



INTRODUCTION TO HIGH-PERFORMANCE COMPUTING (HPC)

April 26th 2021 - Guest lecture for Information Technology and Data Processing course

Kenneth Hoste, HPC-UGent - kenneth.hoste@ugent.be

<https://www.ugent.be/hpc> - hpc@ugent.be





whoami

kenneth.hoste@ugent.be

@boege1 (*GitHub, IRC, Slack*)

@kehoste (*Twitter*)

- Masters (2005) & PhD (2010) in Computer Science from Ghent University
- PhD topic: machine learning applied to software performance, compilers, ...
- **joined HPC-UGent team in October 2010**
- **main tasks: user support & training, software installations**
- likes family, beer, loud music, FOSS, helping people, dad jokes, stickers, ...
- doesn't like CMake, SCons, Bazel, TensorFlow, OpenFOAM, ...

LIVE POLL

- A couple of poll questions will be raised during the presentation
- Please join in (anonymously), no login required
- Use your smartphone or browser
- Visit ahaslides.com/HPCUGENT, or scan this QR code:



OUTLINE

- What are supercomputers? What are they used for?
- Important concepts in HPC: multi-core, MPI, Amdahl's law, ...
- Trends in HPC: GPUs, AI, cloud computing, Intel vs AMD vs Arm vs ...
- HPC in the future: quantum computing



TERMINOLOGY

- Supercomputer, nodes, processors, cores, interconnect, FLOPS
- High-Performance Computing (HPC), scientific computing
- Parallel computing, shared memory vs distributed memory computing
- OpenMP, threads, MPI, processes, CPUs, GPUs, ...
- Accelerated computing, cloud computing
- Quantum computer, qubits



POLL

How familiar are you with supercomputing?

- I have no idea what it is, or why I should care.
- I think I know what a supercomputer is...
- I have used a supercomputer already.
- I frequently use a supercomputer.
- I am an HPC system administrator
- I have my own supercomputer.





WHAT IS A SUPERCOMPUTER?

SUPERCOMPUTER

- A bunch of "normal" computers (a.k.a. servers, (worker)nodes, ...)
- **Fast local network** to link them together ("interconnect")
- Large amount of **shared storage**, various types (tape, disks, SSDs, ...)
- Can be considered one big system (in some sense)
- Used for **resource-intensive workloads** (compute, memory, storage, ...)
- Sometimes with special-purpose accelerator hardware (GPUs, FPGAs)
- Usually in a dedicated datacenter

EXAMPLE: STAMPEDE 2 (TACC, 2017)



> 350,000 cores - 100Gb/sec Omni-Path interconnect - 31PB of storage

<https://www.tacc.utexas.edu/systems/stampede2>

PERFORMANCE IN FLOPS

- Performance of supercomputers can be measured in FLOPS:
(64-bit) floating-point operations per second
- Example FLOP: $1.34 \times 54.97 = 73.6598$
- Used to create Top 500 list of supercomputers since 1993
- 1 megaFLOPS (MFLOPS): 1 million FLOPS (10^6)
- x1000 each generation (gigaFLOPS, teraFLOPS, petaFLOPS, exaFLOPS)
- Current generation of supercomputers: petaFLOPS (10^{15})



POLL

Where is the current fastest supercomputer in the Top 500 located?

- China
- Europe
- Russia
- United States
- Somewhere else





CDC 6600 (1964)

"first supercomputer"

3 MFLOPS, up to 982 kilobytes of memory



Cray-2 (1985)

1.9 GFLOPS (4 vector processors)

256 MWords of memory



ASCI Red, US (1997)

1.6 TFLOPS (first over 1 TFLOPS)

9,298 processors (76 nodes)

1,212 GB of RAM memory



Jaguar, US (2009)

1.75 PFLOPS

18,688 AMD Opteron nodes (224,256 cores)

360TB of memory, 10PB of storage



Titan, US (2013)

17.58 PFLOPS - 18,688 nodes (~300,000 cores)

each node: 16-core AMD Opteron CPU + 1 NVIDIA Tesla GPU

~700TB of memory (CPU+GPU), 40PB of storage



Piz Daint, Lugano - Switzerland (2016)

~25 PFLOPS

~7,500 Intel Xeon nodes, ~133,000 cores + 5,704 NVIDIA P100 GPUs

8.8 PB shared storage, Cray Aries interconnect

#3 in Top500 (June 2017)



MareNostrum 4, Barcelona (2017)

11.15 PFLOPS , 1.3 MW power

3,456 Intel Xeon nodes (165,888 cores)

14 PB shared storage, Omnipath interconnect

#13 in Top500 (June 2017)



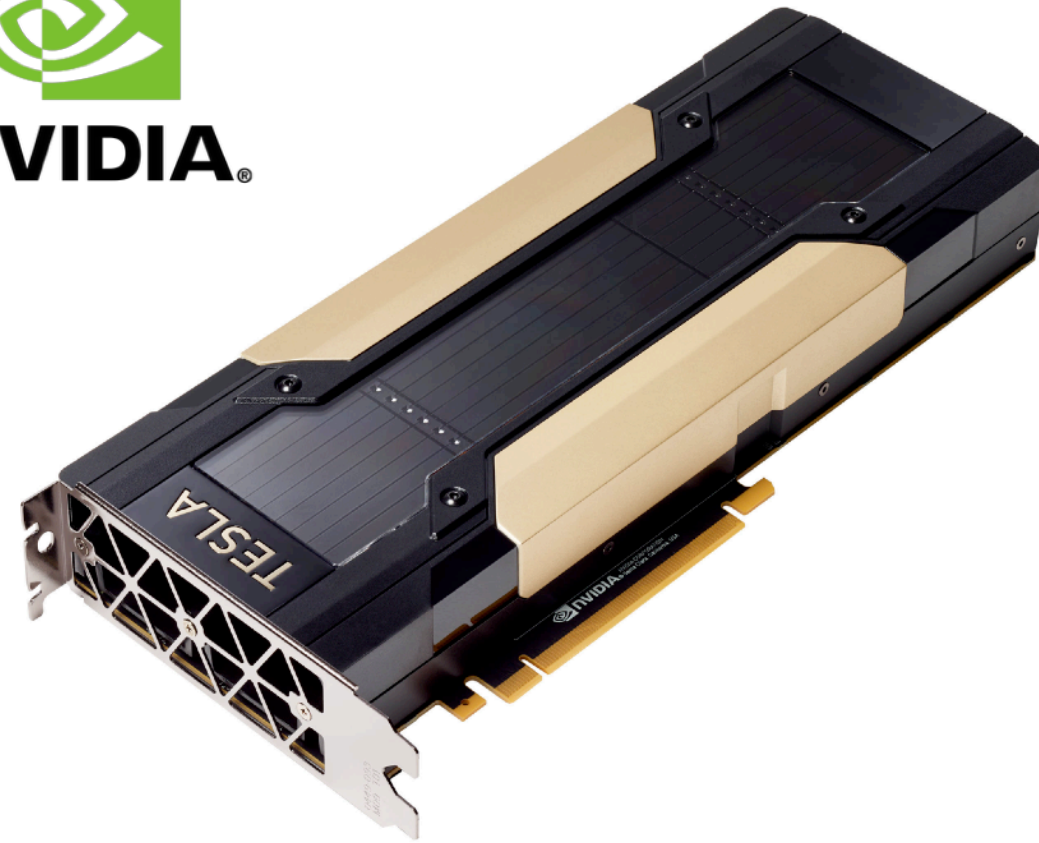
Fugaku, Japan (2020)

Current #1 in Top 500 of supercomputers, since June 2020

158,976 nodes (7.6 million cores), 442 PFLOPS

Fujitsu A64FX CPUs (Arm 64-bit)

Torus fusion interconnect, 150 PB of shared storage



NVIDIA V100 GPU

~\$10,000

7.8 TFLOPS



Equivalent with #1 supercomputer in 2000!



iPhone X (2017)

~500 EUR

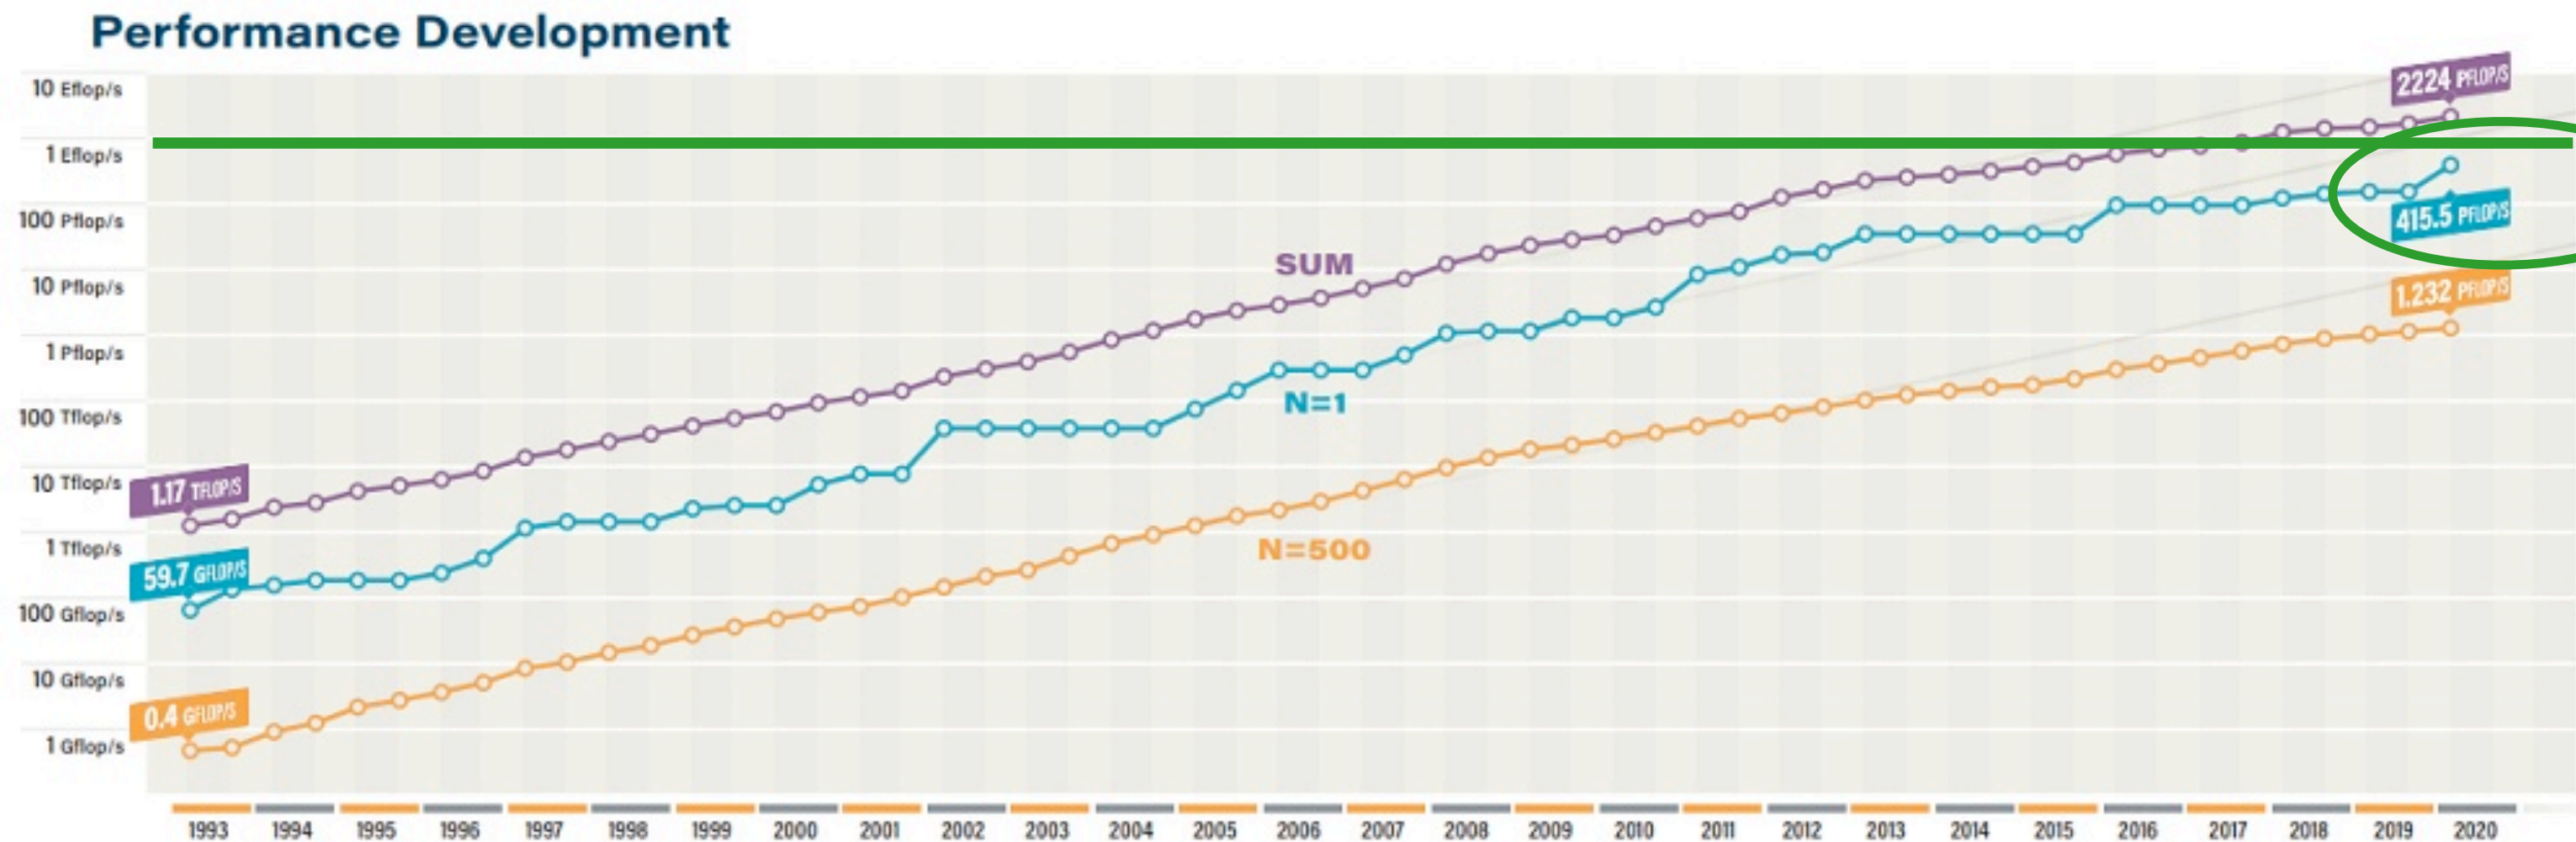
Includes NPU (Neural Processing Unit) of ~600 GFLOPS

