


Shared Memory in Python

Concurrency vs. Parallelism -- Why It Matters

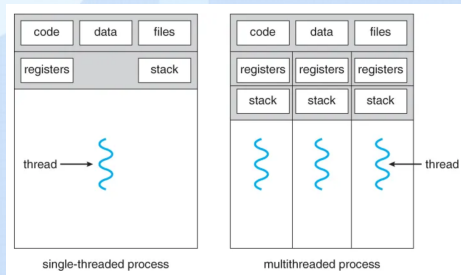
- Concurrency: multiple tasks make progress by interleaving their execution (not necessarily at the same time).
- Parallelism: tasks execute at the same time on different cores/processing units to increase throughput.



Shared Memory in Python

Threads vs. Processes -- Memory Model

- Thread: is a thread of execution in a program. Aka, lightweight process.
- Process: is an instance of a computer program that is being executed.
- Implications:
 - Threads share the memory and state of the parent, process share nothing.
 - Processes use inter-process communication (IPC) to communicate, thread do not.
 - A process can have 1 or several threads.
- Scheduling: the OS kernel schedules threads on CPU cores; a process can host one or multiple threads.



Shared Memory in Python

Multi-threading Basics: Creating Threads with `threading.Thread`

- Two threads (`t1`, `t2`) are created to run the `worker` function in parallel, each simulating work with `time.sleep(2)`.
- Because both threads sleep concurrently, the total execution time is about 2 seconds instead of 4.

```
1 Thread A starting
2 Thread B starting
3 Thread A done
4 Thread B done
5 Timing: 2.002 sec
```

```
1 import threading
2 import time
3
4 def worker(name):
5     print(f"Thread {name} starting")
6     time.sleep(2)
7     print(f"Thread {name} done")
8
9 # Create two threads
10 t1 = threading.Thread(target=worker, args=("A",))
11 t2 = threading.Thread(target=worker, args=("B",))
12
13 start = time.time()
14 t1.start()
15 t2.start()
16 t1.join()
17 t2.join()
18 end = time.time()
19
20 print("Timing: ", end - start, "sec")
```

Shared Memory in Python

Multi-threading Basics: Subclassing `threading.Thread`

- Step 1: Create a class inheriting from `threading.Thread`.
- Step 2: Override the `run()` method with the task logic.
- Step 3: Instantiate the class and call `start()`.
- Step 4: Use `join()` to wait for completion.

```
1 Thread A starting
2 Thread B starting
3 Thread A done
4 Thread B done
```

```
1 import threading
2 import time
3
4 class Worker(threading.Thread):
5     def __init__(self, name):
6         super().__init__()
7         self.name = name
8
9     def run(self):
10         print(f"Thread {self.name} starting")
11         time.sleep(2)
12         print(f"Thread {self.name} done")
13
14 # Create and start threads
15 t1 = Worker("A")
16 t2 = Worker("B")
17 t1.start()
18 t2.start()
19 t1.join()
20 t2.join()
```

Shared Memory in Python

Multi-threading Basics: Prime Calculation with single Thread

- Compute the sum of all primes up to 200,000.
- Using 1 thread only.

```
1 def isPrime(n):
2     if n < 2:
3         return False
4     if n == 2:
5         return True
6     max_val = int(math.ceil(math.sqrt(n)))
7     i = 2
8     while i <= max_val:
9         if n % i == 0:
10            return False
11        i += 1
12    return True
13
14 def sum_primes(n):
15    return sum([x for x in range(2, n) if isPrime(x)])
```

```
1 if __name__ == "__main__":
2
3     for i in range(0, 200000, 500):
4         sum_primes(i)
```



