

- Scheduler runs go routines

- Pauses/resumes when goroutine is blocked on a channel or mutex operation

- Co-ordinates blocking system calls (file I/O)

- Takes care of running tasks

 - ↳ Garbage Collection |

8

go routines are USER-space threads

- Not kernel OS threads

 - ↳ similar though

 - ↳ Lighter + faster ^{to create} + smaller ~~mem.~~ footprint
+ easier to switch context

- OS only knows how to schedule threads on hardware (CPU)

go routines are managed by go runtime

go schedule puts go routines on the kernel threads!

go scheduler wants to utilise ^{hardware} parallelism

ie scale up to more cores!

work stealing ; Contention

Co-operative Preemption

⇒ Sys mon unscheduled GWS running
go-routines when possible

→ Background monitor thread

→ Moves them to a global run queue

↳ A lower priority run queue
checked less frequently than local thread
run queue

Thread Spinning

- Threads without work "spin" and look for work in global runqueue, poll the network and look at other runqueues to perform work stealing
- This uses up more CPU cycles, but leverages parallelism

Runqueues

- Run queues for cores are stored in a heap allocated struct 'p'

↳ Also stores resources the thread needs to run goroutines
⇒ like a memory cache.

- A thread will claim a runqueue to run goroutines, system will handover when thread is blocked.

Limitations of go scheduler

- FIFO runqueues \Rightarrow no way to prioritise goroutines
 - \hookrightarrow Not like linux scheduler which uses priority queues.
- No strong preemption
 - \hookrightarrow No ~~strong~~ fairness/latency guarantee
- Scheduler is not aware of hardware topology
 - \hookrightarrow Use LIFO runqueues for better cache utilisation?