



# High performance web server 高效能 web 伺服器

學生：李祥宇

指導教授：王宏鏜

# Outline

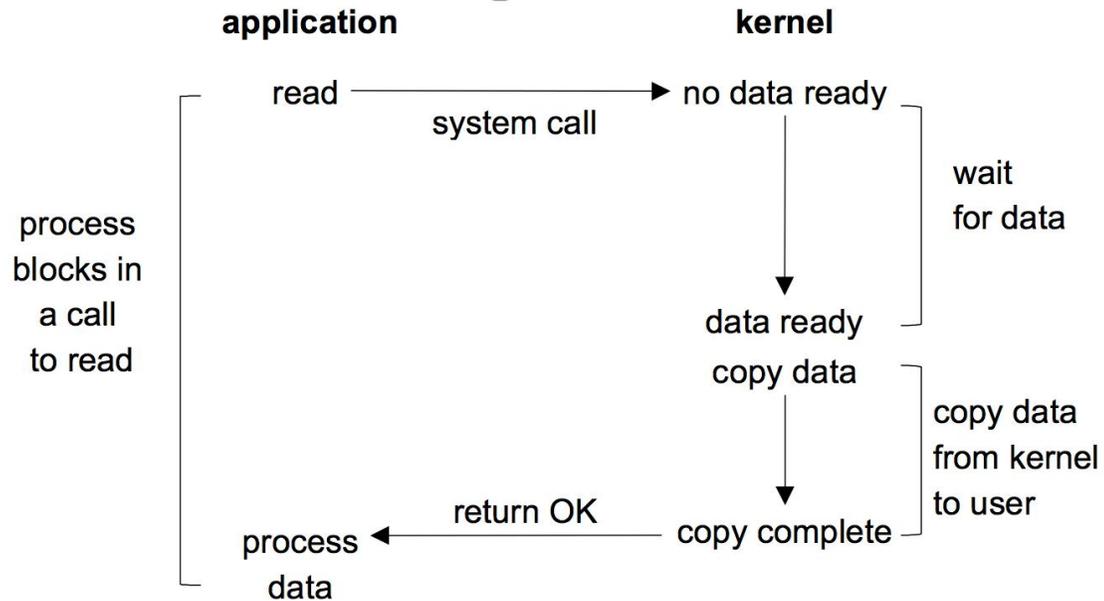
- ▶ 成果概述
- ▶ IO model 簡介
- ▶ Epoll & Io\_uring 簡介
- ▶ Web 伺服器實作
- ▶ 測試結果與 benchmarking

# 成果概述

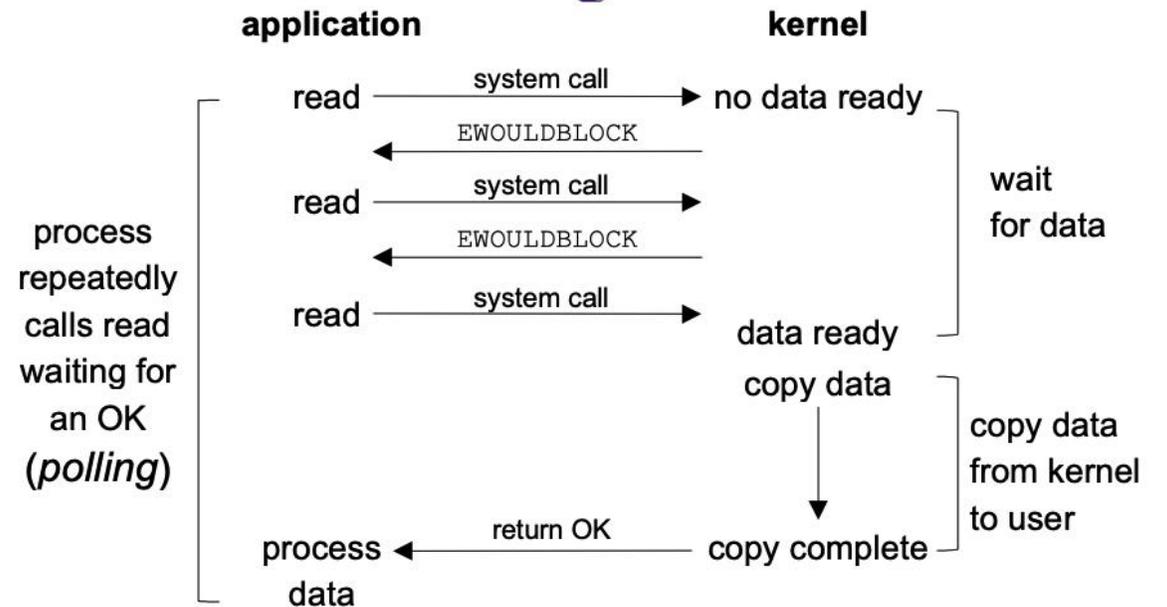
- ▶ 研究 Linux 上的 I/O model
- ▶ 運用近年來嶄新的 io\_uring 系統呼叫，結合多執行緒，打造高性能 web 伺服器
- ▶ 每秒處理約 20 萬個 http requests
- ▶ 比較 epoll 和 io\_uring 的效能差異

