

Threads and Concurrency

— DM510 Operating Systems

— Lars Rohwedder



Windows



macOS



iOS

Disclaimer

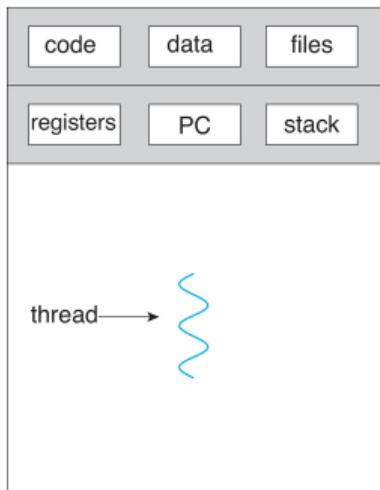
These slides contain (modified) content and media from the official Operating System Concepts slides:
<https://www.os-book.com/OS10/slide-dir/index.html>

Today's lecture

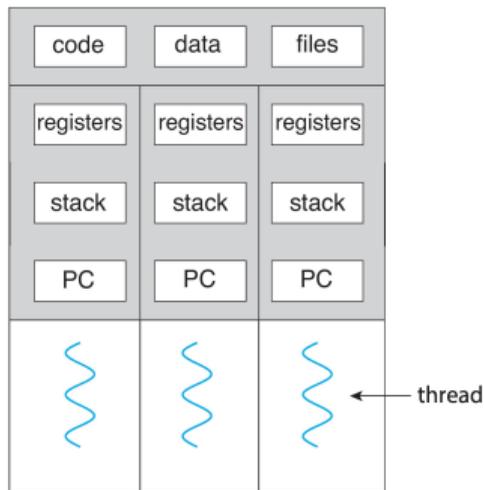
- > Chapter 4 of course book

Multi-threaded process

- > Last lecture: several processes (running simultaneously) can cooperate on tasks
- > Instead of several processes, often several **threads** within one process are used
- > Reasons for using several threads/processes: continuing during a blocking call or CPU intensive task (especially for responsiveness), parallelism (multicore processor)



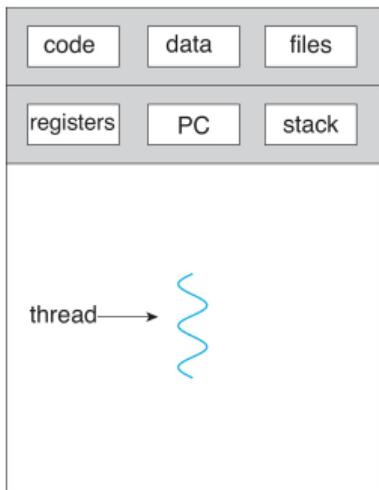
single-threaded process



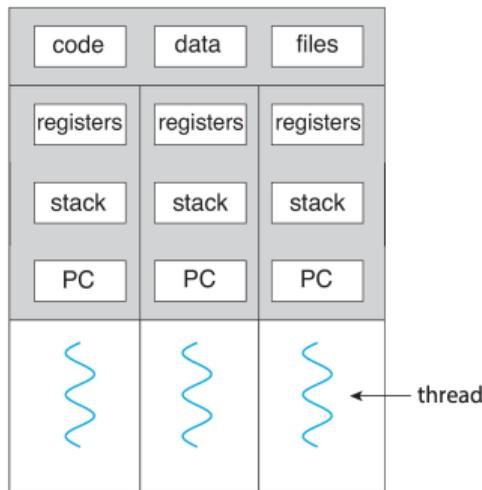
multithreaded process

Multi-threaded process

- > Last lecture: several processes (running simultaneously) can cooperate on tasks
- > Instead of several processes, often several **threads** within one process are used
- > Reasons for using several threads/processes: continuing during a blocking call or CPU intensive task (especially for responsiveness), parallelism (multicore processor)



