



CS 51: Introduction to Computer Programming II
Final Examination
Spring, 2016

We recommend that you read all instructions carefully and look over the exam in its entirety before beginning the exam.

This is an open-book exam: You are free to use books and notes in preparing your solutions for this exam. However, no electronic devices of any kind may be used.

We have provided a separate handout with reference material: a “cheat sheet” on OCaml and the lab code for NativeLazyStreams.

The exam is in five sections comprising 26 problems. Numbers in brackets are the points (out of 100 total) allocated to the problem and may provide an approximate recommendation for allocating time.

Write the answers to all problems in the boxes provided. Anything written outside of the provided boxes will not be graded. If additional room is needed for an answer, make a note in the box provided and write the remainder of your answer in one of the extra boxes at the back of the examination paper.

Many of the problems ask you to define something or write code to do something. Throughout the exam, when we ask you to define a value or function or type or class or module, we mean that you should provide a top-level OCaml definition written in idiomatic OCaml using the appropriate OCaml definitional construct (`let`, `type`, `class`, `module`, etc.).

Please write your name and ID number clearly in the boxes below, but nowhere else on the exam to maintain anonymity of grading.

YOUR NAME:

SOLUTION

YOUR HARVARD ID NUMBER:

1. CIRCUITS AND BOOLEAN STREAMS

A Boolean circuit is a device with one or more inputs and a single output that receives over time a sequence of Boolean values on its inputs and converts them to a corresponding sequence of Boolean values on its output. The building blocks of circuits are called *gates*. For instance, the *and* gate is a Boolean device with two inputs; its output is true when its two inputs are both true, and false if either output is false. The *not* gate is a Boolean device with a single input; its output is true when its input is false and vice versa.

In this problem, you'll generate code for modeling Boolean circuits. The inputs and outputs will be modeled as *lazy Boolean streams*. For your reference, we've provided the pertinent code for the `NativeLazyStreams` module from lab in the separate reference material handout. You can (and should) make use of that module as if it had been opened (so no module name prefix is required).

Let's start with an infinite stream of false values.

Problem 1. [4 points] Define a value `false`s to be an infinite stream of the Boolean value `false`.

Solution:

```
let rec false =
  lazy (Cons(false, false)) ;;
```

Problem 2. [1 points] What is the type of `false`s?

Solution: `false`s: `bool stream`

A useful function is the `trueat` function. The expression `trueat n` generates a stream of values that are all false except for a single true at index `n`:

```
# first 5 (trueat 1) ;;
- : bool list = [false; true; false; false; false]
```

You need not write this function; you should just assume you have access to this function, as it is used in demonstrating the functions described below.

Problem 3. [4 points] Define a function `circnot` : `bool stream -> bool stream` to represent the not gate. It should have the following behavior:

```
# first 5 (circnot falses) ;;
- : bool list = [true; true; true; true; true]
```

Solution:

```
let circnot : bool stream -> bool stream =
  smap (not) ;;
```

Note: use of the already built `NativeLazyStreams smap`; advantage taken of currying.

Problem 4. [4 points] Define a function `circand` to represent the and gate. It should have the following behavior:

```
# first 5 (circand (circnot (trueat 1)) (circnot (trueat 3))) ;;
- : bool list = [true; false; true; false; true]
```

Solution:

```
let circand : bool stream -> bool stream -> bool stream =
  smap2 (&&) ;;
```

Note: use of the already built `NativeLazyStreams smap2`; advantage taken of currying.

A *nand* gate is a gate that computes the negation of an *and* gate. That is, it negates the *and* of its two inputs, so that its output is false only if both of its inputs are true.

Problem 5. [4 points] Succinctly define a function `circnand` using the functions above to represent the nand gate. It should have the following behavior:

```
# first 5 (circnand falses (trueat 3)) ;;
- : bool list = [true; true; true; true; true]
# first 5 (circnand (trueat 3) (trueat 3)) ;;
- : bool list = [true; true; true; false; true]
```

Solution:


```
let circnand (s: bool stream) (t: bool stream) : bool stream =
  circnot (circand s t) ;;
```

Note: use of the already built `circnot` and `circand`.