# Blocking I/O vs Non-blocking I/O

## Blocking I/O Model



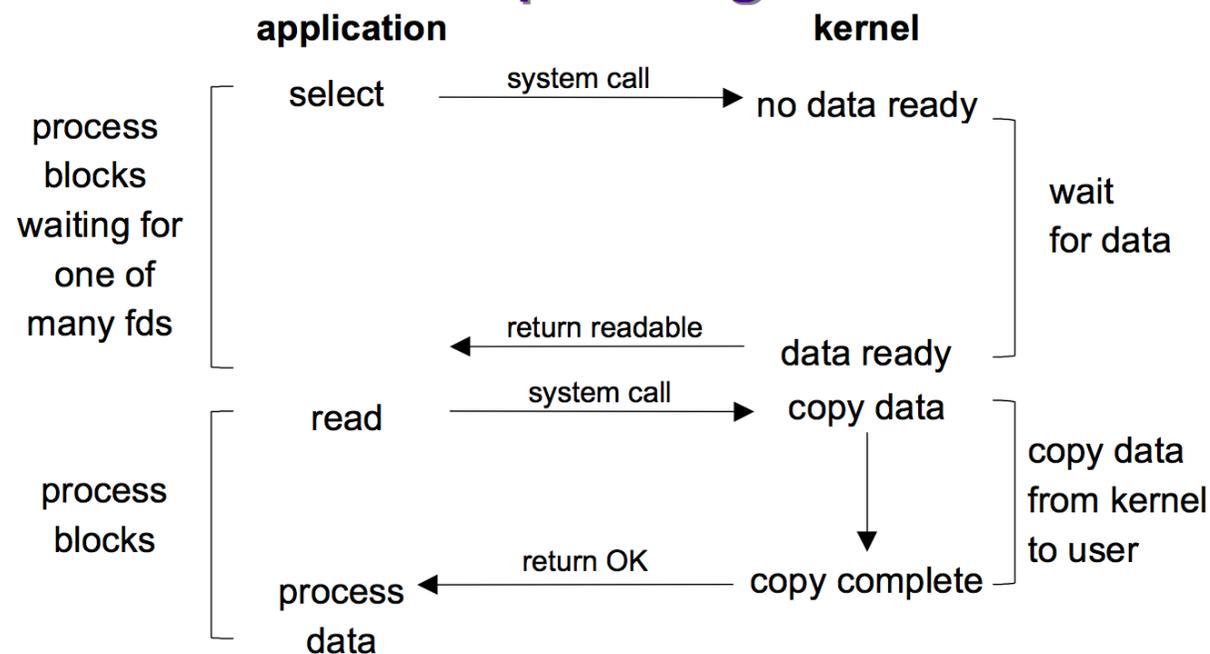
## Nonblocking I/O Model



# I/O Multiplexing

- ▶ I/O multiplexing 其實是 Blocking I/O 的延伸，如 select 和 epoll 的行為都是屬於 I/O multiplexing，差別在於他們能夠同時監事多個 file descriptors 而達到多工。