Shared Memory in Python

Multi-threading Basics: Profiling Single-Thread Prime Calculation

- Profiling with cProfile shows most time spent in:
 - `sum_primes, isPrime`
- $\text{tottime} = \text{percall}_{\text{local}} \times \text{ncalls}$
- $\text{percall}_{\text{cumul}} = \text{cumtime} / \text{ncalls}$

ncalls	tottime	percall	cumtime	percall	filename:lineno(function)
1	0.000	0.000	47.330	47.330	{built-in method builtins.exec}
1	0.000	0.000	47.330	47.330	isprime.py:1(<module>)
400	4.139	0.010	47.330	0.118	isprime.py:16(sum_primes)
39899202	38.625	0.000	43.177	0.000	isprime.py:3(isPrime)
39898803	2.470	0.000	2.470	0.000	{built-in method math.sqrt}
39898803	2.081	0.000	2.081	0.000	{built-in method math.ceil}
400	0.014	0.000	0.014	0.000	{built-in method builtins.sum}
1	0.000	0.000	0.000	0.000	{method 'disable' of '_lsprof.Profiler' objects}

```
Execution time (1 thread) = 47.330 sec
```

$$4.139 + 38.625 + 2.470 + 2.081 + 0.014 = 47.33 \text{ secondes}$$


Shared Memory in Python

Multi-threading Basics: Prime Calculation with Multi Thread (8 threads)

```
1 def do_work(q):
2     while True:
3         try:
4             x = q.get(block=False) # Get an item from the queue (non-blocking)
5             sum_primes(x)          # Compute the sum of primes below x
6         except Empty:
7             break
8
9 if __name__ == "__main__":
10     work_queue = Queue()
11     for i in range(0, 200000, 500):
12         work_queue.put(i)          # inserts the value i into the queue.
13
14     threads = [Thread(target=do_work, args=(work_queue,)) for _ in range(8)]
15
16     for t in threads:
17         t.start()
18     for t in threads:
19         t.join()
```



Shared Memory in Python

Profiling Multi-Thread Prime Calculation

- Profiling with `cProfile` highlights time spent in:
 - `threading.py:join`
 - `isPrime`
 - `_thread.lock.acquire`

tottime	percall	cumtime	percall	filename:lineno(function)
0.000	0.000	171.816	21.477	threading.py:1115(join)
0.000	0.000	171.811	21.476	threading.py:1153('wait_for_tstate_lock')
0.001	0.000	114.813	0.261	{method 'acquire' of 'thread.lock' objects}
50.089	0.000	57.532	0.001	thread_sqrt.py:6(isPrime)
0.000	0.000	57.442	57.442	{built-in method builtins.exec}
0.000	0.000	57.442	57.442	thread_sqrt.py:1(<module>)
0.003	0.000	57.440	57.440	threading.py:1016(_bootstrap)
0.008	0.001	57.440	57.440	threading.py:1056(_bootstrap_inner)
2.978	0.000	2.978	0.000	{built-in method math.sqrt}
2.525	0.000	2.525	0.000	{built-in method math.ceil}

Execution time (8 threads) = 57.442 sec




Shared Memory in Python

Multi-threading Basics: The Global Interpreter Lock (GIL)

- Only one thread can run Python bytecode at a time.
- CPU-bound: No true parallelism \Rightarrow threads wait.
- I/O-bound: GIL is released during blocking I/O \Rightarrow overlap possible.

```
def add(a, b):  
    return a + b
```

```
0 LOAD_FAST 0 (a)  
2 LOAD_FAST 1 (b)  
4 BINARY_OP 0 (+)  
6 RETURN_VALUE
```



Shared Memory in Python

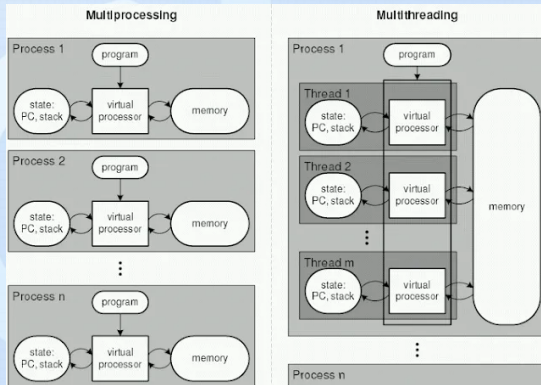
Multi-threading Basics: Why Multi-threading in Python Doesn't Scale for CPU-bound Tasks

- Consequence: Even on multi-core CPUs, threads run in a concurrent but not parallel way.
- In Profiling:
 - Real runtime $\approx 57s$.
 - Cumulative time (e.g., `join` = 171s) adds up waiting + scheduling overhead from all threads.
- Takeaway:
 - For I/O tasks \rightarrow threads can still improve responsiveness.
 - For CPU-heavy work \rightarrow use `multiprocessing` (separate processes, separate GILs).