single-threaded process



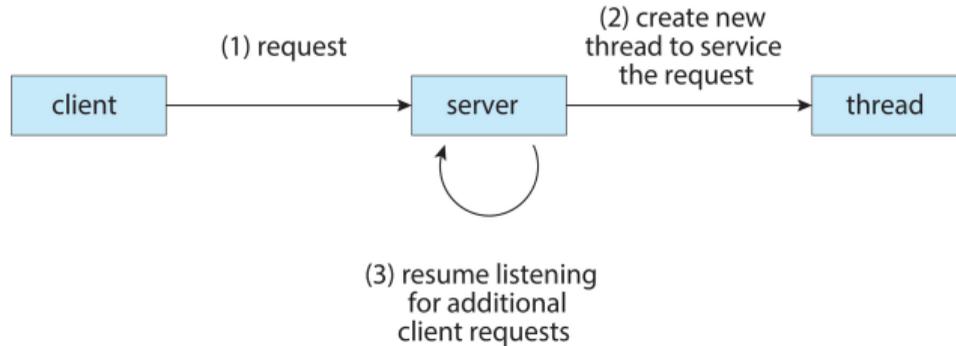
multithreaded process

Advantages of threads

- > Lower resource overhead (especially memory)
- > Faster thread creation and thread switches
- > Communication via shared memory simpler (address space is shared)

Examples

- > A web server responding to HTTP requests:



- > A database needs to search a large table for a specific entry. It scans the first quarter on one thread (running on CPU core 1), the second quarter on another thread (running on CPU core 2), etc.

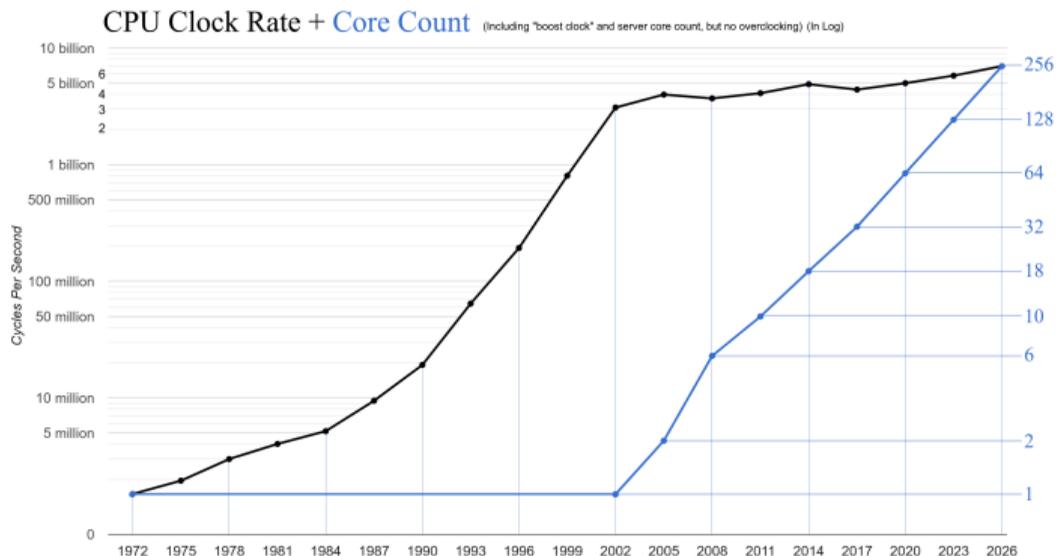
Parallelism

Background on parallelism

- > Speed at which processors execute instructions (in million cycles per second: MHz) is no longer increasing due to physical limits
- > Performance improvements are now mainly due to **parallelism** (more CPU cores)

Challenges

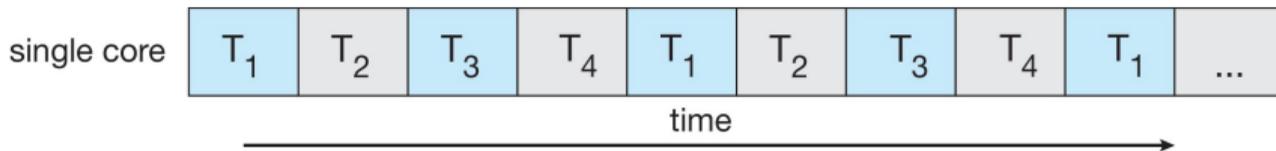
- > Dividing activities
- > Balancing
- > Data splitting
- > Data dependency
- > Testing and debugging