Equivalent with #1 supercomputer in 1996!



SOON: EXASCALE COMPUTING

- Top supercomputer will soon be > 1 exaFLOPS (10^{19})
- US, China, Europe, Japan, ... ?



POLL



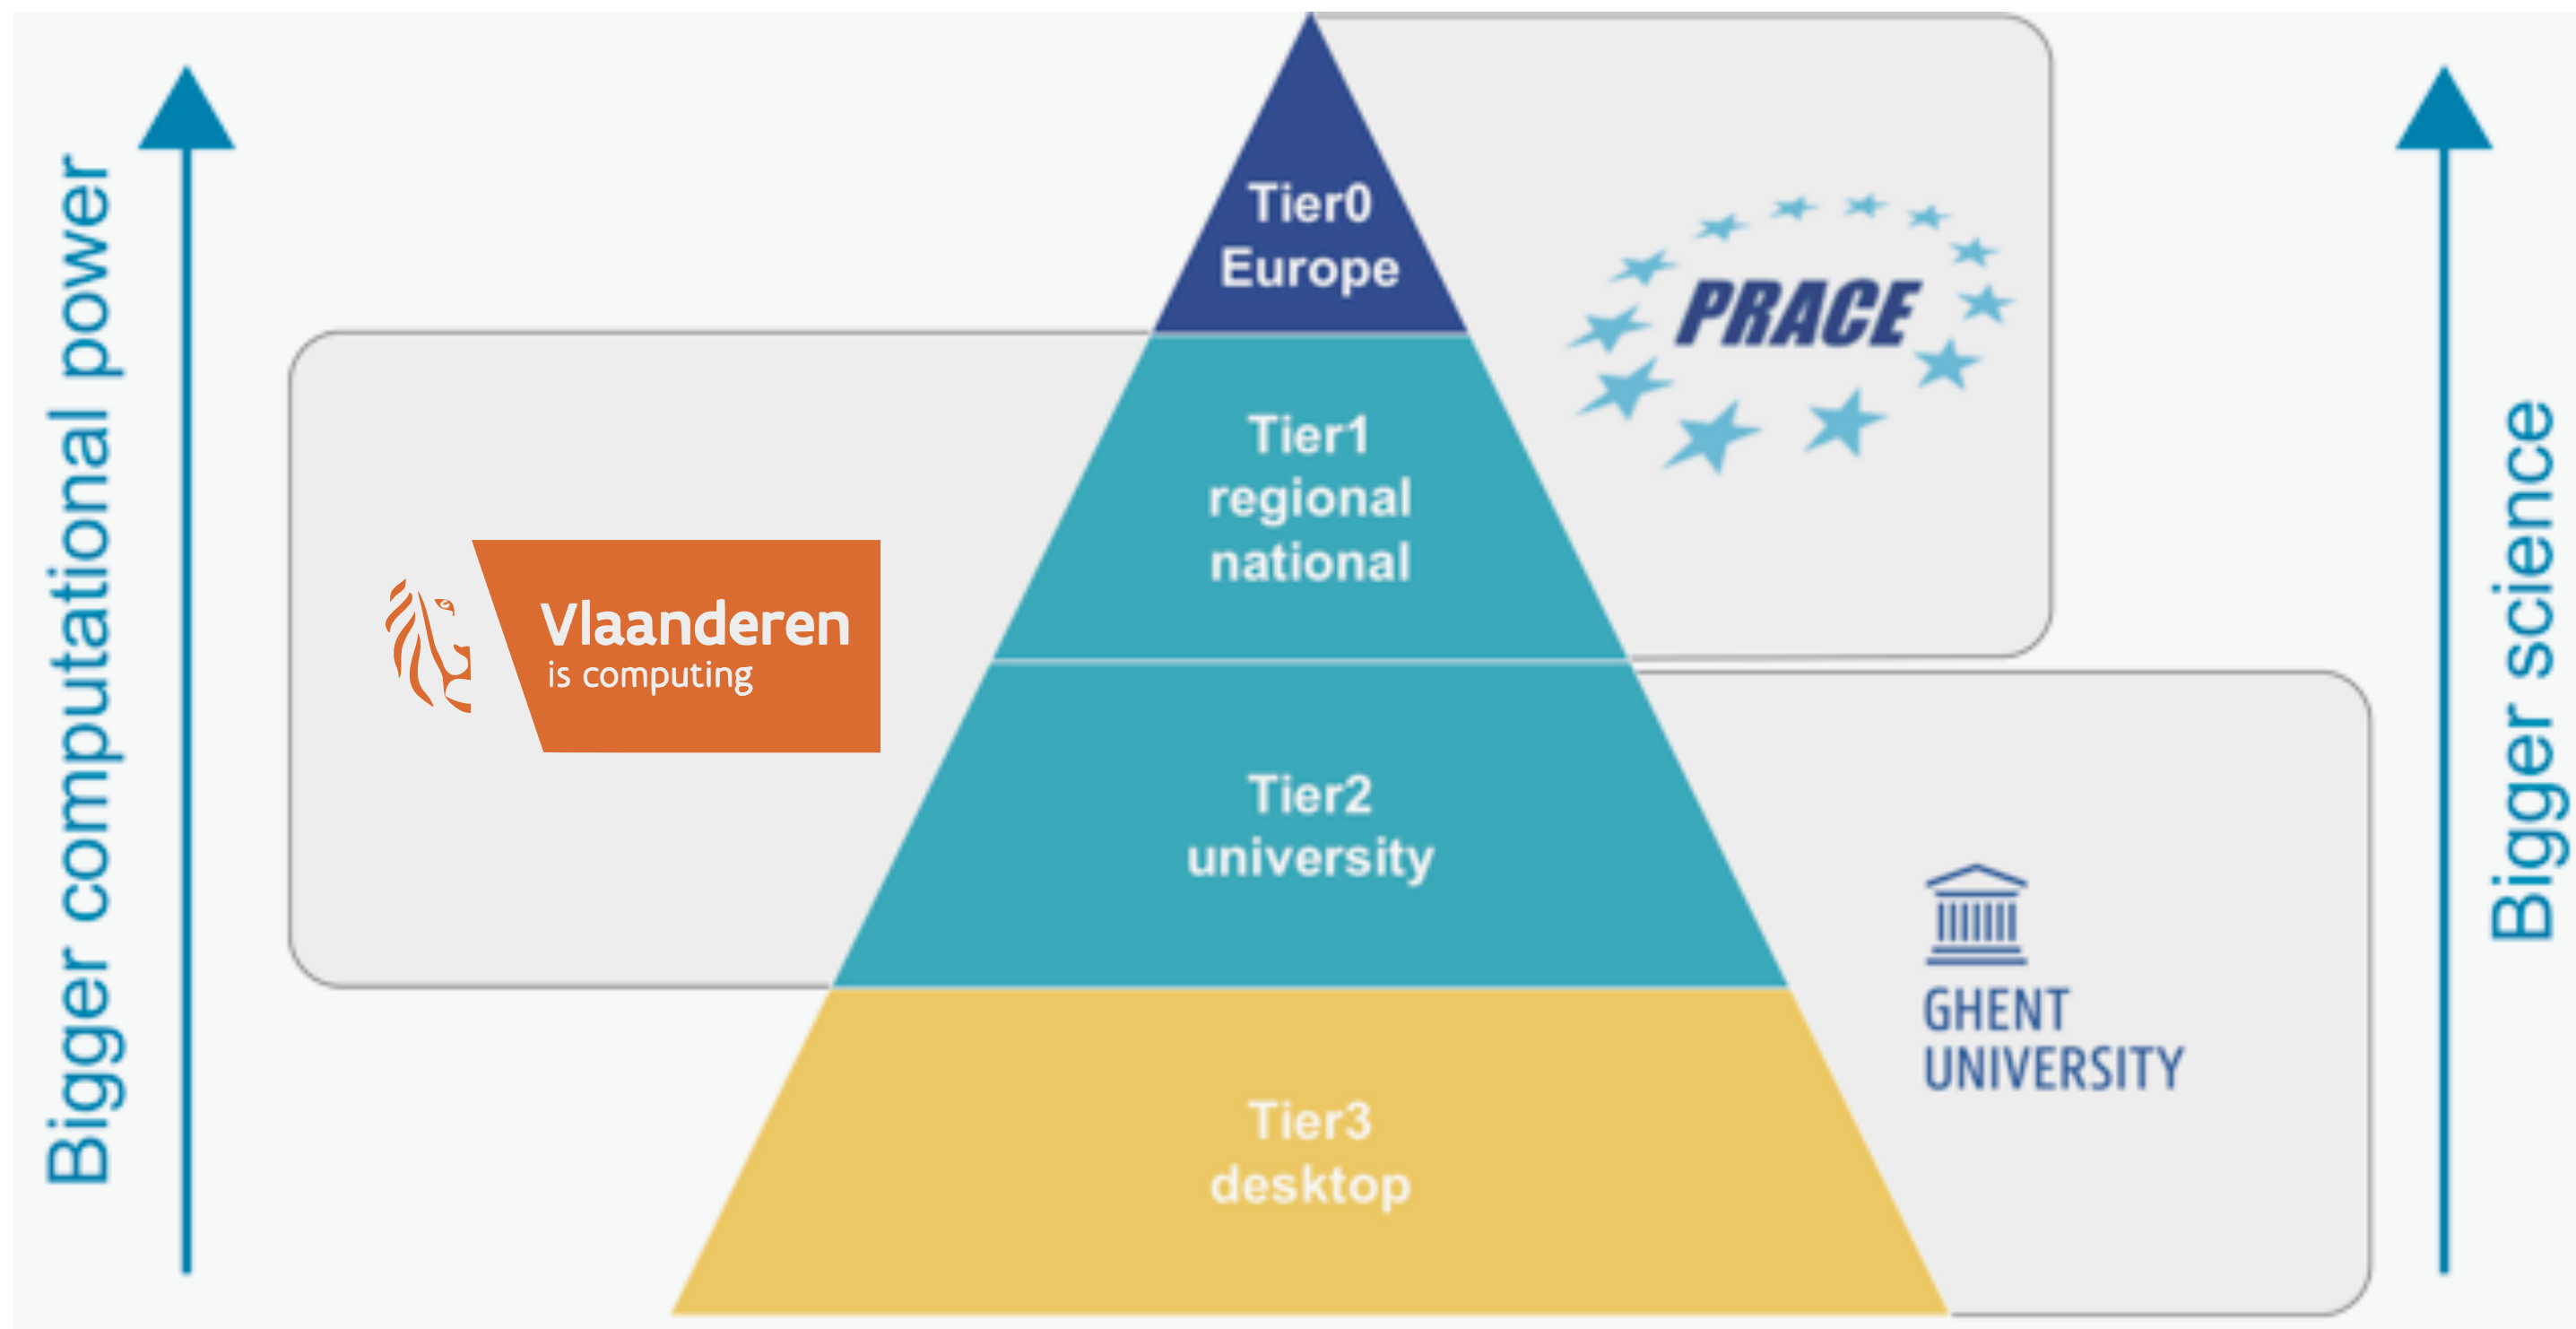
Vlaanderen
is computing

Where is the fastest supercomputer in Flanders?

- There is none
- Antwerp
- Bruges
- Brussels
- Ghent
- Leuven
- Somewhere else



EUROPEAN TIERS OF SCIENTIFIC COMPUTING





UGent datacenter
S10 (campus Sterre)



HPC-UGent (since 2008)

Current infrastructure (Tier-2)