Shared Memory in Python

Multithreading vs. Multiprocessing


- Multithreading \Rightarrow Concurrency
 - Multiple threads share the same memory inside one process.
 - Good for I/O-bound tasks (overlap waiting times).
 - Limited by the GIL: no true parallel CPU execution.
- Multiprocessing \Rightarrow Parallelism
 - Each process has its own memory and interpreter.
 - Achieves true parallelism across CPU cores.
 - Higher cost: inter-process communication (IPC).
- Key takeaway: `threading` = concurrency (I/O-bound).
`multiprocessing` = parallelism (CPU-bound).



Shared Memory in Python

Multi-processing Basics: Elements of Programming

- Memory Isolation
 - Processes do NOT share memory address space
- Fork/Join Execution Model
 - Fundamental way of expressing concurrency within a computation
 - `Fork` creates a new child process
 - Parent continues after the `Fork` operation
 - Child begins operation separate from the parent
 - Parent waits until child `joins` (continues afterwards)



Shared Memory in Python

Multi-processing Basics: Race Conditions

- A Race Condition occurs if:
 - Two or more processes manipulate a shared resource concurrently
 - The outcome depends on the order of access

Process P1:

- (1) MOV SUM, Reg1
- (2) ADD #1, Reg1
- (3) MOV Reg1, SUM

Process P2:

- (1') MOV SUM, Reg1
- (2') ADD #1, Reg1
- (3') MOV Reg1, SUM

Possible interleavings:

- (1')(1)(2)(3)(2')(3') \Rightarrow SUM = SUM+1
- (1)(1')(2')(3')(2)(3) \Rightarrow SUM = SUM+1
- (1)(2)(3)(1')(2')(3') \Rightarrow SUM = SUM+2

Solution

Synchronization needed to prevent race conditions

Mutual Exclusion: prevents simultaneous access to a shared resource.

Shared Memory in Python


Multi-processing Basics: Synchronization

Variable Mutex: S

- Boolean: 0 / 1
- General: Integer ≥ 0

Functions:

- **Lock(S)**
 - If $S == 0$ then wait until $S > 0$
 - If $S > 0$ then $S = S - 1$
- **Unlock(S)**
 - $S = S + 1$



- $(1')(2')(3')(1)(2)(3) \Rightarrow SUM = SUM+2$
- $(1)(2)(3)(1')(2')(3') \Rightarrow SUM = SUM+2$

Shared Memory in Python

Multi-processing Basics: Join/Fork Model

- The `multiprocessing` module provides an easy API for parallelism.
- Steps:
 1. Create a `Process` structure with target function + args
 2. Start processes with `.start()`
 3. Wait for processes to finish with `.join()`
- This model avoids the GIL by using separate processes.

```
1 import multiprocessing
2
3 def print_cube(num):
4     ...
5
6 def print_square(num):
7     ...
8
9 if __name__ == "__main__":
10     p1 = multiprocessing.Process(target=print_square, args=(10,))
11     p2 = multiprocessing.Process(target=print_cube, args=(10,))
12
13     p1.start()
14     p2.start()
15
16     p1.join()
17     p2.join()
18
19     print("Done!")
```

Shared Memory in Python

Multi-processing Basics: Shared Memory

- Each process has its own memory space.
- Global variables are not shared between processes.
- Example: modifying a global list in a process does not affect the parent.

```
In process: [1, 4, 9, 16]
In main: []
```

```
1 import multiprocessing
2
3 result = []
4
5 def square_list(mylist):
6     global result
7     for num in mylist:
8         result.append(num * num)
9     print("In process:", result)
10
11 if __name__ == "__main__":
12     mylist = [1,2,3,4]
13     p1 = multiprocessing.Process(target=square_list, args=(mylist,))
14
15     p1.start()
16     p1.join()
17
18     print("In main:", result)
```


