

21 Jan 2025 - Operating Systems - II - Week 03

Data and Task parallelism \rightsquigarrow distribution among cores

→ a combination is used most of times

→ data parallelism → good for GPUs (task same)

→ Amdahl's law

logical cores
later

speed up \leq $\frac{1}{S + \frac{1-S}{N}}$

① → program runs serially

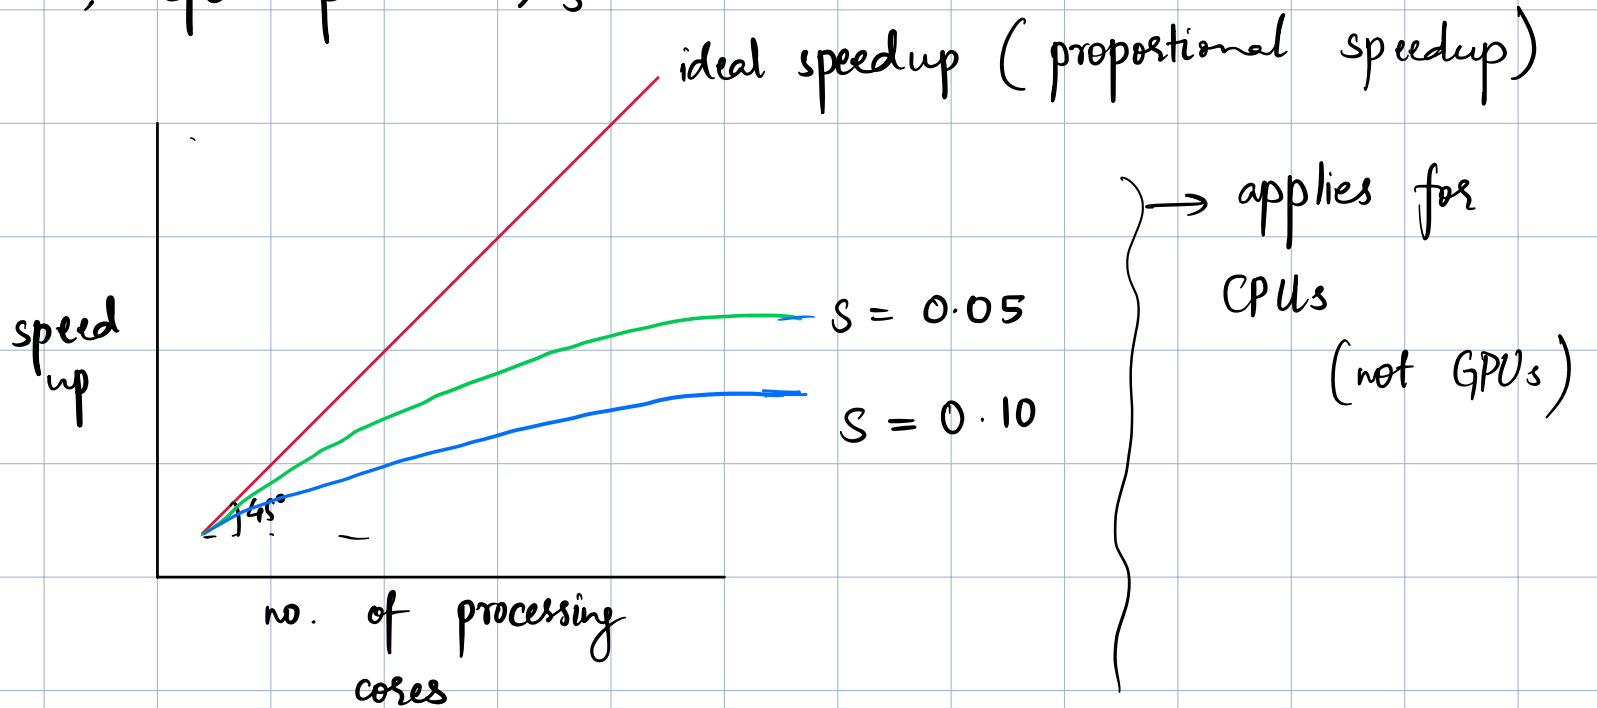
⑤ → serial portion

$\frac{1-S}{N}$ → } → 1

time for parallel portion
no. of cores

→ serial portion creates a disproportionate effect on perf gained by adding additional cores.

