

2 MapReduce/Spark Practice: Optimize Your GPA

2.1 Given the student's name and course taken, output their name and total GPA.

Declare any custom data types here:

CourseData:

```
int courseID
float studentGrade // a number from 0-4
```

```
1 map(_____, _____):      1 reduce(_____, _____):
map(String student, CourseData value):      reduce(String key, Iterable<float> values):
    emit(student, value.studentGrade)        totalPts = 0
                                              totalClasses = 0
                                              for grade in values:
                                                  totalPts += grade
                                                  totalClasses += 1
                                              emit(key, totalPts / totalClasses)
```

3 Coherency and Atomic Operations

The benefits of multi-threading programming come only after you understand concurrency. Here are two of the most common concurrency issues:

1. **Cache-incoherence**: each hardware thread has its own cache, hence data modified in one thread may not be immediately reflected in the other. This can often be solved by bypassing the cache and writing directly to memory, i.e. using volatile keywords in many languages.
2. **Read-modify-write**: Read-modify-write is a very common pattern in programming. In the context of multi-threading programming, the **interleaving** of R, M, W stages often produces a lot of issues.

In order to solve the problems created by Read-modify-write, we have to rely on the idea of **uninterrupted execution**, also known as atomic execution.

In RISC-V, we have two categories of atomic instructions:

1. **Load-reserve, store-conditional**: allows us to have uninterrupted execution across multiple instructions
2. **Amo.swap**: allows for uninterrupted memory operations within a single instruction

Both of these can be used to achieve atomic primitives. Here are examples for each:

Test-and-set	Compare-and-swap
Start: <code>addi t0 x0 1 # Locked = 1</code>	<code># a0 holds address of memory location</code>
<code>amoswap.w.aq t1 t0 (a0)</code>	<code># a1 holds expected value</code>
<code>bne t1 x0 Start</code>	<code># a2 holds desired value</code>
<code># If the lock is not free, retry</code>	<code># a0 holds return value, 0 if successful, !0 otherwise</code>
<code>... # Critical section</code>	<code>cas:</code>
<code>amoswap.w.rl x0 x0 (a0) # Release lock</code>	<code>lr.w t0, (a0) # Load original value.</code>
	<code>bne t0, a1, fail # Doesnt match, so fail.</code>
	<code>sc.w a0, a2, (a0) # Try to update.</code>
	<code>jr ra # Return.</code>
	<code>fail:</code>
	<code>li a0, 1 # Set return to failure.</code>
	<code>jr ra # Return.</code>

Instruction definitions:

1. **Load-reserve**: Loads the four bytes at `M[R[rs1]]`, writes them to `R[rd]`, sign-extending the result and registers a reservation on that word in memory.
2. **Store-conditional rd, rs2, (rs1)**: Stores the four bytes in register `R[rs2]` to `M[R[rs1]]`, provided there exists a load reservation on that memory address. Writes 0 to `R[rd]` if the store succeeded, or a nonzero error code otherwise.
3. **Amoswap rd, rs2, (rs1)**: Atomically, puts the sign-extended word located at `M[R[rs1]]` into `R[rd]` and puts `R[rs2]` into `M[R[rs1]]`.

Explanations for both methodologies:

1. **Test-and-set**: We have a lock stored at the address specified by `a0`. We utilize `amoswap` to put in 1 and get the old value. If the old value was a 1, we would not have changed the value of the lock and we will realize that someone currently has the lock. If the old value was a 0, we will have just "locked" the lock and can continue with the critical section. When we are done, we put a 0 back into the lock to "unlock" it.
2. **Compare-and-swap**: CAS tries to first reserve the memory and gets the value stored and compares it to the expected value. If the expected value and the value that was stored do not match, the entire process fails and we must restart to update based on the new information. Otherwise, we register a reservation on the memory and try to store the new value. If the exit code is nonzero, something went wrong with the store and we must retry the entire LR/SC process. Otherwise with a zero exit code, we continue into the critical section, then release the lock.

3.1 Why do we need special instructions for these operations? Why cant we use normal load and store for `lr` and `sc`? Why cant we expand `amoswap` to a normal load and store?

For `lr` and `sc`, after `lr`, other threads cannot write to the location marked `reserve`, hence the value loaded from memory (`a3` in the above example) will be unchanged between `lr` and `sc`. For `amoswap`, it does load and store in one single CPU cycle, hence the operation is atomic and uninterruptable.

3.2 Now that we have atomic operations, let's try to experiment with them. Let us try to implement an algorithm that enforces ordered thread execution. This means that if we have four threads, thread 0 goes first, thread 1 goes next, etc. For this problem assume that `a1` holds the location of a piece of memory we have access to for the entire duration of our algorithm. Also, we can assume there exists a label `get_thread_num` that returns the thread's number in `a0`. Try to fill in the blanks below. Please use LR/SC for this problem:

```

1      addi t0, x0, 0
2
3      _____ # Setup for the first (0-th) thread
4      ...
5      # Assume we now spawn 4 threads in this code
6      ...
7
8  Check: jal _____ # Get the current thread number
9              # Get the ID of the next thread that should operate
10     _____ # (make sure this can't get interfered with)
11
12     _____
13  Done: addi t0, t0, 1
14
15     _____ # Set which thread is next to run
16
17     bne _____
18     ...
19     # Assume we now join the 4 threads in this code
20     ...
21     jr ra

1      addi t0, x0, 0
2      sw t0, 0(a1)
3      ...
4      # Assume we now spawn 4 threads in this code
5      ...
6  Check: jal ra, get_thread_num
7      lr.w t0, (a1)
8      bne a0, t0, Check
9  Done: addi t0, t0, 1
10     sc.w a0, t0, (a1)
11     bne a0, x0, Check
12     ...

```

```
13     # Assume we now join the 4 threads in this code
14     ...
15     jr ra
```