

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?

1.3 Solution

$$S = \frac{1}{1 - p + \frac{p}{n}}$$
$$p = \frac{0.9 * 2}{0.9 * 2 + 0.1 * 3}$$

With $n = 4$ this results in:

$$S = \frac{14}{5} = 2.8$$

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 do K times	1 do K times
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 ?


2.3 Solution

The final value of n can be any value between $-K$ and K . There are two main cases. Either P executes one loop iteration between the moment Q reads n in line 2 and the moment Q is about to write to n in line 3. In this case Q will overwrite the effect of P 's iteration. In the other case the roles of P and Q are swapped. If the first case occurs over and over again then we end up with $n = -K$. In the second case the execution results in $n = K$. The other possible combinations of the two cases result in the values between $-K$ and K .

3 LTL Models

3.1 Background

We have the following model, \mathcal{M} :



3.2 Task

You are given formulae:

- a) Ga
- b) aUb
- c) $a \wedge Fb$
- d) $aU(X(a \wedge \neg b))$
- e) $X\neg b \wedge G(\neg a \vee \neg b)$
- f) $X(a \wedge b) \wedge F(\neg a \wedge \neg b)$

For each of the above formulae, ϕ :

1. Find a path from q_3 that satisfies ϕ .
2. Determine whether $\mathcal{M}, q_3 \models \phi$.

3.3 Solution

1. Successful paths given:
 - a) $q_3, q_4, q_3, q_4, \dots$
 - b) q_3, q_2, \dots
 - c) q_3, q_4, \dots
 - d) q_3, q_4, \dots
 - e) $q_3, q_1, q_2, q_2, \dots$
 - f) $q_3, q_4, q_3, q_1, \dots$
2. Counter-example paths given, where applicable:
 - a) No, q_3, q_2, \dots
 - b) No, q_3, q_1, \dots
 - c) Yes, a is satisfied immediately, and either state q_4 or q_2 is entered infinitely often.
 - d) No, q_3, q_1, \dots
 - e) No, q_3, q_4, \dots
 - f) No, q_3, q_1, \dots

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.

4.2 Solution

With respect to the order provided in the question, the answers are as follows.

1. This is a liveness property. The good thing to happen is that the thing having gone up is coming down.
2. This is a liveness property. The good thing to happen is that at least one processor succeeds.
3. This is a liveness property. The good thing to happen is the message being printed.
4. This is a safety property. The “bad” thing that will never happen is that the cost of living decreases.
5. This is a liveness property. The “good” thing to happen is death and taxes.
6. This is a safety property. The bad thing that will never happen is that you cannot tell a Harvard man.

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/2013a_spring/ccp/reading_materials/book/.

5.3 Solution

The solution is available in the source code that comes with this solution.

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.