CS/ECE 374 Sec A ✧ Spring 2025

Homework 2

Due Wednesday, Feb 5th, 2025 at 9am

- You can work in a group of up to **three** students. Read the instructions on the course website for additional details.
- **Submit your solutions electronically on the course Gradescope site as PDF files.** Submit a separate PDF file for each numbered problem. If you plan to typeset your solutions, please 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, not just a phone camera).
- Late submissions will be accepted (for 75% credit) until midnight the day of the deadline.

☞ Some important course policies ☜

- **You may use any source at your disposal**—paper, electronic, or human—but you ***must*** cite ***every*** source that you use, and you ***must*** write everything yourself in your own words. See the academic integrity policies on the course web site for more details including the policy on using AI tools.
- **Avoid the Two Deadly Sins!** Any homework or exam solution that breaks any of the following rules will be given an ***automatic zero***, unless the solution is otherwise perfect. Yes, we really mean it. We're not trying to be scary or petty (Honest!), but we do want to break a few common bad habits that seriously impede mastery of the course material.
  - Always give complete solutions, not just examples.
  - Always declare all your variables, in English. In particular, always describe the specific problem your algorithm is supposed to solve.

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 Ed.

---

0 Several problems on creating regular expressions and DFAs can be found in Jeff's notes, Sipser's book, Aho-Motwani-Ullman book, and several others. They cover a wide range from easy to hard. We have placed several optional problem on PrarieLearn that you can also try out.

1. For each of the following languages assume that the alphabet is $\{\mathtt{0}, \mathtt{1}\}$ unless it is specified otherwise. Give a regular expression that describes that language, and **briefly argue why your expression is correct**.

   (a) All strings that end in **010** *and* contain **000** as a substring.

   (b) All strings that do not contain the *subsequence* **010**.

   (c) The complement of the language $\{\mathtt{0}, \mathtt{01}, \mathtt{10}, \mathtt{101}, \mathtt{011}, \mathtt{111}\}$.

   (d) The language $\{\mathtt{0}^a\mathtt{1}^b\mathtt{2}^c \mid a, b, c \geq 0, a \equiv b + c \mod 2\}$ over the alphabet $\{\mathtt{0}, \mathtt{1}, \mathtt{2}\}$.

   (e) Let r be a regular expression and let $L(r)$ be the language denoted by r. We want to remove the empty string from $L(r)$ and create a regular expression r' such that $L(r') = L(r) - \{\epsilon\}$.[1] One might try to "understand" $L(r)$ and write a regular expression r' from scratch for the new language, but this may be difficult if r is complicated and long, and doesn't give us a general procedure. Here we explore a recursive algorithmic procedure to obtain r' from r (in an automatic way) by considering the following ideas.

   - For each of the base cases $\emptyset, \epsilon, \mathtt{0}, \mathtt{1}$ obtain r' assuming that r is one of them.
   - Suppose $r = r_1 + r_2$ and r'_1 and r'_2 are regular expressions for $L(r_1) - \{\epsilon\}$ and $L(r_2) - \{\epsilon\}$ respectively. How would you obtain r' from r'_1, r'_2, r_1, r_2? Note that there may be multiple equivalent expressions; you can give any correct one.
   - Suppose $r = r_1r_2$. How would you obtain r' from r'_1, r'_2, r_1, r_2?
   - Suppose $r = (r_1)^*$. How would you obtain r' from r'_1, r_1?

   Given the ideas above, describe a recursive algorithm to obtain r' from an arbitrary regular expression r. You do not need to justify the constructions explicitly but it is in your own interest to think about the correctness of each case.

2. DFA design. For the first three parts describe a DFA by drawing it explicitly and **briefly explain the *meaning* of each state**. For the last two do not draw the DFA. Instead use formal notation, **still briefly explaining the *meaning* of each state**.

   (a) All strings in $\{\mathtt{0}, \mathtt{1}\}^*$ that end in **11**.

   (b) All strings in $\{\mathtt{0}, \mathtt{1}\}^*$ that start with **01** and whose length is divisible by 4.