$\rightarrow N \rightarrow \infty$, speedup $\rightarrow 1/s$



0 Sudoku assignment

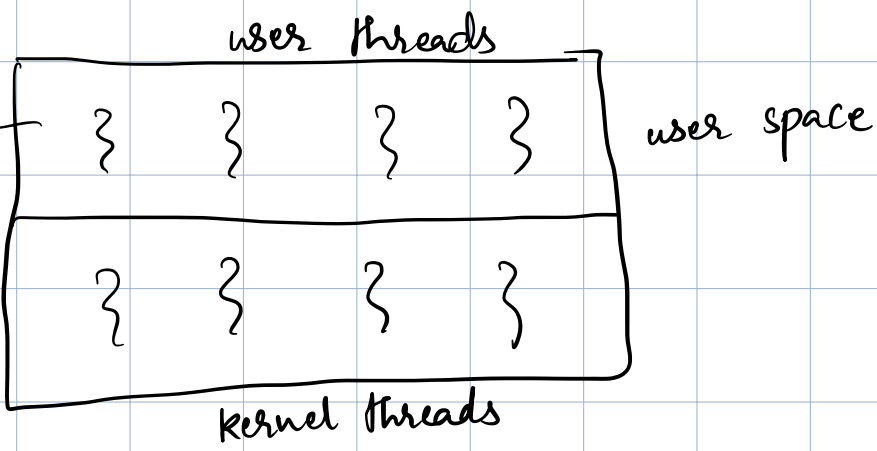
core of O.S → makes system calls on behalf of application programs

User threads and Kernel threads

↓
supported by kernel

different thread libraries for diff O.Ses

- pthreads → Linux
- C++ threads



can only access user space memory

for accessing kernel, make system call e.g printf internally makes system call

Why?
→ security
→ bug in user code can affect others

Multithreading models

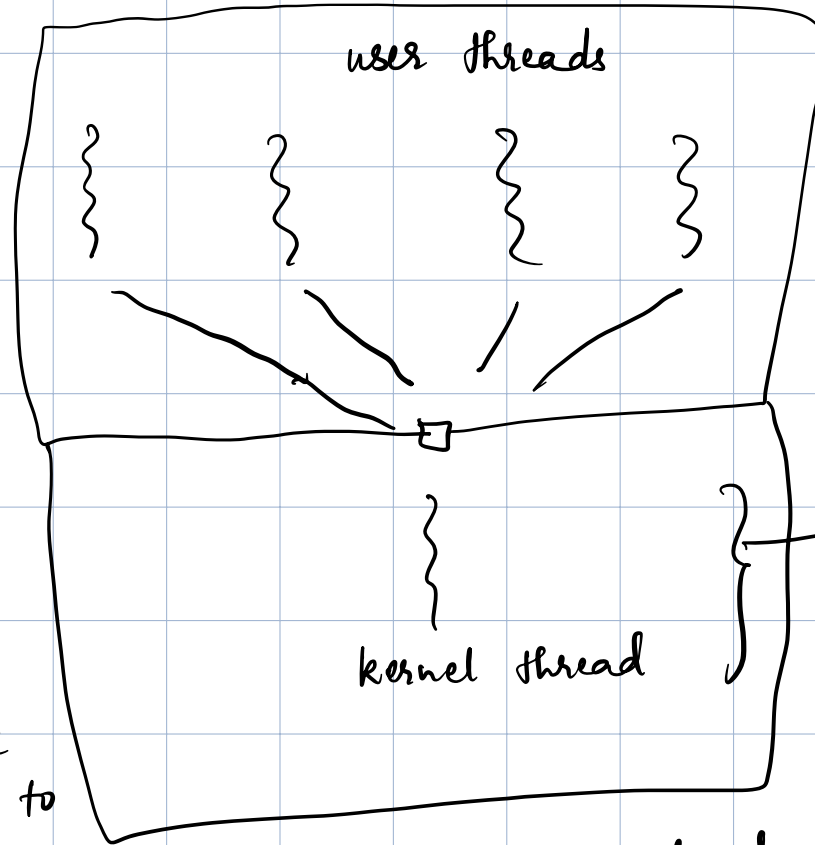
no parallelism } can run even
concurrent } on single core

Many-to-one:

→ Many user-level threads mapped to single kernel thread

→ One thread blocking causes all to block

not able to execute further ^{with} due to lack of data or due to scheduling algorithm



OS recognizes only this

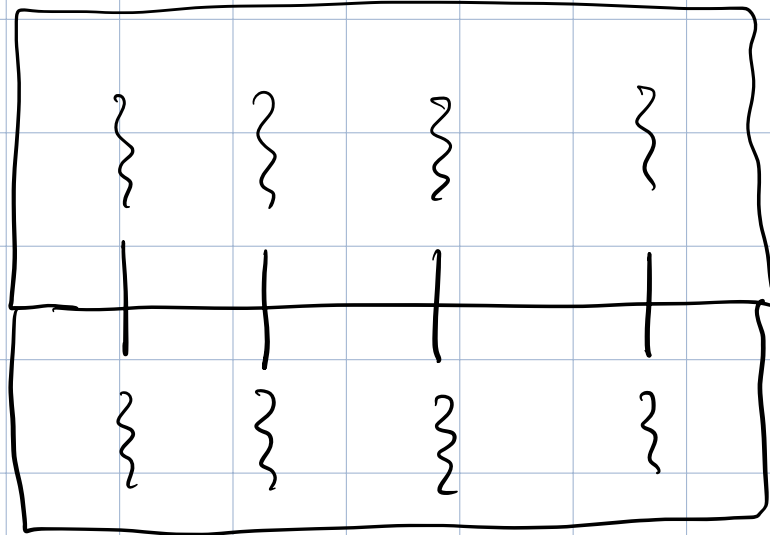
Advantages

??
..

