & \\ \qquad \vee \big(Shuf(i,j-1) \wedge (B[j] = C[i+j])\big) & \text{otherwise} \end{cases}$$
>
> We can memoize this function into a two-dimensional array $Shuf[0..m][0..n]$. Each array entry $Shuf[i,j]$ depends only on the entries immediately above and immediately to the left: $Shuf[i-1,j]$ and $Shuf[i,j-1]$. Thus, we can fill the array in standard row-major order in $O(mn)$ **time.** ∎

> **Solution:** The following algorithm runs in $O(mn)$ **time**.
>
> ```
> ISSHUFFLE?(A[1..m], B[1..n], C[1..m+n]):
>     Shuf[0,0] ← TRUE
>     for j ← 1 to n
>         Shuf[0,j] ← Shuf[0,j-1] ∧ (B[j] = C[j])
>     for i ← 1 to n
>         Shuf[i,0] ← Shuf[i-1,0] ∧ (A[i] = B[i])
>     for j ← 1 to n
>         Shuf[i,j] ← FALSE
>         if A[i] = C[i+j]
>             Shuf[i,j] ← Shuf[i-1,j]
>         if B[i] = C[i+j]
>             Shuf[i,j] ← Shuf[i,j] ∨ Shuf[i,j-1]
>     return Shuf[m,n]
> ```
>
> Here $Shuf(i,j) = $ TRUE if and only if the prefix $C[1..i+j]$ is a shuffle of the

prefixes $A[1..i]$ and $B[1..j]$.   ■

(b) Given three strings $A[1..m]$, $B[1..n]$, and $C[1..m+n]$, describe and analyze an algorithm to determine *the number of different ways* that A and B can be shuffled to obtain C.

Solution: Let $\#Shuf(i,j)$ denote the number of different ways that the prefixes $A[1..i]$ and $B[1..j]$ can be shuffled to obtain the prefix $C[1..i+j]$. We need to compute $\#Shuf(m,n)$.

The $\#Shuf$ function satisfies the following recurrence. Here I am using Iverson bracket notation to convert booleans to integers: For any proposition P, the expression $[P]$ is equal to 1 if P is true and 0 if P is false.

$$\#Shuf(i,j) = \begin{cases} 1 & \text{if } i = j = 0 \\ \#Shuf(0,j-1) \cdot \big[B[j] = C[j]\big] & \text{if } i = 0 \text{ and } j > 0 \\ \#Shuf(i-1,0) \cdot \big[A[i] = C[i]\big] & \text{if } i > 0 \text{ and } j = 0 \\ \big(\#Shuf(i-1,j) \cdot \big[A[i] = C[i]\big]\big) \\ \quad + \big(\#Shuf(i,j-1) \cdot \big[B[j] = C[j]\big]\big) & \text{otherwise} \end{cases}$$

We can memoize this function into a two-dimensional array $\#Shuf[0..m][0..n]$. As in part (a), we can fill this array in standard row-major order in $O(mn)$ *time*.   ■

Solution: The following algorithm runs in $O(mn)$ *time*:

```
NumShuffles(A[1..m], B[1..n], C[1..m+n]):
    #Shuf[0,0] ← 1
    for j ← 1 to n
        #Shuf[0,j] ← 0
        if (B[j] = C[j])
            #Shuf[0,j] ← #Shuf[0,j−1]
    for i ← 1 to n
        #Shuf[0,j] ← 0
        if (A[i] = B[i])
            #Shuf[0,j] ← #Shuf[i−1,0]
        for j ← 1 to n
            #Shuf[i,j] ← 0
            if A[i] = C[i+j]
                #Shuf[i,j] ← #Shuf[i−1,j]
            if B[i] = C[i+j]
                #Shuf[i,j] ← #Shuf[i,j] + #Shuf[i,j−1]
    return Shuf[m,n]
```

Here #*Shuf*[i, j] stores the number of different ways that the prefixes $A[1..i]$ and $B[1..j]$ can be shuffled to obtain the prefix $C[1..i+j]$. ■

Rubric: 5 points, standard dynamic programming rubric. **Again, each of these solutions is separately worth full credit.** These are not the only correct solutions. —½ for reporting running time as $O(n^2)$. 3 points for a slower polynomial-time algorithm; scale partial credit accordingly.

♫ Homework 7 ♫

1. The City Council of Sham-Poobanana needs to partition Purple Street into voting districts. A total of n people live on Purple Street, at consecutive addresses $1, 2, \ldots, n$. Each voting district must be a contiguous interval of addresses $i, i+1, \ldots, j$ for some $1 \le i < j \le n$. By law, each Purple Street address must lie in exactly one district, and the number of addresses in each district must be between k and $2k$, where k is a positive integer parameter.

   Every election in Sham-Poobanana is between two rival factions: Oceania and Eurasia. A majority of the current City Council are from Oceania, so they consider a district to be *good* if more than half the residents of that district voted for Oceania in the previous election. Naturally, the City Council has complete voting records for all n residents.

   For example, the figure below shows a legal partition of 22 addresses (of which 9 are good and 13 are bad) into 4 good districts and 3 bad districts, where $k = 2$ (so each district contains either 2, 3, or 4 addresses). Each **O** indicates a vote for Oceania, and each **X** indicates a vote for Eurasia.

   Describe an algorithm to find the largest possible number of *good* districts in a legal partition. Your input consists of the integer k and a boolean array GoodVote[$1..n$] indicating which residents previously voted for Oceania (True) or Eurasia (False). You can assume that a legal partition exists. Analyze the running time of your algorithm in terms of the parameters n and k. (In particular, do **not** assume that k is a constant.)

2. The StupidScript language includes a binary operator @ that computes the *average* of its two arguments. For example, the StupidScript code `print(3 @ 6)` would print `4.5`, because $(3 + 6)/2 = 4.5$.

   Expressions like 3 @ 7 @ 4 that use the @ operator more than once yield different results when they are evaluated in different orders:

   $$(3 \text{ @ } 7) \text{ @ } 4 = 5 \text{ @ } 4 = 4.5 \qquad \text{but} \qquad 3 \text{ @ } (7 \text{ @ } 4) = 3 \text{ @ } 5.5 = 4.25$$

   Here is a larger example:

   $$((((8 \text{ @ } 6) \text{ @ } 7) \text{ @ } 5) \text{ @ } 3) \text{ @ } (0 \text{ @ } 9) = 4.5$$
   $$((8 \text{ @ } 6) \text{ @ } (7 \text{ @ } 5)) \text{ @ } ((3 \text{ @ } 0) \text{ @ } 9) = 5.875$$
   $$(8 \text{ @ } (6 \text{ @ } (7 \text{ @ } (5 \text{ @ } (3 \text{ @ } 0))))) \text{ @ } 9 = 7.890625$$

   Your goal for this problem is to describe and analyze an algorithm to compute, given a sequence of integers separated by @ signs, the **largest possible** value the expression can take by adding parentheses. Your input is an array $A[1..n]$ listing the sequence of integers.

   For example, if your input sequence is $[3, 7, 4]$, your algorithm should return 4.5, and if your input sequence is $[8, 6, 7, 5, 3, 0, 9]$, your algorithm should return 7.890625. Assume all arithmetic operations (including @) can be performed exactly in $O(1)$ time.

   (a) Tommy Tutone suggests the following natural greedy algorithm: Merge the adjacent pair of numbers with the smallest average (breaking ties arbitrarily), replace them with their average, and recurse. For example:

   $$8 \text{ @ } 6 \text{ @ } 7 \text{ @ } 5 \text{ @ } \underline{3 \text{ @ } 0} \text{ @ } 9$$
   $$8 \text{ @ } 6 \text{ @ } 7 \text{ @ } \underline{5 \text{ @ } 1.5} \text{ @ } 9$$
   $$8 \text{ @ } 6 \text{ @ } \underline{7 \text{ @ } 3.25} \text{ @ } 9$$
   $$8 \text{ @ } \underline{6 \text{ @ } 5.125} \text{ @ } 9$$
   $$\underline{8 \text{ @ } 5.5625} \text{ @ } 9$$
   $$\underline{6.78125 \text{ @ } 9}$$
   $$7.890625$$

   Tommy reasons that with an efficient priority queue, this algorithm will run in $O(n \log n)$ time, which is *way* faster than any dynamic programming algorithm.

   Prove that Tommy's algorithm is incorrect, by describing a specific input array and proving that his algorithm does not yield the largest possible value for that array.

   (b) Describe and analyze a correct algorithm for this problem. Poor, poor Tommy.

3. *Practice only. Do not submit solutions.*

Suppose we need to broadcast a message to all the nodes in a rooted binary tree. Initially, only the root node knows the message. In a single round, any node that knows the message can forward it to at most one of its children. See the figure below for an example.

Design an algorithm to compute the minimum number of rounds required to broadcast the message to every node.

A message being distributed through a binary tree in five rounds.

Solved problems

3. A string w of parentheses $($ and $)$ and brackets $[$ and $]$ is **balanced** if and only if w is generated by the following context-free grammar:

$$S \rightarrow \varepsilon \mid (S) \mid [S] \mid SS$$

For example, the string $w = ([()][]())[()()]()$ is balanced, because $w = xy$, where

$$x = ([()][]()) \qquad \text{and} \qquad y = [()()]().$$

Describe and analyze an algorithm to compute the length of a longest balanced subsequence of a given string of parentheses and brackets. Your input is an array $A[1..n]$, where $A[i] \in \{(,), [,]\}$ for every index i.

Solution: Suppose $A[1..n]$ is the input string. For all indices i and k, let $\textbf{LBS}(\textbf{\textit{i}},\textbf{\textit{k}})$ denote the length of the longest balanced subsequence of the substring $A[i..k]$. We need to compute $LBS(1,n)$. This function obeys the following recurrence:

$$LBS(i,k) = \begin{cases} 0 & \text{if } i \geq k \\[2ex] \max \begin{cases} 2 + LBS(i+1,k-1) \\ \displaystyle\max_{j=1}^{k-1}\big(LBS(i,j) + LBS(j+1,k)\big) \end{cases} & \text{if } A[i] \sim A[k] \\[3ex] \displaystyle\max_{j=1}^{k-1}\big(LBS(i,j) + LBS(j+1,k)\big) & \text{otherwise} \end{cases}$$

Here $\textbf{\textit{A}}[\textbf{\textit{i}}] \sim \textbf{\textit{A}}[\textbf{\textit{k}}]$ indicates that $A[i]$ is a left delimiter and $A[k]$ is the corresponding right delimiter: Either $A[i] = ($ and $A[k] =)$, or $A[i] = [$ and $A[k] =]$.

We can memoize this function into a two-dimensional array $LBS[1..n, 1..n]$. Because each entry $LBS[i,k]$ depends only on entries in later rows or earlier columns (or both), we can fill this array row-by-row from bottom up (decreasing i) in the outer loop, scanning each row from left to right (increasing k) in the inner loop.

We can compute each entry $LBS[i,k]$ in $O(n)$ time, so the resulting algorithm runs in $\textbf{\textit{O}}(\textbf{\textit{n}}^3)$ **time**. ∎

Solution (pseudocode): The following algorithm runs in $O(n^3)$ **time**:

```
LONGESTBALANCEDSUBSEQUENCE(A[1..n]):
    for i ← n down to 1
        LBS[i,i] ← 0
        for k ← i+1 to n
            if (A[i] = ( and A[k] = )) or (A[i] = [ and A[k] = ])
                LBS[i,k] ← LBS[i+1,k−1] + 2
            else
                LBS[i,k] ← 0
            for j ← i to k−1
                LBS[i,k] ← max {LBS[i,k], LBS[i,j] + LBS[j+1,k]}
    return LBS[1,n]
```

4

Here $LBS[i, k[$ stores the length of the longest balanced subsequence of the substring $A[i .. k]$.   ■

4. Oh, no! You've just been appointed as the new organizer of Giggle, Inc.'s annual mandatory holiday party! The employees at Giggle are organized into a strict hierarchy, that is, a tree with the company president at the root. The all-knowing oracles in Human Resources have assigned a real number to each employee measuring how "fun" the employee is. In order to keep things social, there is one restriction on the guest list: An employee cannot attend the party if their immediate supervisor is also present. On the other hand, the president of the company *must* attend the party, even though she has a negative fun rating; it's her company, after all.

Describe an algorithm that makes a guest list for the party that maximizes the sum of the "fun" ratings of the guests. The input to your algorithm is a rooted tree T describing the company hierarchy, where each node v has a field $v.fun$ storing the "fun" rating of the corresponding employee.

Solution (two functions): We define two functions over the nodes of T.

- *MaxFunYes*(v) is the maximum total "fun" of a legal party among the descendants of v, where v is definitely invited.

- *MaxFunNo*(v) is the maximum total "fun" of a legal party among the descendants of v, where v is definitely not invited.

We need to compute *MaxFunYes*$(root)$. These two functions obey the following mutual recurrences:

$$MaxFunYes(v) = v.fun + \sum_{\text{children } w \text{ of } v} MaxFunNo(w)$$

$$MaxFunNo(v) = \sum_{\text{children } w \text{ of } v} \max\{MaxFunYes(w), MaxFunNo(w)\}$$

These recurrences do not require separate base cases, because $\sum \varnothing = 0$.[a]

We can memoize these functions by adding two additional fields $v.yes$ and $v.no$ to each node v in the tree. The values at each node depend only on the vales at its children, so we can compute all $2n$ values using a postorder traversal of T.

The resulting algorithm spends $O(1)$ time at each node of T, and therefore runs in **$O(n)$ time**.   ∎

---

[a]A naïve recursive implementation of these recurrences would run in $O(\phi^n)$ time in the worst case, where $\phi = (1 + \sqrt{5})/2 \approx 1.618$ is the golden ratio. The worst case occurs when T is a single path.

Solution (two functions, pseudocode): The following algorithm runs in **$O(n)$ time**.

```
BestParty(T):
    ComputeMaxFun(T.root)
    return T.root.yes
```

```
ComputeMaxFun(v):
    v.yes ← v.fun
    v.no ← 0
    for all children w of v
        ComputeMaxFun(w)
        v.yes ← v.yes + w.no
        v.no ← v.no + max{w.yes, w.no}
```

We are storing two pieces of information in each node v of the tree:

- $v.yes$ is the maximum total "fun" of a legal party among the descendants of v, assuming v is invited.

- $v.no$ is the maximum total "fun" of a legal party among the descendants of v, assuming v is not invited.

(Yes, this is still dynamic programming; we're only traversing the tree recursively in COMPUTEMAXFUN because that's the most natural way to traverse trees!) ∎

Solution (one function): For each node v in the input tree T, let $MaxFun(v)$ denote the maximum total "fun" of a legal party among the descendants of v, where v may or may not be invited.

The president of the company must be invited, so none of the president's "children" in T can be invited. Thus, the value we need to compute is

$$root.fun + \sum_{\text{grandchildren } w \text{ of } root} MaxFun(w).$$

The function $MaxFun$ obeys the following recurrence:

$$MaxFun(v) = \max \left\{ \begin{array}{c} v.fun + \sum_{\text{grandchildren } x \text{ of } v} MaxFun(x) \\ \sum_{\text{children } w \text{ of } v} MaxFun(w) \end{array} \right\}$$

(This recurrence does not require a separate base case, because $\sum \varnothing = 0$.) We can memoize this function by adding an additional field $v.maxFun$ to each node v in the tree. The value at each node depends only on the values at its children and grandchildren, so we can compute all values using a postorder traversal of T.

The algorithm spends $O(1)$ time at each node (because each node has exactly one parent and one grandparent) and therefore runs in $O(n)$ **time** altogether. ∎

Solution (one function, pseudocode):

BESTPARTY(T):
  COMPUTEMAXFUN($T.root$)
  $party \leftarrow T.root.fun$
  for all children w of $T.root$
    for all children x of w
      $party \leftarrow party + x.maxFun$
  return $party$

COMPUTEMAXFUN(v):
  $yes \leftarrow v.fun$
  $no \leftarrow 0$
  for all children w of v
    COMPUTEMAXFUN(w)
    $no \leftarrow no + w.maxFun$
    for all children x of w
      $yes \leftarrow yes + x.maxFun$
  $v.maxFun \leftarrow \max\{yes, no\}$

Here $v.maxFun$ stores the maximum total "fun" of a legal party among the descendants of v, where v may or may not be invited.

