
Parallel Programming

Final Exam

Wednesday, June 1, 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.

Open-book The exam is open book. You are allowed to refer to any printed or hand written material.

No communication You are not allowed to use internet, mobile phones, laptops, smart watches, etc.

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

Exercise 1: Counting Semaphore using Compare-and-Set (25 points)

A counting semaphore is a synchronization construct that is used to control accesses to a finite number of shared resources (e.g. processors, network ports), when there are multiple threads competing to access the resource.

Assume that we have n units of a shared resource. A semaphore for the resource maintains a counter to track the number of units that are available. It supports two operations: `lock` and `unlock`. The `lock` operation decrements by *one*, and the `unlock` operation increments by one the number of available units of the shared resource. The `unlock` method may cause the number of available units to exceed the initial n . In case there no units available at the time of invoking the `lock` operation, the operation blocks until at least one unit becomes available.

Your task is to implement the counting semaphore, as per the above description, using only an `AtomicReference` that has the methods `compareAndSet` and `get`, and *no* other synchronization primitive. In particular, you *cannot* use `synchronized`, `notify`, `wait`, or `fetchAndIncrement` operations. However, you are *allowed* to perform busy waiting, using loops and/or tail recursion for blocking.

Question (25 points)

Implement a counting semaphore based on the following template.

```
class CountingSemaphore(units: Int) {  
    // declare any private variables needed here
```

```
    def lock() : Unit = {  
        // fill in whatever is needed here
```

```
    }
```

```
def unlock(): Unit = {  
  // fill in whatever is needed here
```

```
}
```

```
}
```

Exercise 2: Asynchronous Histogram (25 points)

A statistics professor from UNIL is doing some ground-breaking work on distributions. His research is almost ready to be published but his proof of concept needs some heavy and time-consuming computations.

As a remarkable concurrency student from EPFL, you have been hired to create a concurrent and non-blocking version of his code. Given the `magicGeneration` function and the starting point he wrote below, your task is to compute the **unit digit histogram** explained later.

```
def magicGeneration(): Future[Int] = /* Some computation */

val asyncData = List(
  magicGeneration(),
  magicGeneration(),
  // ...
  magicGeneration()
)
```

You may refer to the Future API page 15.

Question 1: Combining Futures (10 points)

`asyncData` is currently a `List[Future[Int]]`. Your first task is to convert a list of futures into a future of lists.

```
def sequence(fs: List[Future[Int]]): Future[List[Int]] =
  fs.foldRight(Future.successful(List.empty[Int])) {
    case (f: Future[Int], list: Future[List[Int]]) =>

}
}
```

Question 2: Asynchronous Counting (10 points)

Now that you have a single piece of data, write a function that computes the number of integers falling in each bin. The function `computeBin` returns the bin corresponding to a given integer.

Hint: do not forget that the computation is asynchronous.

```
def countByBin(f: Future[List[Int]], computeBin: Int => Bin): Future[Map[Bin, Int]] =
```

Question 3: Unit Digit Histogram (5 points)

Finally using the starting point, you can glue everything together to compute the unit digit histogram of `asyncData`. The unit digit histogram is a mapping from digits to the number of times they appear in the unit position (rightmost position) in base 10. For example `List(Future(1), Future(14), Future(21))` should print `Success(Map(1 -> 2, 4 -> 1))` once the computation is finished. You can assume that all numbers are positive. Values of type `Bin` are integers representing a digit between 0 and 9. Values of type `Count` are non-negative (≥ 0) integers.

```
val histogram: Future[Map[Bin, Count]] =
```

```
    histogram.onComplete { values =>
      println(values)
    }
```

Exercise 3: Quorum and Vote (25 points)

An important voting must take place at EPFL. Students are asked if *they agree to double the tuition fee*. Each student, referred to as an `AssemblyMember`, has to vote on the matter, and an EPFL `Scrutineer` will handle the voting process. The `Scrutineer` must collect all votes and announce the outcome of the vote. The voting session should start only when more than half of the assembly have gathered to vote. In other words, when the *quorum is met*, To be more precise, an absolute majority of the `AssemblyMember` must have informed the `Scrutineer` that they are present before the voting starts.

