

Exercise 3: Distributed Game of Life (15 points)

Important notice: you are not allowed to use any *var* in this assignment.

In this assignment you will have to implement the missing parts of an actor based Game of Life. The game of life consists of a rectangular grid of cells, in which each cell can be alive or dead. The grid wraps around, which means that the left side of the grid connects to its right side, and similarly for the top and bottom sides.

We model the game of life with an actor system, composed of one `Grid` actor, which spawns a `Cell` actor for each of its cells.

The Game of Life is turn-based; a turn is called a *generation*. The state (alive or dead) of a cell for the next generation is computed based on the state of its neighbours in the current generation, as follows (taken from wikipedia):

1. Any live cell with fewer than two live neighbours dies, as if caused by under-population.
2. Any live cell with two or three live neighbours lives on to the next generation.
3. Any live cell with more than three live neighbours dies, as if by overcrowding.
4. Any dead cell with exactly three live neighbours becomes a live cell, as if by reproduction.

Are considered neighbours the 8 cells surrounding a given cell. Note that even cells on the borders of the grid have exactly 8 neighbours, since the grid wraps around.

Neighbours computation (3 pt)

Your first task is to implement the `computeAndSendNeighbours` method of the `Grid` actor (by filling in the blank in the following code). This method take the `population` argument. Which is a 2d array of `ActorRefs` representing each cell, along with their respective initial state. It should send to each of the `Cell` actors an initialisation message holding the initial state of the cell and a set of its neighbours.

Do not forget to declare classes and/or objects for the messages you intend to send in addition to `Init` here.

```
case object Init

class Grid(width: Int, height: Int) extends Actor {

  override def receive: Receive = {
    case Init =>
      val population = Array.tabulate(height, width) { (row, col) =>
        val initiallyAlive = partOfPattern(row, col)
        val cellActor = context.actorOf(Props(new Cell(row, col)),
          name = "cell" + row + "_" + col)
        (cellActor, initiallyAlive)
      }
      computeAndSendNeighbours(population)
  }
}
```

```
def computeAndSendNeighbours(population: Array[Array[(ActorRef, Boolean)]]): Unit = {  
  
    }  
  
    def partOfPattern(row: Int, col: Int): Boolean = {  
        val pattern = Set((0, 0), (0, 1), (1, 0), (2, 3), (3, 2), (3, 3))  
        pattern((row, col))  
    }  
  
    def neighboursOf(row: Int, col: Int): List[(Int, Int)] = {  
        (for (r <- row - 1 to row + 1; c <- col - 1 to col + 1 if row != r || col != r)  
            yield ((r + height) % height, (c + width) % width)).toList  
    }  
  
}
```

Synchronized cells (4 pt)

Given the `Grid` actor you designed in the previous part, you will now implement the `receive` method of your cells. In this first version you will assume that there is some kind of clock that sends a `Tick` message to all actors at the start of every generation. We also assume that when this message is received, all the processing of the previous generation is done, and hence the system has stabilized (there is no more messages being processed or pending).

First your cell will need to handle the initialisation message sent from the `Grid` (holding informations about its initial state and its neighbours). Upon receiving a `Tick` message, each cell has to

1. Inform its neighbours (through messages) of its state in the current generation.
2. Update its state for the next generation, based on messages received from its neighbours, as soon as those have all been received.

Reminder: no vars allowed

```
class Cell(xPos: Int, yPos: Int) extends Actor {
```

```
  override def receive =
```

```
}
```

Independent cells (8 pt)

It is not possible to know when the system is stable—and thus send `Tick` messages at the appropriate time—without having a centralized actor which would gather status from all the cells.

To fix this, you will now design a truly distributed system where cells only communicate with their neighbours (as soon as they are initialized, initialisation is still taken care of by the `Grid`). You will have to keep track, somehow, of the generation, in order to keep the evolution of the game conforming to the rules.

Compared to the previous part, cells do not wait for a `Tick` message to start the next generation. Instead, as soon as they have updated their state for the next generation, they spontaneously start the next generation, and therefore send their new status to its neighbours immediately.

Pay particular attention to the *races* this would cause if you implement this naively. What happens when a cell has already processed all messages from its neighbours, but one of them is still waiting for messages coming from *its* neighbours?

Reminder: no vars allowed

```
class Cell(xPos: Int, yPos: Int) extends Actor {
```

```
  override def receive =
```

```
}
```

Exercise 4: Parallel Word Count (10 points)

In this exercise we will count the number of words in a piece of text using parallel collections. In our simplified model, words are sequences of non-whitespace characters, and you are given the `isWhitespace` function with the following signature:

```
def isWhitespace(char: Char): Boolean
```

- 1) To warm up, given a sequence of characters (of type `Seq[Char]`), implement a sequential word count. Keep in mind that the string may begin and end with whitespaces. Use a single `foldLeft` traversal:

```
def sequentialWordCount(seq: Seq[Char]): Int =
```

(3 points)

- 2) To enable parallelization, we provide a method which splits a sequence of characters into non-empty sub-sequences, producing a `ParSeq[Seq[Char]]` (you need not implement it!). The number of sub-sequences corresponds to the number of processor cores and the method guarantees they are all non-empty:

```
def segments(text: Seq[Char]): ParSeq[Seq[Char]]
```

Assuming the input sequence is "Hello Scala!", show the four different cases in which it could be split in two sub-sequences and indicate the result of `sequentialWordCount` for each of them:

```
segments("Hello Scala!") =
```

Case 1:

Case 2:

Case 3:

Case 4:

In two of the cases above, the sum of `sequentialWordCount` should be different from the actual number of words. Why is this?

