

Introduction to HPC

content and definitions

Jan Thorbecke, Section of Applied Geophysics

Motivation and Goal

- Get familiar with hardware building blocks, how they operate, and how to make use of them in writing efficient programs.
- Understand the function of an Operating System (Linux).
- Basic skills to write, compile and run numerical code efficiently on standard hardware.
- Using the memory hierarchy to write programs that use computational resources efficiently.
- Start to think how to solve problems in Parallel (try to avoid thinking in sequential loops).
- During the course keep in mind a numerical problem which you would like to implement/improve.

Learning Outcomes

- Understand the architecture of modern CPU's and how this architecture influences the way programs should be written.
- Optimize all aspects in the processes of programming: from compilation, starting and running program by OS, executing (parallel) instructions by CPU, to writing output to disk.
- Write numerical software, that exploits the memory hierarchy of a CPU, to obtain a code with close to optimal performance.
- Analyze an existing program for OpenMP and MPI parallelization possibilities.
- Evaluate the possibilities of accelerators to speed up computational work.

Organization

Dates: 14, 21, 28 April and 12, 19 May 2022

Time: 13:30 - 17:15

Points: 3

Examination: attendance, exercises and discussions during the course. For MSc students there is a multiple-choice exam with discussion afterwards.

Schedule and rooms

Civil Engineering & Geosciences, Building 23:
Stevinweg 1

BouwCampus, Building 26:
Van der Burghweg 1 + 2

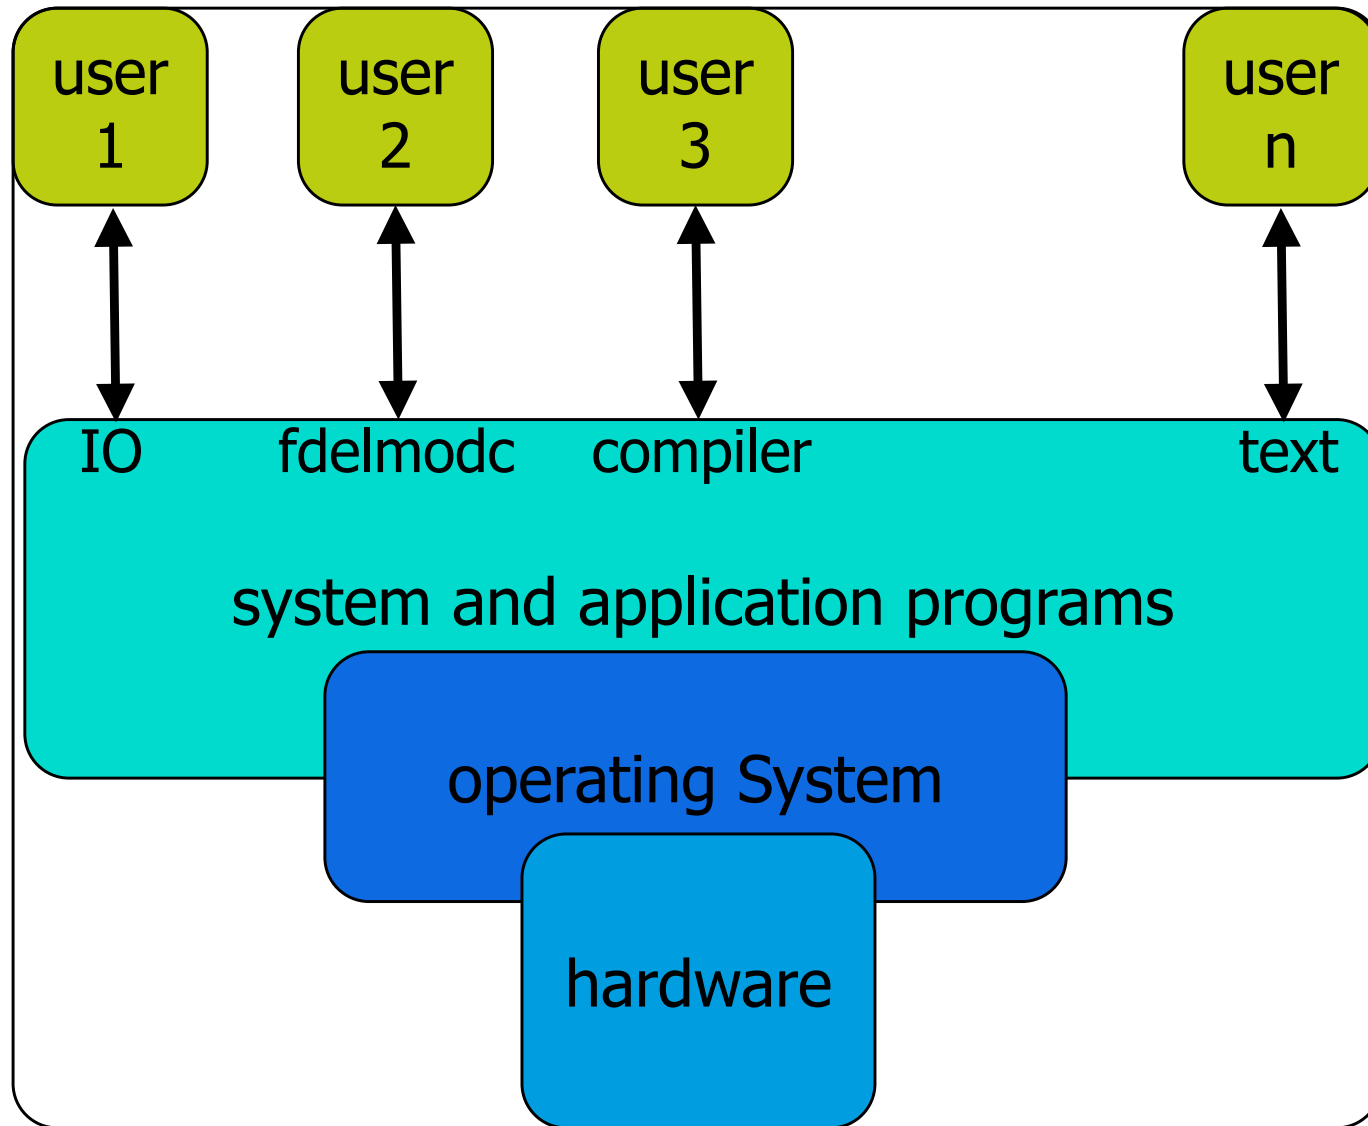
- 14 April room F (building 23 CITG)
- 21 April room 2 (building 26 BouwCampus)
- 28 April room 2 (building 26 BouwCampus))
- 12 May room 2 (building 26 BouwCampus))
- 19 May room 2 (building 26 BouwCampus))