One-to-one \rightsquigarrow pthread_create

\rightarrow Windows, Linux

\rightarrow No. of threads per process sometimes restricted due to overhead (Not in Linux)




\rightarrow disadvantage: one process can create many threads $\left. \begin{array}{l} \rightarrow \text{advantage} \\ \text{for many to one} \end{array} \right\}$

Many-to-Many

best of both but
not very common

→ overhead in
thread management




later
scheduling
across multiple
cores

Two-level model

→ always mapped to same
kernel thread

→ You can create kernel threads (kthread) directly but you cannot test it without an application program

→ drivers as programs attached to kernel
 ↓
 interact with kernel



Thread libraries → slides

pthread
 ↘ specification → like documentation - states behaviour
 ↘ not implementation

ANSI standard
 ↙ gcc supports

Java }
 C# } → platform independent
 JS }
 Python }


C } → a.out .exe


Java } → not an executable intermediate state (byte code)

22 Jan 2025

Assignment

2 plots \rightarrow each plot \rightsquigarrow 3 curves
mixed, chunk, sequential





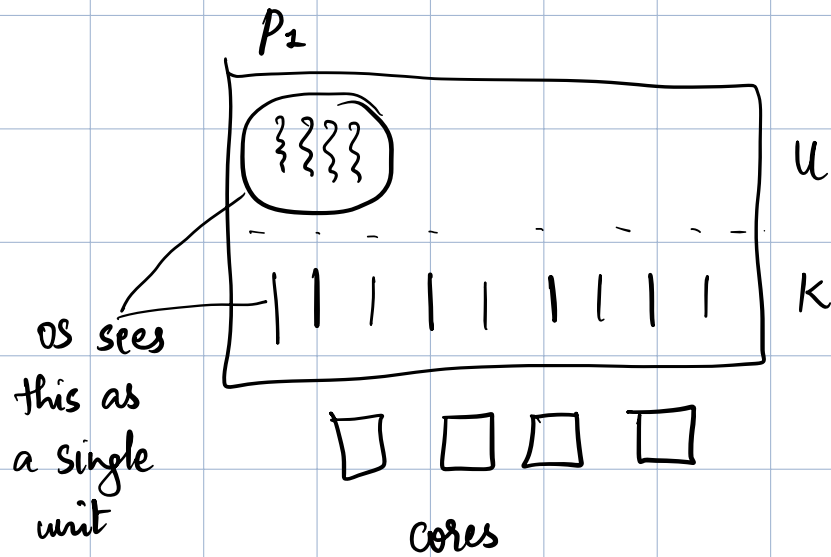
— Thr 1
— Thr 2
⋮
32

} create using for loop
p-thread - create

timer : start bet pthread - create
stop after join

Many - to - one : how does one thread getting blocked
blocks others ?

— OS sees only kernel level threads



Thr 1 : scanf

} time multiplexing

no. of threads \uparrow

\Rightarrow context switching, etc.

perf \downarrow

Implicit Threading \rightsquigarrow not in C++, provided by other languages

\rightsquigarrow they wait in a pool step

\rightarrow Create threads before

\rightarrow faster ; creating threads comes with an overhead

Java threadpool \rightarrow threads remain idle when task is done.

- tasks to be performed independently from creation
separately
- easy for programming
- flexibility reduces

Fork - join parallelism



OpenMP → set of compiler directives

→ compiler manages the thread

```
#include <omp.h>
#include <stdio.h>
C code sequential
```

```
#pragma omp parallel
{ printf("moo")
}
```

...

```
# pragma omp parallel for
```

```
for (i = 0 ; i < N ; i++) {
```

```
    c[i] = a[i] + b[i]
```

```
}
```

caching
??