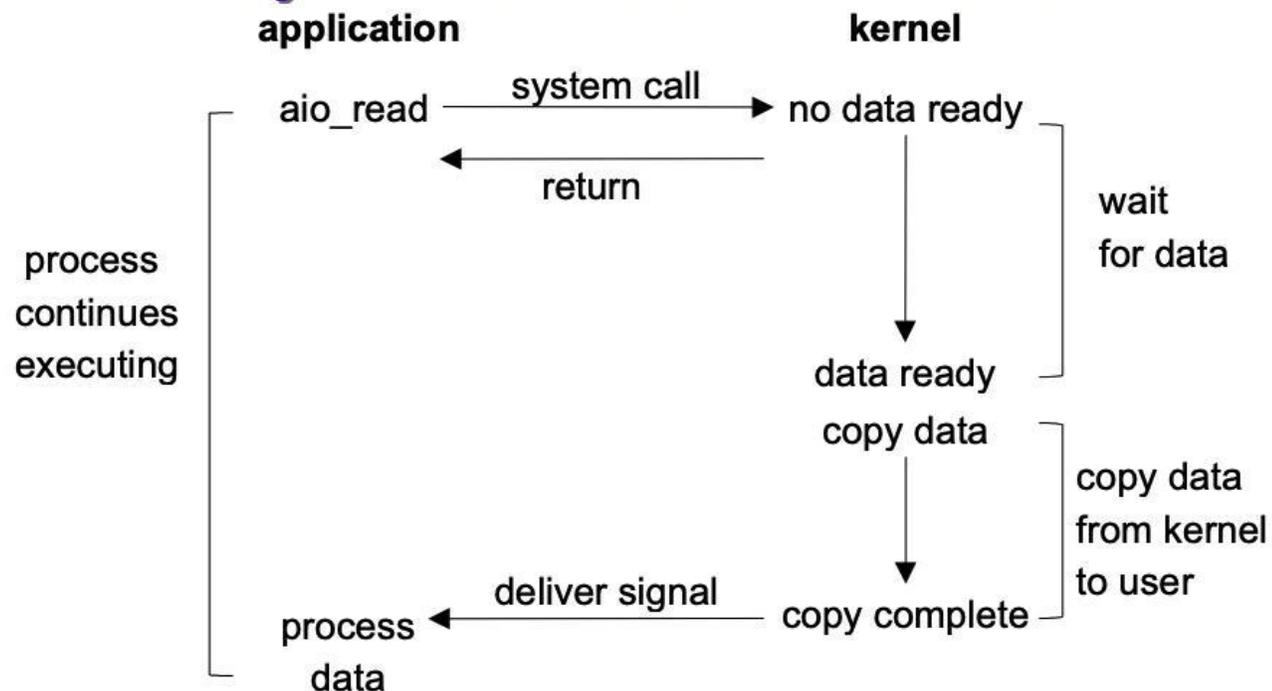
## I/O Multiplexing Model



# Asynchronous I/O

- ▶ 呼叫 `read()` 並切換至 `kernel mode` 之後，不論資料是否就緒，立刻返回 `user level`，方才的工作則留在核心，轉交其他 `kernel thread` 來完成。當完成後，核心會藉由 `signal` 等方式通知 `user level` 的 `process`。