see <https://esviewer.tudelft.nl/> for room locations.

What is HPC

- **High-performance computing (HPC)** uses [supercomputers](#) and [computer clusters](#) to solve advanced computation problems. Today, computer systems in the [teraflops](#)-region are counted as HPC-computers. (Wikipedia)
- A list of the most powerful high performance computers can be found on the [TOP500](#) list. The TOP500 list ranks the world's 500 fastest high performance computers as measured by the HPL benchmark.
- For this course:
Scientific programs that use computational hardware in an optimal way.

Four Components Computer System



Computer System Structure

- Computer system can be divided into four components
 - Hardware – provides basic computing resources
 - CPU, memory, I/O devices
 - Operating system
 - Controls and coordinates use of hardware among various applications and users
 - Application programs – to solve the computing problems of users
 - Word processors, compilers, web browsers, database systems, games
 - Users
 - People, machines, other computers

Contents

- Introduction (1)
- Hardware (1,2)
- OS (2)
- Programming (2/3)
- Optimization (3)
- Multi-processor hardware(3/4)
- Writing Parallel programs (4/5)
- Parallel programming + wrap-up and discussion (5)

Part 1

- Components of a computer system
 - addressing
 - memory hardware
- Memory Hierarchy
 - types of cache mappings
 - cache levels
 - virtual memory, TLB
- Processors
 - out of order
 - pipelining
 - branches
 - multi-core cache-coherency
 - vector instructions: SSE / AVX
- Modern processors
 - Intel, AMD, IBM, ARM, GPU
- Future trends in hardware

Part 2

- Operating System
- IO management
- Processes
- Virtual memory management
- File systems
 - storage, RAID
- Linux, Unix, OSX, XP

Part 3

- User environment
 - login, shells
- Compiling and linking
 - compile stages
 - Makefiles
- Number representations
 - IEEE 754 standard
- Programming languages
 - C, Fortran
 - profiling and debugging
 - numerical libraries