2. CREDIT CARD NUMBERS AND THE LUHN ALGORITHM

Here's an interesting bit of trivia: Not all credit card numbers are well-formed. The final digit in a 16-digit credit card number is in fact a *checksum*, a digit that is computed from the previous 15 by a simple algorithm.

The algorithm used to generate the checksum is called the *Luhn check*. To calculate the correct final checksum digit used in a 16-digit credit card number, you perform the following computation on the first 15 digits of the credit card number:



- (1) Take all of the digits in an odd-numbered position (the leftmost digit being the first, not the zero-th digit, hence an odd-numbered one) and double them, subtracting nine if the doubling is greater than nine (called “casting out nines”).

As an example, we'll use the (partial) credit card number in the sample above:

$$\underline{4}2\underline{7}5 \quad \underline{3}1\underline{5}6 \quad \underline{0}3\underline{7}2 \quad \underline{5}4\underline{9}x$$

Here, the odd-numbered digits (4, 7, 3, 5, 0, 7, 5, and 9) have been underlined. We double them and cast out nines to get 8, 5, 6, 1, 0, 5, 1, and 9.

- (2) Add all of the even position digits and the doubled odd position digits together. For the example above, the sum would be

$$(2 + 5 + 1 + 6 + 3 + 2 + 4) + (8 + 5 + 6 + 1 + 0 + 5 + 1 + 9) = 23 + 35 = 58 \quad .$$

- (3) The checksum is then the digit that when added to this sum makes it a multiple of ten. In the example above the checksum would be 2, since adding 2 to 58 generates 60, which is a multiple of 10. Thus, the sequence 4275 3156 0372 5492 is a valid credit card number, but changing the last digit to any other makes it invalid. (In particular, the final 3 in the image above is not the correct checksum!)

Problem 6. [5 points] Define an explicitly recursive polymorphic function `odds` to extract the elements at odd-numbered indices in a list, where the indices are counted starting with 1, so that

```
# let cc = [4;2;7;5;3;1;5;6;0;3;7;2;5;4;9];;
val cc : int list = ...
# odds cc ;;
- : int list = [4; 7; 3; 5; 0; 7; 5; 9]
```

Solution:

```
let rec odds lst =
  match lst with
  | [] -> []
  | [a] -> [a]
  | a :: _b :: rest -> a :: odds rest ;;
```

Problem 7. [2 points] What is the type of the `odds` function that you defined?

Solution:

```
odds: 'a list -> 'a list
```

Note the polymorphic type.

Problem 8. [3 points] What is an appropriate recurrence equation for defining the time complexity of `odds` in terms of the length of its list argument?

Solution:

$$T_{\text{odds}}(n) = c + T_{\text{odds}}(n - 2)$$

More detail in the equation in terms of constants for different bits is unnecessary, but benign. Note the $n - 2$, though $n - 1$ yields the same complexity.

Problem 9. [3 points] What is the time complexity of odds (in big-O notation)?

Solution: Linear – $O(n)$.

Problem 10. [7 points] If the function $f(n)$ is the time complexity of odds on a list of n elements, which of the following is true? Check all that apply:

- $f \in O(1)$
- $f \in O(\log n)$
- $f \in O(\log n/c)$ for all $c > 0$
- $f \in O(c \cdot \log n)$ for all $c > 0$
- $f \in O(n)$
- $f \in O(n/c)$ for all $c > 0$
- $f \in O(c \cdot n)$ for all $c > 0$
- $f \in O(n^2)$
- $f \in O(n^2/c)$ for all $c > 0$
- $f \in O(c \cdot n^2)$ for all $c > 0$
- $f \in O(2^n)$
- $f \in O(2^n/c)$ for all $c > 0$
- $f \in O(c \cdot 2^n)$ for all $c > 0$

Solution: The O classes are independent of multiplication or division by constants, so each “triplet” of answers after the first are equivalent. Since f is $O(n)$, it is also $O(n^2)$ etc. for all higher classes. Thus, all boxes from the fifth on should be checked.

