

## 02: Concurrent & Network Programming

### Distributed Software Architectures

Stefan Huber <[shuber.lba@fh-salzburg.ac.at](mailto:shuber.lba@fh-salzburg.ac.at)>

March 22, 2019

## Section 1

# Concurrent Programming

## Process:

- ▶ Multi-tasking: An OS can execute multiple programs at a time.
- ▶ A process is an OS abstraction to provide a *virtual processor* with *virtual memory*.
- ▶ The OS isolates processes from each other: It gives the illusion of having exclusive access to the processor and memory.
- ▶ *Process control block* contains information on:
  - ▶ Process Id
  - ▶ CPU states: instruction pointer, registers, memory mapping, ...
  - ▶ Privilege information: user id, group id, ...
  - ▶ Resources: opened files, sockets, pipes, ...
- ▶ Processes hierarchy: Every process has a parent process.

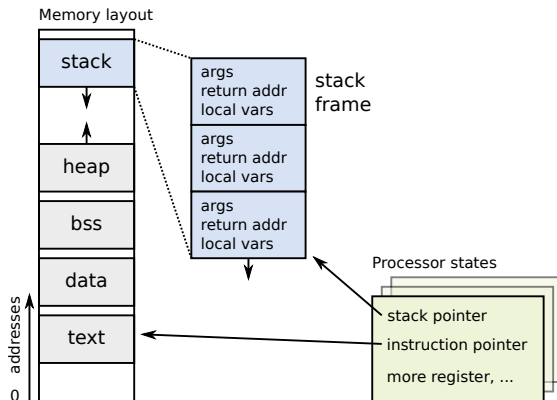
## Threads:

- ▶ Multi-threading: Multiple threads of execution in one process context.
- ▶ Threads operate on the process memory and with the process resources.
- ▶ Thread safety:
  - ▶ Code is thread-safe if, when executed by multiple threads, there is no unintended interference.
  - ▶ Global states<sup>1</sup> may break thread safety; concurrent access must be taken into account.

---

<sup>1</sup> This includes static variables in C functions. Thread-safety requires at least *re-entrancy*.

# Process memory layout



**Figure:** A simplified process memory layout and the relationship to processor registers. A “thread of execution” refers to the processor states.

## Latencies:

- ▶ Communication in a distributed system may involve significant latencies.
- ▶ Idea: Hide latencies by doing something else while waiting.

## Web browser:

- ▶ Has to fetch many resources besides referenced in an HTML document.
- ▶ Fetch them in parallel to avoid summing up latencies.
- ▶ Displaying content before having all resources fetched.
- ▶ Multi-threading simplifies implementation. For instance, each fetch can be implemented with a simple blocking read.
- ▶ Multiple connections in parallel may utilize multiple web servers.

## Example

*Electrolysis*<sup>2</sup> is a multi-process architecture of Firefox. The parent process forks 4 child processes to handle web content.<sup>3</sup> Each child process has a plethora of threads running.

---

<sup>2</sup> <https://wiki.mozilla.org/Electrolysis>

<sup>3</sup> Firefox 66 increases this to 8: [https://bugzilla.mozilla.org/show\\_bug.cgi?id=1470280](https://bugzilla.mozilla.org/show_bug.cgi?id=1470280).

- ▶ A modern web browser may create hundreds of threads.
- ▶ Most are probably not running most of the time but waiting.
- ▶ To what extent are those running in parallel?

## Definition (Thread-level parallelism (TLP))

For a total of  $N$  threads let  $c_i$  be the fraction of time where exactly  $i$  threads executed simultaneously. Then we define

$$TLP = \frac{\sum_{i=1}^N i \cdot c_i}{1 - c_0}.$$


- ▶ Modern web browsers have a TLP of 1.5 to 2.5.
- ▶ Hence, a processor with 2 to 3 cores could be utilized well.
- ▶ Moreover, the hundreds of threads are used with software architecture in mind.

Servers deal with multiple clients, hence multi-threading is typically more important for servers than clients.

- ▶ Simpler code.
- ▶ Higher performance.

## Example: File server

- ▶ Popular architecture: A dispatcher and multiple worker threads.
- ▶ Dispatcher thread receives requests from clients.
- ▶ Dispatcher distributes work to one of the worker threads.
- ▶ Worker thread performs the read with blocking reads on the file system.



## Benefits:

- ▶ Parallel handling of clients, unlike for a single-threaded implementation with blocking file IO.
- ▶ Simpler to implement than managing asynchronous calls with multiple clients.