Part 4

- Computer Performance
 - Amdahl's law
 - standard benchmarks Spec, Linpack
- Optimization
 - cache blocking
 - loops tricks
 - many examples

Part 5

- Programming for parallel systems
- Multi-processor hardware
 - Classification
 - HPC hardware
- Programming for parallel systems
 - OpenMP
 - MPI
 - examples
- Finish the last parts of the course
- Followed by questions, discussion and evaluation of exercises

Exercises

- Exercises are set up to get familiar with concepts and ideas. If you already have enough experience with some parts in the exercise you can skip those parts.
- Exercises are simple and straightforward, no programming skills are needed, and consists of compilation and running of small programs. Every exercise comes with a step by step procedure to follow.
- The exercises are hopefully fun to do and are not obligatory to pass the course, but **highly** recommended.

Getting the Exercises

- git clone <https://github.com/JanThorbecke/HPCourse.git>
- Shall we try that out now?

The slides

- There are quite a few of them...
- Not all of them will be discussed
- Sometimes techniques are explained in two different ways, one will be used in the lectures, the other one can be used to understand the concepts better.
- At the moment there are no lecture notes and the detailed slides try to compensate that.

Definitions and Terms

- bit = binary number with states 0 or 1
 - byte = 8 bits
 - word = 4 or 8 bytes (the number of bits used in a register)
 - Kbit = 1000 bits, used for communication speeds
 - KiB = 1024 (2^{10}) Bytes, used for memory and storage
 - MiB = 1048576 (2^{20}) Bytes
 - flops = floating point operations per second
 - CPU = Central Processing Unit
 - FPU = Floating Point Unit
 - ALU = Arithmetic Logic Unit
 - AGU = Address Generation Unit
-
- **<http://physics.nist.gov/cuu/Units/binary.html>**

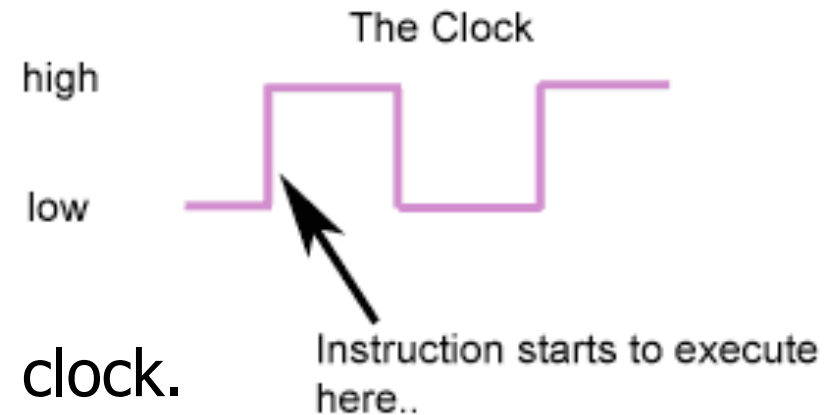
Definitions

Latency: Latency is a time delay between the moment something is initiated, and the moment one of its effects begins or becomes detectable. The word derives from the fact that during the period of latency the effects of an action are latent, meaning "potential" or "not yet observed".

Bandwidth: a data transmission rate; the maximum amount of information (bits/second) that can be transmitted along a channel.

flops: number of floating point operations (addition multiplication) per second.