For this whole exercise, you are given the following setup and a `Timer` actor that notifies a requester after a desired duration. The case classes and objects defined below should be sufficient to solve this exercise. **Think twice before adding a new one. Focus on the part you have to implement and pay attention to the methods signatures.** But you might find useful infos in the given code (regarding code structure and/or logic).

```
// These are the messages sent by the Scrutineer
case object VoteIsOver
case class IssueUnderVote(val question: String)

// These are the messages sent by an AssemblyMember
case object ImHere
sealed trait VoteOption
case object Yes extends VoteOption
case object No extends VoteOption
case object Blank extends VoteOption

// Message related to the Timer
case object TimesUp
case class StartTimer(duration: Long)

object Quorum extends App {
  val system = ActorSystem("Quorum")
  val AssemblySize: Int = 153
  val VotingItem: IssueUnderVote = IssueUnderVote("raise the EPFL tuition fee")

  val scrutineer =
    system.actorOf(Props(new Scrutineer(AssemblySize, VotingItem)))

  val AssemblyMembers = (0 until AssemblySize).map{
    id => system.actorOf(Props(new AssemblyMember(scrutineer,
      (10000 * Math.random()).toInt)), name = s"Assembly_member_$id")
  }
}

class Timer extends Actor {
  override def receive: Receive = {
    case StartTimer(duration: Long) =>
      Thread.sleep(duration)
      context.sender() ! TimesUp
  }
}
```

Question 1: Assembly Member Actor (10 points)

Implement the `AssemblyMember` actor by filling the template shown below. This actor has to do the following:

1. Announce itself to the `Scrutineer` (with an `ImHere` message)
2. Wait for receiving an `IssueUnderVote` message, and then start thinking. The actor must wait for `thinkDelay` ms by starting a `timer`.
3. Then, there are two options:
 - When done with thinking, the actor must reply to the `Scrutineer` with a `VoteOption` given by the helper method `makeUpMind()` and start waiting for an issue again.
 - When receiving a `VoteIsOver`, it means that the member did not vote in time. The actor must start waiting on an issue again directly. You can assume that the next issue will start long after the previous timer sends the ignored `TimesUp` message.

Do not use any var declarations, use multiple Receive methods instead

```
class AssemblyMember(scrutineer: ActorRef, thinkDelay: Int) extends Actor {
  val timer = context.actorOf(Props(new Timer))

  def makeUpMind(): VoteOption = {
    Math.random() match {
      case rng if rng < 0.45 => Yes
      case rng if rng < 0.9  => No
      case _                  => Blank
    }
  }

  // Announce itself to the scrutineer and implement the receive routine

}
```

Question 2: Scrutineer Actor (15 points)

In this part, you are given a partial implementation of the `Scrutineer` actor over the next pages. Implement the missing parts (focus on them). The handling of votes is already implemented but you must ensure those are valid votes by adding a `case` filtering invalid `VoteOption` before the `case v: VoteOption` already implemented. Ensure the following points:

- The voting process must start only when *the quorum is met*, which means that a majority has gathered. Starting the vote means sending an `IssueUnderVote` message to all `presentMembers`.
- The voting process must be open for `VoteDuration` only. Use the given `Timer` actor, the `TimesUp` handling is already implemented for you.
- `AssemblyMember` arriving after the voting process started and before it ended should still have the opportunity to vote.
- No `AssemblyMember` should be able to vote twice even if they are disfunctioning.

Do not use any var declarations, use the different Receive methods instead

```
class Scrutineer(assemblySize: Int, votingItem: IssueUnderVote) extends Actor {
  val AbsoluteMajority = assemblySize/2 + 1
  val VoteDuration = 50000L
  val timer = context.actorOf(Props(new Timer))

  override def receive = waitingToStart()

  def waitingToStart(presentAssemblyMembers: Set[ActorRef] = Set()): Receive = {

    // Wait for the quorum to be met and start the voting process

  }
}
```

```

def voting(alreadyVoted: Set[ActorRef],
  presentMembers: Set[ActorRef], yesNoVotes: (Int, Int)): Receive = {

  /* Take care of AssemblyMember arriving late
   * Make sure no one is voting twice by filtering
   * invalid received VoteOption in a separate case
   */

  case v: VoteOption =>
    val (yesVotes, noVotes) = yesNoVotes
    val votes = v match {
      case Blank => yesNoVotes
      case Yes => (yesVotes + 1, noVotes)
      case No => (yesVotes, noVotes + 1)
    }

    val voteCasted = alreadyVoted + context.sender()

    tallyVotes(votes, voteCasted, presentMembers,
      presentMembers == voteCasted){ () =>
        context.become(voting(voteCasted, presentMembers, votes))
      }

  case TimesUp =>
    tallyVotes(yesNoVotes, alreadyVoted, presentMembers,
      isOver = true){ () => {} }
}

def done(): Receive = {
  case _ => println(s"The voting session is over, it is too late to vote " +
    context.sender())
}