(2 points = 4×0.5 points)

- 3) Now implement the `parallelWordCount` method, using a `map` and a `foldLeft` comprehension. You can use the `segments` and `sequentialWordCount` methods defined previously and can assume none of the sub-sequences is empty:

```
def parallelWordCount(seq: Seq[Char]): Int =
```

(4 points)

- 4) Why did we use `map` + `foldLeft` and not `aggregate`? What would make the aggregation more difficult in this case?

(1 point)

Exercise 3: Efficient Data Science at CFF (11 points)

In this exercise, we will concentrate on doing a number of operations on RDDs *efficiently*. The key to obtaining full points for each problem is to find the most efficient solution. After each part we give you enough space to implement the assignment; please use that space.

Imagine CFF is interested in analyzing the customer data that it has stored and real-time behavior of its customers. To this end, in the company, they collect data about `Customers` and `Tickets` defined as follows:

```
case class Customer(  
  customerId: Int,  
  homeCity: String,  
  purchaseHistory: List[Ticket]  
)  
case class Ticket(  
  customer: Option[Int],  
  origin: String,  
  destination: String,  
  price: Double  
)
```

Assume we have an RDD containing all known customers that use the CFF **mobile app**:

```
val customers: RDD[Customer] = sc.parallelize(List(  
  Customer(1, "Lausanne", List(  
    Ticket(Some(1), "Lausanne", "Zurich", 70.4),  
    Ticket(Some(1), "Zurich", "Lausanne", 70.4)  
  )),  
  Customer(2, "Montreux", List(  
    Ticket(Some(2), "Montreux", "Vevey", 3.50),  
    Ticket(Some(2), "Montreux", "Martigny", 5.40)  
  )),  
  ...  
)
```

Part 1 (1 point) Using `customers`, calculate the total amount of money the CFF has earned in ticket sales from known customers.

```
val sales: Double =
```

Part 2 (1 point) Using `customers`, and knowledge of their `homeCity`, determine how much money each Swiss city grosses in CFF ticket sales of known customers.

```
val cityGross: Seq[(String, Double)] =
```

Part 3 (3 points) Suppose the CFF has a service that collects all ticket purchases done in previous 10 minutes into an RDD. The tickets that are bought on a ticket machine have the `customerId` field set to `None` and tickets bought by known customers have a `customerId` field set to `Some(id)`:

```
// RDD containing all purchases made in the last 10 minutes
// RDD containing all purchases made in the last 10 minutes
val recentPurchases: RDD[Ticket] = sc.parallelize(List(
  Ticket(None, "Bern", "Zurich", 15.20), // paper ticket: no customer info
  Ticket(Some(1), "Lausanne", "Zurich", 70.4),
  Ticket(Some(1), "Zurich", "Lausanne", 70.4),
  Ticket(Some(2), "Montreux", "Vevey", 3.50),
  Ticket(Some(2), "Montreux", "Martigny", 5.40),
  ...
))
```

With this data CFF wants to calculate the ratio of people leaving a city that is their home city versus the overall number of people leaving that city. Assume that `customers` and `recentPurchases` are already partitioned by the option of their `id`.

```
def leavingTheCityRatio(
  customers: RDD[(Option[Int], Customer)],
  recentPurchases: RDD[(Option[Int], Ticket)],
  city: String): Double =
```

Part 4 (3 points) CFF data scientists analyzed their queries and cluster size and realized that they want to persist customer data by the option of their id (`Option[Int]`) in 10 different in-memory partitions. How can we achieve this:

```
val idCustomers: RDD[(Option[Int], Customer)] =  
  
val partitioner: Partitioner =  
  
val persistedCustomers: RDD[(Option[Int], Customer)] =
```

Given the `persistedCustomers`, CFF wants to periodically compute `leavingTheCityRatio` for Lausanne. Their timer process is set up to call the `leavingLausanneRatio` function every 10 minutes with fresh `recentPurchases`. Can you implement the `leavingLausanneRatio` function so that a minimal amount of network traffic happens. Assume that you have access to—and you are allowed to reuse—variables `persistedCustomers`, `partitioner`, and `leavingTheCityRatio`:

```
// assume that persistedCustomers, partitioner, and leavingTheCityRatio are visible  
def leavingLausanneRatio(recentPurchases: RDD[Ticket]): Double =
```

If we call `mapValues` on `persistedCustomers` would the result of `mapValues` still have a `partitioner`? If so, which kind?

Bonus (2 points) After running the `leavingLausanneRatio` service the CFF data analysts realized that their query is taking far more time than expected. What could be the reason?

Question 1 (1 points) Can you name 5 Spark methods that may cause a shuffle?

Question 2 (2 points) For which class of problems is it particularly advantageous to repartition the dataset? (i.e., for which pattern of computation would it save you time if you first created a partitioner and used `partitionBy`?) Why?