- 5 CPU clusters (4x Intel, 1x AMD)
- ~22,000 cores in total
- 1 cluster with 40 NVIDIA V100 GPUs
- Most clusters with Infiniband interconnect
- 4PB of shared storage in total



FLEMISH SUPERCOMPUTER CENTRE (VSC)

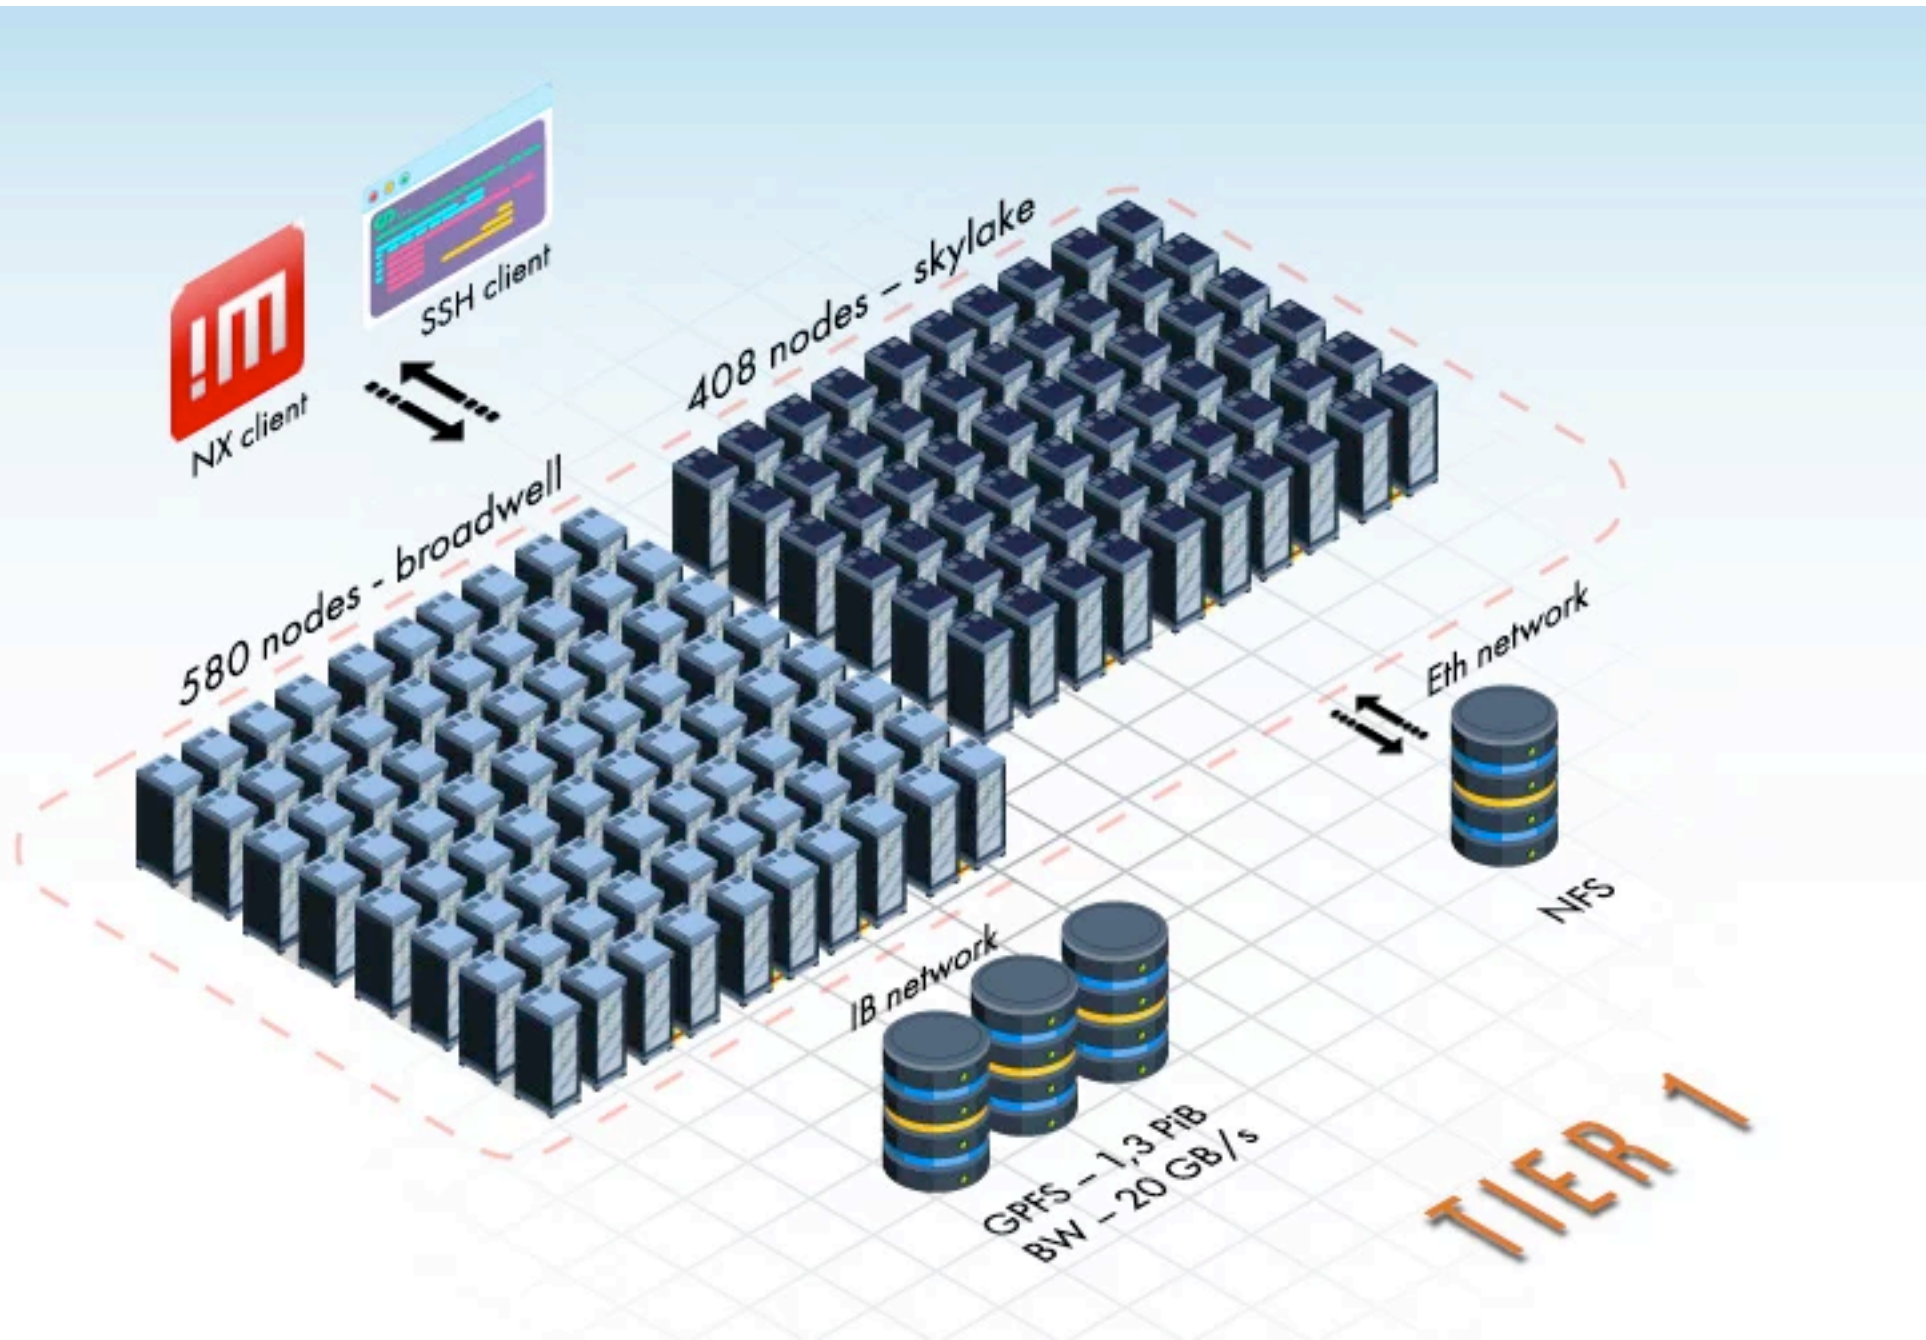
- Partnership between different Flemish universities
- Virtual centre with hubs in Ghent, Leuven, Brussels, Antwerp
- Shared infrastructure: account management, cloud, Tier-1, ...
- Funded by Flemish government + universities
- More info: <https://www.vscentrum.be>