## Performance gains:

- ▶ Parallel execution:
  - ▶ Hyper-threading
  - ▶ Multi-core processors
  - ▶ Multiple processors
- ▶ IO:
  - ▶ Native Command Queuing (NCQ) of hard disks, RAID systems.
  - ▶ Parallel fetch of web page resources to reduce serialized latencies.
  - ▶ (Concurrent IO can easily degrade performance, too!)

## Software architecture gains:

- ▶ Modularization into concurrent modules as architecture.
- ▶ Simpler code, e.g., we can still use blocking IO instead of asynchronous IO.
- ▶ Splitting up a process into multiple can improve security:
  - ▶ Isolation mechanisms
  - ▶ Finer privilege control

*Concurrent programming does not imply parallel (simultaneous) execution.*

Python documentation to concurrency:

- ▶ <https://docs.python.org/3/library/concurrency.html>

Modules in the Python Standard Library:

- ▶ threading
- ▶ multiprocessing
- ▶ and more...

`Thread` class represents a thread:

- ▶ A callable object (e.g., function) can be passed to the constructor.
- ▶ The method `start()` starts the thread execution.
- ▶ The method `join()` waits until thread terminates.

The standard Python interpreter (CPython) has a Global Interpreter Lock (GIL):

- ▶ No two threads can execute Python byte code at the same time.
- ▶ Makes access to Python data structures thread-safe.
- ▶ However, no parallelism for Python threads within a process!

The class `Lock` forms synchronization primitive:

- ▶ Can have two states: locked, unlocked
- ▶ The method `release()` unlocks a `Lock`.
- ▶ The method `acquire()` locks it.
  - ▶ If it is unlocked, it gets locked and continues execution.
  - ▶ If it is already locked, the calling thread gets blocked until another thread releases it.
- ▶ Can be used with `with` statement (context manager).

When dealing with all kind of resources we often produce code like this:

```
1 VAR = EXPR                # Creating file object, threading.Lock(), etc.
2 VAR.some_initialization() # Opening file, acquiring lock, etc.
3 try:
4     some_logic(VAR)        # File I/O, accessing resource, etc.
5 finally:
6     VAR.some_cleanup()     # No matter whether exception has been raised
                             # Closing file, releasing lock, etc.
```

This is a typical situation where the *context manager protocol*<sup>4</sup> can help:

```
1 # Acquire and later release mylock in an exception-safe fashion
2 with mylock:
3     do_something()
```

This translates basically to

```
1 mylock.__enter__()        # Does lock.acquire()
2 try:
3     do_something()
4 finally:
5     mylock.__exit__()     # Does lock.release()
```

<sup>4</sup> See <https://docs.python.org/3/reference/datamodel.html#context-managers> and <https://www.python.org/dev/peps/pep-0343/>.

The class `Barrier` makes all threads to wait for each other.

- ▶ A barrier is created for a number  $N$  of threads.
- ▶ The method `wait()` makes the thread wait until  $N$  threads called `wait()` and then all threads are released simultaneously.

The class `Event` is a one-bit signal (flag).

- ▶ The method `is_set()` returns `True` iff flag is set.
- ▶ The method `set()` raises the flag and `clear()` resets it.
- ▶ The method `wait()` makes the thread wait until flag is raised.



The `multiprocessing` package provides a similar API to `threading`.

- ▶ The class `Process` is like `Thread`.
- ▶ The method `start()` and `join()` start execution and wait until termination.
- ▶ There is also `Lock()`, `Event()`, `Barrier()`, and the like.

The `multiprocessing` package provides a similar API to `threading`.

- ▶ The class `Process` is like `Thread`.
- ▶ The method `start()` and `join()` start execution and wait until termination.
- ▶ There is also `Lock()`, `Event()`, `Barrier()`, and the like.

But there is more:

- ▶ Processes have their own address space; we require *Inter-Process Communication*:
  - ▶ The class `Queue` allows to `put()` and `get()` objects.
  - ▶ The function `Pipe()` creates a pair of connection object.
  - ▶ Shared memory.
- ▶ The class `Pool` allows to easily distribute computational load.

```
1 def f(x):  
2     return x*x  
3  
4 if __name__ == "__main__":  
5     with multiprocessing.Pool(5) as p:  
6         print(p.map(f, [1, 2, 3]))
```


Processes are scheduled by the OS: No GIL between processes.

## Section 2

# Network Programming

Anwesenheitsliste nicht vergessen.

