WORKSHEET

A2.8 Thinking concurrently

Recursion and computational thinking · A level · OCR H446 2.1.5 · about 25 min

BugBotLab

Name _____ Class _____ Date _____

What this lesson is about

Concurrent and parallel processing, what can happen at once, benefits and trade-offs, and pipelining.

- Concurrent and parallel
- Which parts can happen at the same time?
- Benefits and trade-offs
- Pipelining

Questions

5 marks in all

1. What is the difference between concurrent and parallel processing? [1 mark]

- ☐ **A** Concurrent tasks are in progress over the same period, possibly taking turns on one core; parallel tasks run at the same instant on separate cores
- ☐ **B** There is no difference
- ☐ **C** Parallel tasks take turns; concurrent tasks run at the same instant
- ☐ **D** Concurrent processing needs a network

2. Which are trade-offs of concurrent processing? [1 mark]

Tick every answer that is true.

- ☐ **A** Tasks sharing data can give timing-dependent results (race conditions)
- ☐ **B** Tasks can deadlock, each waiting for a resource the other holds
- ☐ **C** Coordinating tasks has an overhead
- ☐ **D** It always makes every task finish sooner

3. Jobs: A takes 20 min, B takes 30 min, C takes 10 min and needs A and B to finish first. With unlimited workers, how many minutes until C finishes? [1 mark]

4. What does this program print? [1 mark]

```
items, stages, minutes = 10, 4, 3
one_at_a_time = items * stages * minutes
pipelined = stages * minutes + (items - 1) * minutes
print(one_at_a_time, pipelined)
```

5. 10% of a program must run in sequence and the rest can be split between processors. What is the most it can ever be sped up, with unlimited processors? [1 mark]

- ☐ **A** 10 times
- ☐ **B** 100 times
- ☐ **C** Unlimited
- ☐ **D** 2 times

The task: getting ready at the same time

The jobs before a match are held in a dictionary. Each job maps to a tuple `(minutes, needs)`: how long it takes, and a list of the jobs that must finish before it can start. - `charge` 40 min, needs nothing; `flash` 10 min, needs nothing; `calibrate` 5 min, needs `flash`; - `print` 30 min, needs nothing; `assemble` 15 min, needs `print` and `flash`; - `test` 10 min, needs `assemble`, `charge` and `calibrate`. With as many workers as you like, a job starts the moment the last job it needs has finished. Write a **recursive** function `finish(job)` that takes a job name and returns the earliest time, in minutes from the start, that the job can be finished. Then print: 1. `one worker: <n> min`, the time with one worker doing the jobs one at a time; 2. `test finishes at <n> min`, using `finish("test")`; 3. `unlimited workers: <n> min`, the latest finish of any job; 4. `speed-up: <s>`, one worker's time divided by unlimited workers' time, to 2 decimal places.

```
# the two lines every program starts with: the commands, then the robot
from bugbot import *
connect()

jobs = {
    "charge": (40, []),
    "flash": (10, []),
    "calibrate": (5, ["flash"]),
    "print": (30, []),
    "assemble": (15, ["print", "flash"]),
    "test": (10, ["assemble", "charge", "calibrate"]),
}

def finish(job):
    minutes, needs = jobs[job]
    return minutes
```

Plan your program here, then type it in and press Run.

Do it on the robot

www.bugbotlab.com/learn/a2-8-thinking-concurrently/

The simulator checks it and tells you when it passes. Nothing to install, no account.

Challenges

1. Add a second worker limit: with exactly two workers, when can the test finish? Work it out by hand first.
2. Which job is on the critical path? Would making `charge` faster finish the test sooner?
3. In the flashing task, what goes wrong if the loop has no `wait`? What if it waits 2 seconds each time round?