```

```

/* Sets the scrutineer's Receive to done() if the voting can be decided,
and calls the callback cont() otherwise. */
def tallyVotes(votes: (Int, Int), alreadyVoted: Set[ActorRef],
  members: Set[ActorRef], isOver: Boolean)(cont: () => Unit): Unit = {

  lazy val IssueUnderVote(question) = votingItem
  val (yesVotes, noVotes) = votes

  def endVote(message: String): Unit = {
    (members -- alreadyVoted) foreach(_ ! VoteIsOver)
    println(message)
    context.become(done())
  }

  if (yesVotes >= AbsoluteMajority)
    endVote(s"The Assembly agreed to $question by an absolute majority!")
  else if (noVotes >= AbsoluteMajority)
    endVote(s"The Assembly refused to $question by an absolute majority!")
  else if (isOver) {
    if (yesVotes > noVotes)
      endVote(s"The Assembly agreed to $question: $yesVotes against $noVotes votes!")
    else if (yesVotes == noVotes)
      endVote(s"The Assembly was unable to choose whether to $question or not")
    else
      endVote(s"The Assembly refused to $question: $noVotes against $yesVotes votes!")
  } else cont()
}
}

```

Exercise 4: MMORPG Data Analysis (25 points)

In this exercise, you will be working with Spark to analyse data on a massively-multiplayer online role playing game (MMORPG). You may refer to the Spark API on page 15-16. In order to obtain full points for each part, your solution not only need to be correct, but also *efficient!* We give you enough space after each part to write down your answer. Please use that space.

In the game, characters are represented by the following case class:

```
case class Character(  
  id: Int,  
  name: String,  
  role: Role,  
  level: Int,  
  guildId: Option[Int]  
)  
  
sealed abstract class Role  
case object Warrior extends Role  
case object Mage extends Role  
case object Priest extends Role
```

Each character has a unique identifier and a name. Each character also has a role, which defines its abilities in battle and a level, which indicates their power. Characters start at level 1 and can reach up to level 60. Characters in this game can join guilds, which are associations of characters. Joining a guild is completely optional.

Assume that we have a RDD containing all characters of the game:

```
val characters = sc.parallelize(List(  
  Character(1, "Alek",    None,    60, Warrior),  
  Character(2, "Beatrix", Some(17), 45, Priest),  
  Character(3, "Cleop",   Some(17), 1,  Mage),  
  Character(4, "Deadra",  Some(41), 60, Priest),  
  ...  
))
```

Assume that the data is too large to process on a single machine.

Question 1: Character Progression (10 points)

In a hurry, Tom, the lead designer of the game, storms your office and asks you to compute some statistics about the game for an important board meeting taking place very soon! Those numbers will help them gauge the popularity of the game.

1.1) Using `characters`, efficiently compute the number of characters which have reached level 60.

```
val numberOfLevel60: Long =
```

1.2) Using `characters`, efficiently compute the average level of all characters in the game depending on their role.

```
val averageLevelPerRole: Map[Role, Double] =
```

Question 2: The Most Numerous Guild Award (5 points)

After his meeting, Tom tells you that the board has decided to give a special award to the guild with the most number of characters in it. It is your job to find out which guild deserves the award.

Using `characters`, efficiently compute the identifier of the guild with the most members. If some guilds have the same number of members, you may choose any of them as the result of your computation.

```
val mostNumerousGuildId: Int =
```

Question 3: Final Boss (5 points)

Very satisfied with your job so far, Tom comes to you with a final request. Some players have been complaining that the final boss fight was too difficult for some characters. Tom suspects that the difficulty of the fight strongly depends on the role of the character. To verify this, the company has been monitoring the boss fight for a week and collected records of the following form:

```
case class FightCompleted(characterId: Int, bossDefeated: Boolean)
```

Where `characterId` is the identifier of the character and `bossDefeated` indicates whether the boss was defeated by the character or not.