   (c) All strings in $\{\mathtt{0}, \mathtt{1}\}^*$ that do *not* contain **100110** as a substring.

   (d) All strings whose 8th-to-last symbol is **1**, or equivalently, the set

   $$\{x\mathtt{1}y \mid x \in \Sigma^* \text{ and } y \in \Sigma^7\}$$

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

[1] It is not immediately obvious that $L(r) - \{\epsilon\}$ is regular but it turns out to be, and in fact we will end up giving a proof of that by solving this problem.

3. **Not for submission:** Consider a fixed string $s = a_1 a_2 \ldots a_k$ of length k. Let $L_s = \{w \in \Sigma^* \mid w \text{ contains string } s \text{ as a substring}\}$. Describe in formal tuple notation a DFA M that accepts L_s. Do not try to optimize number of states. Briefly describe the meaning of the states and your construction. How many states does your DFA have as a function of k? Suppose you want to optimize the number of states. How would you find a smaller DFA? Can you find a DFA with $O(k)$ states? Think of a concrete string such as $aabaab$. Note that a DFA for this problem gives a linear-time algorithm for string matching.

4. **Not for submission.** Consider the strings over the alphabet $\{0, 1, 2\}$ as representing ternary numbers (i.e., numbers in base 3). Let L be the language of strings that represent ternary numbers divisible by 5. For example, `120` would be in the language since $120_3 = 1 \cdot 3^2 + 2 \cdot 3 = 15$, while `200` would not.

   Describe a DFA over the alphabet $\Sigma = \{0, 1, 2\}$ that accepts the language L. Argue that your machine accepts every string in L and nothing else, by explaining what each state in your DFA *means*.

   You may either draw the DFA or describe it formally, but the states Q, the start state s, the accepting states A, and the transition function δ must be clearly specified.

5. **Not for submission.** Give a string w we use the notation w^R to denote its reverse. We can extend this notion to languages as follows. Given a language $L \subseteq \Sigma^*$ we define the reverse of L, denoted by L^R as follows:

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

   In other words L^R is the language consisting of the reverses of the strings in L. In this problem your goal is to develop an algorithm that given a regular expression r, converts it into a regular expression r' such that $L(r')$ is $(L(r))^R$. As a byproduct this also shows that L^R is regular whenever R is regular. Fix Σ to be $\{0, 1\}$ for simplicity.

   - Consider each of the base cases of the regular expressions. For each r corresponding to the base cases define r' appropriately such that $L(r') = (L(r))^R$.
   - Suppose $r = r_1 + r_2$. Assuming that you have found, recursively, regular expressions r_1' and r_2' for the reverses of r_1, r_2 respectively, describe a regular expression r' for the reverse of $L(r)$.
   - Suppose $r = r_1 r_2$. Do the same as in the preceding part.
   - Suppose $r = (r_1)^*$. Find a regular expression r' for the reverse of $L(r)$.
   - Assuming $r = 0^* + (01 + 100)^*(11^* + \epsilon + 0)$ what would your algorithm output for the reverse of $L(r)$?

Solved problem

4. ***C comments*** are the set of strings over alphabet $\Sigma = \{$⋆, /, A, ◇, ↲$\}$ that form a proper comment in the C program language and its descendants, like C++ and Java. Here ↲ represents the newline character, ◇ represents any other whitespace character (like the space and tab characters), and A represents any non-whitespace character other than ⋆ or /.[2] There are two types of C comments:

   - Line comments: Strings of the form //⋯↲.
   - Block comments: Strings of the form /⋆⋯⋆/.

   Following the C99 standard, we explicitly disallow ***nesting*** comments of the same type. A line comment starts with // and ends at the first ↲ 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:

   - /⋆⋆⋆/
   - //◇//◇↲
   - /⋆///◇⋆◇↲⋆⋆/
   - /⋆◇//◇↲◇⋆/

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

   - /⋆/
   - //◇//◇↲◇↲
   - /⋆◇/⋆◇⋆/◇⋆/

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

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

   You must explain *in English* how your DFAs work. Drawings or formal descriptions without English explanations will receive no credit, even if they are correct.