## Asynchronous I/O Model

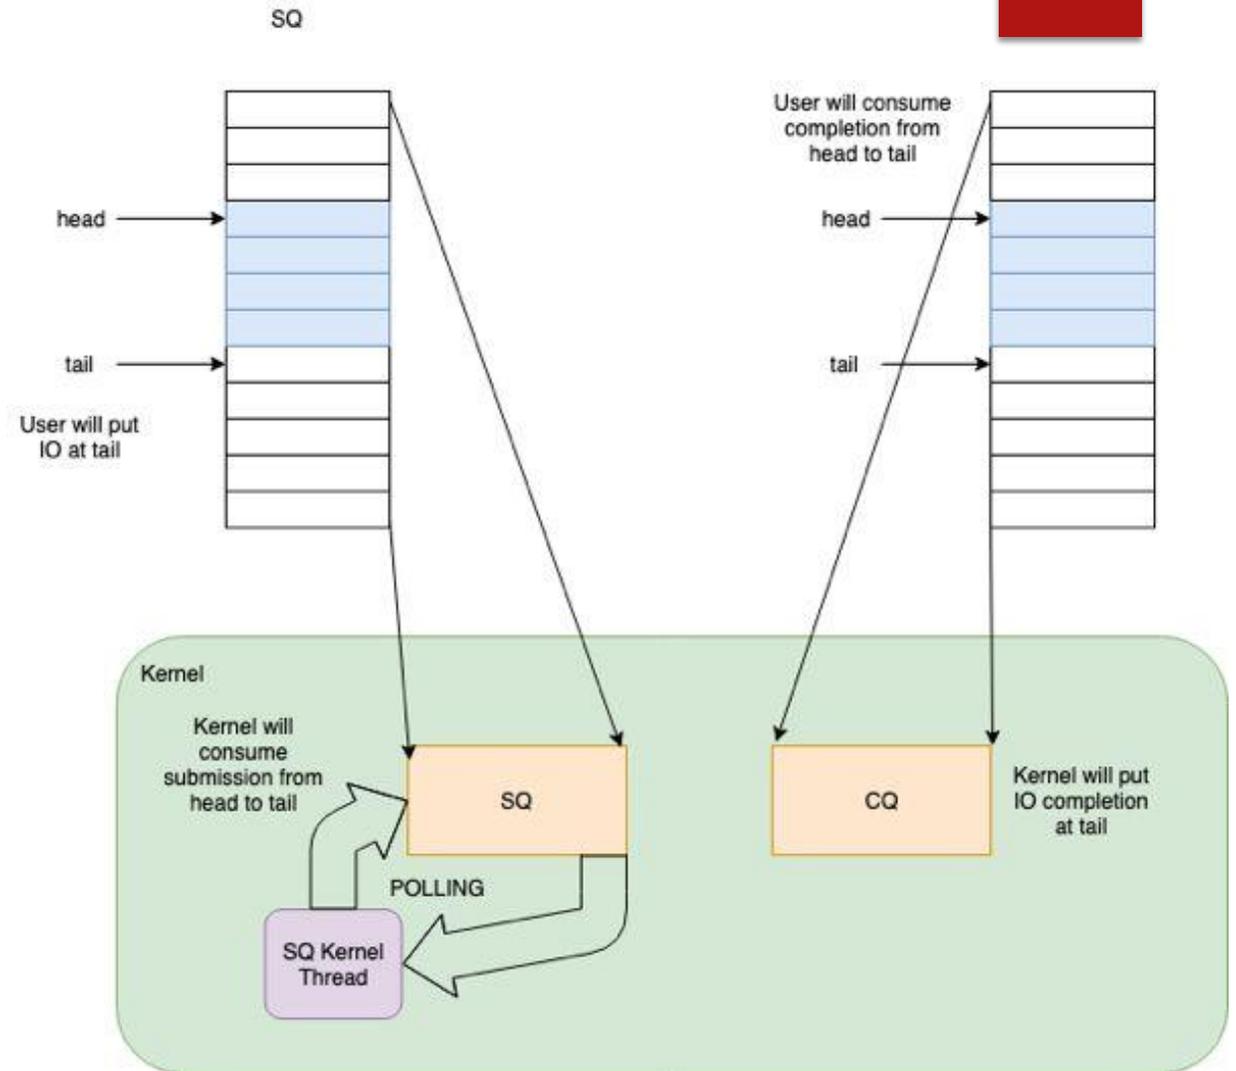


# Epoll

- ▶ IO multiplexing
- ▶ Interest list & ready list (maintained by kernel)
- ▶ Search fds in Rbtree
- ▶ Faster than select