A week's worth of records have been collected so far in the following RDD:

```
val bossEvents = sc.parallelize(List(  
  FightCompleted(1, false),  
  FightCompleted(7, false),  
  FightCompleted(4, true),  
  ...  
))
```

Efficiently compute, using `characters` and `bossEvents`, the ratio (between 0 and 1) of times the boss was defeated per encounter depending on the role of the character.

```
val bossDefeatedRatio: Map[Role, Double] =
```

Question 4: Shuffling (5 points)

During the coffee break, you meet with Tom again. He tells you that, being very impressed with your job, he has started learning to use Spark himself. During the conversation, he tells you that he doesn't quite understand what shuffling is and why it is relevant.

4.1) Using your own words, explain what shuffling is and why it is problematic. Limit yourself to 4 phrases at most.

4.2) What technique can be used to avoid unnecessary shuffling within Spark? Limit yourself to 2 phrases at most.

Future API

Relevant API for `scala.concurrent.Future[+T]`:

- `def map[S](f: T => S): Future[S]`: Creates a new future by applying a function to the successful result of this future.
- `def flatMap[S](f: T => Future[S]): Future[S]`: Creates a new future by applying a function to the successful result of this future, and returns the result of the function as the new future.
- `def onComplete[U](f: Try[T] => U): Unit`: When this future is completed, either through an exception, or a value, apply the provided function.

Spark API

Relevant API for Spark `RDD[T]`:

- `def collect(): Array[T]`: Return an array that contains all of the elements in this RDD.
- `def count(): Long`: Return the number of elements in the RDD.
- `def distinct(): RDD[T]`: Return a new RDD containing the distinct elements in this RDD.
- `def filter(f: (T) => Boolean): RDD[T]`: Return a new RDD containing only the elements that satisfy a predicate.
- `def flatMap[U](f: (T) => TraversableOnce[U]): RDD[U]`: Return a new RDD by first applying a function to all elements of this RDD, and then flattening the results.
- `def fold(zeroValue: T)(op: (T, T) => T): T`: Aggregate the elements of each partition, and then the results for all the partitions, using a given associative and commutative function and a neutral “zero value”.
- `def groupBy[K](f: (T) => K): RDD[(K, Iterable[T])]`: Return an RDD of grouped items.
- `def map[U](f: (T) => U): RDD[U]`: Return a new RDD by applying a function to all elements of this RDD.
- `def reduce(f: (T, T) => T): T`: Reduces the elements of this RDD using the specified commutative and associative binary operator.
- `def sortBy[K](f: (T) => K, ascending: Boolean = true): RDD[T]`: Return this RDD sorted by the given key function.
- `def take(num: Int): Array[T]`: Take the first num elements of the RDD.
- `def union(other: RDD[T]): RDD[T]`: Return the union of this RDD and another one.

Additional methods available to RDDs of type `RDD[(K, V)]`:

- `def collectAsMap(): Map[K, V]`: Return the key-value pairs in this RDD to the master as a Map.
- `def groupByKey(): RDD[(K, Iterable[V])]`: Group the values for each key in the RDD into a single sequence.
- `def join[W](other: RDD[(K, W)]): RDD[(K, (V, W))]`: Return an RDD containing all pairs of elements with matching keys in this and other.
- `def leftOuterJoin[W](other: RDD[(K, W)]): RDD[(K, (V, Option[W]))]`: Perform a left outer join of this and other.
- `def mapValues[U](f: (V) => U): RDD[(K, U)]`: Pass each value in the key-value pair RDD through a map function without changing the keys.
- `def partitionBy(partitioner: Partitioner): RDD[(K, V)]`: Return a copy of the RDD partitioned using the specified partitioner.
- `def reduceByKey(func: (V, V) => V): RDD[(K, V)]`: Merge the values for each key using an associative reduce function.
- `def rightOuterJoin[W](other: RDD[(K, W)]): RDD[(K, (Option[V], W))]`: Perform a right outer join of this and other.

Constructors for `Partitioner`:

- `new HashPartitioner(partitions: Int)`: A partitioner that implements hash-based partitioning using Java's `Object.hashCode`.
- `new RangePartitioner(partitions: Int, rdd: RDD[(K, V)], ascending: Boolean = true)`: A partitioner that partitions sortable records by range into roughly equal ranges.