> Each value *v.maxFun* is read at most three times during the algorithm's execution: Once in ComputeMaxFun(*v.parent*), and once in ComputeMaxFun(*v.parent.parent*), and at most once in the non-recursive part of BestParty. Thus, the entire algorithm runs in $O(n)$ *time*. ∎

Rubric: 10 points: standard dynamic programming rubric. These are not the only correct solutions. Yes, each of these solutions is independently worth full credit.

♫ Homework 8 ♫

Due Tuesday, October 24, 2023 at 9pm Central Time

---

1. A *six-sided die* (plural *dice*) is a cube with each side marked with a different number of dots (called *pips*) from 1 to 6. On a *standard* die, numbers on opposite sides always add up to 7.

   A *rolling die maze* is a puzzle involving a standard six-sided die and a grid of squares. You should imagine the grid lying on a table; the die always rests on and exactly covers one square of the grid. In a single step, you can *roll* the die 90 degrees around one of its bottom edges, moving it to an adjacent square one step north, south, east, or west.

   Some squares in the grid may be *blocked*; the die can never rest on a blocked square. Other squares may be *labeled* with a number; whenever the die rests on a labeled square, the number on the *top* face of the die must equal the label. Squares that are neither labeled nor marked are called *free*. You may not roll the die off the edges of the grid. A rolling die maze is *solvable* if it is possible to place a die on the lower left square and roll it to the upper right square under these constraints.

   Figure 1. Rolling a (right-handed) die

   Figure 2 shows five rolling die mazes. The first two mazes are solvable using any standard die. Specifically, the first maze can be solved by placing the die on the lower left square with 1 on the top face, and then rolling the die east, north, north, east; the second maze can be solved in 12 moves. The third maze is only solvable using a *right-handed* die, where faces 1, 2, 3 appear in counterclockwise order around a common corner.[1] The last two mazes cannot be solved even with non-standard dice.

   Figure 2. Five rolling die mazes.

   Describe and analyze an algorithm that determines whether a given rolling die maze can be solved with a right-handed standard die. Your input is a two-dimensional array $Label[1..n, 1..n]$, where each entry $Label[i, j]$ stores the label of the square in the ith row and jth column, where 0 means the square is free and -1 means the square is blocked.

   [Hint: You have some freedom in how to place the initial die. There are rolling die mazes that can be solved only if the initial placement is chosen correctly. Describe your solution in high-level language; don't get bogged down in grungy case analysis.]

   ---
   [1]Right-handed dice are more common in the Western hemisphere; left-handed dice are more common in east Asia.

2. The Cheery Hells neighborhood of Sham-Poobanana runs a popular and well-regulated Halloween celebration, attended by thousands of costumed children from all across Poobanana County. To regulate the flood of costumed children, the Cheery Hells Neighborhood Association has designated a walking direction for each stretch of sidewalk.

   After paying the $25 entrance fee, each child receives a map of the neighborhood, in the form of a directed graph G, whose vertices represent houses. Each edge $v \rightarrow w$ indicates that one can walk directly from house v to house w following the designated sidewalk directions. (Anyone caught walking backward along a sidewalk will be ejected from Cheery Hells, without their candy. No refunds.) A special vertex s designates the entrance to Cheery Hells. Children can visit houses as many times as they like, but biometric scanners at every house ensure that each child receives candy only at their *first* visit to each house.

   The children of Cheery Hells have published a secret web site listing the amount of candy that each house in Cheery Hells will give to each visitor.

   Describe and analyze an algorithm to compute the maximum amount of candy that a single child can obtain in a walk through Cheery Hells, starting at the entrance node s. The input to your algorithm is the directed graph G, along with a non-negative integer $v.candy$ for each vertex v, describing the amount of candy the correspinding house gives to each first-time visitor.

   [Hint: Think about two special cases first: (1) Cheery Hells is strongly connected, and (2) Cheery Hells is acyclic. Solving only these two special cases is worth half credit.]

3. *Practice only. Do not submit solutions.*

   One of my daughter's elementary-school math workbooks[2] contains several puzzles of the following type:

   Complete each angle maze below by tracing a path from Start to Finish that has only acute angles.

   Describe and analyze an algorithm to solve arbitrary acute-angle mazes.

   Your input is a connected undirected graph G, whose vertices are points in the plane and whose edges are straight line segments. Edges do not intersect, except at their common endpoints. For example, a drawing of the letter X would have five vertices and four edges, and the first maze above has 18 vertices and 21 edges. You are also given two vertices Start and Finish.

   Your algorithm should return TRUE if G contains a walk from Start to Finish that has only acute angles, and FALSE otherwise. Formally, a walk through G is valid if, for any two

---

[2]Jason Batterson and Shannon Rogers, *Beast Academy Math: Practice 3A*, 2012. See https://www.beastacademy.com/resources/printables.php for several more examples.

consecutive edges $u{\to}v{\to}w$ in the walk, either $\angle uvw = \pi$ (straight) or $0 < \angle uvw < \pi/2$ (acute). Assume you have a subroutine that can determine in $O(1)$ time whether the angle between two given segments is straight, obtuse, right, or acute.

Solved problem

4. Professor McClane takes you out to a lake and hands you three empty jars. Each jar holds a positive integer number of gallons; the capacities of the three jars may or may not be different. The professor then demands that you put exactly k gallons of water into one of the jars (which one doesn't matter), for some integer k, using only the following operations:

   (a) Fill a jar with water from the lake until the jar is full.

   (b) Empty a jar of water by pouring water into the lake.

   (c) Pour water from one jar to another, until either the first jar is empty or the second jar is full, whichever happens first.

   For example, suppose your jars hold 6, 10, and 15 gallons. Then you can put 13 gallons of water into the third jar in six steps:

   • Fill the third jar from the lake.
   • Fill the first jar from the third jar. (Now the third jar holds 9 gallons.)
   • Empty the first jar into the lake.
   • Fill the second jar from the lake.
   • Fill the first jar from the second jar. (Now the second jar holds 4 gallons.)
   • Empty the second jar into the third jar.

   Describe and analyze an efficient algorithm that either finds the smallest number of operations that leave exactly k gallons in any jar, or reports correctly that obtaining exactly k gallons of water is impossible. Your input consists of the capacities of the three jars and the positive integer k. For example, given the four numbers 6, 10, 15, and 13 as input, your algorithm should return the number 6 (the length of the sequence of operations listed above).

   Solution: Let A, B, C denote the capacities of the three jars. We reduce the problem to breadth-first search in a directed graph $G = (V, E)$ defined as follows:

   • $V = \{(a, b, c) \mid 0 \le a \le A \text{ and } 0 \le b \le B \text{ and } 0 \le c \le C\}$. Each vertex corresponds to a possible **configuration** of water in the three jars. There are $(A+1)(B+1)(C+1) = O(ABC)$ vertices altogether.

   • G contains a directed edge $(a, b, c){\to}(a', b'c')$ whenever it is possible to change the first configuration into the second in one step. Specifically, G contains an edge from (a, b, c) to each of the following vertices (except those already equal to (a, b, c)):

      – $(0, b, c)$ and $(a, 0, c)$ and $(a, b, 0)$ — dumping a jar into the lake
      – (A, b, c) and (a, B, c) and (a, b, C) — filling a jar from the lake

$$- \begin{cases} (0, a+b, c) & \text{if } a+b \le B \\ (a+b-B, B, c) & \text{if } a+b \ge B \end{cases} \text{— pouring from jar 1 into jar 2}$$

$$- \begin{cases} (0, b, a+c) & \text{if } a+c \le C \\ (a+c-C, b, C) & \text{if } a+c \ge C \end{cases} \text{— pouring from jar 1 into jar 3}$$

$$- \begin{cases} (a+b, 0, c) & \text{if } a+b \le A \\ (A, a+b-A, c) & \text{if } a+b \ge A \end{cases} \text{— pouring from jar 2 into jar 1}$$

$$- \begin{cases} (a, 0, b+c) & \text{if } b+c \le C \\ (a, b+c-C, C) & \text{if } b+c \ge C \end{cases} \text{— pouring from jar 2 into jar 3}$$

$$- \begin{cases} (a+c, b, 0) & \text{if } a+c \le A \\ (A, b, a+c-A) & \text{if } a+c \ge A \end{cases} \text{— pouring from jar 3 into jar 1}$$

$$- \begin{cases} (a, b+c, 0) & \text{if } b+c \le B \\ (a, B, b+c-B) & \text{if } b+c \ge B \end{cases} \text{— pouring from jar 3 into jar 2}$$

Because each vertex has at most 12 outgoing edges, there are at most $12(A+1) \times (B+1)(C+1) = O(ABC)$ edges altogether.

To solve the jars problem, we need to find the **shortest path** in G from the start vertex $(0,0,0)$ to any target vertex of the form $(k, \cdot, \cdot)$ or $(\cdot, k, \cdot)$ or $(\cdot, \cdot, k)$.

We can compute this shortest path by calling **breadth-first search** starting at $(0,0,0)$, and then examining every target vertex by brute force. If BFS does not visit any target vertex, we report that no legal sequence of moves exists. Otherwise, we find the target vertex closest to $(0,0,0)$ and trace its parent pointers back to $(0,0,0)$ to determine the shortest sequence of moves. The resulting algorithm runs in $O(V+E) = \boldsymbol{O(ABC)}$ **time**.

We can speed up this algorithm by observing that every move leaves at least one jar either completely empty or completely full. Thus, we only need vertices (a, b, c) where either $a = 0$ or $b = 0$ or $c = 0$ or $a = A$ or $b = B$ or $c = C$; no other vertices are reachable from $(0,0,0)$. The number of non-redundant vertices and edges is $O(AB + BC + AC)$. Thus, if we only construct and search the relevant portion of G, the algorithm runs in $\boldsymbol{O(AB + BC + AC)}$ **time**. ∎

Rubric: 10 points: standard graph reduction rubric

- Brute force construction is fine.

 − 1 for calling Dijkstra instead of BFS

- max 8 points for $O(ABC)$ time; scale partial credit.

☙ Homework 9 ☙

Due Tuesday, October 31, 2023 at 9pm

This is the last homework before Midterm 2.

---

1. You are planning a hiking trip in Jasper National Park in British Columbia over winter break. You have a complete map of the park's trails, which indicates that hikers on certain trails have a higher chance of encountering a sasquatch. All visitors to the park are required to purchase a canister of sasquatch repellent. You can safely traverse a high-risk trail segment only by *completely* using up a *full* canister. The park rangers have helpfully installed several refilling stations around the park, where you can refill empty canisters at no cost. The canisters themselves are expensive and heavy, so you can only carry one. The trails are narrow, so each trail segment allows traffic in only one direction.

   You have converted the trail map into a directed graph $G = (V, E)$, whose vertices represent trail intersections, and whose edges represent trail segments. A subset $R \subseteq V$ of the vertices indicate the locations of the *R*epellent *R*efilling stations, and a subset $H \subseteq E$ of the edges are marked as *H*igh-risk. Each edge e is labeled with the length $\ell(e)$ of the corresponding trail segment. Your campsite appears on the map as a vertex $s \in V$, and the visitor center is another vertex $t \in V$.

   (a) Describe and analyze an algorithm that finds the shortest *safe* hike from your campsite s to the visitor center t. Assume there is a refill station at your campsite, and another refill station at the visitor center.

   (b) Describe and analyze an algorithm to decide if you can safely hike from *any* refill station to *any* other refill station. In other words, for *every* pair of vertices u and v in R, is there a safe hike from u to v?

2. You are driving through the back-country roads of Tenkucky, desperately trying to leave the state before the state's annual Halloween Purge begins. Every road in the state is patrolled by a Driving Posse who will let you exercise your god-given right to drive as fast as you damn well please, provided you pay the appropriate speed tax. The faster you traverse any road, the more you have to pay. What's the fastest way to escape the state?

   You have an accurate map of the state, in the form of a directed graph $G = (V, E)$, whose vertices V represent small towns and whose edges E represent one-lane dirt roads between towns.[1] One vertex s is marked as your starting location; a subset $X \subset V$ of vertices are marked as exits. Each edge e has an associated value $\$(e)$ with the following interpretation.

   - If you drive from one end of road e to the other in m minutes, for any positive real number m, then you must pay road e's Driving Posse a speed tax of $\lceil \$(e)/m \rceil$ dollars.

---

[1] Paved roads are far too expensive!

- Equivalently, if you pay road e's Driving Posse a speed tax of d dollars, for any positive integer d, you are allowed to drive the entire length of road e in $\$(e)/d$ minutes, but no less.

In particular, any road you drive on *at all* will cost you *at least* one dollar. Anyone who violates this rule (for example, by running out of money) will be thrown in jail, which means almost certain death in the Purge.

The Driving Posses do not accept coins, credit cards, Venmo, Zelle, or any other mobile payment app—only cold hard American paper currency—and they do not give change. Fortunately, you are starting your journey with a pile of D crisp new $1 bills.

Describe and analyze and algorithm to compute the fastest possible driving route from s to any exit node in X. The input to your algorithm consists of the map $G = (V, E)$, the start vertex s, the exit vertices X, and the positive integer D. Report the running time of your algorithm as a function of the parameters V, E, and D.

3. *Practice only. Do not submit solutions.*

   After a grueling midterm at the See-Bull Center for Commuter Silence, you decide to take the bus home. Since you planned ahead, you have a schedule that lists the times and locations of every stop of every bus in Sham-Poobanana. Unfortunately, no single bus visits both the See-Bull Center and your home; you must change buses at least once. There are exactly b different buses. Each bus starts at 12:00:01AM, makes exactly n stops, and finally stops running at 11:59:59PM. Buses always run exactly on schedule, and you have an accurate watch. Finally, you are far too tired to walk between bus stops.

   (a) Describe and analyze an algorithm to determine a sequence of bus rides that gets you home as early as possible. Your goal is to minimize your *arrival time*, not the time you spend traveling.

   (b) Oh, no! The midterm was held on Halloween, and the streets are infested with zombies! Describe how to modify your algorithm from part (a) to minimize *the total time you spend waiting at bus stops*; you don't care how late you get home or how much time you spend on buses. (Assume you can wait inside the See-Bull Center until your first bus is just about to leave.)

   For both questions, your input consists of the exact time when the midterm ends See-Bull and two arrays $Time[1..b, 1..n]$ and $Stop[1..b, 1..n]$, where $Time[i, j]$ is the scheduled time of the ith bus's jth stop, and $Stop[i, j]$ is the location of that stop. Report the running times of your algorithms as functions of the parameters n and b.

Solved Problems

4. Although we typically speak of "the" shortest path from one vertex to another, a single graph could contain several minimum-length paths with the same endpoints.

Four (of many) equal-length shortest paths.

Describe and analyze an algorithm to compute the *number* of shortest paths from a source vertex s to a target vertex t in an arbitrary directed graph G with weighted edges. Assume that all edge weights are positive and that any necessary arithmetic operations can be performed in $O(1)$ time each.

[*Hint: Compute shortest path distances from s to every other vertex. Throw away all edges that cannot be part of a shortest path from s to another vertex. What's left?*]

Solution: We start by computing shortest-path distances $dist(v)$ from s to v, for every vertex v, using Dijkstra's algorithm. Call an edge $u \to v$ **tight** if $dist(u) + w(u \to v) = dist(v)$. Every edge in a shortest path from s to t must be tight. Conversely, every path from s to t that uses only tight edges has total length $dist(t)$ and is therefore a shortest path!

Let H be the subgraph of all tight edges in G. We can easily construct H in $O(V+E)$ time. Because all edge weights are positive, H is a directed acyclic graph. It remains only to count the number of paths from s to t in H.