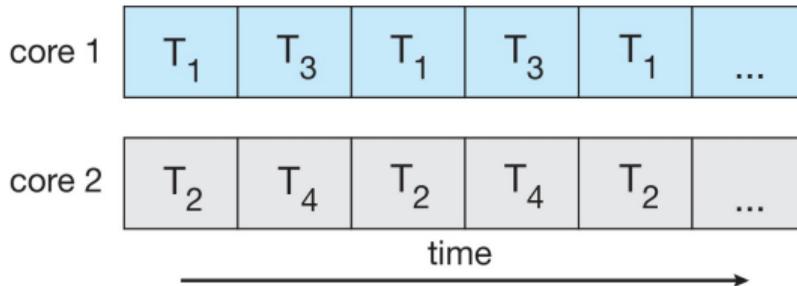
Source: <https://www.singularity.com/charts/page6f1.html> 1976-1998 "Microprocessor clockspeed" (modified) | 1999-2026 Manufacturer specifications
Method: Fastest / highest (not coupled) core count CPU released (1999-2024) 2026 - Leaks + conservative prediction

Concurrency and parallelism

Note that **concurrency** (simultaneous progress on different tasks) is also possible on single core via a preemptive scheduler:

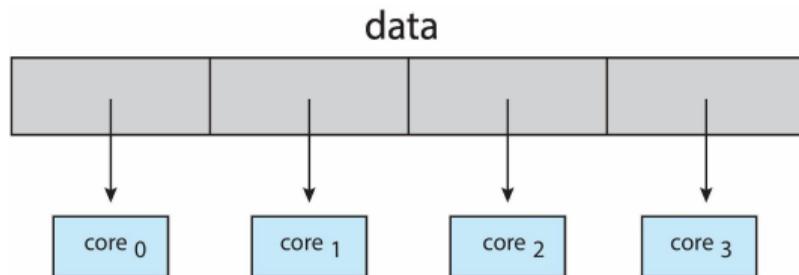


True **parallelism** on several cores:

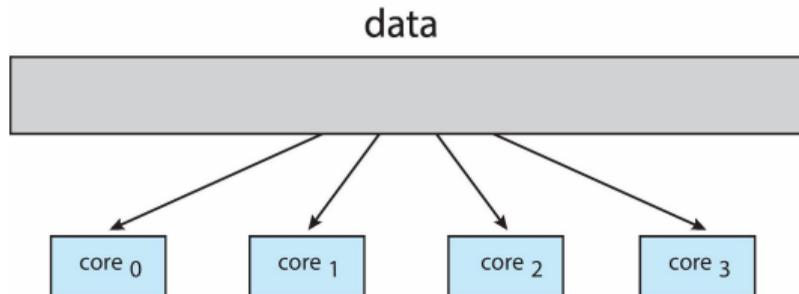


Forms of parallelism

Data parallelism: distribute data to cores performing the same task on each batch