7	Application	High-level APIs, remote file access, resource sharing
6	Presentation	Data translation including encoding, compression, encryption
5	Session	Communication sessions with back-and-forth transmission
4	Transport	Sending data segments and datagrams
3	Network	Routing packets, addressing, traffic control
2	Data link	Data frames between nodes
1	Physical	Bit stream over physical medium




- ▶ A *reference model*.
- ▶ Each layer dedicated to a specific task.
- ▶ Each layer builds upon the layer below, provides a service to the layer above
- ▶ Software engineering aspects:
  - ▶ *Separation of concerns*: Dedication to specific tasks, inner cohesion.
  - ▶ A concrete layer implementation must be *exchangeable*: It must only rely on the *contract* with the layer below and fulfills the *contract* to the layer above.

---

Application	HTTP, FTP, DNS, IMAP, POP, SMTP, SSH, Telnet, ...
Transport	TCP, UDP, ...
Internet	IPv4, IPv6, ICMP, ...
Link	Ethernet, WiFi, PPP, FDDI, ...

---



`host` DNS lookup utility

`netstat` Prints network connections, and more

`nc` Netcat is the TCP/IP swiss army knife

`socat` Socket cat, a mix of nc and cat

`curl` A command line data transfer tool by URL.

`tcpdump` Dumps traffic on a network

`wireshark` An extensive network analyzer for all kind of protocols

## Berkley sockets:

- ▶ *The* API for communication, not only for Internet network communication:
  - ▶ A socket is a communication endpoint.
  - ▶ The communication happens between two sockets.
- ▶ A typical sequence of socket calls looks as follows:

server	client
s = socket() s.bind() s.listen()	
a = s.accept()	b = socket() b.connect()
a.send()/recv() a.close()	b.send()/recv() b.close()
s.close()	

- ▶ The middle column is typically done repeatedly:
  - ▶ In a single-threaded loop.
  - ▶ In parallel threads.
  - ▶ In a forked process.



# Sockets in Python

```
1 import socket
2
3 # See https://docs.python.org/3/library/socket.html
4 # AF_INET = "Internet protocol", SOCK_STREAM = TCP, SOCK_DGRAM = UDP
5 s = socket.socket(socket.AF_INET, socket.SOCK_STREAM)
6 s.bind(("127.0.0.1", 1200))
7 s.listen()
8 a, addrinfo = s.accept()
9 print("recv: ", a.recv(1024).decode("utf-8"))
10 a.send(bytes("greetings from server", "utf-8"))
11 a.close()
12 s.close()
```

```
1 import socket
2
3 b = socket.socket(socket.AF_INET, socket.SOCK_STREAM)
4 b.connect(("127.0.0.1", 1200))
5 b.send(bytes("greetings from client", "utf-8"))
6 print("recv: ", b.recv(1024).decode("utf-8"))
7 b.close()
```

```
1 import socket
2
3 # With SOCK_DGRAM we do not use listen() and accept().
4 s = socket.socket(socket.AF_INET, socket.SOCK_DGRAM)
5 s.bind(("localhost", 1200))
6 print("recv:", s.recv(1024))
```

```
1 import socket
2
3 s = socket.socket(socket.AF_INET, socket.SOCK_DGRAM)
4 # We can use sendto() instead of connect() and send().
5 s.sendto(bytes("hi!", "utf-8"), ("localhost", 1200))
```

### Timeouts:

- ▶ Operations are by default blocking operations.
- ▶ A timeout of  $x$  (float) seconds can be set by `socket.socket.settimeout(x)`.
- ▶ If  $x$  is `None` then timeout is infinity.

### Privileged ports:

- ▶ Ports below 1024 are *privileged*.
- ▶ On Unix-like operating systems superuser rights are required.

The module `http.server` implements a simple (non-production) HTTP server.

► <https://docs.python.org/3/library/http.server.html>

```
1 import http.server
2 import threading
3
4 # A customized Handler class to process HTTP requests
5 class Handler(http.server.BaseHTTPRequestHandler):
6     def do_GET(self):
7         self.send_response(200)
8         self.send_header("Content-type", "text/plain")
9         self.end_headers()
10        self.wfile.write("Path was {}\n".format(self.path).encode())
11
12 if __name__ == "__main__":
13     server = http.server.HTTPServer(('localhost', 1080), Handler)
14     server_th = threading.Thread(target=server.serve_forever)
15     server_th.start()
16
17     input("Press <return> to exit")
18     server.shutdown()
```

The module `urllib.request` implements access to URLs in all detail.

- ▶ Authentication, SSL/TLS, redirections, cookies, ...
- ▶ <https://docs.python.org/3/library/urllib.request.html>

```
1 import urllib.request
2
3 if __name__ == "__main__":
4     # Does actually a 303 Moved Permanently to https://...
5     with urllib.request.urlopen("http://www.sthu.org/monalisa.txt") as f:
6         print(f.read().decode("utf-8"))
```