Vlaanderen
is computing



CURRENT VSC TIER-1: BRENIAC (@ KUL)



Vlaanderen
is computing



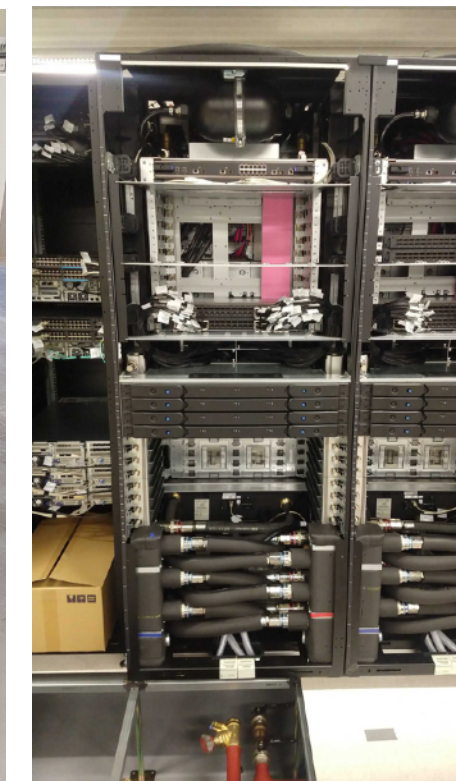
NEXT VSC TIER-1: HORTENSE (@ UGENT)

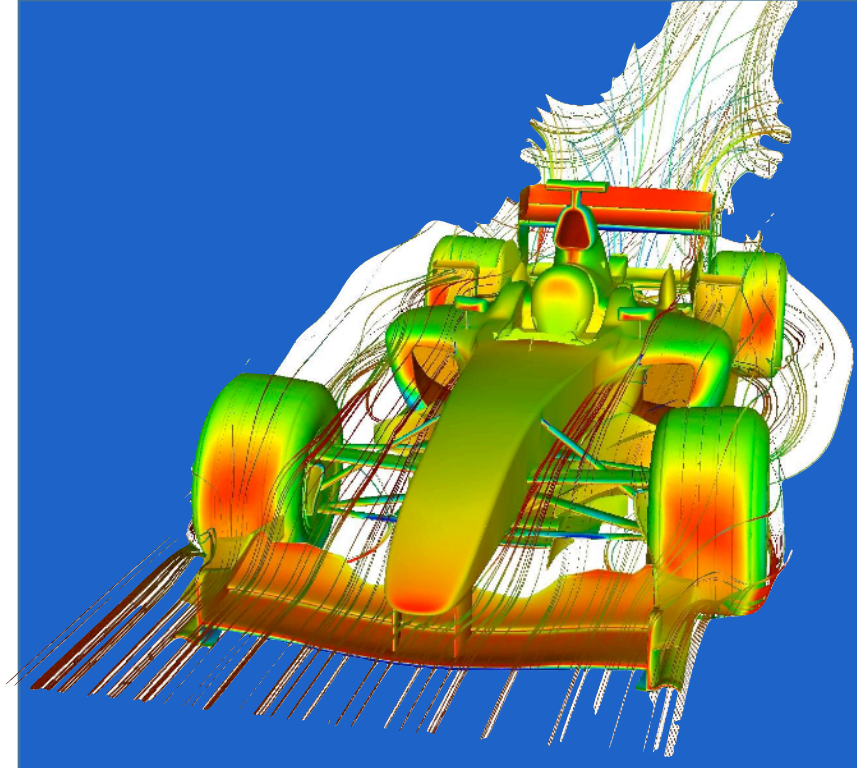
<https://www.ugent.be/hpc/en/news-events/news/hortense>



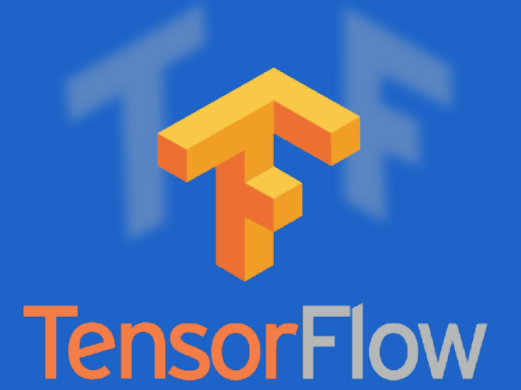
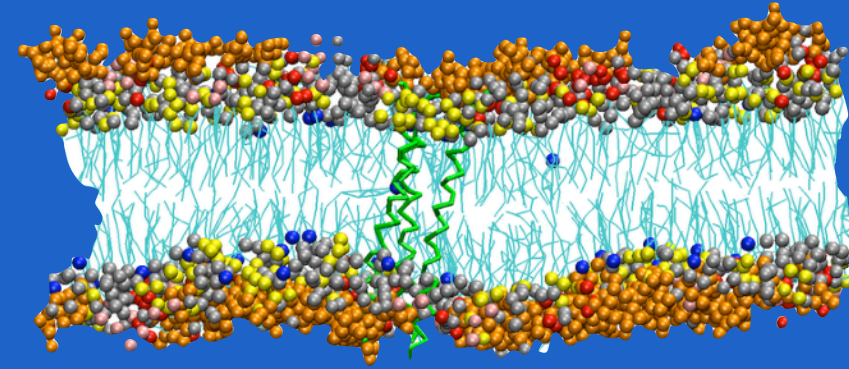
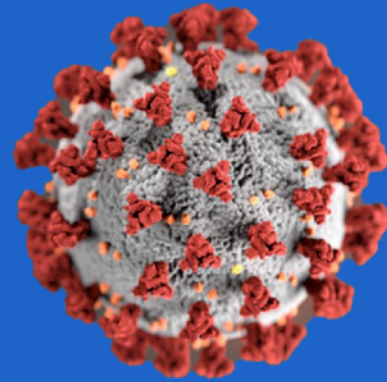
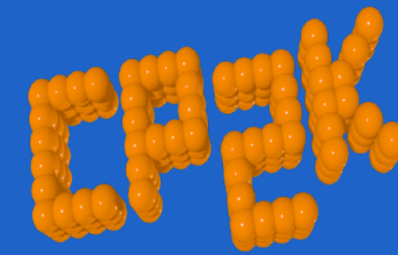
Vlaanderen
is computing

- 356 nodes (AMD EPYC 7H12), water-cooled
- 128 cores per node (~44,000 cores in total)
- 80 NVIDIA A100 GPUs (4 per node)
- 3PB of dedicated storage (Lustre)
- InfiniBand HDR-100 interconnect
- ~3.3 PFLOPS (likely Top 500)
- Should be fully operational in 2020...





Open  FOAM



 PyTorch

SUPERCOMPUTING

SOFTWARE & APPLICATIONS



HIGH-PERFORMANCE COMPUTING (HPC)

- Running **large scale** calculations (on a supercomputer)
- Example: running a physics simulation on 100,000 cores
- Typically heavily relying on fast interconnect, shared storage, ...
- Related:
 - High-Throughput Computing (HTC):
running **a lot** of (small & short) calculations
 - Scientific computing: simulations, data analysis, machine learning, ...

EVOLUTION IN SCIENTIFIC COMPUTING: BIGGER, FASTER



Margaret Hamilton (1969)

Standing next to source code she developed for the Apollo moon mission

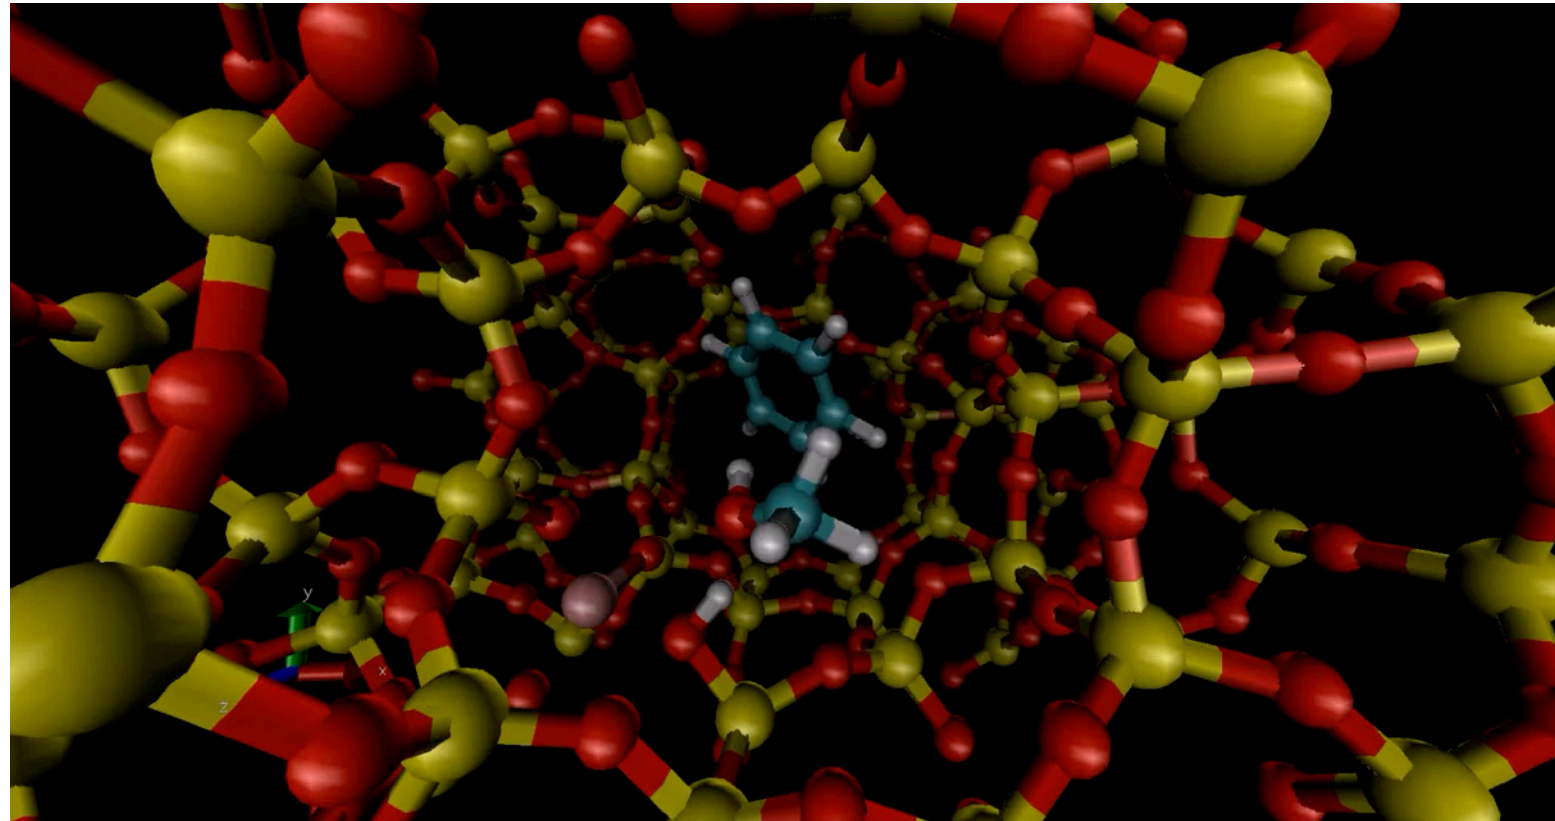


Katie Bouman (2019)