Grand central dispatch → macOS, iOS

```
→ ^ } printf ( ... ); }
```

```
^ { }
```

Blocks

↓
assigned to a
thread pool

→ Two types of dispatch queue

- serial

- concurrent

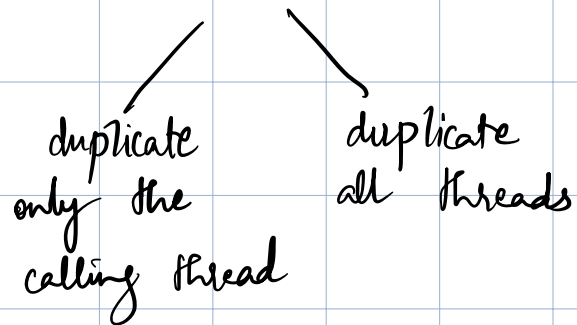
Intel Thread Building Block (TBB) \rightarrow open sourced

\rightarrow threads are implicitly created

\rightarrow very helpful for p-threads

Threading issues

Semantics of fork() and exec()



like interrupts
Signal handling

→ default sigle

→ user defined signal handler

standard UNIX fn for delivering a signal

→ kill

int kill (pid_t p1, int sig)

↓

if p₂ is multithreaded

which thread should

receive

* pthread_kill

→ windows does not explicitly provide support for

signals. → APC

* thread - cancellation

↳ searching done in 1st thread

↳ quitting a tab in browser while it's loading

→ deprecated → can leave the system in an inconsistent state.

- asynchronous - immediate termination

- deferred - target thread checks periodically

↙ ↘ better

default

↙ cancellation

occurs only when thread reaches cancellation

points → system defined

bank transaction

file moving

flag is raised ⇒ do cleanup and terminate

→ You can also disable thread cancellation

- linux implements cancel internally using signals
kill

→ cleanup handler ~ user defined

Thread Local Storage

↳ similar to static

static
within
function

local

global

```
Thr 1() {  
    TLS t1;  
    ...  
    i;  
    fn1()  
}
```

```
Thr 2() {  
    ...  
    t2  
}
```

```
Main {  
    ...  
    thr1();  
    thr2();  
    ...  
}
```

fn 1()

fn 2()

local to thread 1 and all its functions
ll → visible to fn 1, not to fn 2, main
i → not visible to fn 1

Why? Thread local functionalities

```
static __thread int threadID; // gcc
```

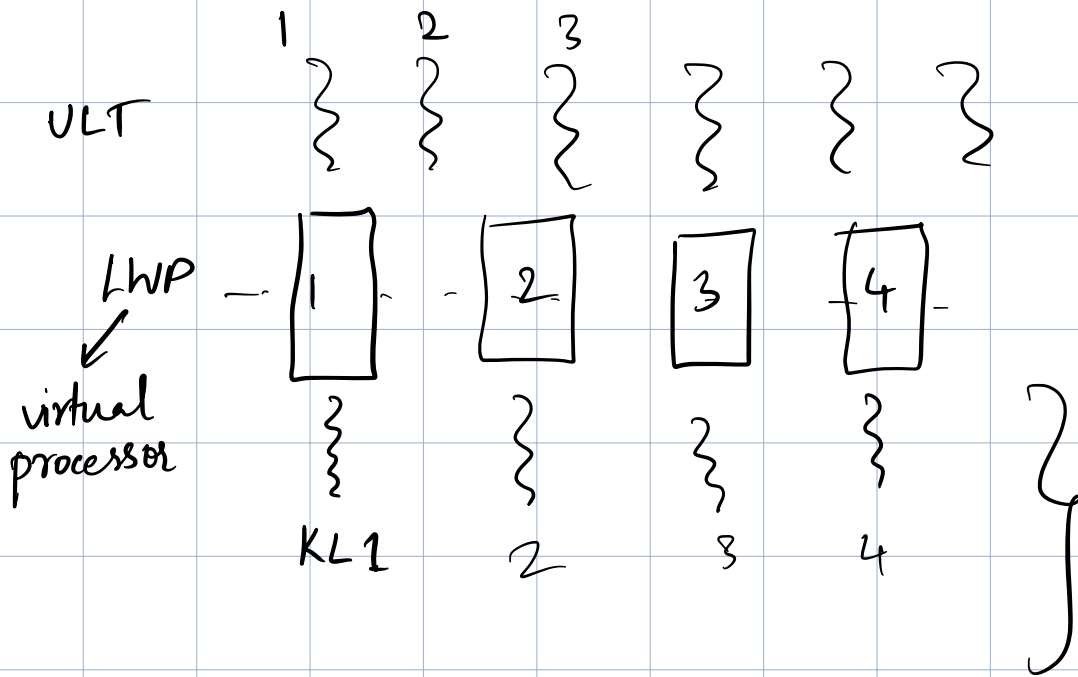
Scheduler Activation

M:M and
Two-level models
require communication
to maintain

← uses thread

LWP data structure

} kernel thread



TML → thread management library

} OS takes care only on KLT

Linux threads

- Calls threads as task
↓
can be a process or a task

- thread creation is done using `clone()` `fork` `pthread_create`

- flags for clone