For any vertex v, let $NumPaths(v)$ denote the number of paths in H from v to t; we need to compute $NumPaths(s)$. This function satisfies the following simple recurrence:

$$NumPaths(v) = \begin{cases} 1 & \text{if } v = t \\ \displaystyle\sum_{v \to w} NumPaths(w) & \text{otherwise} \end{cases}$$

In particular, if v is a sink but $v \neq t$ (and thus there are no paths from v to t), this recurrence correctly gives us $NumPaths(v) = \sum \varnothing = 0$.

We can memoize this function into the graph itself, storing each value $NumPaths(v)$ at the corresponding vertex v. Since each subproblem depends only on its successors in H, we can compute $NumPaths(v)$ for all vertices v by considering the vertices in reverse topological order, or equivalently, by performing a depth-first search of H starting at s. The resulting algorithm runs in $O(V + E)$ time.

The overall running time of the algorithm is dominated by Dijkstra's algorithm in the preprocessing phase, which runs in **$O(E \log V)$ time**. ∎

---

5. After moving to a new city, you decide to choose a walking route from your home to your new office. Your route must consist of an uphill path (for exercise) followed by a downhill path (to cool down), or just an uphill path, or just a downhill path. But you also want the *shortest* path that satisfies these conditions, so that you actually get to work on time.

   Your input consists of an undirected graph G, whose vertices represent intersections and whose edges represent road segments, along with a start vertex s and a target vertex t. Every vertex v has a value $h(v)$, which is the height of that intersection above sea level, and each edge uv has a value $\ell(uv)$, which is the length of that road segment.

   (a) Describe and analyze an algorithm to find the shortest uphill–downhill walk from s to t. Assume all vertex heights are distinct.

   > **Solution:** We define a new directed graph $G' = (V', E')$ as follows:
   >
   > - $V' = \{v^{\uparrow}, v^{\downarrow} \mid V \in V\}$. Vertex $v^{\uparrow}$ indicates that we are at intersection v moving uphill, and vertex $v^{\downarrow}$ indicates that we are at intersection v moving downhill.
   > - E' is the union of three sets:
   >   - Uphill edges: $\{u^{\uparrow} \rightarrow v^{\uparrow} \mid uv \in E \text{ and } h(u) < h(v)\}$. Each uphill edge $u^{\uparrow} \rightarrow v^{\uparrow}$ has weight $\ell(uv)$.
   >   - Downhill edges: $\{u^{\downarrow} \rightarrow v^{\downarrow} \mid uv \in E \text{ and } h(u) > h(v)\}$. Each downhill edge $u^{\downarrow} \rightarrow v^{\downarrow}$ has weight $\ell(uv)$.
   >   - Switch edges: $\{v^{\uparrow} \rightarrow v^{\downarrow} \mid v \in V\}$; each switch edge has weight 0.
   >
   > We need to compute three shortest paths in this graph:
   >
   > - The shortest path from $s^{\uparrow}$ to $t^{\downarrow}$ gives us the best uphill-then-downhill route.
   > - The shortest path from $s^{\uparrow}$ to $t^{\uparrow}$ gives us the best uphill-only route.
   > - The shortest path from $s^{\downarrow}$ to $t^{\downarrow}$ gives us the best downhill-only route.
   >
   > G' is a directed **acyclic** graph; we can get a topological ordering by listing the up vertices $v^{\uparrow}$, sorted by increasing height, followed by the down vertices $v^{\downarrow}$, sorted by decreasing height. Thus, we can compute the shortest path in G' from any vertex to any other in $O(V' + E') = O(V + E)$ time by dynamic programming. (The algorithm is the same as the longest-path algorithm in the notes, except we use "min" instead of "max" in the recurrence, and define $\min \emptyset = \infty$.)
   >
   > Our overall algorithm runs in $O(V + E)$ **time**. ∎

(b) Suppose you discover that there is no path from s to t with the structure you want. Describe an algorithm to find a path from s to t that alternates between "uphill" and "downhill" subpaths as few times as possible, and has minimum length among all such paths. (There may be even shorter paths with more alternations, but you don't care about them.) Again, assume all vertex heights are distinct.

> **Solution (Dijkstra, 5/5):** Let $L = 1 + \sum_{u \to v} \ell(u \to v)$. Define a new graph $G' = (V', E')$ as follows:
>
> - $V' = \{v^\uparrow, v^\downarrow \mid V \in V\} \cup \{s, t\}$. Vertex $v^\uparrow$ indicates that we are at intersection v moving uphill, and vertex $v^\downarrow$ indicates that we are at intersection v moving downhill.
> - E' contains four types of edges:
>   - Uphill edges: $\{u^\uparrow \to v^\uparrow \mid uv \in E \text{ and } h(u) < h(v)\}$. Each uphill edge $u^\uparrow \to v^\uparrow$ has weight $\ell(uv)$.
>   - Downhill edges: $\{u^\downarrow \to v^\downarrow \mid uv \in E \text{ and } h(u) > h(v)\}$. Each downhill edge $u^\downarrow \to v^\downarrow$ has weight $\ell(uv)$.
>   - Switch edges: $\{v^\uparrow \to v^\downarrow \mid v \in V\} \cup \{v^\downarrow \to v^\uparrow \mid v \in V\}$. Each switch edge has weight L.
>   - Start and end edges $s \to s^\uparrow$, $s \to s^\downarrow$, $t^\uparrow \to t$, and $t^\downarrow \to t$, each with weight 0,
>
> We need to compute the shortest path from s to t in G'; the large weight L on the switch edges guarantees that this path with have the minimum number of switches, and the minimum length among all paths with that number of switches. Dijkstra's algorithm finds this shortest path in $O(E' \log V') = O(E \log V)$ *time*.
>
>   (Because G' includes switch edges in both directions, G' is not a dag, so we can't use dynamic programming directly.) ∎

> **Rubric:** 5 points, standard graph-reduction rubric. This is not the only correct solution with running time $O(E \log V)$.

> **Solution (clever, extra credit):** Our algorithm works in two phases: First we determine the minimum number of switches required to reach every vertex, and then we compute the shortest path from s to t with the minimum number of switches. The first phase is can be solved in $O(V + E)$ time by a modification of breadth-first search; the second by computing shortest paths in a dag.
>
> For the first phase, we define a new graph $G' = (V', E')$ as follows:
>
> - $V' = \{v^\uparrow, v^\downarrow \mid V \in V\} \cup \{s, t\}$. Vertex $v^\uparrow$ indicates that we are at intersection v moving uphill, and vertex $v^\downarrow$ indicates that we are at intersection v moving downhill.
> - E' contains four types of edges:
>   - Uphill edges: $\{u^\uparrow \to v^\uparrow \mid uv \in E \text{ and } h(u) < h(v)\}$. Each uphill edge has weight 0.
>   - Downhill edges: $\{u^\downarrow \to v^\downarrow \mid uv \in E \text{ and } h(u) > h(v)\}$. Each downhill

edge has weight 0.
  – Switch edges: $\{v^\uparrow \to v^\downarrow \mid v \in V\} \cup \{v^\downarrow \to v^\uparrow \mid v \in V\}$. Each switch edge has weight 1.
  – Start and end edges $s \to s^\uparrow$, $s \to s^\downarrow$, $t^\uparrow \to t$, and $t^\downarrow \to t$, each with weight 0.

Now we compute the shortest path distance from s to every other vertex in G'. We could use Dijkstra's algorithm in $O(E \log V)$ time, but the structure of the graph supports a faster algorithm.

Intuitively, we break the shortest-path computation into phases, where in the kth phase, we mark all vertices at distance k from the source vertex s. During the kth phase, we may also discover vertices at distance $k + 1$, but no further. So instead of using a binary heap for the priority queue, it suffices to use two bags: one for vertices at distance k, and one for vertices at distance $k + 1$.

$\underline{\text{ZeroOneDijkstra}(G, \ell, s):}$
  $s.dist \leftarrow 0$
  for all vertices $v \neq s$
    $v.dist \leftarrow \infty$
  $curr \leftarrow$ new empty bag
  add s to $curr$
  for $k \leftarrow 0$ to V
    $next \leftarrow$ new empty bag
    while $curr$ is not empty
      take v from $curr$ ⟨⟨$v.dist = k$⟩⟩
      for all edges $v \to w$
        if $w.dist > v.dist + \ell(v \to w)$
          $w.dist \leftarrow v.dist + \ell(v \to w)$
          if $\ell(v \to w) = 0$
            add w to $curr$
          else ⟨⟨$if \; \ell(v \to w) = 1$⟩⟩
            add w to $next$
    $curr \leftarrow next$

This phase of the algorithm runs in $O(V' + E') = O(V + E)$ time.

Once we have computed distances in G', we construct a second graph $G'' = (V', E'')$ with the same vertices as G', but only a subset of the edges:

$$E'' = \left\{ u' \to v' \in E' \mid u'.dist + \ell(u' \to v') = v'.dist \right\}$$

Equivalently, an edge $u' \to v'$ belongs to E'' if and only if that edge is part of at least one shortest path in G' from s to another vertex. It follows (by induction, of course), that every path in G'' from s to another vertex v' is a shortest path in G', and therefore a minimum-switch path in G.

We also reassign the edge weights in G''. Specifically, we assign each uphill edge $u^\uparrow \to v^\uparrow$ and downhill edge $u^\downarrow \to v^\downarrow$ in G'' weight $\ell(uv)$, and we assign every switch edge, start edge, and end edge weight 0. **Now we need to compute the shortest path from s to t in G'', with respect to these new edge weights.**

We can expand the definition of E'' in terms of the original input graph as follows:

$$
\begin{aligned}
E'' = \ & \big\{ u^{\uparrow} {\to} v^{\uparrow} \ \big| \ uv \in E \text{ and } h(u) < h(v) \text{ and } u^{\uparrow}.dist = v^{\uparrow}.dist \big\} \\
& \cup \ \big\{ u^{\downarrow} {\to} v^{\downarrow} \ \big| \ uv \in E \text{ and } h(u) > h(v) \text{ and } u^{\downarrow}.dist = v^{\downarrow}.dist \big\} \\
& \cup \ \big\{ v^{\uparrow} {\to} v^{\downarrow} \ \big| \ v \in V \text{ and } v^{\uparrow}.dist < v^{\downarrow}.dist \big\} \\
& \cup \ \big\{ v^{\downarrow} {\to} v^{\uparrow} \ \big| \ v \in V \text{ and } v^{\downarrow}.dist < v^{\uparrow}.dist \big\}
\end{aligned}
$$

We can topologically sort G'' by first sorting the vertices by increasing $v'.dist$, and then within each subset of vertices with equal $v'.dist$, listing the up-vertices by increasing height, followed by the down vertices by decreasing height. It follows that G'' is a dag! Thus, we can compute shortest paths in G'' in $O(V'' + E'') = O(V + E)$ time, using the same dynamic programming algorithm that we used in part (a).

The overall algorithm runs in $O(V + E)$ **time**. ∎

Rubric: max 10 points =

- 5 for computing minimum-switch paths = 1 for vertices + 1 for edges (including weights) + 2 for 0/1 shortest path algorithm + 1 for running time.
- 5 for computing shortest minimum-switch paths = 1 for vertices + 1 for edges (including weights) + 1 for proving dag + 1 for dynamic programming algorithm + 1 for running time

---

1. This problem asks you to describe polynomial-time reductions between two closely related problems:

   - SUBSETSUM: Given a set S of positive integers and a target integer T, is there a subset of S whose sum is T?

   - PARTITION: Given a set S of positive integers, is there a way to partition S into two subsets S_1 and S_2 that have the same sum?

   (a) Describe a polynomial-time reduction from SUBSETSUM to PARTITION.

   (b) Describe a polynomial-time reduction from PARTITION to SUBSETSUM.

   Don't forget to to prove that your reductions are correct.

2. A subset S of vertices in an undirected graph G is called *triangle-free* if, for every triple of vertices $u, v, w \in S$, at least one of the three edges uv, uw, vw is *absent* from G. Prove that finding the size of the largest triangle-free subset of vertices in a given undirected graph is NP-hard.

A triangle-free subset of 7 vertices and its induced edges.
This is **not** the largest triangle-free subset in this graph.

Solved Problem

4. **RedBlue** is a puzzle that consists of an $n \times m$ grid of squares, where each square may be empty, occupied by a red stone, or occupied by a blue stone. The goal of the puzzle is to remove some of the given stones so that the remaining stones satisfy two conditions:

   (1) Every row contains at least one stone.

   (2) No column contains stones of both colors.

   For some RedBlue puzzles, reaching this goal is impossible; see the example below.

   Prove that it is NP-hard to determine whether a given RedBlue puzzle has a solution.

A solvable RedBlue puzzle and one of its many solutions.   An unsolvable RedBlue puzzle.

Solution: We show that RedBlue is NP-hard by describing a reduction from 3SAT.

Let Φ be a 3CNF boolean formula with m variables and n clauses. We transform this formula into a RedBlue instance X in polynomial time as follows. The size of the board is $n \times m$. The stones are placed as follows, for all indices i and j:

- If the variable x_j appears in the ith clause of Φ, we place a blue stone at (i, j).
- If the negated variable $\overline{x_j}$ appears in the ith clause of Φ, we place a red stone at (i, j).
- Otherwise, we leave cell (i, j) blank.

To prove that RedBlue is NP-hard, it suffices to prove the following claim:

> Φ is satisfiable
> if and only if
> RedBlue puzzle X is solvable.

$\implies$ First, suppose Φ is satisfiable; consider an arbitrary satisfying assignment. For each index j, remove stones from column j according to the value assigned to x_j:

- If $x_j = $ True, remove all red stones from column j.
- If $x_j = $ False, remove all blue stones from column j.

In other words, remove precisely the stones that correspond to False literals. Because every variable appears in at least one clause, each column now contains stones of only one color (if any). On the other hand, each clause of Φ must contain at least one True literal, and thus each row still contains at least one stone. We conclude that RedBlue puzzle X is solvable.

⟸ On the other hand, suppose RedBlue puzzle X is solvable; consider an arbitrary solution. For each index j, assign a value to x_j depending on the colors of stones left in column j:

- If column j contains blue stones, set $x_j = $ TRUE.
- If column j contains red stones, set $x_j = $ FALSE.
- If column j is empty, set x_j arbitrarily.

In other words, assign values to the variables so that the literals corresponding to the remaining stones are all TRUE. Each row still has at least one stone, so each clause of Φ contains at least one TRUE literal, so this assignment makes $\Phi = $ TRUE. We conclude that Φ is satisfiable.

This reduction clearly requires only polynomial time.   ∎

Standard NP-hardness rubric.   10 points =

+ 1 point for choosing a reasonable NP-hard problem X to reduce from.
  - The Cook-Levin theorem implies that *in principle* one can prove NP-hardness by reduction from *any* NP-complete problem. What we're looking for here is a problem where a simple and direct NP-hardness proof seems likely.
  - You can use any of the NP-hard problems listed on the next page or in the textbook (except the one you are trying to prove NP-hard, of course).
+ 2 points for a *structurally sound* polynomial-time reduction. Specifically, the reduction must:
  - take an *arbitrary* instance of the declared problem X **and nothing else** as input,
  - transform that input into a corresponding instance of Y (the problem we're trying to prove NP-hard),
  - transform the output of the magic algorithm for Y into a reasonable output for X, and
  - run in polynomial time.

  (The output transformation is usually trivial.) This is strictly about the structure of the reduction algorithm, not about its correctness. **No credit for the rest of the problem if this is wrong.**
+ 2 points for a *correct* polynomial-time reduction. That is, assuming a black-box algorithm that solves Y in polynomial time, the proposed reduction actually solves problem X in polynomial time.
+ 2 points for the "if" proof of correctness. (Every good instance of X is transformed into a good instance of Y.)
+ 2 points for the "only if" proof of correctness. (Every bad instance of X is transformed into a bad instance of Y.)
+ 1 point for writing "polynomial time"

- An incorrect but structurally sound polynomial-time reduction that still satisfies half of the correctness proof is worth at most 6/10.
- A reduction in the wrong direction is worth at most 1/10.

Some useful NP-hard problems. You are welcome to use any of these in your own NP-hardness proofs, except of course for the specific problem you are trying to prove NP-hard.