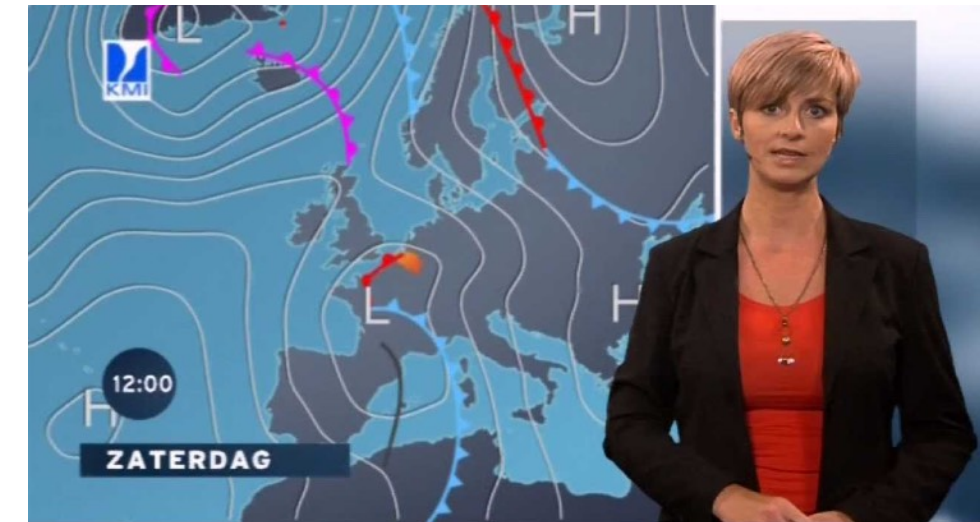
Showing hard drives that store the data collected to make the first image of a black hole (5PBs = 5000 TBs)



APPLICATIONS ACROSS ALL SCIENTIFIC DOMAINS



Material research ([CMM](#) @ UGent)



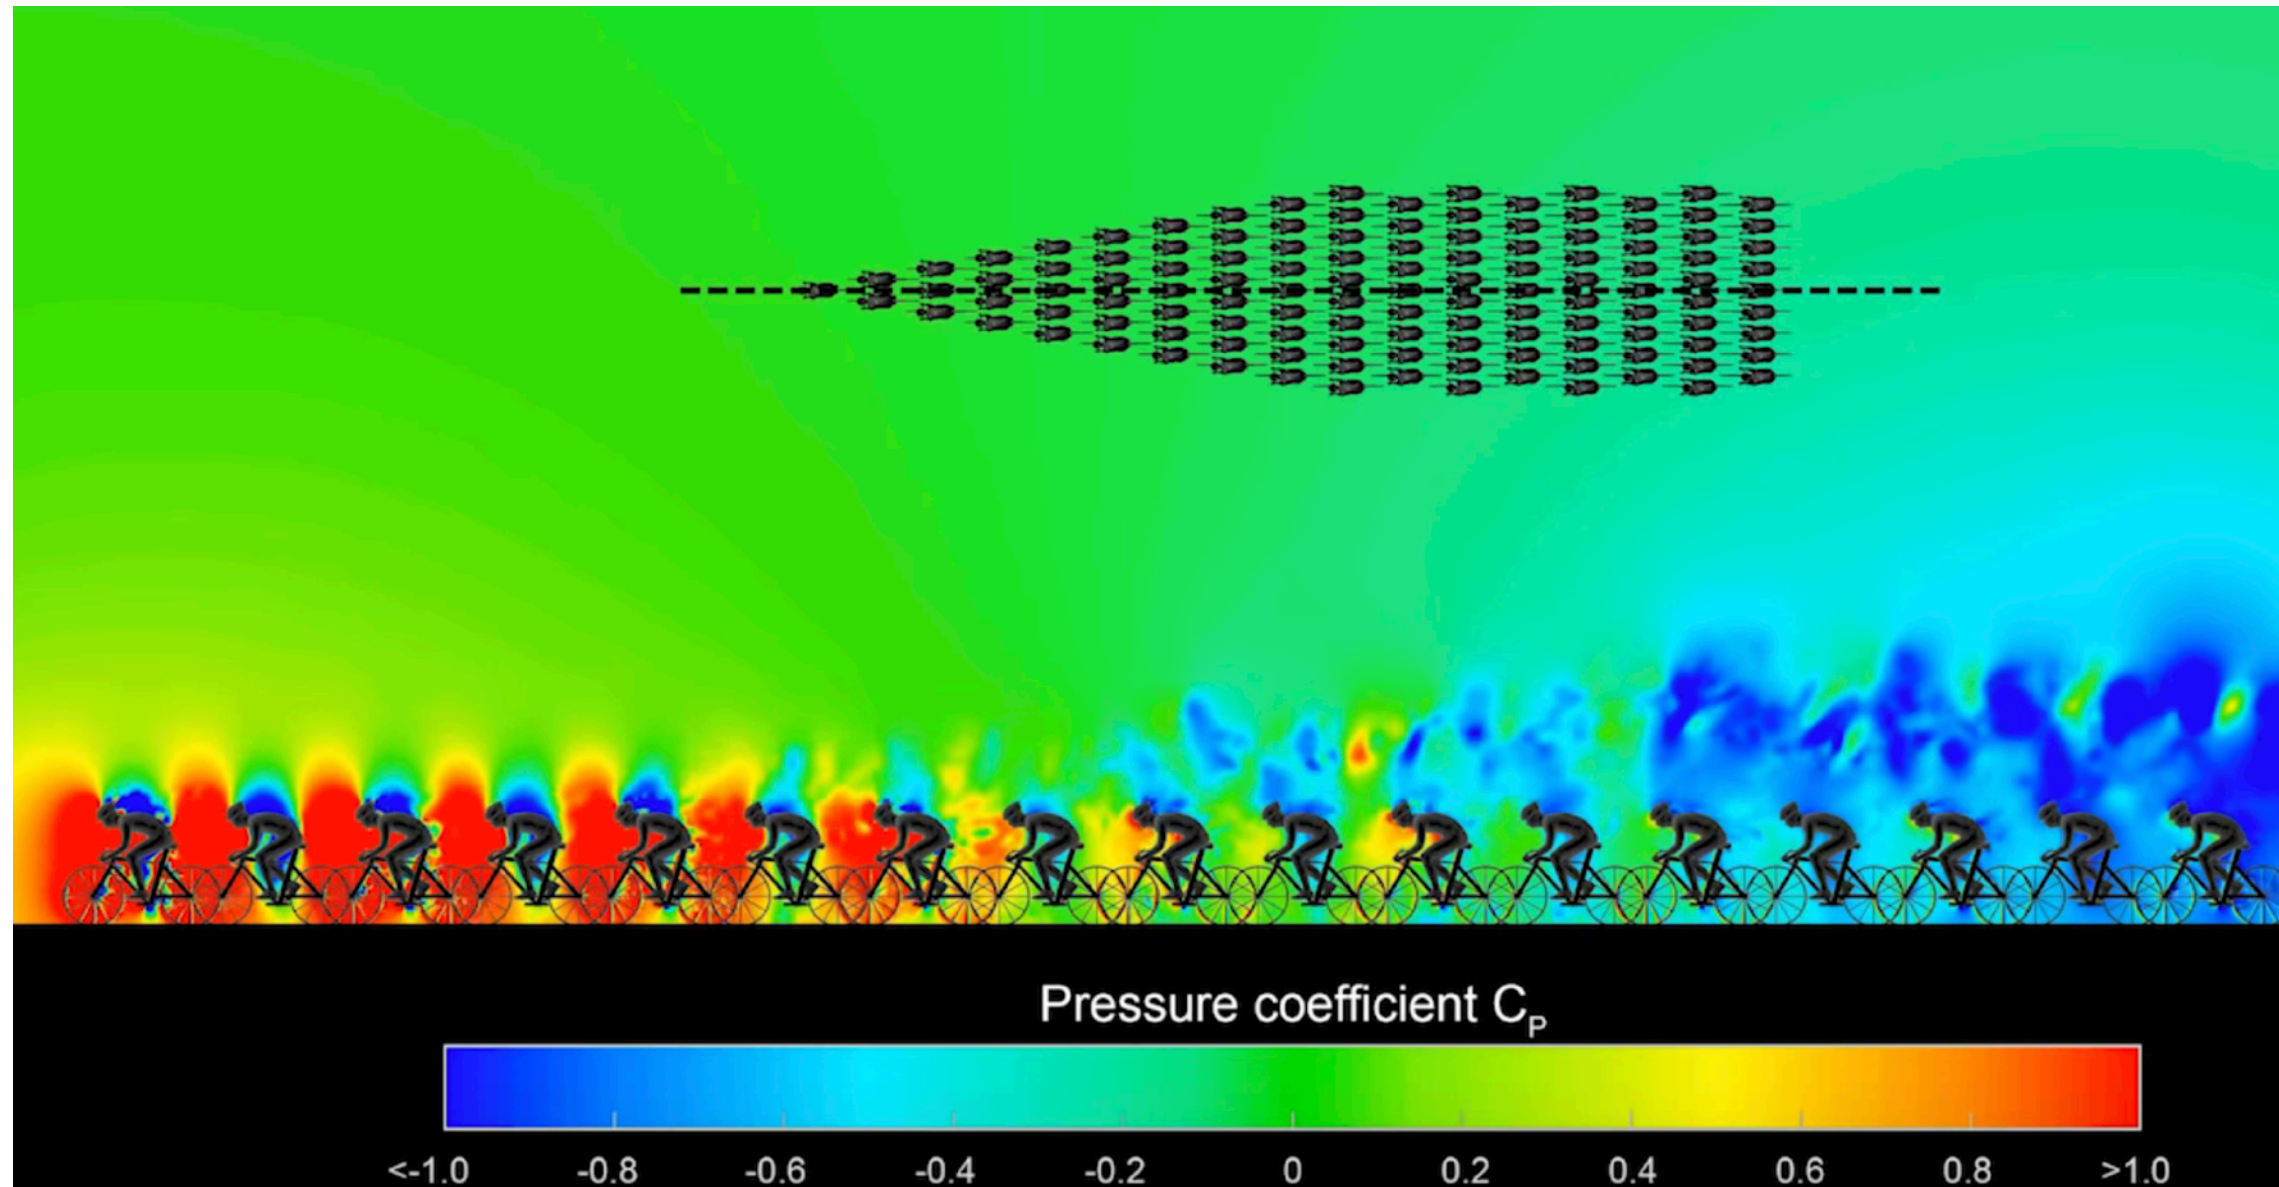
Weather predictions & climate simulations ([RMI](#))



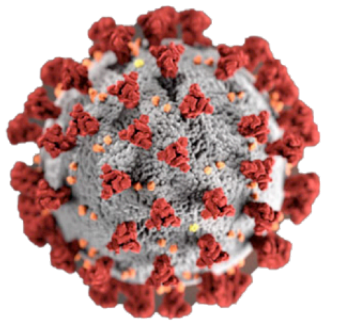
DNA analysis, for example
in cancer research ([CRIG](#) @ UGent)