---

[2]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 (“⋯“) or a (multi-)character literal (’⋯’).
- 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 (\")
- The opening single-quote of a character literal must not be inside a string literal ("⋯’⋯") or a comment.
- The closing single-quote of a character literal must not be escaped (\’)
- A backslash escapes the next symbol if and only if it is not itself escaped (\\) or inside a comment.

For example, the string "/⋆\\\"⋆/"/⋆"/⋆\"/⋆"⋆/ is a valid string literal (representing the 5-character string /⋆\"\⋆/, which is itself a valid block comment!) followed immediately by a valid block comment. ***For this homework question, just pretend that the characters ’, ", and \ 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.

Solution:

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

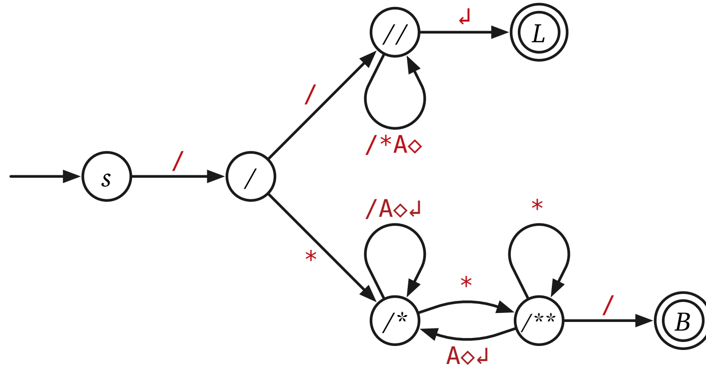

The states are labeled mnemonically as follows:

- s — We have not read anything.
- / — We just read the initial /.
- // — We are reading a line comment.
- L — We have 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 read a complete block comment.

(b) 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.

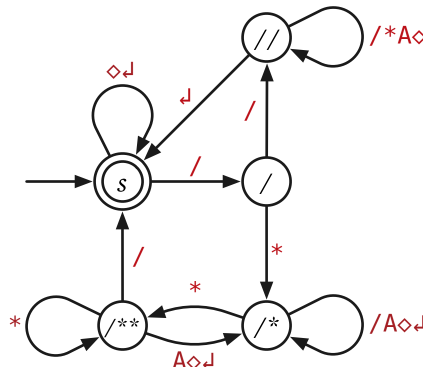

The states are labeled mnemonically as follows:

- 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: 10 points = 5 for each part, using the standard DFA design rubric (scaled)

Rubric (DFA design): For problems worth 10 points:

- 2 points for an unambiguous description of a DFA, including the states set Q, the start state s, the accepting states A, and the transition function δ.
  - **For drawings:** Use an arrow from nowhere to indicate s, and doubled circles to indicate accepting states A. If $A = \varnothing$, say so explicitly. If your drawing omits a reject state, say so explicitly. **Draw neatly!** If we can't read your solution, we can't give you credit for it,.
  - **For text descriptions:** You can describe the transition function either using a 2d array, using mathematical notation, or using an algorithm.
  - **For product constructions:** You must give a complete description of the states and transition functions of the DFAs you are combining (as either drawings or text), together with the accepting states of the product DFA.
- **Homework only:** 4 points for *briefly* and correctly explaining the purpose of each state *in English*. This is how you justify that your DFA is correct.
  - For product constructions, explaining the states in the factor DFAs is enough.
  - **Deadly Sin:** ("Declare your variables.") No credit for the problem if the English description is missing, *even if the DFA is correct*.
- 4 points for correctness. (8 points on exams, with all penalties doubled)
  - -1 for a single mistake: a single misdirected transition, a single missing or extra accept state, rejecting exactly one string that should be accepted, or accepting exactly one string that should be accepted.
  - -2 for incorrectly accepting/rejecting more than one but a finite number of strings.
  - -4 for incorrectly accepting/rejecting an infinite number of strings.
- DFA drawings with too many states may be penalized. DFA drawings with *significantly* too many states may get no credit at all.
- Half credit for describing an NFA when the problem asks for a DFA.