CIRCUITSAT: Given a boolean circuit, are there any input values that make the circuit output TRUE?

3SAT: Given a boolean formula in conjunctive normal form, with exactly three distinct literals per clause, does the formula have a satisfying assignment?

MAXINDEPENDENTSET: Given an undirected graph G, what is the size of the largest subset of vertices in G that have no edges among them?

MAXCLIQUE: Given an undirected graph G, what is the size of the largest complete subgraph of G?

MINVERTEXCOVER: Given an undirected graph G, what is the size of the smallest subset of vertices that touch every edge in G?

MINSETCOVER: Given a collection of subsets $S_1, S_2, \ldots, S_m$ of a set S, what is the size of the smallest subcollection whose union is S?

MINHITTINGSET: Given a collection of subsets $S_1, S_2, \ldots, S_m$ of a set S, what is the size of the smallest subset of S that intersects every subset S_i?

3COLOR: Given an undirected graph G, can its vertices be colored with three colors, so that every edge touches vertices with two different colors?

CHROMATICNUMBER: Given an undirected graph G, what is the minimum number of colors required to color its vertices, so that every edge touches vertices with two different colors?

HAMILTONIANPATH: Given graph G (either directed or undirected), is there a path in G that visits every vertex exactly once?

HAMILTONIANCYCLE: Given a graph G (either directed or undirected), is there a cycle in G that visits every vertex exactly once?

TRAVELINGSALESMAN: Given a graph G (either directed or undirected) with weighted edges, what is the minimum total weight of any Hamiltonian path/cycle in G?

LONGESTPATH: Given a graph G (either directed or undirected, possibly with weighted edges), what is the length of the longest simple path in G?

STEINERTREE: Given an undirected graph G with some of the vertices marked, what is the minimum number of edges in a subtree of G that contains every marked vertex?

SUBSETSUM: Given a set X of positive integers and an integer k, does X have a subset whose elements sum to k?

PARTITION: Given a set X of positive integers, can X be partitioned into two subsets with the same sum?

3PARTITION: Given a set X of $3n$ positive integers, can X be partitioned into n three-element subsets, all with the same sum?

INTEGERLINEARPROGRAMMING: Given a matrix $A \in \mathbb{Z}^{n \times d}$ and two vectors $b \in \mathbb{Z}^n$ and $c \in Z^d$, compute $\max\{c \cdot x \mid Ax \leq b, x \geq 0, x \in \mathbb{Z}^d\}$.

FEASIBLEILP: Given a matrix $A \in \mathbb{Z}^{n \times d}$ and a vector $b \in \mathbb{Z}^n$, determine whether the set of feasible integer points $\max\{x \in \mathbb{Z}^d \mid Ax \leq b, x \geq 0\}$ is empty.

DRAUGHTS: Given an $n \times n$ international draughts configuration, what is the largest number of pieces that can (and therefore must) be captured in a single move?

SUPERMARIOBROTHERS: Given an $n \times n$ Super Mario Brothers level, can Mario reach the castle?

This is the last graded homework before the final exam.

---

1. A *balloon* of size ℓ is an undirected graph consisting of a (simple) cycle of length ℓ and a (simple) path of length ℓ, where one endpoint of the path lies on the cycle, and otherwise the cycle and the path are disjoint. Every balloon of size ℓ has exactly 2ℓ vertices and 2ℓ edges. For example, the 4×4 grid graph shown below contains a balloon subgraph of size 8.

   Prove that it is NP-hard to find the size of the the largest balloon subgraph of a given undirected graph.

2. Recall that a *3-coloring* of a graph assigns each vertex one of three colors, say red, yellow, and blue. A 3-coloring is *proper* if every edge has endpoints with different colors. The 3COLOR problem asks, given an arbitrary undirected graph G, whether G has a proper 3-coloring.

   Call a 3-coloring of a graph G *slightly improper* if each vertex has *at most one neighbor* with the same color. The SLIGHTLYIMPROPER3COLOR problem asks, given an arbitrary undirected graph G, whether G has a slightly improper 3-coloring.

   (a) Consider the following attempt to prove that SLIGHTLYIMPROPER3COLOR is NP-hard, using a reduction from 3COLOR.

   > **Non-solution:** We reduce from 3COLOR. Given an arbitrary input graph G, we construct a new graph H by attaching a clique of 4 vertices to every vertex of G. Specifically, for each vertex v in G, the graph H contains three new vertices v_1, v_2, v_3, along with edges $vv_1, vv_2, vv_3, v_1v_2, v_1v_3, v_2v_3$. I claim that

> *G* has a proper 3-coloring
> if and only if
> *H* has a slightly improper 3-coloring.
>
> $\Longrightarrow$ Suppose *G* has a proper 3-coloring, using the colors red, yellow, and blue.
> Extend this color assignment to the vertices of *H* by coloring each vertex v_1
> red, each vertex v_2 yellow, and each vertex v_3 blue. With this assignment,
> each vertex of *H* has at most one neighbor with the same color. Specifically,
> each vertex of *G* has the same color as one of the vertices in its gadget, and
> the other two vertices in *v*'s gadget have no neighbors with the same color.
>
> $\Longleftarrow$ Now suppose *H* has a slightly improper 3-coloring. Then *G* must have a
> proper 3-coloring because. . . um. . .
>
> ♣

Describe a graph *G* that does not have a proper 3-coloring, such that the graph *H*
constructed by this reduction does have a slightly improper 3-coloring.

(b) Describe a small graph *X* with the following property: In every slightly improper
3-coloring of *X*, every vertex of *X* has *exactly* one neighbor with the same color.

(c) Describe a correct polynomial-time reduction from 3Color to SlightlyImproper-
3Color. *[Hint: Use your graph from part (b) as a gadget.]* This reduction will prove
that SlightlyImproper3Color is indeed NP-hard.

Solved Problem

3. A *double-Hamiltonian tour* in an undirected graph G is a closed walk that visits every vertex in G exactly twice. Prove that it is NP-hard to decide whether a given graph G has a double-Hamiltonian tour.

This graph contains the double-Hamiltonian tour $a \to b \to d \to g \to e \to b \to d \to c \to f \to a \to c \to f \to g \to e \to a$.

Solution: We prove the problem is NP-hard with a reduction from the standard Hamiltonian cycle problem. Let G be an arbitrary undirected graph. We construct a new graph H by attaching a small gadget to every vertex of G. Specifically, for each vertex v, we add two vertices $v^\sharp$ and $v^\flat$, along with three edges $vv^\flat$, $vv^\sharp$, and $v^\flat v^\sharp$.

A vertex in G and the corresponding vertex gadget in H.

Now I claim that

> G has a Hamiltonian cycle
> if and only if
> H has a double-Hamiltonian tour.

$\implies$ Suppose G contains a Hamiltonian cycle $C = v_1 \to v_2 \to \cdots \to v_n \to v_1$. We can construct a double-Hamiltonian tour of H by replacing each vertex v_i in C with the following walk:

$$\cdots \to v_i \to v_i^\flat \to v_i^\sharp \to v_i^\flat \to v_i^\sharp \to v_i \to \cdots$$

$\impliedby$ Conversely, suppose H has a double-Hamiltonian tour D. Consider any vertex v in the original graph G; the tour D must visit v exactly twice. Those two visits split D into two closed walks, each of which visits v exactly once. Any walk from $v^\flat$ or $v^\sharp$ to any other vertex in H must pass through v. Thus, one of the two closed walks visits only the vertices v, $v^\flat$, and $v^\sharp$. Thus, if we remove the vertices and edges in $H \setminus G$ from D, we obtain a closed walk in G that visits every vertex in G exactly once.

Given any graph G, we can clearly construct the corresponding graph H in polynomial time by brute force.


  With more effort, we can construct a graph H that contains a double-Hamiltonian tour ***that traverses each edge of H at most once*** if and only if G contains a Hamiltonian

cycle. For each vertex v in G we attach a more complex gadget containing five vertices and eleven edges, as shown on the next page.

A vertex in G, and the corresponding modified vertex gadget in H.

■

Rubric: 10 points, standard polynomial-time reduction rubric. This is not the only correct solution.

Non-solution (self-loops): We attempt to prove the problem is NP-hard with a reduction from the Hamiltonian cycle problem. Let G be an arbitrary undirected graph. We construct a new graph H by attaching a self-loop every vertex of G. Given any graph G, we can clearly construct the corresponding graph H in polynomial time.


An incorrect vertex gadget.

Now I claim that

> G has a Hamiltonian cycle
> if and only if
> H has a double-Hamiltonian tour.

$\implies$ Suppose G has a Hamiltonian cycle $v_1 \rightarrow v_2 \rightarrow \cdots \rightarrow v_n \rightarrow v_1$. We can construct a double-Hamiltonian tour of H by alternating between edges of the Hamiltonian cycle and self-loops: $v_1 \rightarrow v_1 \rightarrow v_2 \rightarrow v_2 \rightarrow v_3 \rightarrow \cdots \rightarrow v_n \rightarrow v_n \rightarrow v_1$.

$\overset{}{\not\Longleftarrow}$ Um...

Unfortunately, if H has a double-Hamiltonian tour, we *cannot* conclude that G has a Hamiltonian cycle, because we cannot guarantee that a double-Hamiltonian tour in H uses *any* self-loops. The graph G shown below is a counterexample; it has a double-Hamiltonian tour (even before adding self-loops!) but no Hamiltonian cycle.

This graph has a double-Hamiltonian tour.

♣

Some useful NP-hard problems. You are welcome to use any of these in your own NP-hardness proofs, except of course for the specific problem you are trying to prove NP-hard.

CIRCUITSAT: Given a boolean circuit, are there any input values that make the circuit output TRUE?

3SAT: Given a boolean formula in conjunctive normal form, with exactly three distinct literals per clause, does the formula have a satisfying assignment?

MAXINDEPENDENTSET: Given an undirected graph G, what is the size of the largest subset of vertices in G that have no edges among them?

MAXCLIQUE: Given an undirected graph G, what is the size of the largest complete subgraph of G?

MINVERTEXCOVER: Given an undirected graph G, what is the size of the smallest subset of vertices that touch every edge in G?

MINSETCOVER: Given a collection of subsets $S_1, S_2, \ldots, S_m$ of a set S, what is the size of the smallest subcollection whose union is S?

MINHITTINGSET: Given a collection of subsets $S_1, S_2, \ldots, S_m$ of a set S, what is the size of the smallest subset of S that intersects every subset S_i?

3COLOR: Given an undirected graph G, can its vertices be colored with three colors, so that every edge touches vertices with two different colors?

CHROMATICNUMBER: Given an undirected graph G, what is the minimum number of colors required to color its vertices, so that every edge touches vertices with two different colors?

HAMILTONIANPATH: Given graph G (either directed or undirected), is there a path in G that visits every vertex exactly once?

HAMILTONIANCYCLE: Given a graph G (either directed or undirected), is there a cycle in G that visits every vertex exactly once?

TRAVELINGSALESMAN: Given a graph G (either directed or undirected) with weighted edges, what is the minimum total weight of any Hamiltonian path/cycle in G?

LONGESTPATH: Given a graph G (either directed or undirected, possibly with weighted edges), what is the length of the longest simple path in G?

STEINERTREE: Given an undirected graph G with some of the vertices marked, what is the minimum number of edges in a subtree of G that contains every marked vertex?

SUBSETSUM: Given a set X of positive integers and an integer k, does X have a subset whose elements sum to k?

PARTITION: Given a set X of positive integers, can X be partitioned into two subsets with the same sum?

3PARTITION: Given a set X of $3n$ positive integers, can X be partitioned into n three-element subsets, all with the same sum?

INTEGERLINEARPROGRAMMING: Given a matrix $A \in \mathbb{Z}^{n \times d}$ and two vectors $b \in \mathbb{Z}^n$ and $c \in Z^d$, compute $\max\{c \cdot x \mid Ax \leq b, x \geq 0, x \in \mathbb{Z}^d\}$.

FEASIBLEILP: Given a matrix $A \in \mathbb{Z}^{n \times d}$ and a vector $b \in \mathbb{Z}^n$, determine whether the set of feasible integer points $\max\{x \in \mathbb{Z}^d \mid Ax \leq b, x \geq 0\}$ is empty.

DRAUGHTS: Given an $n \times n$ international draughts configuration, what is the largest number of pieces that can (and therefore must) be captured in a single move?

SUPERMARIOBROTHERS: Given an $n \times n$ Super Mario Brothers level, can Mario reach the castle?

CS/ECE 374 A ✧ Fall 2023

ൟ "Homework" 12 ൟ

"Due" Monday, December 4, 2023

This homework is *not* for submission. However, we are planning to ask a few (true/false, multiple-choice, or short-answer) questions about undecidability on the final exam, so we still strongly recommend treating these questions as regular homework. Solutions will be released next Monday.

---

1. Let $\langle M \rangle$ denote the encoding of a Turing machine M (or if you prefer, the Python source code for the executable code M). Recall that w^R denotes the reversal of string w. Prove that the following language is undecidable.

$$\text{SELFREVACCEPT} := \left\{ \langle M \rangle \mid M \text{ accepts the string } \langle M \rangle^R \right\}$$

Note that Rice's theorem does *not* apply to this language.

2. Let M be a Turing machine, let w be a string, and let s be an integer. We say that M *accepts w in space s* if, given w as input, M accesses **at most** the first s cells on its tape and eventually accepts. (If you prefer to think in terms of programs instead of Turing machines, "space" is how much memory your program needs to run correctly.)

   Prove that the following language is undecidable:

   $$\text{SOMESQUARESPACE} = \left\{ \langle M \rangle \mid M \text{ accepts at least one string } w \text{ in space } |w|^2 \right\}$$

   Note that Rice's theorem does *not* apply to this language.

   [Hint: The only thing you actually need to know about Turing machines for this problem is that they consume a resource called "space".]

3. Prove that the following language is undecidable:

   $$\text{PICKY} = \left\{ \langle M \rangle \;\middle|\; \begin{array}{l} M \text{ accepts at least one input string} \\ \text{and } M \text{ rejects at least one input string} \end{array} \right\}$$

   Note that Rice's theorem does *not* apply to this language.

Solved Problem

4. Consider the language SOMETIMESHALT $= \{\langle M \rangle \mid M$ halts on at least one input string$\}$. Note that $\langle M \rangle \in$ SOMETIMESHALT does not imply that M *accepts* any strings; it is enough that M *halts* on (and possibly rejects) some string.

   (a) Prove that SOMETIMESHALT is undecidable.

   > **Solution (Rice):** Let $\mathscr{L}$ be the family of all non-empty languages. Let N be any Turing machine that never halts, so HALT$(N) = \varnothing \notin \mathscr{L}$. Let Y be any Turing machine that always halts, so HALT$(Y) = \Sigma^* \in \mathscr{L}$. Rice's Halting Theorem immediately implies that SOMETIMESHALT $=$ HALTIN$(\mathscr{L})$ is undecidable. ∎

   > **Solution (closure):** Let ENCODINGS be the language of all Turing machine encodings (for some fixed universal Turing machine); this language is decidable. We immediately have ENCODINGS $=$ NEVERHALT $\cup$ SOMETIMESHALT, or equivalently, NEVERHALT $=$ ENCODINGS $\setminus$ SOMETIMESHALT.
   >
   > The lectures notes include a proof that NEVERHALT is undecidable. On the other hand, the existence of a universal Turing machine implies that ENCODINGS is decidable. So Corollary 3(d) in the undecidability notes implies that SOMETIMESHALT is undecidable. ∎

   > **Solution (reduction from HALT):** We can reduce the standard halting problem to SOMETIMESHALT as follows:
   >
   > <u>DECIDEHALT$(\langle M, w \rangle)$:</u>
   >   Encode the following Turing machine M':
   >
   > > <u>$M'(x)$:</u>
   > >   (ignore x)
   > >   run M on input w
   >
   >   return DECIDESOMETIMESHALT$(\langle M' \rangle)$
   >
   > We prove this reduction correct as follows:
   >
   > $\Longrightarrow$ Suppose M halts on input w.
   >   Then M' halts on *every* input string x.
   >   So DECIDESOMETIMESHALT must accept the encoding $\langle M' \rangle$.
   >   We conclude that DECIDEHALT correctly accepts the encoding $\langle M, w \rangle$.
   >
   > $\Longleftarrow$ Suppose M does not halt on input w.
   >   Then M' diverges on *every* input string x.
   >   So DECIDESOMETIMESHALT must reject the encoding $\langle M' \rangle$.
   >   We conclude that DECIDEHALT correctly rejects the encoding $\langle M, w \rangle$.
   >
   > ∎

> You have 120 minutes to answer five questions.
> **Write your answers in the separate answer booklet.**
> Please return this question sheet and your cheat sheet with your answers.

