

# Assignment 1: Introduction and challenges of concurrency

ETH Zurich

## 1 Amdahl's Law

### 1.1 Background

Consider a program where multiple threads operate on a buffer. Some threads only read from the buffer and other threads only write to the buffer. Any number of readers can simultaneously operate on the buffer. While a writer is operating on the buffer, no other writer or reader can be active on the buffer.

Assume a pool of  $N$  threads where each reader and writer is a thread. Hereby, 90 % of the threads are readers and 10 % of the threads are writers. Each reader thread takes 2 seconds to execute and each writer thread takes 3 seconds to execute. The program terminates when all threads in the thread pool terminated.

### 1.2 Task

According to Amdahl's Law, what is an upper bound for the speedup of the above implementation on a 4-core processor?

## 2 Interleavings

### 2.1 Background

This exercise is taken from the book *Principles of Concurrent and Distributed Programming* [2]. Imagine two threads  $P$  and  $Q$  that share the variables  $K$  and  $n$ .

$n := 0$			
$P$		$Q$	
1	<b>do <math>K</math> times</b>	1	<b>do <math>K</math> times</b>
2	$temp := n$	2	$temp := n$
3	$n := temp + 1$	3	$n := temp - 1$

### 2.2 Task

What are the possible final values of  $n$  for a given positive value of  $K$ ?

## 3 LTL Models

### 3.1 Background

The transition system  $\mathcal{M}$  in Figure 3.1 models a barbershop with one barber, two customers, and one chair for waiting. Each customer  $i$  can be in four different states:  $entering_i$  upon

entering the shop,  $waiting_i$  when waiting for the barber,  $haircut_i$  when getting a haircut, and  $gone_i$  upon leaving the shop. A waiting customer can leave the shop when the barber is busy, and a customer on the barber chair can switch with a waiting customer.




Figure 1: Barbershop model

The following notation is used in the formulas:

$G$ Globally	$F$ Future
$X$ Next	$U$ Until
$\rightarrow$ Implication	$\wedge$ Conjunction
$\vee$ Disjunction	$\neg$ Negation

### 3.2 Task

Answer the following questions about the formulas given in linear-time temporal logic (LTL). Justify your answers in each case.

1. Let  $\phi$  be  $(F gone_1) \wedge (F gone_2)$ .

- (a) Is  $\phi$  satisfied at state  $q_8$ , i.e.,  $\mathcal{M}, q_8 \models \phi$ ?
  - (b) Is  $\phi$  satisfied globally, i.e.,  $\mathcal{M}, q \models \phi$  for all states  $q$ ?
2. Let  $\phi$  be *waiting<sub>1</sub> U haircut<sub>1</sub>*.
- (a) Is  $\phi$  satisfied at state  $q_6$ , i.e.,  $\mathcal{M}, q_6 \models \phi$ ?
3. Let  $\phi$  be *(entering<sub>1</sub>  $\wedge$  XX haircut<sub>1</sub>)  $\rightarrow$  X haircut<sub>1</sub>*.
- (a) Is  $\phi$  satisfied at state  $q_1$ , i.e.,  $\mathcal{M}, q_1 \models \phi$ ?
  - (b) Is  $\phi$  satisfied globally, i.e.,  $\mathcal{M}, q \models \phi$  for all states  $q$ ?

Translate each of the following properties from natural language into LTL.

- 4. No two customers get a haircut at the same time.
- 5. If customer 1 enters the barber shop, (s)he will leave the shop eventually.
- 6. Before leaving the barber shop, customer 2 is always either waiting or getting a haircut.
- 7. Once customer 1's haircut is done, (s)he leaves immediately.

## 4 Safety vs. liveness

### 4.1 Task

Consider the following properties.

- 1. What goes up must come down.
- 2. If two or more processes are waiting to enter their critical sections, at least one succeeds.
- 3. If an interrupt occurs, then a message is printed.
- 4. The cost of living never decreases.
- 5. Two things are certain: death and taxes.
- 6. You can always tell a Harvard man.

For each of the above properties, state whether it is a safety or liveness property. Identify the bad or good thing of interest.

## 5 Interleavings in practice

### 5.1 Background

We know that the interleavings in a concurrent program may give rise to different behavior. This exercise is designed to give a way to see how unpredictable these effects may be.

## 5.2 Task

Your task is to design a Haiku composer. A Haiku is a Japanese form of poetry with 17 syllables in three lines, where the first line must contain 5 syllables, the second must contain 7, and the third line must contain 5 (this is the traditional layout). The lines may contain any number of words, as long as the syllable restrictions are followed. The Haiku composer will have a small (20-30 should be enough) list of words, and will spawn 3 threads to compose a Haiku poem. Each thread is responsible for a single line of the Haiku.

For this task, you must use a single shared store of words. Once a thread has used a word, it must be removed from the store. You may find the usage of the `java.util.concurrent` package helpful here. The store should have a reasonable number of 1-3 syllable words. It is also perfectly OK to keep removing words until you find the one that “fits” your syllable requirement. You may wish to define a **Word** class which can model a word, including syllable count.

This should be done without using concurrency operations such as `synchronized` and the `wait/notify` capabilities of objects.

To spawn threads and the basics of java concurrency, you may refer to the chapter on Java concurrency in the course book available at

[http://se.inf.ethz.ch/courses/2014a\\_spring/ccp/reading-materials/book/](http://se.inf.ethz.ch/courses/2014a_spring/ccp/reading-materials/book/).

## References

- [1] Andrei Voronkov. Script to Logic in Computer Science. 2009.
- [2] Mordechai Ben-Ari. Principles of Concurrent and Distributed Programming (2nd Edition). Addison-Wesley, 2006.