Task parallelism: distribute different tasks to cores, working on same data



Amdahl's law

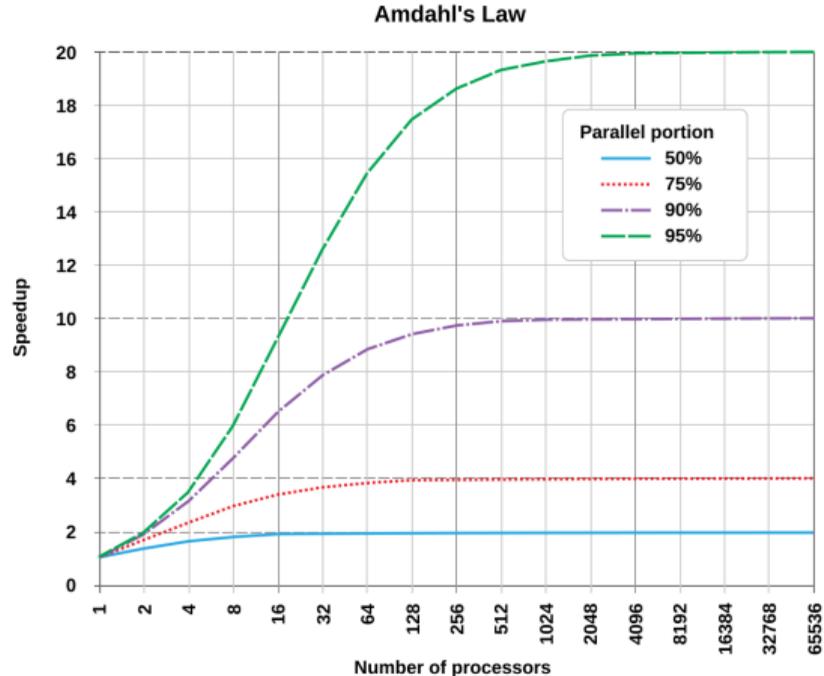
Possible speedup due to parallelism (even with many cores) is limited by how sequential the program is

Theoretical speedup

Let S be the ratio of sequential operations to all operations. Then

$$\text{speedup} \leq \frac{1}{S + \frac{1-S}{\text{cores}}}$$

This estimate is very optimistic since it does not account for communication overhead



Thread internals

Overview

- > Typically, libraries handle internals of thread implementation
- > **pthread**s (POSIX threads) is a widespread API specification implemented in many libraries, especially on UNIX systems, see exercises

Overview

- > Typically, libraries handle internals of thread implementation
- > **pthread** (POSIX threads) is a widespread API specification implemented in many libraries, especially on UNIX systems, see exercises

We distinguish:

Kernel threads

The kernel creates, schedules, terminates threads like processes, but resources are shared. In particular:

- > threads may run on different cores
- > other threads can continue when one is in a blocking system call

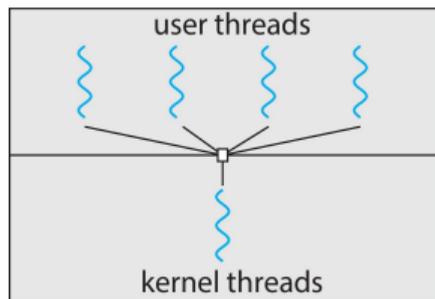
User threads

Threads that are visible to programmer. Sometimes implemented purely in user code without kernel.

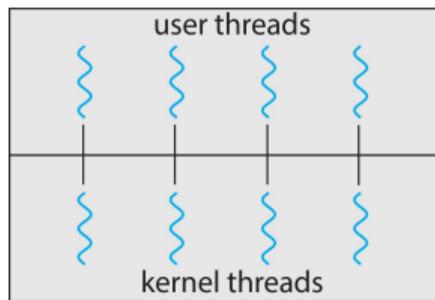
- > Limited: process uses single core, blocking call suspends all threads
- > Low resource overhead, applicable also to embedded systems without OS or with very limited OS

User-kernel thread mapping

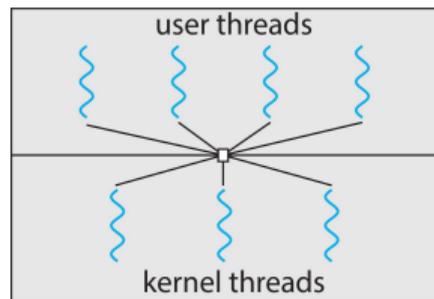
> We can map user threads (of one process) to a smaller number of kernel threads in different ways:



Many-to-one



One-to-one



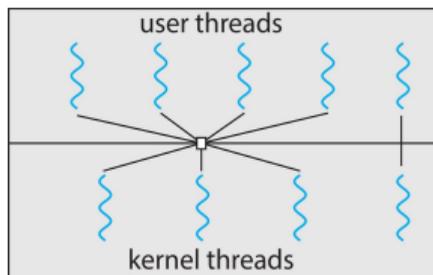
Many-to-many

user
space

kernel
space

- > **one-to-one** model is by far the most popular and default option
- > Possible motivation for not using one-to-one is that kernel threads are more resource heavy than user threads and we may not have perfect control over scheduling of kernel threads

Light-weight processes

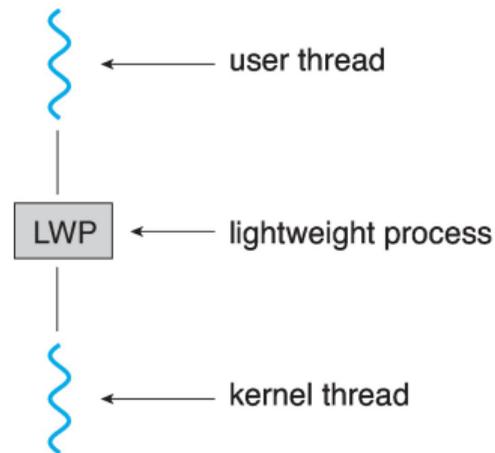


Two-level

- > A variant of many-to-many is **two-level** model that allows user process to carefully decide how user threads are allocated to kernel threads
- > This provides fine-grained control, but also requires a sophisticated kernel interface, usually achieved by **lightweight processes (LWP)**

Lightweight process

- > A LWP is a wrapper for a kernel thread
- > threads are allocated to LWPs by user process
- > Kernel informs process via **upcall** when: user thread enters blocking call (so that LWP can be used for other user thread) or thread is no longer blocked



Cancelling threads

- > **asynchronous cancellation**: The simplest way to cancel a thread is to stop its execution at the point it currently is, but this might leave process in an inconsistent state
- > **deferred cancellation**: A safer alternative is that a thread specifies points where it can be cancelled. Then it continues to run until it reaches the next such a point

Implicit Multi-Threading

Motivation of implicit multi-threading

Optimizing performance and ensuring correctness in multithreaded programs can be difficult:

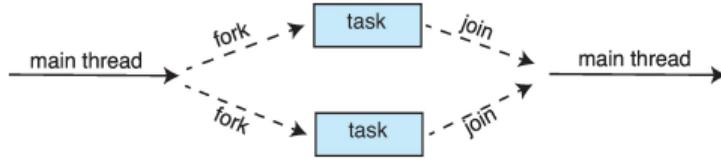
- > how many kernel threads should be used? *Depends on CPU architecture*
- > How to pass data between different threads? *Requires synchronization, which causes overhead*
- > In general, writing optimized code for all architectures may be difficult.

At the same time, many similar programs face these challenges.

There are various libraries and frameworks, which take a computational task and parallelize it almost automatically

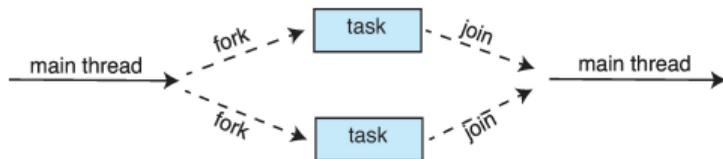
Fork-join model

- > When task is large, split into subtasks, do both in parallel (fork), then wait for results (join)

