

Please note the following: To pass the course you need at least 50% of the home assignment points. Please contact the Lecturer after the exam if you've not completed the home assignments successfully.

Assignment 1 Please write a short essay (1/2 page max) explaining what are weak memory models? Why are they being used as the semantics of programming languages instead of sequential consistency? Give some examples of weak memory models from hardware level, programming language level, and from databases. Discuss what kind of challenges weak memory models impose to compiler writers, and to programmers in general. (4p)

Assignment 2 Consider the Java Monitors notification mechanisms and answer using 2-4 sentences for each item:

- a) What happens when a thread executes a `wait()` method on an object? (1p)
- b) What are the three main reasons under which a Java `waiting` thread will continue execution? (1p)
- c) What is the difference between `notify()` and `notifyAll()` methods? (1p)
- d) What are spurious wakeups and how to defend against them in Java? (1p)

Assignment 3 Explain the following notions and terms using 2-4 sentences for each item:

- a) RDD (1p)
- b) Safety property (1p)
- c) Liveness property (1p)
- d) X86-TSO (1p)
- e) Monitor (1p)
- f) Critical section (1p)
- g) Deadlock (1p)
- h) Race condition (1p)

Assignment 4 Please write a short essay (1/2 page max) of the Akka Actor programming framework. What is the programming framework like to the programmer, what are its main uses, and what its benefits and drawbacks compared to other programming models for concurrent programs. (4p)

Assignment 5 Consider the Scala Parallel Collections programming framework, and answer using 2-4 sentences for each item:

- a) What is Work Stealing, and how is it related to Scala Parallel collections implementation? (1p)

The name of the course, the course code, the date, your name, your student id, and your signature must appear on every sheet of your answers. All calculators are allowed in this exam.

- b) What are the two core abstractions used by Scala Parallel collections implementation to allow for parallelization of computation and obtaining a joint result? (1p)
- c) What are the requirements imposed on the function applied by the Scala Parallel Collections `reduce` operation? (1p)
- d) How do the requirements for the Apache Spark function applied by parallel `reduce` differ from the requirements imposed by Scala parallel collections? How does this affect portability between these frameworks? (1p)