COMPUTATIONAL FLUID DYNAMICS

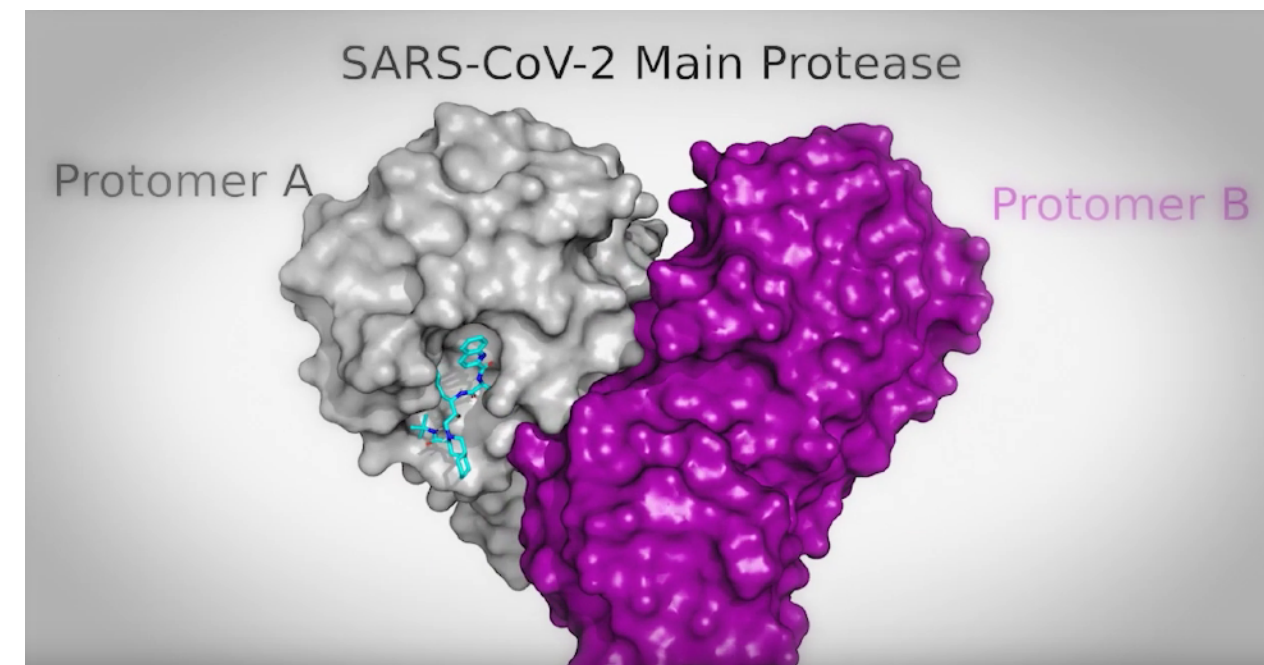
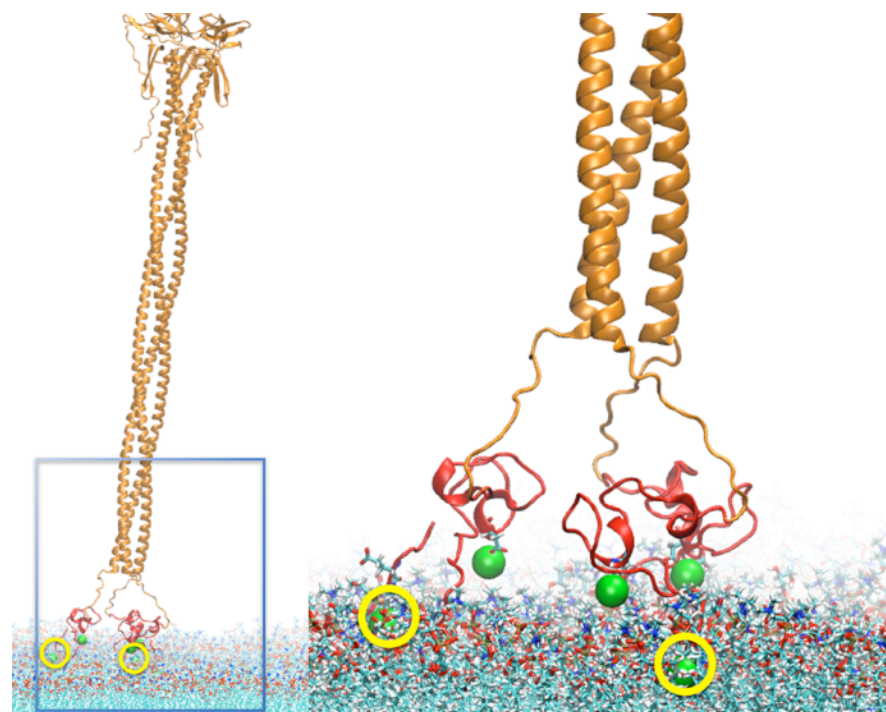
- Example: determine reduce of drag on riders well embedded in peloton (only ~5-10% compared to drag of isolated rider!)
- "Largest sports simulation to date" (July 2018)



FIGHTING COVID-19 WITH HPC



- **COVID-19 High Performance Computing Consortium** (industry + academia)
- Free access for researchers to some of the largest supercomputers in the world
- ~600 PFLOPS of available compute resources in total
- CFD simulations of airflow in rooms & airplanes to evaluate distribution of particles
- Analysis of virus RNA to trace back how virus was spread, identify variants, ...
- No doubt had positive impact on time-to-market for effective vaccines, drugs, etc.



POP QUIZ

Most popular programming languages in scientific computing?

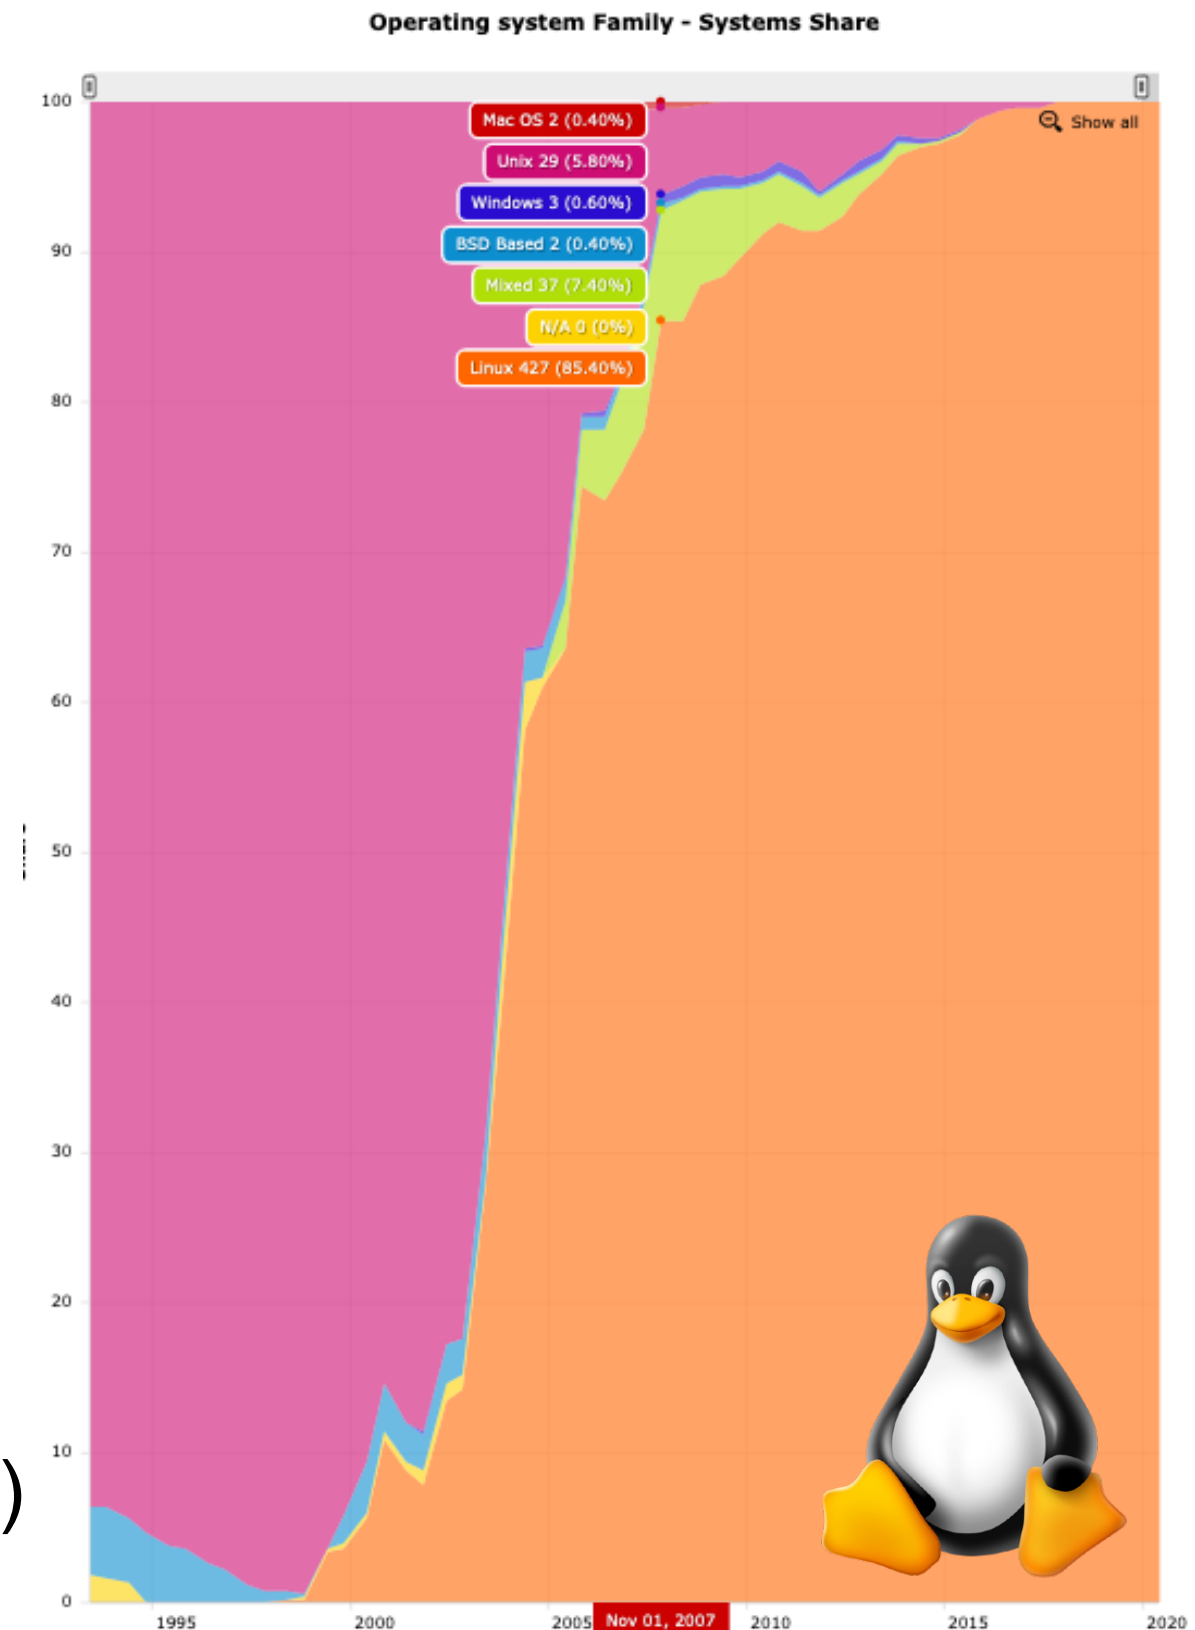
- Assembly
- C
- C++
- C#
- COBOL
- Delphi
- Fortran
- Go
- Java
- Javascript
- MATLAB
- Perl
- PHP
- Python
- R
- SQL
- Swift
- Ruby
- Visual Basic
- ...



