
Programming Principles

Midterm Solution

Friday, April 17 2015

Exercise 1: Associativity (10 points)

Definition (1 pt)

Give the definition of associativity. What does it mean for an operator f to be associative?

Answer

$$\forall x, y, z : f(x, f(y, z)) = f(f(x, y), z)$$

Associativity and commutativity (1 pt)

Give an example operator that is associative but that is *not* commutative.

Answer

Possible answers are:

- Matrix multiplication
- String/list concatenation
- $f(x, y) = x$

An associative operator (8 pt)

Given two types A and B , and two *associative* operators f and g with the following signatures:

- $f: (A, A) \Rightarrow A$
- $g: (B, B) \Rightarrow B$

Define a third operator h with the following signature:

- $h: ((A, B), (A, B)) \Rightarrow (A, B)$

such that h is associative. (1 pt)

Then, prove that h is indeed associative. (7 pt)

Be **very precise** in your proof. At each step, specify which axiom, hypothesis, or previous result you use.

Answer

We define h as:

$$h((x_1, x_2), (y_1, y_2)) = (f(x_1, y_1), g(x_2, y_2))$$

We know that

$$\forall x, y, z : f(x, f(y, z)) = f(f(x, y), z)$$

$$\forall x, y, z : g(x, g(y, z)) = g(g(x, y), z)$$

We have to prove that

$$\forall x, y, z : h(x, h(y, z)) \stackrel{?}{=} h(h(x, y), z)$$

Replacing x by (x_1, x_2) , and similarly for y and z :

$$\forall x_1, x_2, y_1, y_2, z_1, z_2 : h((x_1, x_2), h((y_1, y_2), (z_1, z_2))) \stackrel{?}{=} h(h((x_1, x_2), (y_1, y_2)), (z_1, z_2))$$

From here on, I will assume implicitly the \forall quantification.

Replacing h by its definition, in the inner invocation, on both sides, yields:

$$h((x_1, x_2), (f(y_1, z_1), g(y_2, z_2))) \stackrel{?}{=} h((f(x_1, y_1), g(x_2, y_2)), (z_1, z_2))$$

Replacing h by its definition a second time, on both sides, yields:

$$(f(x_1, f(y_1, z_1)), g(x_2, g(y_2, z_2))) \stackrel{?}{=} (f(f(x_1, y_1), z_1), g(g(x_2, y_2), z_2))$$

Separating the equality of tuples into equality of their members:

$$f(x_1, f(y_1, z_1)) \stackrel{?}{=} f(f(x_1, y_1), z_1)$$

$$g(x_2, g(y_2, z_2)) \stackrel{?}{=} g(g(x_2, y_2), z_2)$$

Which is true because of f and g are associative.

QED

Exercise 2: Parallel Merge Sort (10 points)

In this exercise, you are required to design and analyze a parallel version of the merge sort algorithm. Consider the function `msort` shown below that implements a *sequential* merge sort algorithm for sorting a list of objects of type `T`. Recall that a merge sort algorithm partitions the input list in two parts (whose sizes differ by at most one), sorts each of the parts independently, and finally assembles the two sorted parts into a single sorted list.

```
def msort[T](less: (T, T) => Boolean)(l: List[T]): List[T] = {
  if (l.length <= 1) {
    l
  } else {
    val (first, second) = partition(l)
    merge(less)(msort(less)(first), msort(less)(second))
  }
}
```

The parameter `less` defines an ordering on the instances of type `T`, and takes *constant time*. Given below are the implementations of the function `merge` that combines two sorted lists into a single sorted list, and the function `partition`.

```
def merge[T](less: (T, T) => Boolean)(xs: List[T], ys: List[T]): List[T] = {
  (xs, ys) match {
    case (Nil, _) => ys
    case (_, Nil) => xs
    case (x :: xtail, y :: ytail) =>
      if (less(x, y))
        x :: merge(less)(xtail, ys)
      else
        y :: merge(less)(xs, ytail)
  }
}

def partition[T](l: List[T]): (List[T], List[T]) = {
  val n = l.length / 2;
  (l take n, l drop n)
}
```

You are given a `parallel` construct that can execute two function calls in parallel:

```
def parallel[A, B](taskA: => A, taskB: => B): (A, B)
```

- (a) Which *one* of the above three functions can you parallelize to obtain a span of $O(n)$ for the function `msort`, where n is the size of the input list? You are allowed to introduce one or more `parallel` constructs but you are not allowed to use a different data structure (other than the linked list). Recall that *span* is (informally) the time taken by an algorithm when the number of parallel processors is unlimited. See below for a more formal definition.
- (b) Implement the parallel version of the function that you identified in part (a).

- (c) Prove that the span of the merge sort algorithm is $O(n)$ when the implementation you gave for part (b) is used ?

For your reference, a brief formal definition of span, which was explained in the lectures, is given below. For an expression $f(e_1, \dots, e_n)$, where f is a function or a primitive operation, $S(f(e_1, \dots, e_n)) = S(e_1) + \dots + S(e_n) + S(f)(v_1, \dots, v_n)$, where v_i denotes the values of e_i . If f is a primitive operation on integers, then $S(f)$ is constant, regardless of v_i . If f is a user-defined function, then $S(f)(v_1, \dots, v_n)$ is equal to the span of the body of f when the values of the parameters are v_1, \dots, v_n . For the parallel construct, $S(\text{parallel}(c1, c2)) = c_p + \max(S(c1), S(c2))$, where c_p is a constant. For the if-then-else construct, $S(\text{if}(e1) \ e2 \ \text{else} \ e3) = S(e1) + S(e2) + c$ if the condition $e1$ holds, otherwise is equal to $S(e1) + S(e3) + c$, where c is a constant. Since the match construct can be reduced to if-then-else, its span can be derived using the span of if-then-else.

Answers

(a) `msort`

(b)

```
def msort[T](less: (T, T) => Boolean)(l: List[T]): List[T] = {
  if (l.length <= 1) {
    l
  } else {
    val (first, second) = partition(l)
    val (sortedFirst, sortedSecond) =
      parallel(msort(less)(first), msort(less)(second))
    merge(less)(sortedFirst, sortedSecond)
  }
}
```

(c)

The `then` part of the `if` is trivially $O(1)$.

For the `else` part, informally we have:

$$S(\text{else part}) = S(\text{partition}) + \max(S(\text{msort}(\text{first})), S(\text{msort}(\text{second}))) + S(\text{merge}) + c$$

Defining n to be the length of the list l :

$$S(n) = O(n) + \max(S(\lceil n/2 \rceil), S(\lfloor n/2 \rfloor)) + O(n) + c$$

$$S(n) = O(n) + \max(S(\lceil n/2 \rceil), S(\lfloor n/2 \rfloor))$$

$$S(n) = O(n) + S(\lceil n/2 \rceil)$$

$$S(n) = O(n)$$