

AP-23 – Programming Assignment #2- v1.0

December 6, 2023

This assignment is made of two parts, consisting of exercises on the Java Stream API and Frameworks, and on Decorators in Python, respectively. It is distributed with an archive `aux_files.zip` containing some auxiliary files.

This document is subject to changes. Check on the course web page if there is a newer version.

Premise: the “ciao” of a string

This definition will be used below. Given a string `str`, we define its *ciao* (*characters in alphabetical order*) as the string having the same length of `str` and containing all the characters of `str` in lower case and alphabetical order. As an example, the *ciao* of “Hello” is “ehllo”. A *ciao string* is a string that is equal to its *ciao*. Clearly, two strings have the same *ciao* if and only if each one is an anagram of the other.


Part 1 (Java Stream API and Frameworks) A job scheduler

In this assignment, students are required to implement a simple software framework providing the functionalities of a job scheduler, but ignoring the aspects of parallelism and distribution. More precisely, the framework includes an *emitter* of jobs, a *compute* phase executing the jobs, a *collect* stage grouping them, and an *output* action printing the results in a suitable format. As a proof of concept, a simple working instance of the framework should be implemented as well.

Exercise 1: The framework

Following the guidelines presented in the lesson of November 21st, 2023, *On Designing Software Frameworks*, (see <http://pages.di.unipi.it/corradini/Didattica/AP-23/index.html#framework>), and more specifically the *Strategy design pattern*, implement in Java a `JobScheduler` software framework, respecting the following specifications:

1. The framework must be generic, using type variables `K` and `V` for the types of **keys** and **values** respectively.
2. For **key/value** pairs, the framework must use the class `Pair.java` from `aux_files.zip` (you can change its package, but nothing else).
3. Jobs will be instances of (subclasses) of the abstract class `AJob.java`, also enclosed, containing the abstract method `execute` with no parameter and returning a stream of **key/value** pairs.



4. The framework must include the following methods, conceptually composed as in the picture:
 - `emit`, which generates a stream of jobs;
 - `compute`, which executes the jobs received from `emit` by invoking `execute` on them, and returns a single stream of **key/value** pairs obtained by concatenating the output of the jobs;

- **collect**, which takes as input the output of **compute** and groups all the pairs with the same keys in a single pair, having the same key and the list of all values;
 - **output**, which prints the result of **collect** in a convenient way.
5. Methods **compute**, **collect** and **main** must be frozen spots of the framework, while **emit** and **output** must be hot spots.

Solution format: An archive **JobScheduler.zip** containing the Java files implementing Exercises 3 and 4, suitably commented. If you use **NetBeans**, please send in the archive the entire project.

Exercise 2: Counting anagrams

Write a program that given the absolute path of a directory prints the number of anagrams of all the words contained in a set of documents in that directory, by developing an instance of the framework of the previous point. You should ignore all words of less than four characters, and those containing non-alphabetic characters. Also, uppercase and lowercase letters should not be distinguished.

Here are some guidelines:

1. Create a subclass of **AJob** having a constructor that accepts the name of a file as parameter; the **execute** method must read the file, and it must return a stream containing all pairs of the form **(ciao(w), w)** where **w** is a word of the file satisfying the above properties.
2. **Emit** asks the user for the absolute path of a directory where documents are stored. It visits the directory and creates a new job for each file ending with **.txt** in that directory.
3. Output should write the list of ciao keys and the number of words associated with each key, one per line, in file **count_anagrams.txt** (or **count_anagrams.csv**, in CSV format).

For testing the program you can use the files in the enclosed archive **Books.zip** which contains parts of some famous books as downloaded from the pages of the [Gutenberg Project](#). (before the site became inaccessible from Italy, see [Raffaele Angius, Perché il Progetto Gutenberg sarà sotto sequestro per sempre](#)).

Part 2: (Decorators in Python) Benchmarking functions in Python, with Multithreading

In the last lesson on Python and the GIL I reported a claim from the literature: *“Two threads calling a function may take twice as much time as a single thread calling the function twice”*.

This assignment requires to write a parametric decorator that can be used to check if this claim is true, at least on your machine. The decorator should measure the execution time of invoking several times a function in parallel on a given number of threads.

Solution format: A single Python file called **twoThreadsClaim.py** containing the solutions to the following two exercises.

Exercise 3: A decorator for multi-threaded benchmarking

Define in Python a *parametric* decorator called **bench** (see the section “Decorators with Arguments” of the “Primer on Decorators”, <https://realpython.com/primer-on-python-decorators/#decorators-with-arguments>). When invoking a function **fun** decorated by **bench**, **fun** is executed several times in parallel on several threads (discarding the results), and the whole is repeated a few times returning the average time of execution and the variance.

The exact behaviour of the **bench** decorator is ruled by the following optional parameters:

- **n_threads**: The number of threads (default: 1)
- **seq_iter**: The number of times **fun** must be invoked in each thread (default: 1)
- **iter**: The number of times the whole execution of the **n_threads** threads, each invoking **seq_iter** times **fun**, is repeated. For each execution, the execution time has to be computed (default: 1)

The decorated function must return a dictionary of the following shape, containing the function name and the tuple of arguments, the values of the three parameters just described, as well as the mean value and the variance of the execution times of the **iter** iterations:

```
{'fun': 'fib_run', 'args': (28,), 'n_threads': 2, 'seq_iter': 8, 'iter': 5, 'mean': 3.382881900004577, 'variance': 0.0032192657625461194}
```

Use the [threading](#) module to exploit Python multithreading, `perf_counter()` from module [time](#) for computing the execution time, and module [statistics](#) for computing mean value and variance.

Exercise 4: Exploiting the decorator

Exploit the **bench** decorator to evaluate the effectiveness of multithreading in Python. Write a function **test** that has three parameters: an integer **iter**, a function **fun**, and a tuple of arguments **args**. Function **test** must execute **fun** on **args** with varying numbers of iterations and degrees of parallelism. More precisely, **test** must invoke **bench** a first time to execute 16 times **fun** on **args** on a single thread (passing **iter** as further parameter), then, similarly, to run **fun** 8 times on two threads, then 4 times on 4 threads, and finally 2 times on 8 threads. The program must write the information returned by the four invocations of **bench** in a file named `<fun>_<args>_<n_threads>_<seq_iter>`, or something similar.

Run the test on simple functions with different degrees of usage of the CPU, for example here are two extreme cases you can use:

```
def just_wait(n):          # NOOP for n/10 seconds
    time.sleep(n * 0.1)

def grezzo(n):            # CPU intensive
    for i in range(2**n):
        pass
```

Discuss the results of the experimentation in a brief comment at the end of the Python file.