1. Let compress0s(w) be a function that takes a string w as input, and returns the string formed by compressing every run of 0s in w by half. Specifically, every run of $2n$ 0s is compressed to length n, and every run of $2n + 1$ 0s is compressed to length $n + 1$. For example:

$$\text{compress0s}(\underline{00000}110\underline{001}) = \underline{000}11\underline{001}$$
$$\text{compress0s}(11\underline{000001}0) = 11\underline{0010}$$
$$\text{compress0s}(11111) = 11111$$

   Let L be an arbitrary regular language.

   (a) **Prove** that $\{w \in \Sigma^* \mid \text{compress0s}(w) \in L\}$ is regular.
   (b) **Prove** that $\{\text{compress0s}(w) \mid w \in L\}$ is regular.

2. For each of the following languages L over the alphabet $\Sigma = \{0, 1\}$, describe a DFA that accepts L **and** give a regular expression that represents L. You do not need to justify your answers.

   (a) All strings in which at least one run has length divisible by 3.
   (b) All strings that do not contain either 100 or 011 as a substring.

3. Consider the following recursive function Bond, which doubles the length of any run of 0s in its input string.

$$\text{Bond}(w) := \begin{cases} \varepsilon & \text{if } w = \varepsilon \\ 00 \cdot \text{Bond}(x) & \text{if } w = 0 \cdot x \text{ for some string } x \\ 1 \cdot \text{Bond}(x) & \text{if } w = 1 \cdot x \text{ for some string } x \end{cases}$$

   (a) **Prove** that $|\text{Bond}(w)| \geq |w|$ for all strings w.
   (b) **Prove** that $\text{Bond}(x \bullet y) = \text{Bond}(x) \bullet \text{Bond}(y)$ for all strings x and y.

   As usual, you can assume any result proved in class, in the lecture notes, in labs, or in homework solutions.

4. Let L be the language $\left\{0^a 1^b 0^c \mid a = b \text{ or } a = c \text{ or } b = c\right\}$

   (a) **Prove** that L is *not* a regular language.

   (b) Describe a context-free grammar for L.

5. For each statement below, check "True" if the statement is always true and check "False" otherwise, and give a brief (one short sentence) explanation of your answer. Read these statements very carefully—small details matter!

   (a) If $2 + 2 = 5$, then zero is odd.

   (b) $\{0^n 1 \mid n > 0\}$ is the only infinite fooling set for the language $\{0^n 1 0^n \mid n > 0\}$.

   (c) $\{0^n 1 0^n \mid n > 0\}$ is a context-free language.

   (d) The context-free grammar $S \rightarrow 00S \mid S11 \mid 01$ generates the language $0^n 1^n$.

   (e) Every regular language is recognized by a DFA with exactly one accepting state.

   (f) Any language that can be decided by an NFA with ε-transitions can also be decided by an NFA without ε-transitions.

   (g) If L is a regular language over the alphabet $\{0, 1\}$, then $\{x y^C \mid x, y \in L\}$ is also regular.

   (h) If L is a regular language over the alphabet $\{0, 1\}$, then $\{w w^C \mid w \in L\}$ is also regular.

   (i) The regular expression $(00 + 11)^*$ represents the language of all strings over $\{0, 1\}$ of even length.

   (j) Let L_1, L_2 be two regular languages. The language $(L_1 + L_2)^*$ is also regular.

1. For each of the following languages L over the alphabet $\Sigma = \{0, 1\}$, describe a DFA that accepts L **and** give a regular expression that represents L. You do not need to justify your answers.

   (a) All strings in which the number of runs is divisible by 3. (Recall that a *run* is a maximal non-empty substring where all symbols are equal.)

   (b) All strings that do not contain the substring 0110.

2. Let take2skip2(w) be a function takes an input string w and returns the subsequence of symbols at positions $1, 2, 5, 6, 9, 10, \ldots 4i+1, 4i+2, \ldots$ in w. In other words, take2skip2(w) takes the first two symbols of w, skip the next two, takes the next two, skips the next two, and so on. For example:

$$\text{take2skip2}(\underline{1}) = 1$$
$$\text{take2skip2}(\underline{01}0) = 01$$
$$\text{take2skip2}(\underline{01}00\underline{11}11\underline{00}011) = 0111001$$

   Let L be an arbitrary regular language.

   (a) **Prove** that the language $\{w \in \Sigma^* \mid \text{take2skip2}(w) \in L\}$ is regular.

   (b) **Prove** that the language $\{\text{take2skip2}(w) \mid w \in L\}$ is regular.

3. Consider the following recursive function censor, which deletes all 1s in its input string.

$$\text{censor}(w) := \begin{cases} \varepsilon & \text{if } w = \varepsilon \\ \text{censor}(x) & \text{if } w = 0 \cdot x \text{ for some string } x \\ 1 \cdot \text{censor}(x) & \text{if } w = 1 \cdot x \text{ for some string } x \end{cases}$$

   (a) **Prove** that $|\text{censor}(w)| \leq |w|$ for all strings w.

   (b) **Prove** that $\text{censor}(\text{censor}(w)) = \text{censor}(w)$ for all strings w.

   As usual, you can assume any result proved in class, in the lecture notes, in labs, or in homework solutions.

4. Consider the language $L = \left\{ 0^a 1^b \mid a > 2b \text{ or } 2a < b \right\}$

    (a) ***Prove*** that L is *not* a regular language.

    (b) Describe a context-free grammar for L.

5. For each statement below, check "True" if the statement is always true and check "False" otherwise, and give a brief (one short sentence) explanation of your answer. Read these statements very carefully—small details matter!

    (a) For every language L, the language L^* is infinite.

    (b) If a language L is finite, the complement of L is context-free.

    (c) The language $\{0^{374n} \mid n \geq 374\}$ is regular.

    (d) The language $\left\{ wxw^R \mid w, x \in \Sigma^* \right\}$ is regular.

    (e) The context-free grammar $S \rightarrow 0S1S \mid S1S0 \mid \varepsilon$ generates the set of all binary strings with the same number of 0s and 1s.

    (f) Every regular language is recognized by a DFA with at least 374 states.

    (g) If the languages L and L' are regular, their intersection $L \cap L'$ is also regular.

    (h) If a language has an infinite fooling set, then it is context-free.

    (i) Let M be a **DFA** over the alphabet Σ. Let M' be identical to M, except that accepting states in M are non-accepting in M' and vice versa. Each string in Σ^* is accepted by exactly one of M and M'.

    (j) Let M be an **NFA** over the alphabet Σ. Let M' be identical to M, except that accepting states in M are non-accepting in M' and vice versa. Each string in Σ^* is accepted by exactly one of M and M'.

> **Write your answers in the separate answer booklet.**
>
> You have 120 minutes (after you get the answer booklet) to answer five questions.
>
> Please return this question sheet and your cheat sheet with your answers.

Questions 2 and 4 were swapped in the question sheet—but not in the answer booklet(!)—that was distributed during the exam.

1. For each statement below, check "True" if the statement is always true and check "False" otherwise, and give a brief (one short sentence) explanation of your answer. Read these statements very carefully—small details matter!

   (a) Every irregular language is infinite.

   (b) The language $(0 + 1(01^*0)^*1)^*$ is context-free.

   (c) Every subset of a regular language is regular.

   (d) The language $\{0^a 1^b \mid a + b \text{ is divisible by } 374\}$ is regular.

   (e) If language L is regular, then L has a finite fooling set.

   (f) For every language L, if for every string $w \in L$ there is a DFA that accepts w, then L is regular.

   (g) If language L is accepted by an NFA with n states, then its complement $\Sigma^* \setminus L$ is also accepted by an NFA with n states.

   (h) 0^*1^* is a fooling set for the language $\left\{0^i 1^j 0^{i+j} \mid i, j \geq 0\right\}$.

   (i) Every regular language is accepted by a DFA with an odd number of accepting states.

   (j) The context-free grammar $S \rightarrow 1T \mid T1 \mid \varepsilon$; $T \rightarrow 0S \mid S0$ generates all strings in which the number of 0s equals the number of 1s.

2. Recall that a *run* in a string w is a maximal non-empty substring of w in which all symbols are equal. For example, the string $0\,111111\,000\,1\,0000$ consists of five runs.

   Let L be the set of all strings in $\{0, 1\}^*$ in which every run of 0s is followed immediately by a longer run of 1s. For example, the strings 000111111011 and 11101100011111 and 11111 are in L, but the strings 00000111 and 0111000000 are not.

   (a) **Prove** that L is *not* a regular language.

   (b) Describe a context-free grammar for L.

3. For any string $w \in \{0, 1\}^*$, let sortpairs(w) denote the string obtained by dividing w into pairs of symbols, sorting each pair into non-decreasing order, and leaving the last symbol if w has odd length. We can define sortpairs recursively as follows:

$$\text{sortpairs}(w) := \begin{cases} w & \text{if } w = \varepsilon \text{ or } w = 0 \text{ or } w = 1 \\ 01 \cdot \text{sortpairs}(x) & \text{if } w = 10x \text{ for some string } x \\ ab \cdot \text{sortpairs}(x) & \text{if } w = abx \text{ for some string } x \text{ and bits } a, b \text{ where } ab \neq 10 \end{cases}$$

For example,

$$\text{sortpairs}(\underbrace{00}\,\underbrace{10}\,\underbrace{11}\,\underbrace{10}\,\underbrace{01}\,1) = \underbrace{00}\,\underbrace{01}\,\underbrace{11}\,\underbrace{01}\,\underbrace{01}\,1$$

Recall that $\#(1, w)$ denotes the number 1s in the string w. For example, $\#(1, \varepsilon) = 0$ and $\#(1, 00101100011) = 5$.

  (a) **Prove** that $\#(1, \text{sortpairs}(w)) = \#(1, w)$ for every string w.

  (b) **Prove** that sortpairs(sortpairs(w)) = sortpairs(w) for every string w.

As usual, you can assume any result proved in class, in the lecture notes, in labs, in lab solutions, or in homework solutions. In particular, you may use the fact that $\#(1, xy) = \#(1, x) + \#(1, y)$ for all strings x and y.

4. Recall that a *run* in a string w is a maximal non-empty substring of w in which all symbols are equal. For example, the string $\underbrace{0}\,\underbrace{111111}\,\underbrace{000}\,\underbrace{1}\,\underbrace{0000}$ consists of five runs.

   For any string $w \in \{0, 1\}^*$, let compact(w) denote the string obtained by replacing each run with a single symbol from that run. For example, compact(ε) = ε and

$$\text{compact}(011111100010000) = 01010.$$

Let L be an arbitrary regular language.

  (a) **Prove** that the language $\text{COMPACT}(L) = \{\text{compact}(w) \mid w \in L\}$ is regular.

  (b) **Prove** that the language $\text{UNCOMPACT}(L) = \{w \in \Sigma^* \mid \text{compact}(w) \in L\}$ is regular.

5. For each of the following languages L over the alphabet $\Sigma = \{0, 1\}$, describe a DFA that accepts L **and** give a regular expression that represents L. You do not need to justify your answers.

  (a) Strings that do not contain the substring 01110.

  (b) Strings that contain at least one odd-length run of 0s that is followed immediately by an even-length run of 1s.

> **Write your answers in the separate answer booklet.**
>
> You have 120 minutes (after you get the answer booklet) to answer five questions.
>
> Please return this question sheet and your cheat sheet with your answers.

1. For each statement below, check "Yes" if the statement is always true and check "No" otherwise, and give a brief (one short sentence) explanation of your answer. Read these statements very carefully—small details matter!

   (a) Every infinite language is regular.

   (b) The language $(0 + 1(01^*0)^*1)^*$ is not context-free.

   (c) Every subset of an irregular language is irregular.

   (d) The language $\left\{0^a 1^b \mid a - b \text{ is divisible by } 374\right\}$ is regular.

   (e) If language L is not regular, then L has a finite fooling set.

   (f) If there is a DFA that rejects every string in language L, then L is regular.

   (g) If language L is accepted by an DFA with n states, then its complement $\Sigma^* \setminus L$ is also accepted by a DFA with n states.

   (h) 1^*0^* is a fooling set for the language $\left\{1^i 0^{i+j} 1^j \mid i, j \geq 0\right\}$.

   (i) Every regular language is accepted by a DFA with an odd number of accepting states.

   (j) The context-free grammar $S \rightarrow \varepsilon \mid 0S1S \mid 1S0S$ generates all strings in which the number of 0s equals the number of 1s.

2. For any string w, let cycleleft(w) denote the string obtained by moving the first symbol of w (if any) to the end. More formally:

$$\text{cycleleft}(w) = \begin{cases} \varepsilon & \text{if } w = \varepsilon \\ x \bullet a & \text{if } w = ax \text{ for some symbol } a \text{ and string } x \end{cases}$$

   For example, cycleleft(001111) = 011110.

   Let L be an arbitrary regular language.

   (a) Prove that $\text{CycleLeft}(L) = \{\text{cycleleft}(w) \mid w \in L\}$ is a regular language.

   (b) Prove that $\text{CycleRight}(L) = \{w \in \Sigma^* \mid \text{cycleleft}(w) \in L\}$ is a regular language.

3. For any string $w \in \{0,1\}^*$, let squish(w) denote the string obtained by dividing w into pairs of symbols, replacing each pair with 0 if the symbols are equal and 1 otherwise, and keeping the last symbol if w has odd length. We can define sortpairs recursively as follows:

$$\text{squish}(w) := \begin{cases} w & \text{if } w = \varepsilon \text{ or } w = 0 \text{ or } w = 1 \\ 0 \cdot \text{squish}(x) & \text{if } w = 00x \text{ or } w = 11x \text{ for some string } x \\ 1 \cdot \text{squish}(x) & \text{if } w = 01x \text{ or } w = 10x \text{ for some string } x \end{cases}$$

For example,

$$\text{squish}(\underbrace{00}\,\underbrace{10}\,\underbrace{11}\,\underbrace{10}\,\underbrace{01}\,1) = 010111$$

(a) **Prove** that $\#(1, \text{squish}(w)) \le \#(1, w)$ for every string w.

(b) **Prove** that $\#(1, \text{squish}(w))$ is even if and only if $\#(1, w)$ is even (or equivalently, that $\#(1, \text{squish}(w)) \bmod 2 = \#(1, w) \bmod 2$) for every string w.

As usual, you can assume any result proved in class, in the lecture notes, in labs, in lab solutions, or in homework solutions. In particular, you may use the fact that $\#(1, xy) = \#(1, x) + \#(1, y)$ for all strings x and y.

4. Let L be the set of all strings in $\{0,1\}^*$ in which every run of 0s is followed immediately by a *shorter* run of 1s. For example, the strings 0001100000111 and 11100100000111 and 11111 are in L, but the strings 00011111 and 000110000 are not.

(a) **Prove** that L is *not* a regular language.

(b) Describe a context-free grammar for L.

5. For each of the following languages L over the alphabet $\Sigma = \{0,1\}$, describe a DFA that accepts L **and** give a regular expression that represents L. You do not need to justify your answers.

(a) Strings that do not contain the subsequence 01110.

(b) Strings that contain at least two even-length runs of 1s.

You have 120 minutes to answer five questions.

Write your answers in the separate answer booklet.

Please return this question sheet and your cheat sheet with your answers.

