



# GIL, Executors and Callables

- [GIL, Executors and Callables](#)
  - [Global Interpreter Lock \(GIL\)](#)
    - [Key Characteristics of the GIL:](#)
  - [Executors](#)
    - [ThreadPoolExecutor](#)
    - [ProcessPoolExecutor](#)
  - [Callables and Futures](#)

## Global Interpreter Lock (GIL)

The Global Interpreter Lock, commonly known as the GIL, is a mutex (lock) in Python that protects access to Python objects, preventing multiple native threads from executing Python bytecodes at once. It is a mechanism implemented in the CPython interpreter to simplify memory management and avoid conflicts in the execution of Python code.

## Key Characteristics of the GIL:

- 1. GIL and CPython:**
  - The GIL is specific to the CPython implementation, the reference implementation of Python.
  - Other implementations like Jython (Python on the Java Virtual Machine) and IronPython (Python on the .NET Framework) do not have a GIL.
- 2. Impact on Multithreading:**
  - The GIL allows only one thread to execute Python bytecodes at a time in a single process.
  - While multiple threads can exist, they cannot execute Python code in true parallel fashion due to the GIL.
- 3. Limitation for CPU-Bound Tasks:**
  - The GIL significantly impacts the performance of CPU-bound tasks that involve intensive computations.
  - In such cases, even if a system has multiple CPU cores, CPython

threads cannot fully utilize the available hardware in parallel.

#### 4. **Less Impact on I/O-Bound Tasks:**

- For I/O-bound tasks (tasks involving waiting for external resources, such as file I/O or network operations), the GIL has less impact.
- In I/O-bound scenarios, threads can release the GIL while waiting for external operations, allowing other threads to execute Python code.

#### 5. **Memory Management Simplification:**

- The GIL simplifies memory management by ensuring that memory operations within the interpreter are thread-safe.
- Without the GIL, extensive use of locks and synchronization mechanisms would be necessary, leading to increased complexity.

### **Why does Python have a GIL?**

The GIL is present in Python to simplify memory management and avoid conflicts in the execution of Python code. It also avoids conflicts between threads that may arise due to the use of Python's reference counting mechanism for memory management.

Reference counting is a technique used by Python to manage memory. It keeps track of the number of references to an object in memory. When the reference count of an object reaches zero, the object is deleted from memory.

## Executors

Executors in Python, part of the `concurrent.futures` module, provide a high-level interface for managing and executing concurrent tasks. Executors abstract away the details of managing threads or processes, making it easier to work with concurrent programming constructs.

The `concurrent.futures` module provides two types of executors:

1. `ThreadPoolExecutor` - manages a pool of threads
2. `ProcessPoolExecutor` - manages a pool of processes

Both executors implement the same interface but differ in how they execute tasks.

# ThreadPoolExecutor

The `ThreadPoolExecutor` class manages a pool of threads. It is initialized with a fixed number of threads. When a task is submitted to the executor, it is executed by one of the threads in the pool. If all the threads in the pool are busy, the task is queued until a thread becomes available.

It is well-suited for I/O-bound tasks and situations where threads can run concurrently without interference.

Let's take a look at an example of using the `ThreadPoolExecutor` class:

```
from concurrent.futures import ThreadPoolExecutor

def task(n: int):
    print(f"Processing {n}")

with ThreadPoolExecutor(max_workers=3) as executor:
    executor.submit(task, 1)
    executor.submit(task, 2)
    executor.submit(task, 3)
```

Some key points to note in the above example:

- The `ThreadPoolExecutor` class is initialized with a maximum of 3 threads.
- The `task` function is submitted to the executor using the `submit` method.
- We wrap the executor in a `with` statement to ensure that the executor is properly shutdown after use.

# ProcessPoolExecutor

The `ProcessPoolExecutor` creates a pool of worker processes, and tasks submitted to this executor are executed concurrently within these processes. This executor is suitable for CPU-bound tasks where parallelism is achieved by running tasks in separate processes.

Let's take a look at an example of using the `ProcessPoolExecutor` class:

```
from concurrent.futures import ProcessPoolExecutor

def task(n: int):
    print(f"Processing {n}")

with ProcessPoolExecutor(max_workers=3) as executor:
    executor.submit(task, 1)
    executor.submit(task, 2)
    executor.submit(task, 3)
```

Both `ThreadPoolExecutor` and `ProcessPoolExecutor` implement the same high-level executor interface, which includes methods like `submit()`, `map()`, and `shutdown()`. This common interface makes it easy to switch between thread-based and process-based concurrency depending on the nature of the tasks.

### Common Executor Interface Methods:

- `submit(fn, *args, **kwargs)` : Submits a callable for execution and returns a `Future` object representing the execution of the callable.
- `map(func, *iterables, timeout=None, chunksize=1)` : Applies the function to each item in the iterable(s) and returns an iterator over the results. It blocks until the result for each item is available.
- `shutdown(wait=True)` : Initiates an orderly shutdown of the executor, allowing previously submitted tasks to complete. If `wait` is `True` (default), it blocks until all tasks have completed.

## Callables and Futures

A callable is an object that can be called like a function. In Python, any object that implements the `__call__()` method is considered callable. This includes functions, methods, and classes. In the context of concurrent programming, callables are often tasks or operations that can be executed asynchronously.

The `concurrent.futures` module provides the `Future` class to represent the result of an asynchronous operation. The `Future` class encapsulates the asynchronous execution of a callable and provides methods to check the status of the operation and retrieve the result (or

exception) when it is complete.

Let's try to calculate the sum of two lists using the `ThreadPoolExecutor` class:

```
from concurrent.futures import ThreadPoolExecutor, as_completed

def calculate_sum(data):
    return sum(data)

data_list = [list(range(10)), list(range(10, 20))]

# Creating a ThreadPoolExecutor
with ThreadPoolExecutor() as executor:
    # Submitting tasks for asynchronous execution
    futures = [executor.submit(calculate_sum, data) for data in data_list]

    # Wait for all tasks to complete and retrieve results
    results = [future.result() for future in as_completed(futures)]

print(results)
```

Things to note in the above example:

- We submit the tasks for asynchronous execution using the `submit` method of the `ThreadPoolExecutor` class. This method returns a `Future` object representing the execution of the callable.
- We use the `as_completed` function to wait for the tasks to complete and retrieve the results. This function returns an iterator over the futures as they complete.

The `Future` class provides the following methods to check the status of the operation and retrieve the result (or exception) when it is complete:

- `done()` : Returns True if the future is done executing.
- `result(timeout=None)` : Returns the result of the operation. If the operation has not completed, this method blocks until the result is available or the optional timeout occurs.