
Parallel Programming

Midterm Exam

Wednesday, April 27, 2016

Your points are *precious*, don't let them go to waste!

Your Time All points are not equal. Note that we do not think that all exercises have the same difficulty, even if they have the same number of points.

Your Attention The exam problems are precisely and carefully formulated, some details can be subtle. Pay attention, because if you do not understand a problem, you can not obtain full points.

Exercise	Points	Points Achieved
1	25	
2	25	
3	25	
4	25	
Total	100	

Exercise 1: Associativity (25 points)

For the this exercise, you will look at multiplication of 2 by 2 matrices containing 32-bits integers. Remember that multiplication of such matrices is defined as follows:

$$AB = C \iff \forall 1 \leq i \leq 2, 1 \leq k \leq 2. C_{ik} = \sum_{j=1}^2 A_{ij} \cdot B_{jk}$$

Also note that 32-bits integers form a commutative ring, meaning that both addition and multiplication are associative and commutative, and that multiplication distributes over addition.

Question 1: Definition (5 points)

State the associative property for matrix multiplication.

Question 2: Proof of associativity (15 points)

Prove that multiplication of 32-bits integer matrices of dimension 2 by 2 is associative.

Hint: To show that two matrices X and Y are equal, you may show that for all indexes i and j that X_{ij} and Y_{ij} are equal.

Question 3: Associativity in parallel programming (5 points)

Why is associativity interesting in the context of parallel programming? Circle the best answer from the list below:

1. Associative operations always have efficient lock-free implementations.
2. The order of parameters of an associative operation can be freely swapped, giving better latency.
3. Associativity opens up possibilities to balance work.
4. Associative operations are more efficiently executed on Intel processors.

Exercise 2: Parallelizing Fold with Minus (25 points)

Let `a` be an `Array[Int]`. We are interested in the value of `acc` computed by the following piece of code:

```
var acc = 0
for(i <- (a.length - 1) to 0 by (-1))
  acc = a(i) - acc
```

For example, with `a = Array(4, 3, 6, 2)` we obtain `acc = 4-(3-(6-(2-0))) = 5`.

Question 1 (10 points)

Complete the template shown below so that the value of `acc` returned by the template is same as that computed by the above code snippet. For your reference, `a.zipWithIndex` returns a new array consisting of all elements of `a` paired with their index. Eg. `Array(1, 2, 5).zipWithIndex = Array((1, 0), (2, 1), (5, 2))`

```
val acc = a.zipWithIndex
  .map{ case (value, index) =>
```

```
  }.sum
```

Question 2 (10 points)

We are now interested by a more efficient way of computing `acc`, which combines the two traversals of the array and make the traversal in parallel. For that, we will use the `parallel` construct we saw in the lectures. For your reference, you will below an implementation of the `parallel` construct. Note that `t = Task { op }` creates a new thread `t` to execute `op`, and `t.join()` waits and returns the result of `op` when `t` finishes.

```
def parallel[A, B](op1: =>A, op2: =>B): (A, B) = {
  val res1 = Task { op1 }
  val res2 = op2
  (res1.join(), res2)
}
```

Complete the code snippet shown below so that `computeAcc(a, 0, a.length)` returns the value of `acc`. The function may recursively call itself on an array segment between indices `start` (inclusive) and `end` (exclusive). Some input/output examples: `computeAcc(Array(1,2,7,13,5), 0, 5) = -2`
`computeAcc(Array(1,2,7,13,5), 1, 4) = -8`

```
def computeAcc(a: Array[Int], startIndex: Int, end: Int): Int = {  
  if(end - start == 1) {  
  
    } else {  
  
  
  
  
  
  
  
  
  }  
}
```

Question 3 (5 points)

Suppose we call `computeAcc` with the following arguments:

```
computeAcc(Array(112, 97, 114, 97, 108, 108, 101, 108), 0, 8)
```

How many `Tasks` will be created when the above described implementation of `parallel` is used?

Exercise 3: Parallel Quick Sort (25 points)

In this exercise, you are required to design and analyze a parallel version of the quick sort algorithm. Consider the function `qsort` shown below that implements a *sequential* quick sort algorithm for sorting (in ascending order) a *list* of `Int`.

```
def qsort(list: List[Int]): List[Int] = {
  if (list.size <= 1) list
  else {
    val (smaller, eqs, bigger) = partition(list, choosePivot(list))
    qsort(smaller) ++ eqs ++ qsort(bigger)
  }
}
```

A brief explanation of the quick sort algorithm

- (a) If the input list has zero or one element, the algorithm returns the list as it is.
- (b) Otherwise, the algorithm chooses an element of the input list referred to as the *pivot*.
- (c) It then partitions the input list into three parts such that the first part (**smaller**) contains all elements smaller than the pivot, the second part (**eqs**) contains all elements equal to the pivot, and the last part (**bigger**) contains all elements greater than the pivot. Note that the size of the parts may depend on the choice of the pivot, and may not necessarily be equal.
- (d) Finally, the algorithm recursively sorts the first and the last parts, and concatenates the results.