In addition to `odds`, it will be useful to have a function `evens` that extracts the elements at even-numbered indices in a list. Rather than have you write that, we'll just assume that it's available.

The process of doubling a number and "casting out nines" is easy to implement as well. Here is some code to do that:

```
let double9 n =  
  (n * 2 - 1) mod 9 + 1;;
```

Finally, it will be useful to have a function to sum a list of integers.

Problem 11. [3 points] Implement the function `sum` using the tail-recursive `List` module function `fold_left`.

Solution:

```
let sum =  
  fold_left (+) 0 ;;
```

Note: advantage taken of currying.

All the parts are now in place to implement the Luhn check algorithm.

Problem 12. [8 points] Implement a function `luhn` that takes a list of integers and returns the check digit for that digit sequence. (You can assume that it is called with a list of 15 integers.) For instance, for the example above

```
# luhn cc ;;
- : int = 2
```

You should feel free to use the functions `evens`, `odds`, `double9`, `sum`, and any other OCaml library functions that you find useful and idiomatic.

Solution:

```
let luhn nums =
  let s = sum ((List.map doublemod9 (odds nums)) @ (evens nums)) in
  10 - (s mod 10) ;;
```

Note: prefix on map (though a local or global open is acceptable).

Problem 13. [5 points] What is the time complexity of the `luhn` function in terms of the length n of its list argument? Use big- O notation. Explain why your implementation has that complexity.

Solution: $O(n)$ – linear. The `odds` and `evens` function are both linear and return a list of length linear in n . The `append` is linear in the length of the `odds` list, so also linear in n . The `sum` is linear in length of its argument, which is identical in length to (and thus linear in) n . The `let` body is constant time.

Problem 14. [1 points] Now that you know how to generate valid credit card numbers not your own, would it be legal for you to use these numbers on an e-commerce web site to purchase goods? Would it be ethical for you to do so?

Solution: No and no.

3. EFFICIENT ALGORITHMS FOR EUCLID CORPORATION

Problem 15. [9 points] Your friend, who works at EuclidCo, tells you that he's looking for a fast algorithm to solve a problem he's working on. So far, he's come across two algorithms, algorithm A, which has time complexity $O(n^3)$, and algorithm B, which is $O(2^n)$. He prefers algorithm A, and uses three different arguments to convince you of his decision. For each, evaluate the truth of the bolded statement, checking the appropriate box, and justify your answer in the provided space.

- (1) "We're all about speed at EuclidCo, and **A will always be faster than B.**"
- (2) "In a high stakes industry like ours, we can't afford to have more than a finite number of inputs that run slower than polynomial time, and **we can avoid this if we go with A.**"
- (3) "We work with big data at EuclidCo. **For suitably large inputs, A will be faster on average than B.**"

Solution:

- (1) Big O notation only gives you information about the worst-case performance as the input size becomes very large. Because of this, it ignores lower-order terms and constants that may have a large effect for small inputs, and so A may be slower than B for small inputs, and the statement is *false*.
- (2) Since Big O notation provides worst-case performance, and A is polynomial in Big O, they can be guaranteed that for any input (except for a finite set), A will run in polynomial time, so the statement is *true*.
- (3) As a worst-case bound, Big O doesn't say anything about average-case performance, so the statement is *false*.

4. MUTABLE LIST MISCELLANY

Warning: The definition of mutable lists defined and used in this section is different from that in the version of the textbook and labs as of spring 2023.

Recall the definition of the mutable list type from Lab 5:

```
type 'a mlist =
  | Nil
  | Cons of 'a * ('a mlist ref) ;;
```

The following functions extract the head and the (dereferenced) tail from a mutable list.

```
let mhead (Cons(hd, _tl)) = hd ;;
let mtail (Cons(_hd, tl)) = !tl ;;
```

The first function returns a list (immutable) of the first n elements of a mutable list `mlst`:

```
let rec first (n: int) (mlst: 'a mlist) : 'a list =
  if n = 0 then []
  else match mlst with
    | Nil -> []
    | Cons(hd, tl) -> hd :: first (n-1) !tl ;;
```

Problem 16. [5 points] Write code to define a mutable integer list *alternating* such that for all integers n , the expression `first n alternating` returns a list of alternating 1s and 2s, for example,

```
# first 5 alternating ;;
- : int list = [1; 2; 1; 2; 1]
# first 10 alternating ;;
- : int list = [1; 2; 1; 2; 1; 2; 1; 2; 1; 2]
```

Solution:

```
let rec alternating =
  Cons(1, ref (Cons(2, ref alternating))) ;;
```

Note: the refs.