Shared Memory in Python

Multi-processing Basics: Shared Memory with Array and Value

- The multiprocessing module provides objects to share data:
 - Array: a ctypes array allocated in shared memory.
 - Value: a ctypes variable allocated in shared memory.
- These objects must be passed as arguments to processes.
- Enables efficient communication between processes.


```
Result array: [1, 4, 9, 16]
Sum of squares: 30
```

```
1 import multiprocessing
2
3 def square_list(mylist, result, square_sum):
4     for idx, num in enumerate(mylist):
5         result[idx] = num * num
6         square_sum.value += result[idx]
7
8 if __name__ == "__main__":
9     mylist = [1,2,3,4]
10    result = multiprocessing.Array('i', 4)
11    square_sum = multiprocessing.Value('i')
12
13    p1 = multiprocessing.Process(target=square_list, args=(mylist, ↵
14                                result, square_sum))
15
16    p1.start()
17    p1.join()
18
19    print("Result array:", result[:])
20    print("Sum of squares:", square_sum.value)
```

Shared Memory in Python

Multi-processing Basics: Shared Memory with `Manager` (Advanced)

- `multiprocessing.Manager` allows sharing complex objects:
 - `list`, `dict`, `Queue`, `Array`, etc.
- A single `Manager` can be used by multiple processes, even across different machines.
- Slower than direct shared memory (`Array` / `Value`).



Process Communication

Multi-processing Basics: Shared Memory with multiprocessing.Manager

```
1 from multiprocessing import Manager, Process
2
3 def worker(shared_list, idx):
4     # each worker squares its index
5     shared_list[idx] = shared_list[idx] **2
6
7 if __name__ == "__main__":
8     with Manager() as manager:
9         data = manager.list([i for i in range(10)])
10
11         processes = [Process(target=worker, args=(data, i)) for i in range(len(data))]
12
13         for p in processes:
14             p.start()
15
16         for p in processes:
17             p.join()
18
19         print("Final result:", list(data))
```

```
1 Final result: [0, 1, 4, 9, 16, 25, 36, 49, 64, 81]
```

Shared Memory in Python

Multi-processing Basics: Using `multiprocessing.Queue`

- Queue is a simple way to communicate between processes.
- Can pass any Python object.
- Key functions:
 - `put()` : insert a value into the queue.
 - `get()` : read/remove a value from the queue.
 - `empty()` : check if the queue is empty.
- Useful for synchronizing and sharing results between workers.

```
1 import multiprocessing
2
3 def square_list(mylist, q):
4     for num in mylist:
5         q.put(num * num)
6
7 def print_queue(q):
8     print("Queue elements:")
9     while not q.empty():
10        print(q.get())
11
12 if __name__ == "__main__":
13     mylist = [1, 2, 3, 4]
14     q = multiprocessing.Queue()
15
16     p1 = multiprocessing.Process(target=square_list, args=(mylist, q))
17     p2 = multiprocessing.Process(target=print_queue, args=(q,))
18
19     p1.start(); p2.start()
20     p1.join(); p2.join()
```

Shared Memory in Python

Multi-processing Basics: Prime Calculation Queue (8 processes)

- Overhead comes from process management (fork/join, waiting).
- Different from threads: no acquire or GIL contention.

```
if __name__ == "__main__":
    work_queue = Queue()
    for i in range(0, 200000, 500):
        work_queue.put(i)

    processes = [Process(target=do_work, args=(work_queue,))
                 for _ in range(8)]

    # start_time = time.time()
    for p in processes:
        p.start()
    for p in processes:
        p.join()
```

tottime	percall	cumtime	percall	filename:lineno(function)
0.000	0.000	11.893	11.893	{built-in method builtins.exec}
0.000	0.000	11.893	11.893	process_sqrt.py:1(<module>)
0.000	0.000	11.879	1.485	process.py:142(join)
0.000	0.000	11.879	0.330	popen_fork.py:24(poll)
0.000	0.000	11.879	1.485	popen_fork.py:36(wait)
11.879	0.330	11.879	0.330	{built-in method posix.waitpid}
0.000	0.000	0.019	0.005	<frozen importlib._bootstrap>:1349(_find_and_load)
0.000	0.000	0.019	0.005	<frozen importlib._bootstrap>:1304(_find_and_load_unlocked)
0.000	0.000	0.018	0.005	<frozen importlib._bootstrap>:911(_load_unlocked)
0.000	0.000	0.018	0.004	<frozen importlib._bootstrap_external>:989(exec_module)