1. Short answers:

   (a) Solve the following recurrences:
      - $A(n) = \mathbf{3}A(n/2) + O(n^2)$
      - $B(n) = \mathbf{7}B(n/2) + O(n^2)$
      - $C(n) = \mathbf{4}C(n/2) + O(n^2)$

   (b) Draw a directed acyclic graph with at most ten vertices, exactly one source, exactly one sink, and more than one topological order.

   (c) Draw a directed graph with at most ten vertices, with distinct positive edge weights, that has more than one shortest path from some vertex s to some other vertex t.

   (d) Describe an appropriate memoization structure and evaluation order for the following (meaningless) recurrence, and give the running time of the resulting iterative algorithm to compute $Huh(1, n)$.

$$Huh(i, k) = \begin{cases} 0 & \text{if } i > n \text{ or } k < 0 \\ \min\begin{cases} Huh(i+1, k-2) \\ Huh(i+2, k-1) \end{cases} + A[i, k] & \text{if } A[i, k] \text{ is even} \\ \max\begin{cases} Huh(i+1, k-2) \\ Huh(i+2, k-1) \end{cases} - A[i, k] & \text{if } A[i, k] \text{ is odd} \end{cases}$$

2. You and your eight-year-old nephew Elmo decide to play a simple card game. At the beginning of the game, the cards are dealt face up in a long row. Each card is worth a different number of points, which could be positive, negative, or zero. After all the cards are dealt, you and Elmo take turns removing either the leftmost or rightmost card from the row, until all the cards are gone. At each turn, you can decide which of the two cards to take. The winner of the game is the player that has collected the most points when the game ends.

   Having never taken an algorithms class, Elmo follows the obvious greedy strategy—when it's his turn, Elmo *always* takes the card with the higher point value. Your task is to find a strategy that will beat Elmo whenever possible. (It might seem mean to beat up on a little kid like this, but Elmo absolutely *hates* it when grown-ups let him win.)

   (a) **Prove** that you should not also use the greedy strategy. That is, show that there is a game that you can win, but only if you do *not* follow the same greedy strategy as Elmo. Assume Elmo plays first.

   (b) Describe and analyze an algorithm to determine, given the initial sequence of cards, the maximum number of points that you can collect playing against Elmo.

3. Suppose you are given a directed graph $G = (V, E)$, whose vertices are either red, green, or blue. Edges in G do not have weights, and G is not necessarily a dag. The **remoteness** of a vertex v is the *maximum* of three shortest-path lengths:

   - The length of a shortest path to v from the closest red vertex
   - The length of a shortest path to v from the closest blue vertex
   - The length of a shortest path to v from the closest green vertex

   In particular, if v is not reachable from vertices of all three colors, then v is infinitely remote.

   Describe and analyze an algorithm to find a vertex of G with *minimum* remoteness.

4. Suppose you are given an array $A[1..n]$ of integers such that $A[i] + A[i+1]$ is even for *exactly one* index i. In other words, the elements of A alternate between even and odd, except for exactly one adjacent pair that are either both even or both odd.

   Describe and analyze an efficient algorithm to find the unique index i such that $A[i] + A[i+1]$ is even. For example, given the following array as input, your algorithm should return the integer 6, because $A[6] + A[7] = 88 + 62$ is even. (Cells containing even integers are shaded blue.)

17	40	23	72	39	88	62	13	40	53	92	21	10	73	68

5. A *zigzag walk* in in a directed graph G is a sequence of vertices connected by edges in G, but the edges alternately point forward and backward along the sequence. Specifically, the first edge points forward, the second edge points backward, and so on. The *length* of a zigzag walk is the sum of the weights of its edges, both forward and backward.

   For example, the following graph contains the zigzag walk $a \to b \leftarrow d \to f \leftarrow c \to e$. Assuming every edge in the graph has weight 1, this zigzag walk has length 5.


   Suppose you are given a directed graph G with non-negatively weighted edges, along with two vertices s and t. Describe and analyze an algorithm to find the shortest zigzag walk from s to t in G.

> You have 120 minutes to answer five questions.
> **Write your answers in the separate answer booklet.**
> Please return this question sheet and your cheat sheet with your answers.

1. **Clearly** indicate the following structures in the directed graph below, or write NONE if the indicated structure does not exist. Don't be subtle; to indicate a collection of edges, draw a heavy black line along the entire length of each edge.

   (a) A depth-first search tree rooted at vertex a.

   (b) A breadth-first tree rooted at vertex c.

   (c) The strong components of G. (Circle each strong component.)

   (d) Draw the strong-component graph of G.

2. Suppose we are given an n-digit integer X. Repeatedly remove one digit from either end of X (your choice) until no digits are left. The *square-depth* of X is the maximum number of perfect squares that you can see during this process. For example, the number 32492 has square-depth 3, by the following sequence of removals:

$$32492 \xrightarrow{57^2} 3249 \xrightarrow{18^2} 324 \to 24 \xrightarrow{2^2} 4 \to \varepsilon.$$

   Describe and analyze an algorithm to compute the square-depth of a given integer X, represented as an array $X[1..n]$ of n decimal digits. Assume you have access to a subroutine IsSquare that determines whether a given k-digit number (represented by an array of digits) is a perfect square *in $O(k^2)$ time*.

3. Suppose you are given a directed graph $G = (V, E)$, each of whose edges are colored red, green, or blue. Edges in G do not have weights, and G is not necessarily a dag. A *rainbow walk* is a walk in G that does *not* contain two consecutive edges with the same color.

   Describe and analyze an algorithm to find all vertices in G that are reachable from a given vertex s through a rainbow walk.

4. Suppose you are given k sorted arrays $A_1[1..n], A_2[1..n], \ldots, A_k[1..n]$, all with the same length n. Describe an algorithm to merge the given arrays into a single sorted array. Analyze the running time of your algorithm as a function of n and k.

5. After moving to a new city, you decide to walk from your home to your new office. To get a good daily workout, you want to reach the highest possible altitude during your walk (to maximize exercise), while keeping the total length of your walk below some threshold (to get to your office on time). Describe and analyze an algorithm to compute the best possible walking route.

   Your input consists of an undirected graph G, where each vertex v has a height $h(v)$ and each edge e has a positive length $\ell(e)$, along with a start vertex s, a target vertex t, and a maximum length L. Your algorithm should return the maximum height reachable by a walk from s to t in G, whose total length is at most L.

> You have 120 minutes to answer five questions.
> **Write your answers in the separate answer booklet.**
> Please return this question sheet and your cheat sheet with your answers.

1. Short answers:

   (a) Solve the following recurrences:
   - $A(n) = A(5n/11) + O(\sqrt{n})$
   - $B(n) = 8B(n/2) + O(n^2)$
   - $C(n) = C(n/2) + C(n/3) + C(n/6) + O(n)$

   (b) Describe an appropriate memoization structure and evaluation order for the following (meaningless) recurrences, and state the running time of the resulting iterative algorithm to compute the requested function value.

   - Compute $Foo(1, n)$ where

   $$Foo(i,k) = \begin{cases} 0 & \text{if } i \geq k-1 \\ \max\left\{ \begin{array}{l} Foo(i,j) \\ + Foo(j,k) \end{array} \,\middle|\, i < j < k \right\} + \sum_{j=i}^{k} A[j] & \text{otherwise} \end{cases}$$

   - Compute $Bar(n, 1)$ where

   $$Bar(i,s) = \begin{cases} \infty & \text{if } i < 0 \text{ or } s > n \\ 0 & \text{if } i = 0 \\ \min\left\{ \begin{array}{l} Bar(i, 2s), \\ X[i] \cdot s + Bar(i-s, s) \end{array} \right\} & \text{otherwise} \end{cases}$$

2. Chandler is moving to Tulsa, Oklahoma, to start a new job, and he wants to plan a walk from home to work. He wants to continue his habit of buying a ~~hot milkshake~~ salted caramel latte from a local coffee shop and a paper from a local newsstand every morning. (Yes, an actual *paper* paper. Could he *be* any more retro?)

   Chandler has a map of his new neighborhood in the form of an undirected graph G, whose vertices represent intersections and whose edges represent roads between them. Every edge e has a positive length $\ell(e)$. A subset of the vertices are marked as newsstands; another disjoint subset of vertices are marked as coffee shops. The graph has two special vertices s and t, which represent Chandler's home and work, respectively.

   Describe an algorithm that computes the shortest route that Chandler can follow from home to work that visits both a coffee shop and a newsstand, or correctly reports that no such route exists (which means Chandler should move back to New York).

   Problems 3 and 4 are on the back of this page.

3. Recall that an *arithmetic progression* is any sequence of real numbers $x_1, x_2, \ldots, x_n$ such that $x_{i+1} - x_i = x_i - x_{i-1}$ for every index $2 \le i \le n-1$.

   Suppose we are given a sorted array $X[1..n]$ that contains an arithmetic sequence **with one element deleted**. Describe and analyze an algorithm to find the deleted element as quickly as possible. (If there are multiple correct answers, your algorithm can return any one of them.)

   For example, given the input array $X = [2, 4, 8, 10, 12]$, your algorithm should return 6, and given the array $X = [21, 18, 15, 12]$, your algorithm should return either 9 or 24.

4. **Ink-deck** is a solitaire puzzle game played on a row of n squares, each marked with either $+$ or $-$. Your goal is to move a token from a specified start square s to a specified target square t using a sequence of *moves*. Each move translates the token either left or right along the row to a new square. The *length* of a move is the distance that the token moves; for example, a move from square 5 to square 12 has length 7. Moves are subject to the following rules:

   - The first move must have length 1.
   - If the token is on a square marked $+$, your next move must be one square longer than your previous move.
   - If the token is on a square marked $-$, your next move must be one square shorter than your previous move. (In particular, if your previous move had length 1, then you cannot move at all!)
   - You are never allowed to move the token off either end of the row.

   The following figure shows an example ink-deck puzzle, along with two solutions (which might not be optimal).

   An ink-deck puzzle with two nine-move solutions.

   (a) Describe an algorithm that either finds a solution *with the minimum number of moves* for a given ink-deck puzzle, or correctly reports that the given puzzle has no solution.

   (b) Describe an algorithm that either finds a solution *whose final move is as long as possible* for a given ink-deck puzzle, or correctly reports that the given puzzle has no solution.

   Your input to both algorithms consists of an array $ID[1..n]$, where $ID[i] \in \{-1, +1\}$ for each index i, along with two indices $1 \le s \le n$ and $1 \le t \le n$.

   Problem 5 is on the next page.

5. Suppose you are given a string of symbols, representing a message in some foreign language that you do not understand, in an array $T[1..n]$. You have access to a black-box subroutine IsWord that takes a string w as input and decides in $O(|w|)$ time whether w is a word.

   You eagerly implement and run the text-splitting algorithm we saw in class, only to discover that the given string *cannot* be split into words! Apparently, as a crude form of cryptography, the message has been corrupted by adding extra symbols between words.

   So you decide instead to look for as many non-overlapping words in T as possible. A *verbal subsequence* of T is a sequence of non-overlapping substrings of T, each of which is a word. The *length* of a verbal subsequence is the number of words it contains. Describe and analyze an algorithm to find the length of the longest verbal subsequence of a given string T.

   For example, suppose IsWord(w) returns True if and only if w is an English word with at least four letters. Then (STUDY, MICE, TRAP, RAMEN) and (DYNAMIC, EXTRA, PROGRAM) are verbal subsequences of the string STUDYNAMICEXTRAPROGRAMEN:

   | STUDY | NA | MICE | X | TRAP | ROG | RAMEN | | STU | DYNAMIC | EXTRA | PROGRAM | EN |

   Given the input string STUDYNAMICEXTRAPROGRAMEN, your algorithm should return the integer 4, which is the length of the verbal subsequence (STUDY, MICE, TRAP, RAMEN).

Nothing to see here.

You have 120 minutes to answer five questions.

Write your answers in the separate answer booklet.

Please return this question sheet and your cheat sheet with your answers.

1. (a) Solve the following recurrences:
   - $A(n) = A(2n/3) + O(\sqrt{n})$
   - $B(n) = 8B(n/4) + O(n^{3/2})$
   - $C(n) = C(n/2) + C(n/3) + O(n)$

   (b) Describe an appropriate memoization structure and evaluation order for the following (meaningless) recurrences, and state the running time of the resulting iterative algorithm to compute the requested function value.

   - Compute $Pleb(n, 1)$ where

   $$Pleb(i, k) = \begin{cases} i & \text{if } k \leq 0 \\ k & \text{if } i > n \\ \max \begin{cases} i + k + Pleb(i-1, k+1) \\ i + Pleb(i-1, k) \\ k + Pleb(i, k+1) \end{cases} & \text{otherwise} \end{cases}$$

   - Compute $Nom(1, n)$ where

   $$Nom(i, k) = \begin{cases} 0 & \text{if } k = i \\ (X[k] - X[i]) + \min \begin{cases} Nom(i, j) \\ + Nom(j, k) \end{cases} \Big| \, i < j < k \Big\} & \text{otherwise} \end{cases}$$

2. Suppose you are given a directed graph G in which every edge is either red or blue, and a subset of the vertices are marked as *special*. A walk in G is *legal* if color changes happen only at special vertices. That is, for any two consecutive edges $u{\to}v{\to}w$ in a legal walk, if the edges $u{\to}v$ and $v{\to}w$ have different colors, the intermediate vertex v must be special.

   Describe and analyze an algorithm that either returns the length of the shortest legal walk in G from vertex s to vertex t, or correctly reports that no such walk exists.

   For example, if you are given the following graph as input (where single arrows are "red" and double arrows are "blue"), with special vertices x and y, your algorithm should return the integer 8, which is the length of the shortest legal walk $s{\to}x{\to}a{\to}b{\to}x{\Rightarrow}y{\Rightarrow}b{\Rightarrow}c{\Rightarrow}t$. The shorter walk $s{\to}a{\to}b{\Rightarrow}c{\Rightarrow}t$ is not legal, because vertex b is not special.


   Problem 3 is on the back of this page.

3. **Ink-deck** is a solitaire puzzle game played on a row of n squares, each marked with either + or −. Your goal is to move a token from a specified start square s to a specified target square t using a sequence of *moves*. Each move translates the token either left or right along the row to a new square. The *length* of a move is the distance that the token moves; for example, a move from square 5 to square 12 has length 7. Moves are subject to the following rules:

   - The first move must have length 1.
   - If the token is on a square marked +, your next move must be one square longer than your previous move.
   - If the token is on a square marked −, your next move must be one square shorter than your previous move. (In particular, if your previous move had length 1, then you cannot move at all!)
   - You are never allowed to move the token off either end of the row.

   The following figure shows an example ink-deck puzzle, along with two solutions.

An ink-deck puzzle with two solutions.

   (a) The *total length* of a solution is the sum of the lengths of the moves. For example, the top solution in the figure above has total length $1+2+3+4+3+4+5+4+3 = 29$, and the bottom solution has total length $1+2+3+4+3+4+3+2+1 = 23$.

   Describe an algorithm that, given an ink-deck puzzle, either finds a solution whose *total* length is as *small* as possible, or correctly reports that there is no solution.

   (b) The *maximum length* of a solution is length of its longest move. For example, the top solution above has maximum length 5, and the bottom solution has maximum length 4.

   Describe an algorithm that, given an ink-deck puzzle, either finds a solution whose *maximum* length is as *large* as possible, or correctly reports that there is no solution.

   Your input to both algorithms consists of an array $ID[1..n]$, where $ID[i] \in \{-1,+1\}$ for each index i, along with two indices $1 \leq s \leq n$ and $1 \leq t \leq n$.

Problems 4 and 5 are on the next page.

4. Recall that an *arithmetic progression* is any sequence of real numbers $x_1, x_2, \ldots, x_n$ such that $x_{i+1} - x_i = x_i - x_{i-1}$ for every index $2 \le i \le n-1$.

   Suppose we are given a sorted array $X[1 .. n]$ containing an arithmetic sequence *with one element repeated once*. Describe and analyze an algorithm to find the repeated element as quickly as possible.