SYSTEM SOFTWARE

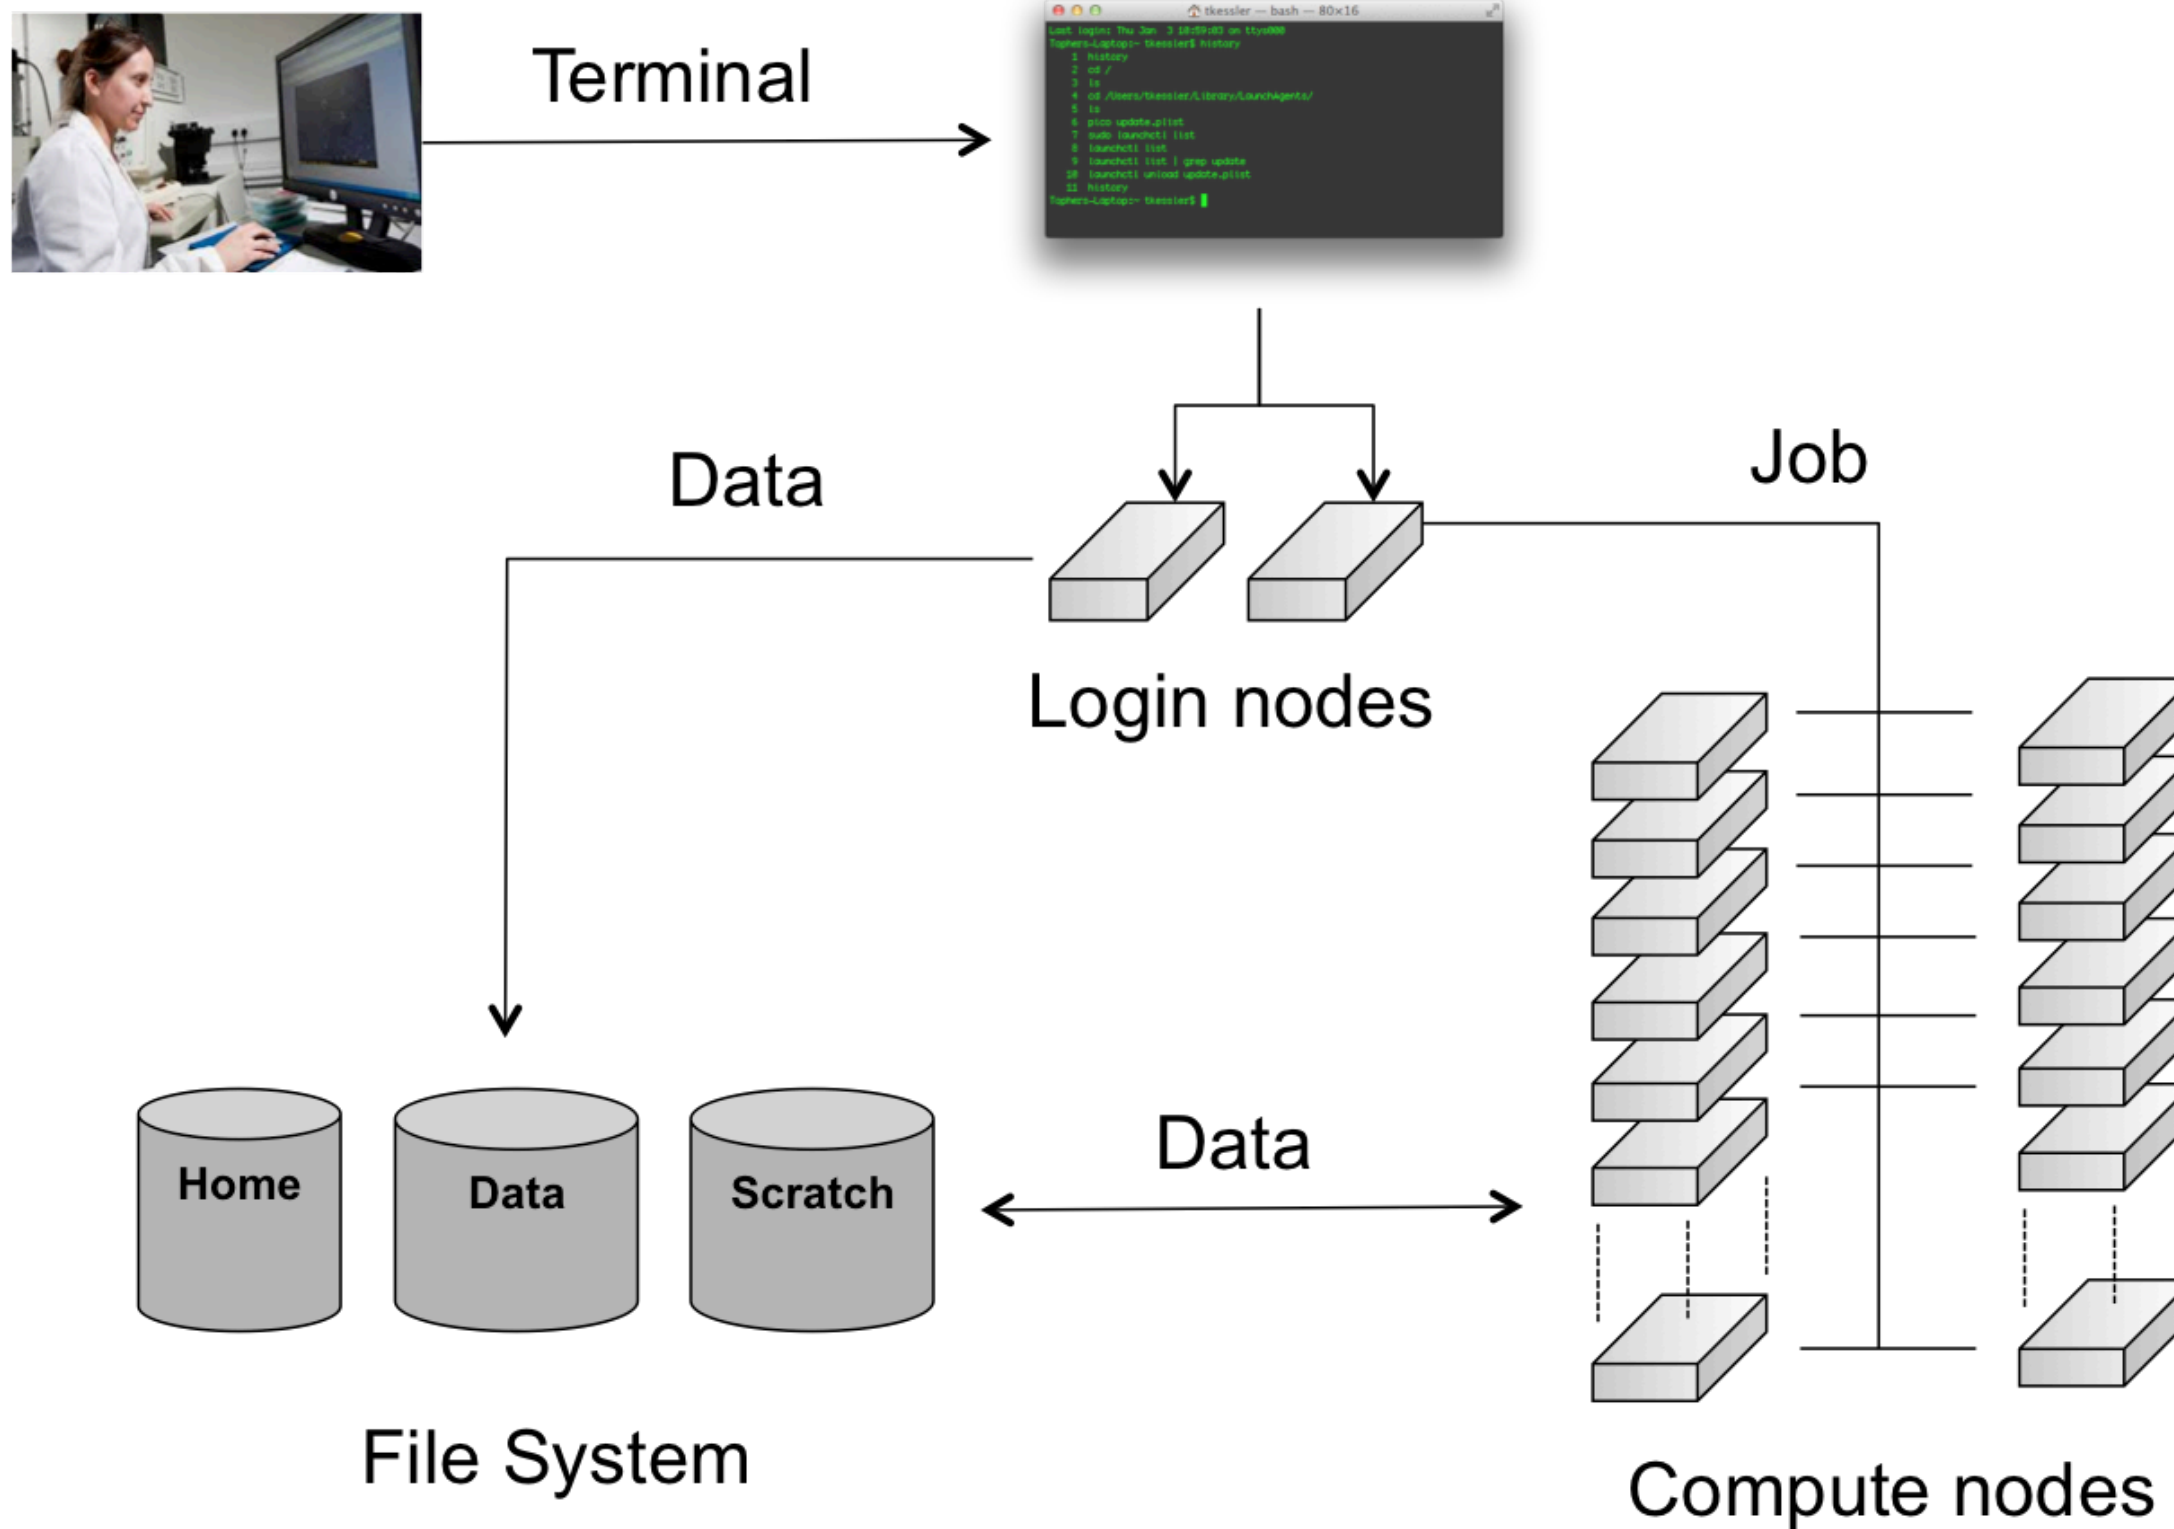
<https://top500.org/statistics/overtime>

- 100% of supercomputers in Top 500 list runs Linux!
- Mix of commercial and open source software on top
- Specific tools & services
 - Shared filesystems (Spectrum Scale, Lustre, ..)
 - Job scheduling (Slurm, ...)
 - Environment modules (Lmod)
 - Testing (ReFrame)
 - Custom scripting (Python, ...)
 - Monitoring of systems and services (nagios, Kibana, ...)