Clock



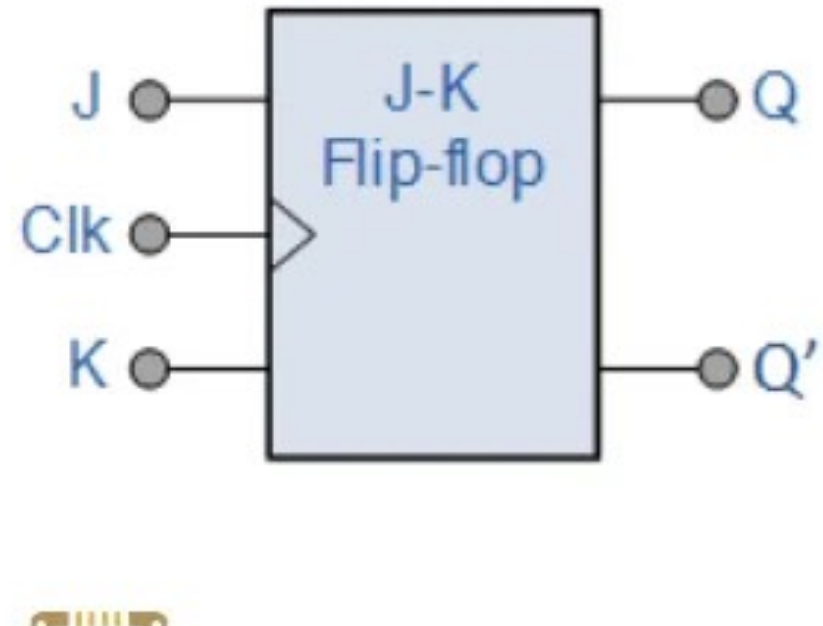
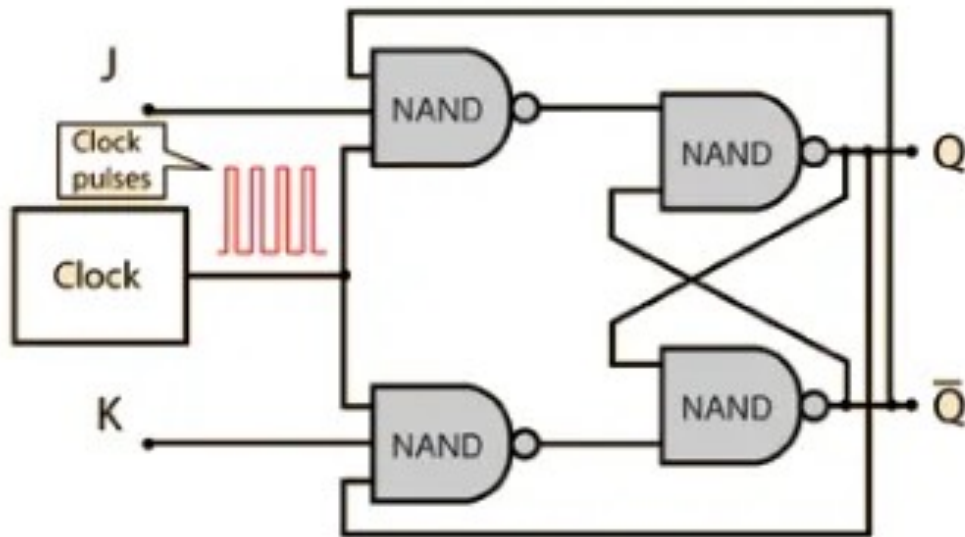
Most modern CPUs are driven by a clock.

The CPU consists internally of logic and memory (flip flops).

When the clock signal arrives, the flip flops take their new value and the logic then requires a period of time to decode the new values.

Then the next clock pulse arrives and the flip flops again take their new values, and so on.

Clock : JK flip-flop



Feedback

Feedback about the course is highly appreciated.

You can always send me an e-mail or come to my desk in room 3.03 (CITG Building 23 on Thursdays only).

Book References

Computer Architecture , A Quantitative Approach

Sixth Edition, John L. Hennessy, David A. Patterson, 2017, Morgan Kaufmann

The C Programming Language (2nd edition)

Kernighan & Ritchie, 1988, Prentice Hall

Modern Fortran Explained

Michael Metcalf, John Reid , Malcolm Cohen, 2018, Oxford University Press

Using OpenMP: Portable Shared Memory Parallel Programming

Barbara Chapman; Gabriele Jost; Ruud van der Pas, 2008, MIT Press

Parallel Programming With MPI

Peter Pacheco, 1996, Morgan Kaufmann

References on internet

<https://www.hpc-carpentry.org>

<http://www.tomshardware.co.uk/>

<http://insidehpc.com/>

<http://www.gpgpu.org/>

<wikipedia.org>

TU-Delft DCSE

Delft Institute for Computational Science and Engineering

- Courses at TU-Delft from DCSE

<https://www.tudelft.nl/cse/education/courses>

- Introduction MPI
- linux bash
- GPU CUDA programming
- OpenFOAM

Linux on Windows

- WSL (Windows 10)
<https://docs.microsoft.com/en-us/windows/wsl/install>
Windows Subsystem for Linux lets developers run a GNU/Linux environment without a Virtual Machine.
- Cygwin (older Windows)
<https://www.cygwin.com/>
Windows Subsystem for Linux lets developers run a GNU/Linux environment
- putty
<https://www.putty.org/>
ssh client on Windows to login to linux server.