   For example, given the input array $X = [2, 4, 6, 6, 8, 10, 12]$, your algorithm should return 6, and given the input array $X = [1, 1, 1, 1]$, your algorithm should return 1.

5. Suppose you are given a string of symbols, representing a message in some foreign language that you do not understand, in an array $T[1 .. n]$. You have access to a black-box subroutine IsWord that takes a string w as input and decides in $O(|w|)$ time whether w is a word.

   You eagerly implement and run the text-splitting algorithm we saw in class, only to discover that the given string *cannot* be split into words! Apparently, as a crude form of cryptography, the author of the message added extra symbols at the beginning and end.

   Describe and analyze an algorithm to find the length of the *longest substring* of T that can be split into words.

   For example, suppose IsWord(w) returns True if and only if w is an English word with at least four letters. Given the input string STURDYNAMICEXTRAPROGRAMBLE, your algorithm should return the integer 19, which is the length of the substring DYNAMICEXTRAPROGRAM, which can be split into the words DYNAMIC, EXTRA, and PROGRAM:

   $$\boxed{\text{STUR}}\ \boxed{\text{DYNAMIC}}\ \boxed{\text{EXTRA}}\ \boxed{\text{PROGRAM}}\ \boxed{\text{BLE}}$$

   The words STURDY, MICE, TRAP, and RAMBLE have larger total length 20, but there are gaps between them; so they can't be formed by splitting a *substring* of T.

   $$\boxed{\text{STURDY}}\ \boxed{\text{NA}}\ \boxed{\text{MICE}}\ \boxed{\text{X}}\ \boxed{\text{TRAP}}\ \boxed{\text{ROG}}\ \boxed{\text{RAMBLE}}$$

Nothing to see here.

> You have 180 minutes to answer six numbered questions.
> **Write your answers in the separate answer booklet.**
> Please return this question sheet and your cheat sheet with your answers.

1. For each statement below, there are two boxes in the answer booklet labeled "Yes" and "No". Check "Yes" if the statement is *always* true and "No" otherwise, and give a *brief* (at most one short sentence) explanation of your answer. **Assume $P \neq NP$.** If there is any other ambiguity or uncertainty about an answer, check "No". For example:

    - $x + y = 5$

        [X̶ Yes] [No]   Suppose $x = 3$ and $y = 4$.

    - 3SAT can be solved in polynomial time.

        [Yes] [X̶ No]   3SAT is NP-hard.

    - If P = NP then Jeff is the Queen of England.

        [X̶ Yes] [No]   The hypothesis is false, so the implication is true.

    Read each statement *very* carefully; some of these are deliberately subtle!

    (a) Which of the following statements are true?

    - The solution to the recurrence $T(n) = \mathbf{8}T(n/\mathbf{2}) + O(n^2)$ is $T(n) = O(n^2)$.

    - The solution to the recurrence $T(n) = \mathbf{2}T(n/\mathbf{8}) + O(n^2)$ is $T(n) = O(n^2)$.

    - Every directed acyclic graph contains at least one sink.

    - Given *any* undirected graph G, we can compute a spanning tree of G in $O(V + E)$ time using whatever-first search.

    - Suppose $A[1 .. n]$ is an array of integers. Consider the following recursive function:

    $$
    What(i, j) = \begin{cases}
    0 & \text{if } i < 0 \text{ or } i > n \\
    0 & \text{if } j < 0 \text{ or } j > n \\
    \max \left\{ \begin{array}{c} What(i, j-1) \\ What(i-1, j) \\ A[i] \cdot A[j] + What(i+1, j+1) \end{array} \right\} & \text{otherwise}
    \end{cases}
    $$

    We can memoize this function into an array $What[0 .. n, 0 .. n]$ in $O(n^2)$ time, by increasing i in the outer loop and increasing j in the inner loop.

Problem 1 continues onto the next page.

1. *[continued]*

   (b) Which of the following statements are true for ***at least one*** language $L \subseteq \{0,1\}^*$?

      • $L^* = (L^*)^*$

      • L is decidable, but L^* is undecidable.

      • L is neither regular nor NP-hard.

      • L is in P, and L has an infinite fooling set.

      • The language $\{\langle M \rangle \mid M \text{ accepts } L\}$ is undecidable.

   (c) Consider the following pair of languages:

      • DIRHAMPATH $:= \{G \mid G \text{ is a directed graph with a Hamiltonian path}\}$

      • ACYCLIC $:= \{G \mid G \text{ is a directed acyclic graph}\}$

      (For concreteness, assume that in both of these languages, graphs are represented by their adjacency matrices.) Which of the following statements are true, assuming $P \neq NP$?

      • ACYCLIC $\in$ NP

      • ACYCLIC $\cap$ DIRHAMPATH $\in$ P

      • DIRHAMPATH is decidable.

      • A polynomial-time reduction from DIRHAMPATH to ACYCLIC would imply P=NP.

      • A polynomial-time reduction from ACYCLIC to DIRHAMPATH would imply P=NP.

   (d) Suppose there is a ***polynomial-time*** reduction from some language A over the alphabet $\{0,1\}$ to some other language B over the alphabet $\{0,1\}$. Which of the following statements are ***always*** true, assuming $P \neq NP$?

      • A is a subset of B.

      • If $B \in P$, then $A \in P$.

      • If B is NP-hard, then A is NP-hard.

      • If B is regular, then A is regular.

      • If B is regular, then A is decidable.

2. Describe and analyze an algorithm to determine whether the language accepted by a given DFA is finite or infinite. You can assume the input alphabet of the DFA is $\{0,1\}$. *[Hint: DFAs are directed graphs.]*

3. Suppose you are asked to tile a $2 \times n$ grid of squares with dominos (1×2 rectangles). Each domino must cover exactly two grid squares, either horizontally or vertically, and each grid square must be covered by exactly one domino.

   Each grid square is worth some number of points, which could be positive, negative, or zero. The **value** of a domino tiling is the sum of the points in squares covered by vertical dominos, **minus** the sum of the points in squares covered by horizontal dominos.

   Describe an algorithm to compute the largest possible value of a domino tiling of a given $2 \times n$ grid. Your input is an array $Points[1 .. 2, 1 .. n]$ of point values.

   As an example, here are three domino tilings of the same 2×6 grid, along with their values. The third tiling is optimal; no other tiling of this grid has larger value. Thus, given this 2×6 grid as input, your algorithm should return the integer 16.

4. Submit a solution to *exactly one* of the following problems. Don't forget to tell us which problem you've chosen!

   (a) Let Φ be a boolean formula in conjunctive normal form, with exactly three literals per clause (or in other words, an instance of 3SAT). **Prove** that it is NP-hard to decide whether Φ has a satisfying assignment in which *exactly half* of the variables are TRUE.

   (b) Let $G = (V, E)$ be an arbitrary directed graph whose edges have colors. A *rainbow Hamiltonian cycle* in G is a cycle that visits every vertex of G exactly once, in which no pair of consecutive edges have the same color. **Prove** that it is NP-hard to decide whether G has a rainbow Hamiltonian cycle.

   (In fact, both of these problems are NP-hard, but we only want a proof for one of them.)

5. Suppose you are given a height map of a mountain, in the form of an $n \times n$ grid of evenly spaced points, each labeled with an elevation value. You can safely hike directly from any point to any neighbor immediately north, south, east, or west, but only if the elevations of those two points differ by at most Δ. (The value of Δ depends on your hiking experience and your physical condition.)

   Describe and analyze an algorithm to determine the longest hike from some point s to some other point t, where the hike consists of an uphill climb (where elevations must increase at each step) followed by a downhill climb (where elevations must decrease at each step). Your input consists of an array $Elevation[1 .. n, 1 .. n]$ of elevation values, the starting point s, the target point t, and the parameter Δ.

6. Recall that a **run** in a string $w \in \{0, 1\}^*$ is a maximal substring of w whose characters are all equal. For example, the string 00011111110000 is the concatenation of three runs:

$$00011111110000 = 000 \bullet 1111111 \bullet 0000$$

(a) Let L_a denote the set of all strings in $\{0, 1\}^*$ where every 0 is followed immediately by at least one 1.

For example, L_a contains the strings 010111 and 1111 and the empty string ε, but does not contain either 001100 or 1111110.

- Describe a DFA or NFA that accepts L_a **and**
- Give a regular expression that describes L_a.

(You do not need to prove that your answers are correct.)

(b) Let L_b denote the set of all strings in $\{0, 1\}^*$ whose run lengths are increasing; that is, every run except the last is followed immediately by a *longer* run.

For example, L_b contains the strings 0110001111 and 1100000 and 000 and the empty string ε, but does not contain either 000111 or 100011.

Prove that L_b is not a regular language.

> You have 180 minutes to answer six numbered questions.
> **Write your answers in the separate answer booklet.**
> Please return this question sheet and your cheat sheet with your answers.

1. Recall that a ***run*** in a string $w \in \{0,1\}^*$ is a maximal substring of w whose characters are all equal. For example, the string `00011111110000` is the concatenation of three runs:

   $$00011111110000 = 000 \bullet 1111111 \bullet 0000$$

   (a) Let L_a denote the set of all non-empty strings in $\{0,1\}^*$ where the length of the first run is equal to the number of runs. For example, L_a contains the strings `0` and `1100000` and `0001110`, but does not contain `000111` or `100011` or the empty string ε (because it has no first run).

       Prove that L_a is not a regular language.

   (b) Let L_b denote the set of all strings in $\{0,1\}^*$ that contain an even number of odd-length runs. For example, L_b contains the strings `010111` and `1111` and the empty string ε, but does not contain either `0011100` or `11110`.

   - Describe a DFA or NFA that accepts L_b ***and***
   - Give a regular expression that describes L_b.

   (You do not need to prove that your answers are correct.)

2. Aladdin and Badroulbadour are playing a cooperative game. Each player has an array of positive integers, arranged in a row of squares from left to right. Each player has a token, which starts at the leftmost square of their row; their goal is to move *both* tokens onto the rightmost squares at the same time.

   On each turn, *both* players move their tokens *in the same direction*, either left or right. The distance each token travels is equal to the number under that token at the beginning of the turn. For example, if a token starts on a square labeled 5, then it moves either five squares to the right or five squares to the left. If *either* token moves past either end of its row, then both players immediately lose.

   For example, if Aladdin and Badroulbadour are given the arrays

A:	7	5	4	1	2	3	3	2	3	1	4	2
B:	5	1	2	4	7	3	5	2	4	6	3	1

   they can win the game by moving right, left, left, right, right, left, right. On the other hand, if they are given the arrays

A:	2	3	5	1	3
B:	3	4	1	2	1

   they cannot win the game. (The first move must be to the right; then Aladdin's token moves out of bounds on the second turn.)

   Describe and analyze an algorithm to determine whether Aladdin and Badroulbadour can solve their puzzle, given the input arrays $A[1..n]$ and $B[1..n]$.

3. Submit a solution to **exactly one** of the following problems. Don't forget to tell us which problem you've chosen!

   (a) Let $G = (V, E)$ be an arbitrary undirected graph. A subset $S \subseteq V$ of vertices is *mostly independent* if more than half the vertices of S have no neighbors in S. **Prove** that finding the largest mostly independent set in G is NP-hard.

   (b) **Prove** that the following problem is NP-hard: Given an undirected graph G, find the largest integer k such that G contains *two disjoint* independent sets of size k.

   (In fact, both of these problems are NP-hard, but we only want a proof for one of them.)

4. Recall that a *palindrome* is any string that is equal to its reversal, like REDIVIDER or POOP.

   (a) Describe and analyze an algorithm to find the length of the longest subsequence of a given string that is a palindrome.

   (b) A *double palindrome* is the concatenation of two *non-empty* palindromes, like REFEREE = REFER • EE or POOPREDIVIDER = POOP • REDIVIDER. Describe and analyze an algorithm to find the length of the longest subsequence of a given string that is a *double* palindrome. *[Hint: Use your algorithm from part (a).]*

   For both algorithms, the input is an array $A[1..n]$, and the output is an integer. For example, given the input string MAYBEDYNAMICPROGRAMMING, your algorithm for part (a) should return 7 (for the subsequences NMRORMN and MAYBYAM, among others), and your algorithm for part (b) should return 12 (for the subsequence MAYBYAMIRORI).

5. You have a collection of n lockboxes and m gold keys. Each key unlocks *at most* one box. Without a matching key, the only way to open a box is to smash it with a hammer. Your baby brother has locked all your keys inside the boxes! Luckily, you know which keys (if any) are inside each box.

   (a) Your baby brother has found the hammer and is eagerly eyeing one of the boxes. Describe and analyze an algorithm to determine if it is possible to retrieve all the keys without smashing any box except the one your brother has chosen.

   (b) Describe and analyze an algorithm to compute the minimum number of boxes that must be smashed to retrieve all the keys.

Problem 6 begins on the next page.

6. For each statement below, there are two boxes in the answer booklet labeled "Yes" and "No". Check "Yes" if the statement is **always** true and "No" otherwise, and give a **brief** (at most one short sentence) explanation of your answer. **Assume $P \neq NP$.** If there is any other ambiguity or uncertainty about an answer, check "No". For example:

- $x + y = 5$

  [Y̶e̶s̶ ✗]   [No]   Suppose $x = 3$ and $y = 4$.

- 3SAT can be solved in polynomial time.

  [Yes]   [N̶o̶ ✗]   3SAT is NP-hard.

- If $P = NP$ then Jeff is the Queen of England.

  [Y̶e̶s̶ ✗]   [No]   The hypothesis is false, so the implication is true.

Read each statement *very* carefully; some of these are deliberately subtle!

(a) Which of the following statements are true for **all** languages $L \subseteq \{0, 1\}^*$?

- $L^* = (L^*)^*$

- If L is decidable, then L^* is decidable.

- L is either regular or NP-hard.

- If L is undecidable, then L has an infinite fooling set.

- The language $\{\langle M \rangle \mid M$ decides $L\}$ is undecidable.

(b) Which of the following statements are true?

- The solution to the recurrence $T(n) = 4T(n/4) + O(n)$ is $T(n) = O(n \log n)$.

- The solution to the recurrence $T(n) = 4T(n/4) + O(n^2)$ is $T(n) = O(n^2 \log n)$.

- Every directed acyclic graph contains at most one source and at most one sink.

- depth-first search explores every path from the source vertex s to every other vertex in the input graph.

- Suppose $A[1 .. n]$ is an array of integers. Consider the following recursive function:

$$Huh(i, j) = \begin{cases} 0 & \text{if } i < 0 \text{ or } j > n \\ \max \begin{cases} Huh(i, j+1) \\ Huh(i-1, j) \\ A[i] \cdot A[j] + Huh(i-1, j+1) \end{cases} & \text{otherwise} \end{cases}$$

  We can compute $Huh(n, 0)$ by memoizing this function into an array $Huh[0 .. n, 0 .. n]$ in $O(n^2)$ time, increasing i in the outer loop and increasing j in the inner loop.

Problem 6 continues onto the next page.

1. *[continued]*

    (c) Suppose we want to prove that the following language is undecidable.

    $$\text{Muggle} := \big\{\langle M\rangle \,\big|\, M \text{ accepts } \texttt{SCIENCE} \text{ but rejects } \texttt{MAGIC}\big\}$$

    Professor Potter, your instructor in Defense Against Models of Computation and Other Dark Arts, suggests a reduction from the standard halting language

    $$\text{Halt} := \big\{(\langle M\rangle, w) \,\big|\, M \text{ halts on input } w\big\}.$$

    Specifically, suppose there is a Turing machine DetectoMuggletum that decides Muggle. Professor Potter claims that the following algorithm decides Halt.

    DecideHalt(⟨M⟩, w):
      Write code for the following algorithm:

        RubberDuck(x):
          run M on input w
          ⟨⟨ignore the output of M⟩⟩
          if x = MAGIC
              return False
          else
              return True

      return DetectoMuggletum(⟨RubberDuck⟩)

    Which of the following statements **must be** true **for all** inputs ⟨M⟩#w?

    - If M accepts w, then RubberDuck accepts MAGIC.

    - If M diverges on w, then RubberDuck rejects MAGIC.