# io\_uring

- ▶ 2019, linux 5.1, by Jens Axboe
- ▶ A “true” asynchronous I/O
- ▶ Shared ring buffer between user space and kernel space.
- ▶ Submission queue & Completion queue
- ▶ Faster than epoll



# Web 伺服器實作 - epoll

- ▶ 使用 Epoll 同時監聽多個 client fd
- ▶ Thread pool (pthread)
- ▶ Memory pool
- ▶ 將 request 交給 thread pool 裡的 workers 處理

# Web 伺服器實作 - io\_uring

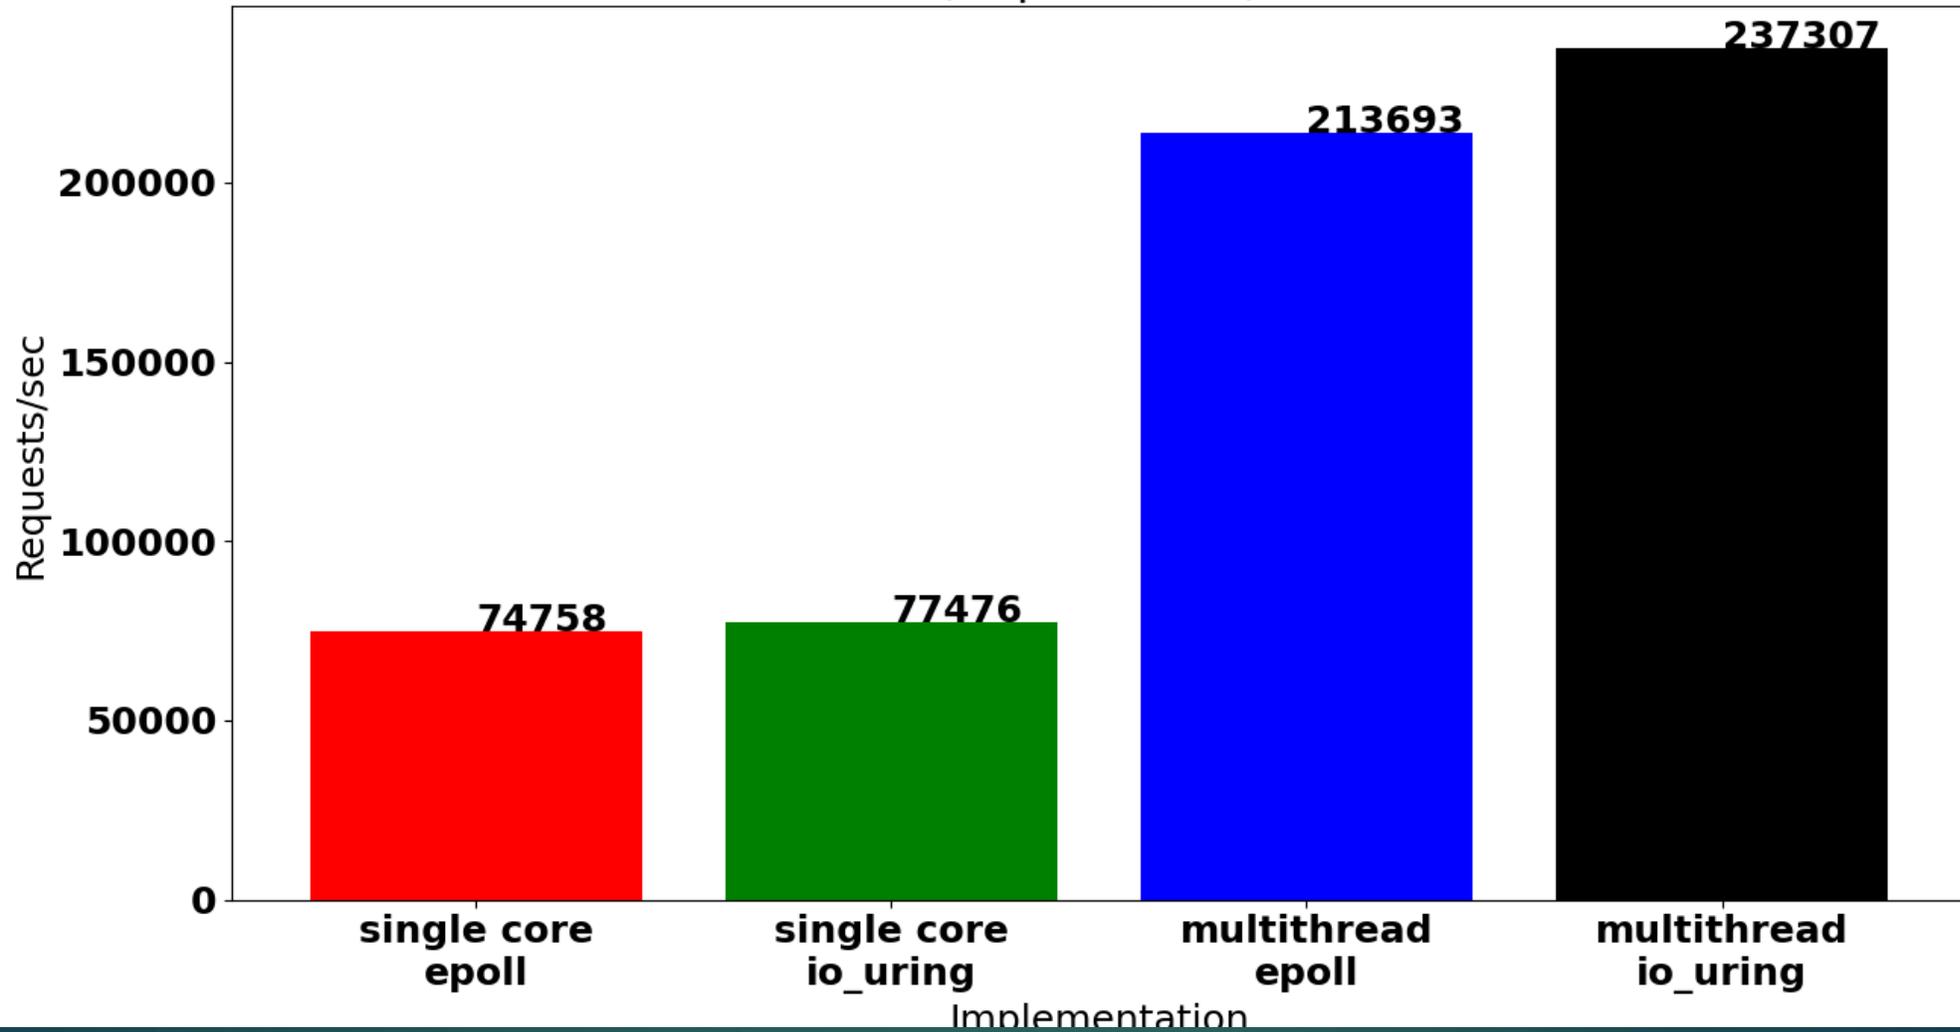
- ▶ Asynchronous I/O
- ▶ Multithreaded (pthread)
- ▶ One ring buffer per thread
- ▶ Each thread handles its own requests
- ▶ Memory pool