If the *pivot* is chosen arbitrarily, the *depth* of the `qsort` algorithm is $O(n^2)$ in the worst case, for a list of size n . Now say, instead of arbitrarily choosing the pivot, we always choose the *median* of the list as the pivot. The median of a list of integers is the middle element in the sorted order of the integers. We can compute the *median* of an arbitrary list `l` (of size n) in worst case *linear time*: $O(n)$ by using an algorithm called *median of medians*. Note that when the median is chosen as the pivot, at most $\lfloor \frac{(n-1)}{2} \rfloor$ elements are smaller than the pivot, and at most $\lceil \frac{(n-1)}{2} \rceil$ elements are greater than the pivot.

Question 1 (15 points) Show a parallel implementation of the `qsort` function whose *depth* is $O(l.size)$, given that `choosePivot` uses the median. You can use the `parallel` construct explained in the previous question. You don't have to show the implementations of the `partition` or the `choosePivot` functions.

Question 2 (10 points) Prove that the depth of the parallel `qsort` function is $O(l.size)$. Note that for computing the depth of `qsort` you need to come up with a suitable *depth* for the `partition` function. You need not prove the *depth* of the `partition`, or `choosePivot` functions.

Hint: In your proof you can first construct a *recurrence relation*, and come up with a solution for it.

For your reference, a brief formal definition of depth, 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 value of e_i . If f is a primitive operation on integers e.g. $+$, $-$, $*$ etc., then $S(f)$ is a constant regardless of v_i . If f is a user-defined function, then $S(f)(v_1, \dots, v_n)$ is equal to the depth of the body of f when the values of the parameters are v_1, \dots, v_n . For the parallel construct, $S(\text{parallel}(c1, c2))$ is equal to $c_p + \max(S(c1), S(c2))$, 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 depth can be derived using the depth of if-then-else.

Exercise 4: Read/Write Lock (25 points)

In this question, you have to implement a class `ReadWrite` that implements a *read/write* lock. Unlike a traditional monitor that has a single operation to build mutual exclusion blocks (`synchronized`), read/write locks have two methods:

```
read(op) // performs 'op' as a reader
write(op) // performs 'op' as a writer
```

The `read` and `write` methods should satisfy the following constraints:

- Only one writer may own the lock at any point in time. That is, no two `write` operations can execute concurrently.
- Multiple readers can own the lock concurrently. That is, multiple `read` operations can happen concurrently.
- A writer may own the lock only if no reader owns it. That is, a `write` operation can happen only if there are no concurrent `read` operations.

Question 1 (15 points)

Implement both `read` and `write` in the class stub on the *next page* and write a small explanation about it just below: (declare any vars you'd need at the beginning of the class). You may use `synchronized`, `wait`, and `notifyAll` methods to implement the `read` and `write` methods.

Explain your solution in at most six sentences here

Hints:

- Since multiple concurrent reads should be allowed to happen, it is important that the `op` operation in `read` is itself executed outside of a `synchronized` block.
- You may want to keep track of the number of reader executing at the same time.

```

class ReadWrite extends Monitor {
  // declare any private variables needed here

  def read[T](op: => T): T = {
    synchronized {
      // fill in what is needed here

    }
    try {
      op
    }
    finally {
      synchronized {
        // fill in what is needed here

      }
    }
  }

  def write[T](op: => Unit) = synchronized {
    // fill in what is needed here

  }
}

```

Question 2 (10 points) In most real world application, reads are far more frequent than writes; the lock implemented in *question 1* would be highly inefficient as reads may block writes from happening.

Implement a read/write lock, based on the previous question's implementation, ensuring that any write operation that entered the synchronized block (i.e. that own the lock) takes precedence over any new reader. That is, if a writer has entered the synchronized block of a `write` no new reader is permitted to enter a `read`, but existing readers are allowed to complete.

```
class ReadWrite extends Monitor {
  // declare any private variables needed here

  def read[T](op: => T): T = {
    // fill in what is needed here

  }

  def write[T](op: => Unit): Unit = synchronized {
    // fill in what is needed here

  }
}
```