Answer the questions about the following interaction with the OCaml REPL, which makes use of the mutable list type and functions previously defined. (We've redacted portions of the OCaml responses.)

```
# let a = Cons(1, ref Nil);;
val a : ...redaction 1...
# let b = Cons(a, ref (Cons(a, ref Nil)));;
val b : ...redaction 2...
# let q1 = (mhead b == mhead (mtail b));;
val q1 : ...redaction 3...
# let q2 = first 5 a;;
val q2 : ...redaction 4...
# let _ =
  match mtail b with
  Cons(hd, tl) -> tl := hd;;
...redaction 5...
# let _ =
  match mhead (mtail b) with
  Cons(hd, tl) -> tl := a;;
...redaction 6...
# let q3 = first 5 a;;
val q3 : ...redaction 7...
```

Problem 17. [2 points] What is the type of *a* at redaction 1?

Solution: `int mlist`

Problem 18. [2 points] What is the type of *b* at redaction 2?

Solution: `int mlist mlist.`

Problem 19. [2 points] What is the value of *q1* at redaction 3?

Solution: `true`

Problem 20. [2 points] What is the value of *q2* at redaction 4?

Solution: `[1]`

Problem 21. [2 points] Briefly explain what happens at redaction 5.

Solution: The expression fails to type-check. The *hd* is of type `int mlist`, but the *tl* is a `ref` to an `int mlist mlist`.

Problem 22. [2 points] *Briefly explain what happens at redaction 6.*

Solution: The second element of the list `b`, that is, the list `a`, is modified so that its tail points to itself, so `a` is mutated to become a cyclic list of 1s.

Problem 23. [2 points] *What is the value of `q3` at redaction 7?*

Solution: `[1,1,1,1,1]`

5. OBJECT-ORIENTED COUNTERS

Here is a class type and class definition for counter objects. Each object maintains an integer state that can be “bumped” by adding an integer. The interface guarantees that only the two methods are revealed.

```
class type counter_interface =
  object
    method bump : int -> unit
    method get_state : int
  end

class counter : counter_interface =
  object
    val mutable state = 0
    method bump n = state <- state + n
    method get_state = state
  end
```

Problem 24. [6 points] Write a class definition for a class `loud_counter` obeying the same interface that works identically, except that it also prints the resulting state of the counter each time the counter is bumped.

Solution:

```
class loud_counter : counter_interface =
  object (this)
    inherit counter as super
    method bump n =
      super#bump n;
      Printf.printf "State is now %d\n" this#get_state
  end
```

Note: the use of inheritance rather than reimplementing; the overwriting of `bump`; the use of `super#bump` rather than reimplementing the old `bump`; the use of `this#get_state`.

Problem 25. [6 points] Write a class type definition for an interface `reset_counter_interface`, which is just like `counter_interface` except that it has an additional method of no arguments intended to reset the state back to zero.

Solution:

```
class type reset_counter_interface =
  object
    inherit counter_interface
    method reset : unit
  end
```

Note: the use of inheritance rather than reimplementation.

Problem 26. [6 points] Write a class definition for a class `loud_reset_counter` satisfying the `reset_counter_interface` that implements a counter that both allows for resetting and is “loud” (printing the state whenever a bump or reset occurs).

Solution:

```
class loud_reset_counter : reset_counter_interface =
  object (this)
    inherit loud_counter as super
    method reset =
      this#bump (-this#get_state)
  end ;;
```

Note: the use of inheritance rather than reimplementation; the use of `bump` to change the state rather than trying to use the unavailable mutable variable directly (`state <- 0`, for instance).

End of the exam!