# Benchmarking

- ▶ Tools : wrk, ab, htstress
- ▶ 測量 20 次取平均
- ▶ Result : io\_uring wins!

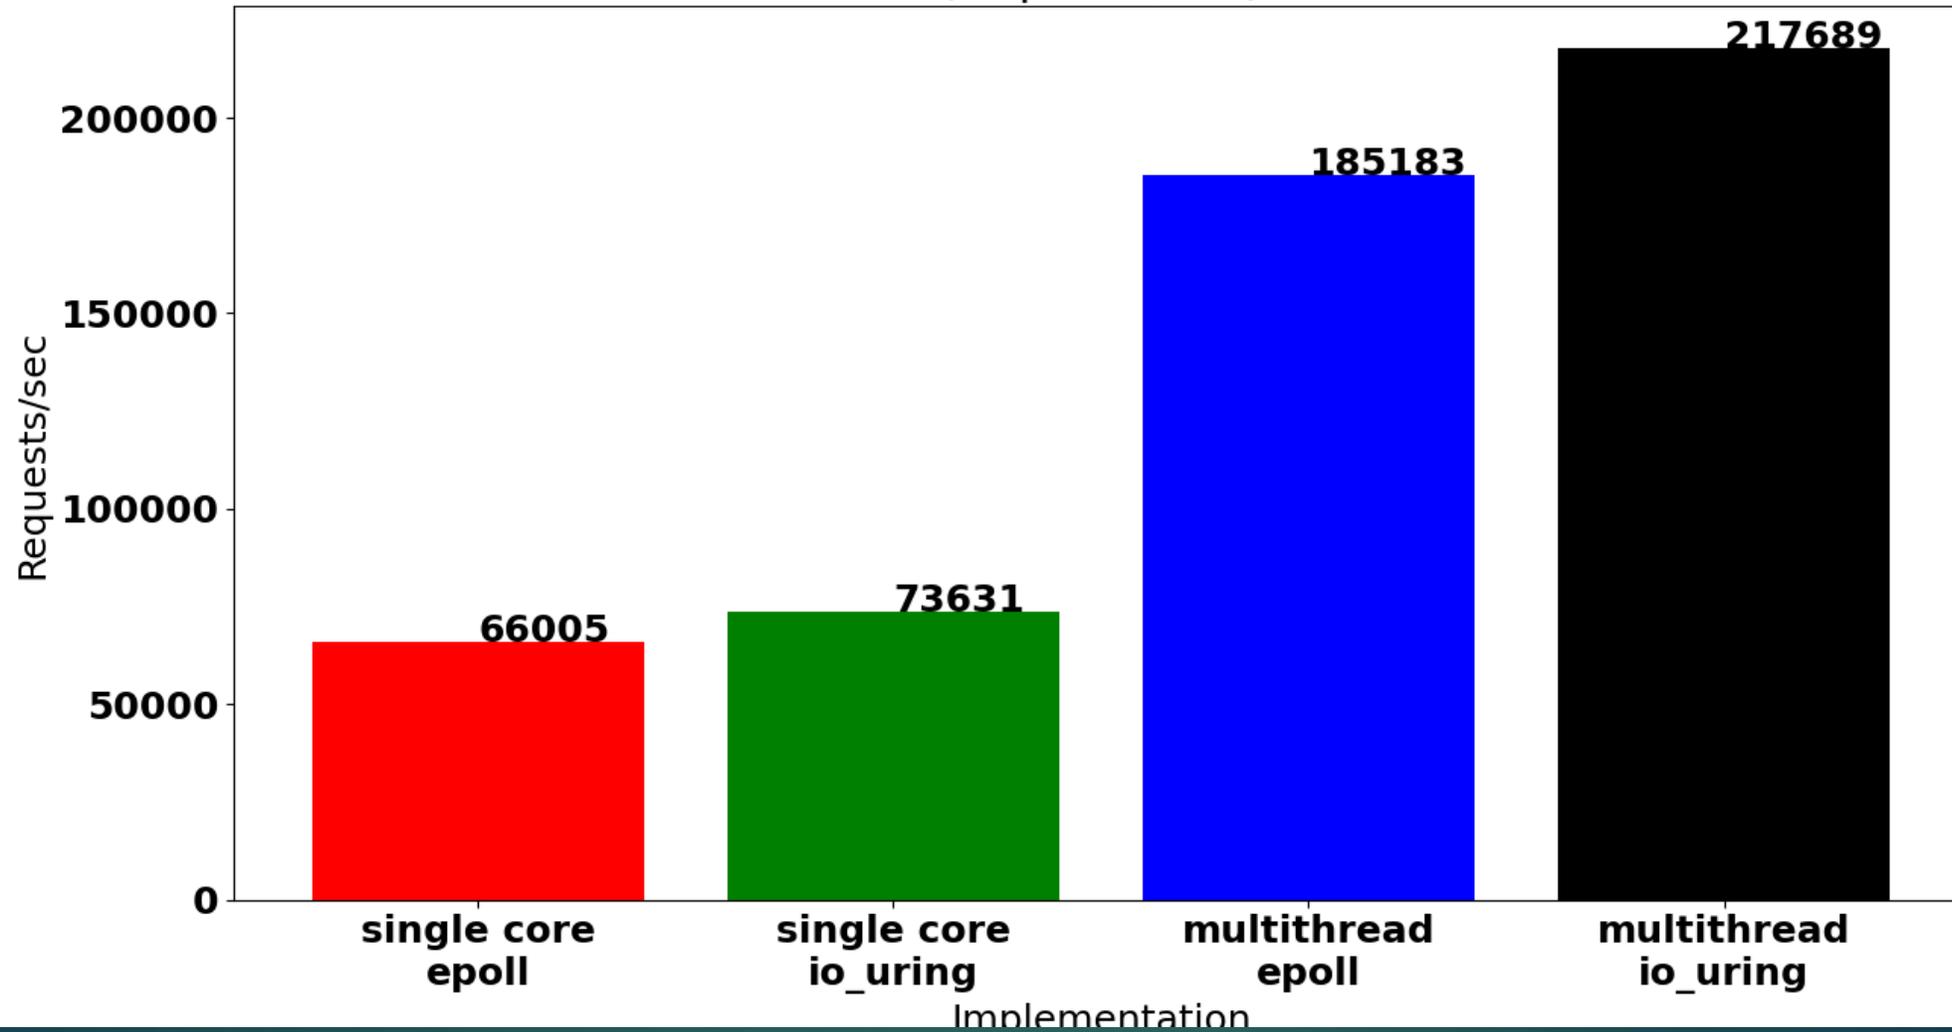
Tool : Wrk

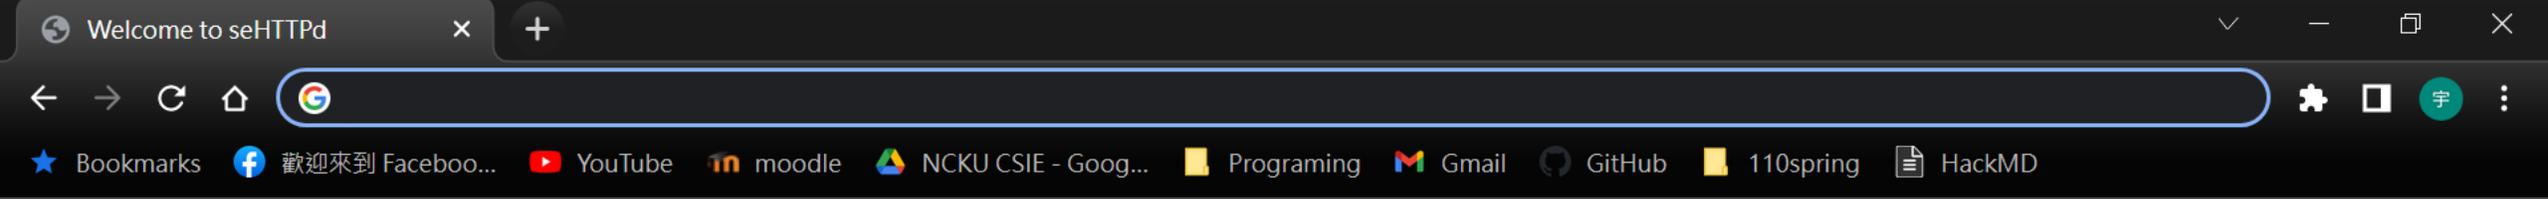
Performance with keep-alive  
(requests/sec)



Tool : htstress

Performance without keep-alive  
(requests/sec)





# Welcome!

If you see this page, the [seHTTPd](#) web server is successfully working.

DEMO