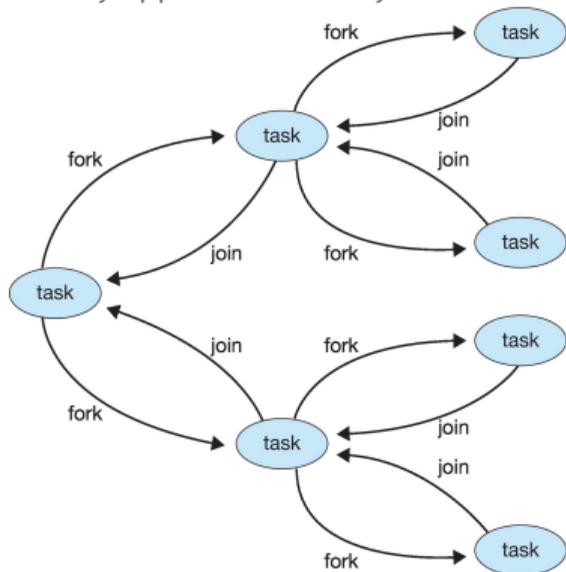


Fork-join model

- > When task is large, split into subtasks, do both in parallel (fork), then wait for results (join)

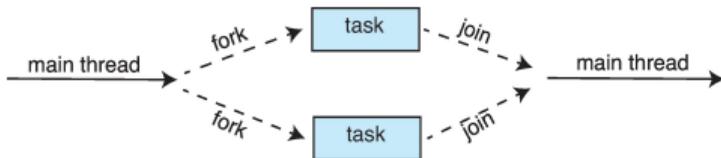


- > Usually applied recursively:

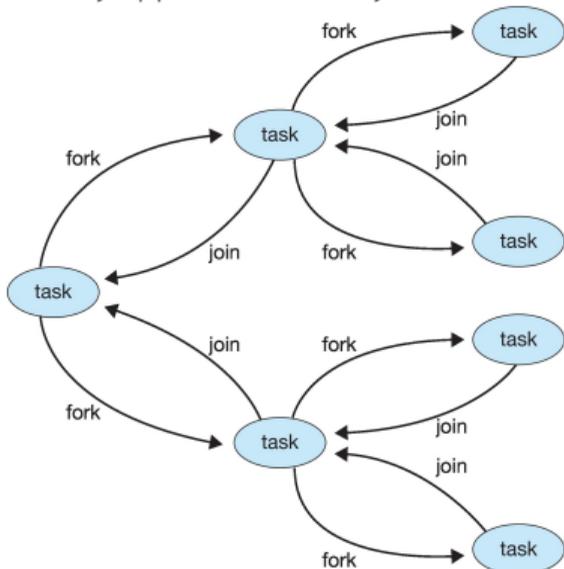


Fork-join model

- > When task is large, split into subtasks, do both in parallel (fork), then wait for results (join)



- > Usually applied recursively:



```
import java.util.concurrent.*;

public class SumTask extends RecursiveTask<Integer>
{
    static final int THRESHOLD = 1000;

    private int begin;
    private int end;
    private int[] array;

    public SumTask(int begin, int end, int[] array) {
        this.begin = begin;
        this.end = end;
        this.array = array;
    }

    protected Integer compute() {
        if (end - begin < THRESHOLD) {
            int sum = 0;
            for (int i = begin; i <= end; i++)
                sum += array[i];

            return sum;
        }
        else {
            int mid = (begin + end) / 2;

            SumTask leftTask = new SumTask(begin, mid, array);
            SumTask rightTask = new SumTask(mid + 1, end, array);

            leftTask.fork();
            rightTask.fork();

            return rightTask.join() + leftTask.join();
        }
    }
}
```

OpenMP

- > Library for parallelizing C/C++ programs
- > Parallelization of for-loops requires only minimal changes: adding `#pragma omp parallel for` before the loop
- > See demo/course resources

Other parallelization

Single-instruction-multiple-data (SIMD)

- > Modern processors come with special registers that can hold vectors, e.g. 4 integers (of 4 bytes each)
- > SIMD instructions can perform parallel operations on vectors.
- > They perform data-parallelism
- > **Note:** this is not threading or an operating system aspect

General-purpose graphics processors (GPGPU)

- > Graphic processors come with very many, sometimes 1000s, of (very simple) cores
- > These traditionally perform computer graphics tasks, but are being used increasingly for other purposes (e.g. linear algebra) as well
- > Accessed as a device via driver
- > Not easy to program