Execution time (8 process) = 11.893 sec

Shared Memory in Python

Multi-processing Basics: Race Conditions

- Race condition occurs when multiple processes access a shared variable concurrently.
- Without synchronization, final result is unpredictable.
- Example: Withdraw and deposit modify the same Value.

```
1 def withdraw(balance):
2     for _ in range(10000):
3         balance.value -=1
4
5 def deposit(balance):
6     for _ in range(10000):
7         balance.value +=1
8
9 if __name__ == "__main__":
10
11     # initial balance (in shared memory)
12     balance =multiprocessing.Value('i', 100)
13
14     p1 =multiprocessing.Process(target=withdraw, args=(balance,))
15     p2 =multiprocessing.Process(target=deposit, args=(balance,))
16
17     p1.start(); p2.start()
18
19     p1.join(); p2.join()
```

Shared Memory in Python

Multi-processing Basics: Locks to Prevent Race Conditions

- `multiprocessing.Lock()` ensures mutual exclusion.
- Only one process can access the shared resource at a time.
- Prevents data corruption, ensures consistent results.

Final balance: 100

```
1 def withdraw(balance, lock):
2     for _ in range(10000):
3         with lock:
4             balance.value -= 1
5 def deposit(balance, lock):
6     for _ in range(10000):
7         with lock:
8             balance.value += 1
9
10 if __name__ == "__main__":
11     balance = multiprocessing.Value('i', 100)
12     # creating a lock object
13     lock = multiprocessing.Lock()
14
15     # creating new processes
16     p1 = multiprocessing.Process(target=withdraw, args=(balance, lock))
17     p2 = multiprocessing.Process(target=deposit, args=(balance, lock))
18
19     p1.start(); p2.start()
20
21     p1.join(); p2.join()
```

Shared Memory in Python

Multi-processing Basics: Using multiprocessing.Pool

- Pool represents a pool of worker processes.
- Allows tasks to be distributed automatically.
- Methods:
 - map(func, iterable) – apply function to list of inputs.
 - apply() – run function once.
 - apply_async() – asynchronous execution.
- Efficient for data-parallel computations.

```
1 import multiprocessing
2 import os
3
4 def square(n):
5     print(f"Id for {n}: {os.getpid()}")
6     return n * n
7
8 if __name__ == "__main__":
9     mylist = [1,2,3,4,5]
10
11     # creating a pool object
12     p = multiprocessing.Pool()
13
14     # map list to target function
15     result = p.map(square, mylist)
16
17     print(result)
```


Shared Memory in Python

Multi-processing Basics: Prime Calculation using Pool (8 processes)

```
if __name__ == "__main__":  
    LIMIT = 200000  
    STEP = 500  
    N_PROCESSES = 8  
  
    # list of jobs  
    tasks = list(range(0, LIMIT, STEP))  
  
    with multiprocessing.Pool(processes=N_PROCESSES) as pool:  
        results = pool.map(sum_primes, tasks) # apply sum_primes for ↔  
                                              each job
```

```
tottime  percall  cuntime  percall  filename:lineno(function)  
0.000    0.000    16.164    0.505    connection.py:246(recv)  
0.000    0.000    16.163    0.505    connection.py:429(_recv_bytes)  
0.000    0.000    16.161    0.253    connection.py:390(_recv)  
5.486    0.081    16.161    0.253    {built-in method posix.read}  
0.000    0.000    10.704    10.704    {built-in method builtins.exec}  
0.000    0.000    10.704    10.704    pool_sqrt.py:1(<module>)  
0.000    0.000    10.692    10.692    pool.py:738(_exit_)  
0.000    0.000    10.686    10.686    pool.py:654(terminate)  
0.000    0.000    10.678    10.678    util.py:208(_call_)  
0.000    0.000    10.678    10.678    pool.py:680(_terminate_pool)  
:3
```

Execution time (8 process) = 10.704 sec