    - If M accepts w, then DetectoMuggletum accepts ⟨RubberDuck⟩.

    - If M diverges on w, then DecideHalt rejects (⟨M⟩, w).

    - DecideHalt decides the language Halt. (That is, Professor Potter's reduction is actually correct.)

    (d) Suppose there is a **polynomial-time** reduction from some language $A \subseteq \{0, 1\}$ reduces to some other language $B \subseteq \{0, 1\}$. Which of the following statements are true, assuming P ≠ NP?

    - $A \cap B \neq \varnothing$.

    - There is an algorithm to transform any Python program that solves B in polynomial time into a Python program that solves A in polynomial time.

    - If B is NP-hard, then A is NP-hard.

    - If B is decidable, then A is decidable.

    - If a Turing machine M accepts every string in B, the *same* Turing machine M also accepts every string in A.

You have 180 minutes to answer six numbered questions.

Write your answers in the separate answer booklet.

Please return this question sheet and your cheat sheet with your answers.

1. For each statement below, there are two boxes in the answer booklet labeled "Yes" and "No". Check "Yes" if the statement is **always** true and "No" otherwise, and give a **brief** (at most one short sentence) explanation of your answer. **Assume $P \neq NP$.** If there is any other ambiguity or uncertainty about an answer, check "No". For example:

   - $x + y = 5$

     ☐ Yes   ☒ No   Suppose $x = 3$ and $y = 4$.

   - 3SAT can be solved in polynomial time.

     ☐ Yes   ☒ No   3SAT is NP-hard.

   - If P = NP then Jeff is the Queen of England.

     ☒ Yes   ☐ No   The hypothesis is false, so the implication is true.

   Read each statement *very* carefully; some of these are deliberately subtle!

   (a) Which of the following statements are true?

   - The solution to the recurrence $T(n) = 2T(n/2) + O(\sqrt{n})$ is $T(n) = O(\sqrt{n}\log n)$.

   - The solution to the recurrence $T(n) = 2T(n/4) + T(n/3) + O(n)$ is $T(n) = O(n\log n)$.

   - There is a forest with 374 vertices and 374 edges. (Recall that a *forest* is an undirected graph with no cycles.)

   - Given any directed graph G whose edges have positive weights, we can compute shortest paths from one vertex s to every other vertex of G in $O(VE)$ time using Bellman-Ford.

   - Suppose $A[1..n]$ is an array of integers. Consider the following recursive function:

     $$Oops(i,k) = \begin{cases} 0 & \text{if } i > k \\ 1 & \text{if } i = k \\ \max\left\{A[i] \cdot A[j] \cdot A[k] + Oops(i,j) + Oops(j,k) \mid i \leq j \leq k\right\} & \text{otherwise} \end{cases}$$

     We can compute $Oops(1,n)$ by memoizing this function into a two-dimensional array $Oops[1..n, 1..n]$, which we fill by decreasing i in the outer loop and increasing k in the inner loop, in $O(n^2)$ time.

Problem 1 continues onto the next page.

1. *[continued]*

    (b) Which of the following statements are true for *every* language $L \subseteq \{0, 1\}^*$?

    - Either L is regular or L is infinite.

    - L^* is regular.

    - L contains arbitrarily long strings.

    - If L is decidable, then its complement $\bar{L}$ is also decidable.

    - If L is undecidable then L^* is undecidable.

    (c) Consider the following pair of languages:

    - 3COLOR $= \{G \mid G$ is an undirected graph with a proper 3-coloring$\}$
    - FOREST $= \{G \mid G$ is an undirected graph with no cycles$\}$

    (For concreteness, assume that in both of these languages, graphs are represented by their adjacency matrices.) Which of the following statements are true, assuming $P \neq NP$?

    - FOREST $\in P$

    - FOREST $\cap$ 3COLOR $\in P$

    - 3COLOR is undecidable.

    - 3COLOR is regular.

    - A polynomial-time reduction from 3COLOR to FOREST would imply $P = NP$.

    (d) Suppose there is a *polynomial-time* reduction R from some language $A \in \{0, 1\}^*$ to some other language $B \in \{0, 1\}^*$. Which if the following statements are *always* true, assuming $P \neq NP$?

    - The reduction transforms every string in A into a string in B.

    - The reduction transforms every string in B into a string in A.

    - If A is infinite, then B is infinite.

    - If B is undecidable, then A is undecidable.

    - If B is undecidable, then A is NP-hard.

Problem 2 appears on the next page.

2. Several years after graduating from Sham-Poobanana University, you decide to open a one-day pop-up art gallery selling NFTs, using the following dynamic pricing strategy.

All NFTs art your gallery have the same advertised price, which you set at the start of the day, but which you can *decrease* later. Customers visit your gallery one at a time. If a customer is willing to pay your current advertised price, they buy one NFT at that price. On the other hand, if your advertised price is too high, the customer will suggest a lower price that they are willing to pay. If you refuse to lower your advertised price, the customer will leave without buying anything. If you agree to lower your advertised price to match their offer, the customer will buy one NFT at the new lower price. Whenever you lower your advertised price, your new lower price stays in effect until you lower it again, or until the end of the day. *You can never increase your advertised price.*

You know your customers extremely well, so you can accurately predict both when each customer will come to the gallery, and how much each customer is willing to pay for one of your NFTs.

Describe and analyze an algorithm that computes the maximum amount of money you can earn using this dynamic pricing strategy. Your input consists of an array $Value[1..n]$, where $Value[i]$ is the amount that the ith customer (in chronological order) is willing to pay for one NFT.

For example, if the input array is $[5, 3, 1, 4, 2]$, your algorithm should output 13, because you can earn $5 + 3 + 0 + 3 + 2 = 13$ dollars using the prices $[5, 3, 3, 3, 2]$, and this is optimal.

3. Submit a solution to **exactly one** of the following problems. Don't forget to tell us which problem you've chosen!

   (a) A *mostly-3-coloring* of a graph $G = (V, E)$ is any function $C : V \rightarrow \{red, yellow, blue, none\}$ such that $C(v) = none$ for less than half the vertices in V. A mostly-3-coloring C is *proper* if, for every edge $uv \in E$, either $C(u) \neq C(v)$ or $C(u) = C(v) = none$.

   Prove that it is NP-hard to determine whether a given graph G has a proper mostly-3-coloring.

   (b) A *Hamiltonian bicycle* in a graph G is a pair of simple cycles in G, with identical lengths, such that every vertex of G lies on exactly one of the two cycles.

A Hamiltonian bicycle in the 6×6 grid graph.

   Prove that it is NP-hard to determine whether a given graph G has a Hamiltonian bicycle.

(In fact, both of these problems are NP-hard, but we only want a proof for one of them.)

4. (a) Let L_a denote the set of all strings in $\{0,1,2\}^*$ that do not contain any symbol twice in a row. For example, this language includes the strings 0101201202, 1010101, 2, and the empty string ε, but it does not include the strings 01210220 or 11.

   - Describe a DFA or NFA that accepts L_a **and**
   - Give a regular expression that describes L_a.

   (You do not need to prove that your answers are correct.)

   (b) Let L_b denote the set of all strings $w \in \{0,1,2\}^*$ such that $\#(0,w)+\#(1,w)=\#(2,w)$. For example, this language includes the strings 01012222 and 20221020 and the empty string ε, but it does not include the string 01212 or 2120210.

   Prove that L_b is not a regular language.

5. Let G be a directed **acyclic** graph, in which every edge $e \in E$ has a weight $w(e)$, which could be positive, negative, or zero. We define the ***alternating length*** of any path in G to be the weight of the first edge, ***minus*** the weight of the second edge, ***plus*** the weight of the third edge, and so on. More formally, for any path $P = v_0 \to v_1 \to v_2 \to \cdots \to v_\ell$ in G, we define

$$AltLen(P) = \sum_{i=0}^{\ell-1}(-1)^i \cdot w(v_i \to v_{i+1}).$$

Describe an algorithm to find a path from s to t with the largest alternating length, given the graph G, the edge weights $w(e)$, and vertices s and t as input.

For example, given the graph shown below, your algorithm should return 5, which is the alternating length of the path $s \to u \to t$.

6. Recall that the *depth* of a vertex v in a binary tree T is the length of the unique path in T from v to the root of T. A binary tree is *complete* if every internal node has two children, and every leaf has exactly the same depth. An *internal subtree* of a binary tree T is any connected subgraph of T.

Describe and analyze a recursive algorithm to compute the *largest complete internal subtree* of a given binary tree. Your algorithm should return both the root and the depth of this internal subtree.

The largest complete internal subtree of this binary tree has depth 2.

> You have 180 minutes to answer six numbered questions.
> **Write your answers in the separate answer booklet.**
> Please return this question sheet and your cheat sheet with your answers.

1. For each statement below, there are two boxes in the answer booklet labeled "Yes" and "No". Check "Yes" if the statement is **always** true and "No" otherwise, and give a **brief** (at most one short sentence) explanation of your answer. **Assume $P \neq NP$.** If there is any other ambiguity or uncertainty about an answer, check "No". For example:

   - $x + y = 5$

     ☐ Yes   ☒ No   Suppose $x = 3$ and $y = 4$.

   - 3SAT can be solved in polynomial time.

     ☐ Yes   ☒ No   3SAT is NP-hard.

   - If P = NP then Jeff is the Queen of England.

     ☒ Yes   ☐ No   The hypothesis is false, so the implication is true.

   Read each statement *very* carefully; some of these are deliberately subtle!

   (a) Which of the following statements are true?

   - The solution to the recurrence $T(n) = 3T(n/3) + O(n^2)$ is $T(n) = O(n^2)$.

   - The solution to the recurrence $T(n) = 9T(n/3) + O(n)$ is $T(n) = O(n^2)$.

   - There is a forest with 374 vertices and 225 edges. (Recall that a *forest* is an undirected graph with no cycles.)

   - Given any directed graph G whose edges have positive weights, we can compute shortest paths from one vertex s to every other vertex of G in $O(VE)$ time using Bellman-Ford.

   - Suppose $A[1..n]$ is an array of integers. Consider the following recursive function:

   $$Rizz(i,k) = \begin{cases} 0 & \text{if } i > k \\ 1 & \text{if } i = k \\ \max \left\{ \left. \begin{array}{c} Rizz(i, j-1) + Rizz(j+1, k) \\ + A[i] \cdot A[j] \cdot A[k] \end{array} \right| i \le j \le k \right\} & \text{otherwise} \end{cases}$$

   We can compute $Rizz(1, n)$ by memoizing this function into a two-dimensional array $Rizz[1..n, 1..n]$, which we fill by decreasing i in the outer loop and increasing k in the inner loop, in $O(n^2)$ time.

   Problem 1 continues onto the next page.

1. *[continued]*

   (b) Which of the following statements are true for **at least one** language $L \subseteq \{0, 1\}^*$?

      - $(L^*)^*$ is finite.

      - L is decidable but its complement $\overline{L}$ is undecidable.

      - $\{\langle M \rangle \mid M$ accepts $L\}$ is undecidable.

      - L is the intersection of two NP-hard languages and L is finite.

      - There is a polynomial-time reduction from L to the halting problem.

   (c) Consider the following pair of languages:

      - TREE $= \{G \mid G$ is a connected undirected graph with no cycles$\}$
      - HAMPATH $= \{G \mid G$ is an undirected graph that contains a Hamiltonian path$\}$

      (For concreteness, assume that in both of these languages, graphs are represented by their adjacency matrices.) Which of the following statements are true, assuming $P \neq NP$?

      - TREE is NP-hard.

      - TREE $\cap$ HAMPATH is NP-hard.

      - TREE $\cup$ HAMPATH is NP-hard.

      - HAMPATH is undecidable.

   - A reduction from TREE to HAMPATH would imply $P = NP$.

   (d) Suppose there is a **polynomial-time** reduction R from some language $A \in \{0, 1\}^*$ to some other language $B \in \{0, 1\}^*$. Which if the following statements are **always** true, assuming $P \neq NP$?

      - Problem B is NP-hard.

      - If A is finite, then B is finite.

      - If A is NP-hard, then B is NP-hard.

      - If A is undecidable, then B is undecidable.

      - If $A \in P$, then $B \in P$.

2. Submit a solution to **exactly one** of the following problems.

    (a) A *theta-graph* is a connected undirected graph in which two vertices have degree 3, and all other vertices have degree 2. Equivalently, a theta-graph is the union of three undirected paths that have the same endpoints, but no other vertices in common. The *size* of a theta-graph is the total number of vertices.

    The 5×5 grid grad contains a theta-subgraph of size 24.

    Prove that is is NP-hard to compute the size of the largest theta-graph that is a subgraph of a given undirected graph G.

    (b) A *clique-partition* of a graph $G = (V, E)$ is a partition of the vertices V into disjoint subsets $V_1 \cup V_2 \cup \cdots \cup V_k$, such that for each index i, every pair of vertices in subset V_i is connected by an edge in G. The *size* of a clique partition is the number of subsets V_i.

    Prove that it is NP-hard to compute the minimum-size clique partition of a given undirected graph G.

    In fact, both of these problems are NP-hard, but we only want a proof for one of them. Don't forget to tell us which problem you've chosen!

3. A *triumph* in a sequence of integers (from the Latin *tri-* meaning "three" and *-umph* meaning "bodacious") is a consecutive triple of sequence elements whose sum is a multiple of 3. For example, the sequence

$$\langle 3, \underline{1, 4, \overline{1}, 5, 9}, 6, 2, 3, 5, \underline{8, 9, 7}, 9, 3, 2, 3, \underline{8, \overline{4, 6, 2}}, 6 \rangle$$

contains five triumphs (indicated by lines above and below).

    We say that one sequence A is *more triumphant* (or *less heinous*) than another sequence B if there are more triumphs in A than in B.

    Describe and analyze an algorithm to compute the number of triumphs in the most triumphant (or equivalently, least heinous) subsequence of a given array $A[1 \mathbin{..} n]$ of integers.

    For example, given the input array $\langle 0, 1, 1, 2, 3, 5, 8, 13, 21 \rangle$, your algorithm should return the integer 4, which is the number of triumphs in the most triumphant subsequence $\langle 0, 1, 2, 3, 8, 13, 21 \rangle$. Excellent!

Problems 4–6 appear on the next page.

4. Suppose we are given a directed graph $G = (V, E)$, where every edge $e \in E$ has a *positive* weight $w(e)$, along with two vertices s and t.

   (a) Suppose each *vertex* of G is colored either orange, green, or purple. Describe and analyze an algorithm to find the shortest walk from s to t in G that never visits two consecutive *vertices* with the same color.

   (b) Now suppose each *edge* of G is colored either orange, green, or purple. Describe and analyze an algorithm to find the shortest walk from s to t in G that never traverses two consecutive *edges* with the same color.

5. let T be a *full* binary tree, meaning that every node has either two children or no children.

   • Recall that the *height* of a vertex v in T is the length of the longest path in T from v down to a leaf. In particular, every leaf of T has height zero.

   • A vertex v is *AVL-balanced* if v is a leaf, or if the heights of v's children differ by at most 1. (You might recall from CS 225 that an **AVL-tree** is a binary search tree in which *every* vertex is AVL-balanced.)

   Describe and analyze an algorithm to compute the number of AVL-balanced vertices in T.

6. (a) Let L_a denote the set of all strings $w \in \{0, 1, 2\}^*$ such that $\#(1, w) + 2 \cdot \#(2, w)$ is divisible by 3. For example, L_a contains the strings `0012` and `20210202` and the empty string ε, but L_a does not include the strings `121` or `0122210`.

   Describe a DFA of NFA that accepts L_a. (You do not need to prove that your answer is correct.)

   (b) Let L_b denote the set of all strings $w \in \{0, 1, 2\}^*$ such that no two symbols appear the same number of times, or in other words, the integers $\#(0, w)$ and $\#(1, w)$ and $\#(2, w)$ are all different. For example, L_b contains the strings `110212` and `20220`, but L_b does not include the string `01212` or `2120210` or the empty string ε.

   Prove that L_b is not a regular language.