GETTING ACCESS

<https://www.ugent.be/hpc/en/access/policy/access>

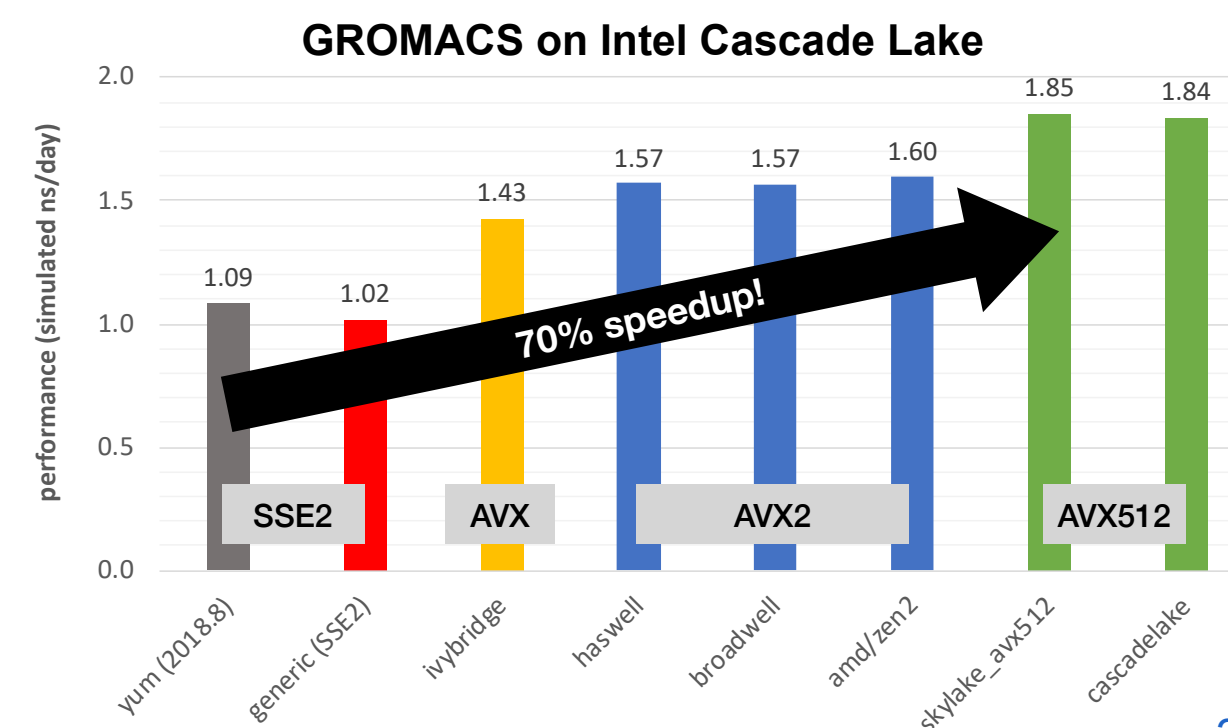


GETTING SCIENTIFIC SOFTWARE INSTALLED

- Scientific software can be difficult to install (scientists are often not trained software developers...)
- Ideally built from source code (results in better performance)
- Used to be a very time-consuming, tedious and manual task...
- EasyBuild was created by the HPC-UGent team to **automate** this
- Open source software: <https://github.com/easybuilders>
- Implemented in Python
- Now used by hundreds of HPC sites worldwide!



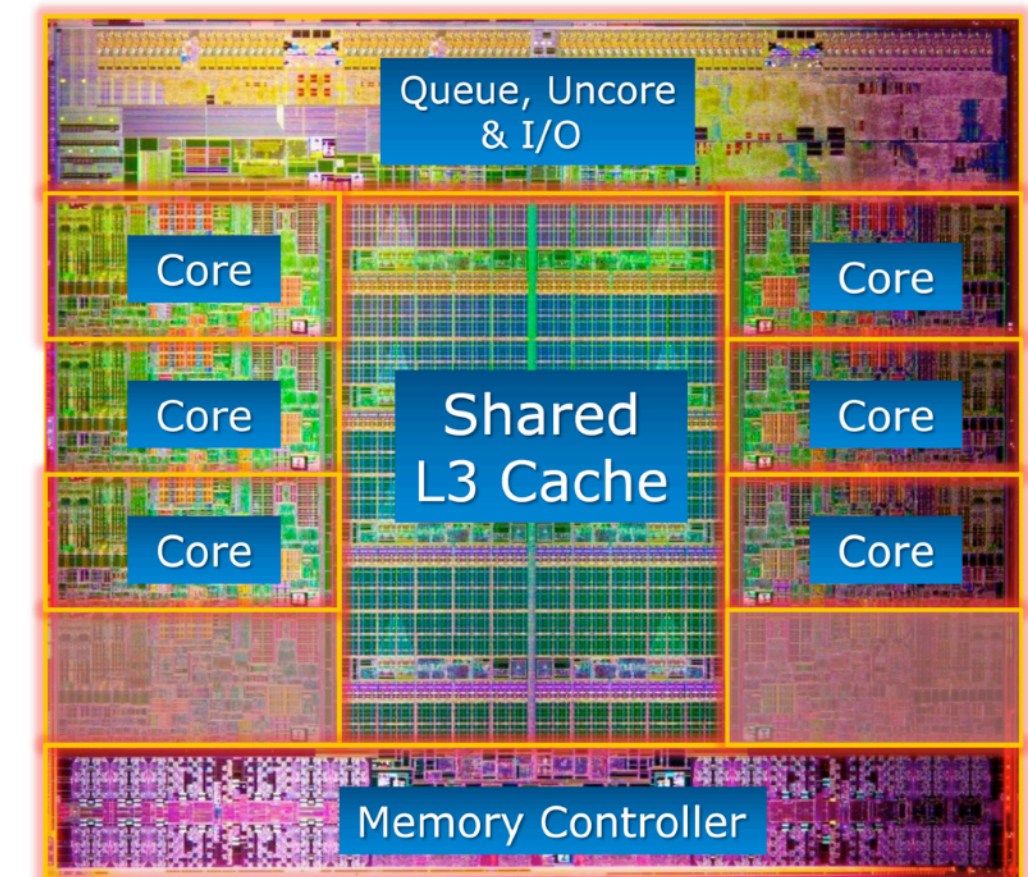
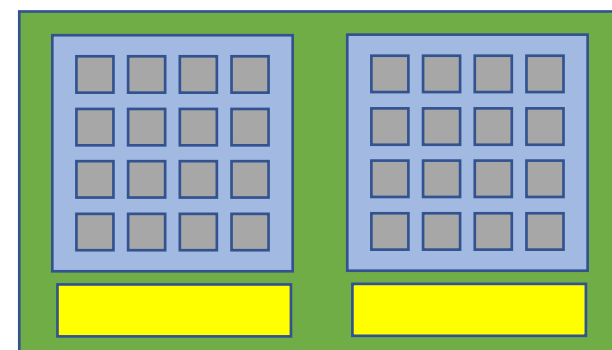
<https://easybuild.io>



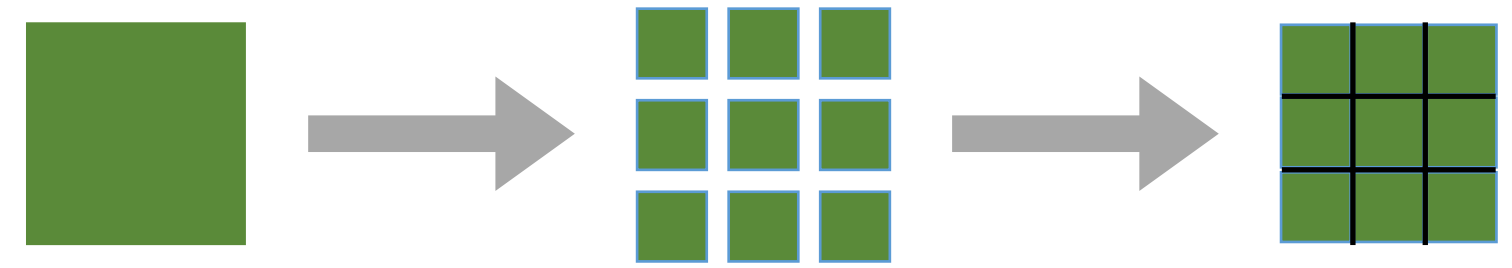
MODERN MULTI-CORE PROCESSORS



- CPU (central processing unit) a.k.a. (micro)processor
- **Multiple cores** to run multiple computations in parallel
- Hierarchy of small (~MBs) but (very) fast on-chip *cache* memory
- Channels to access memory & co
- Typically 2 multi-core processors per node (each housed in a *socket*)
- Simplified view (2x 16 cores):



PARALLEL COMPUTING

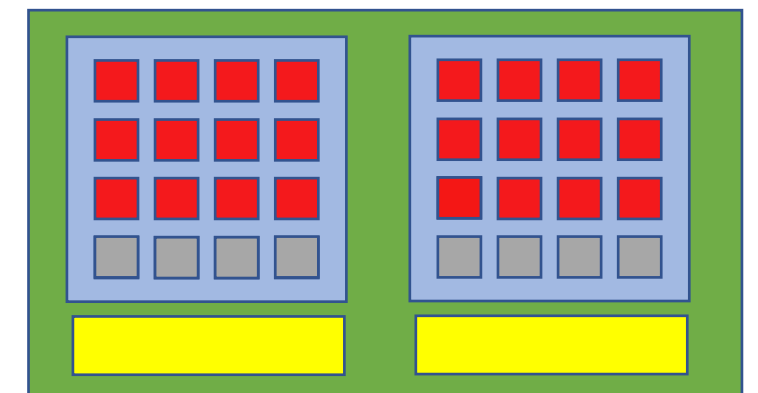


- Long-running calculations can be done faster through **parallelization**
- Split up large problem into smaller problems, divide and conquer
- Challenge is to **keep all cores busy** (doing useful work)
 - Avoid bottlenecks: slow memory/disk, network latency, limited bandwidth
 - Avoid load imbalance (waiting for longest running part to be done)
- Avoid overhead in breaking up problem, and composing final result

SHARED MEMORY PROGRAMMING

- **Intra-node** parallel computing: using multiple *threads* on a **single node**
- Each thread runs calculations on 1 core
- Communication between threads via shared data structures in memory
- Limited to resources available in a single node (cores, memory, ...)
- Care must be taken when writing to shared data structures (use of locks)
- Typical programming paradigm: OpenMP

```
#pragma omp parallel for  
for(int x=0; x < width; x++){  
    ...  
}
```



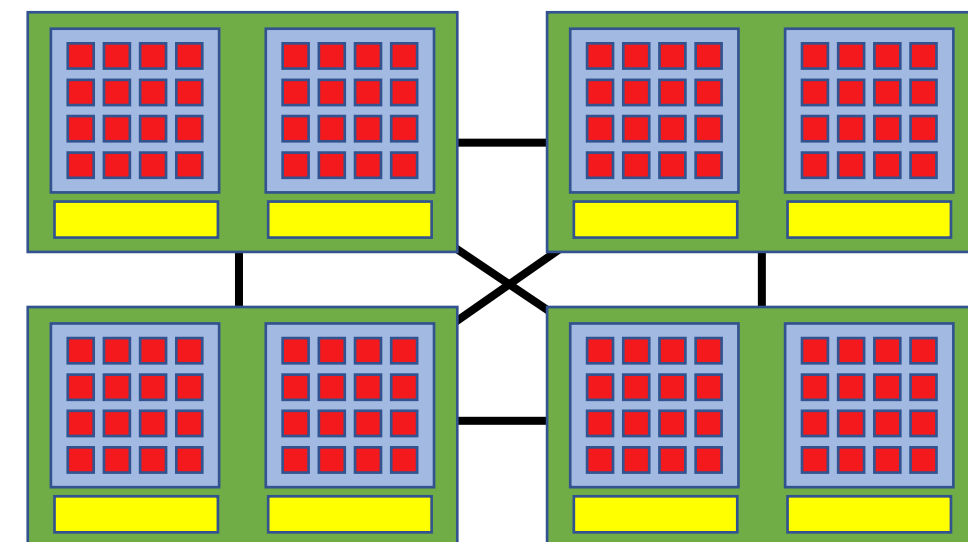
DISTRIBUTED MEMORY PROGRAMMING

- **Inter-node** parallel programming: using cores in **multiple nodes**
- Multiple instances of same software are started (*processes, ranks*)
- Coordination by exchanging messages (who does what, partial results, ...)
- Fast interconnect is *crucial* if a lot of *communication* is needed to coordinate

- Typical programming paradigm: Message Passing Interface (MPI)

```
MPI_Send(...)  
MPI_Recv(...)
```

- *Hybrid* parallel programming:
multiple processes, each with multiple threads



AMDAHL'S LAW

$$speedup(c) = \frac{1}{(1 - p) + \frac{p}{c}}$$

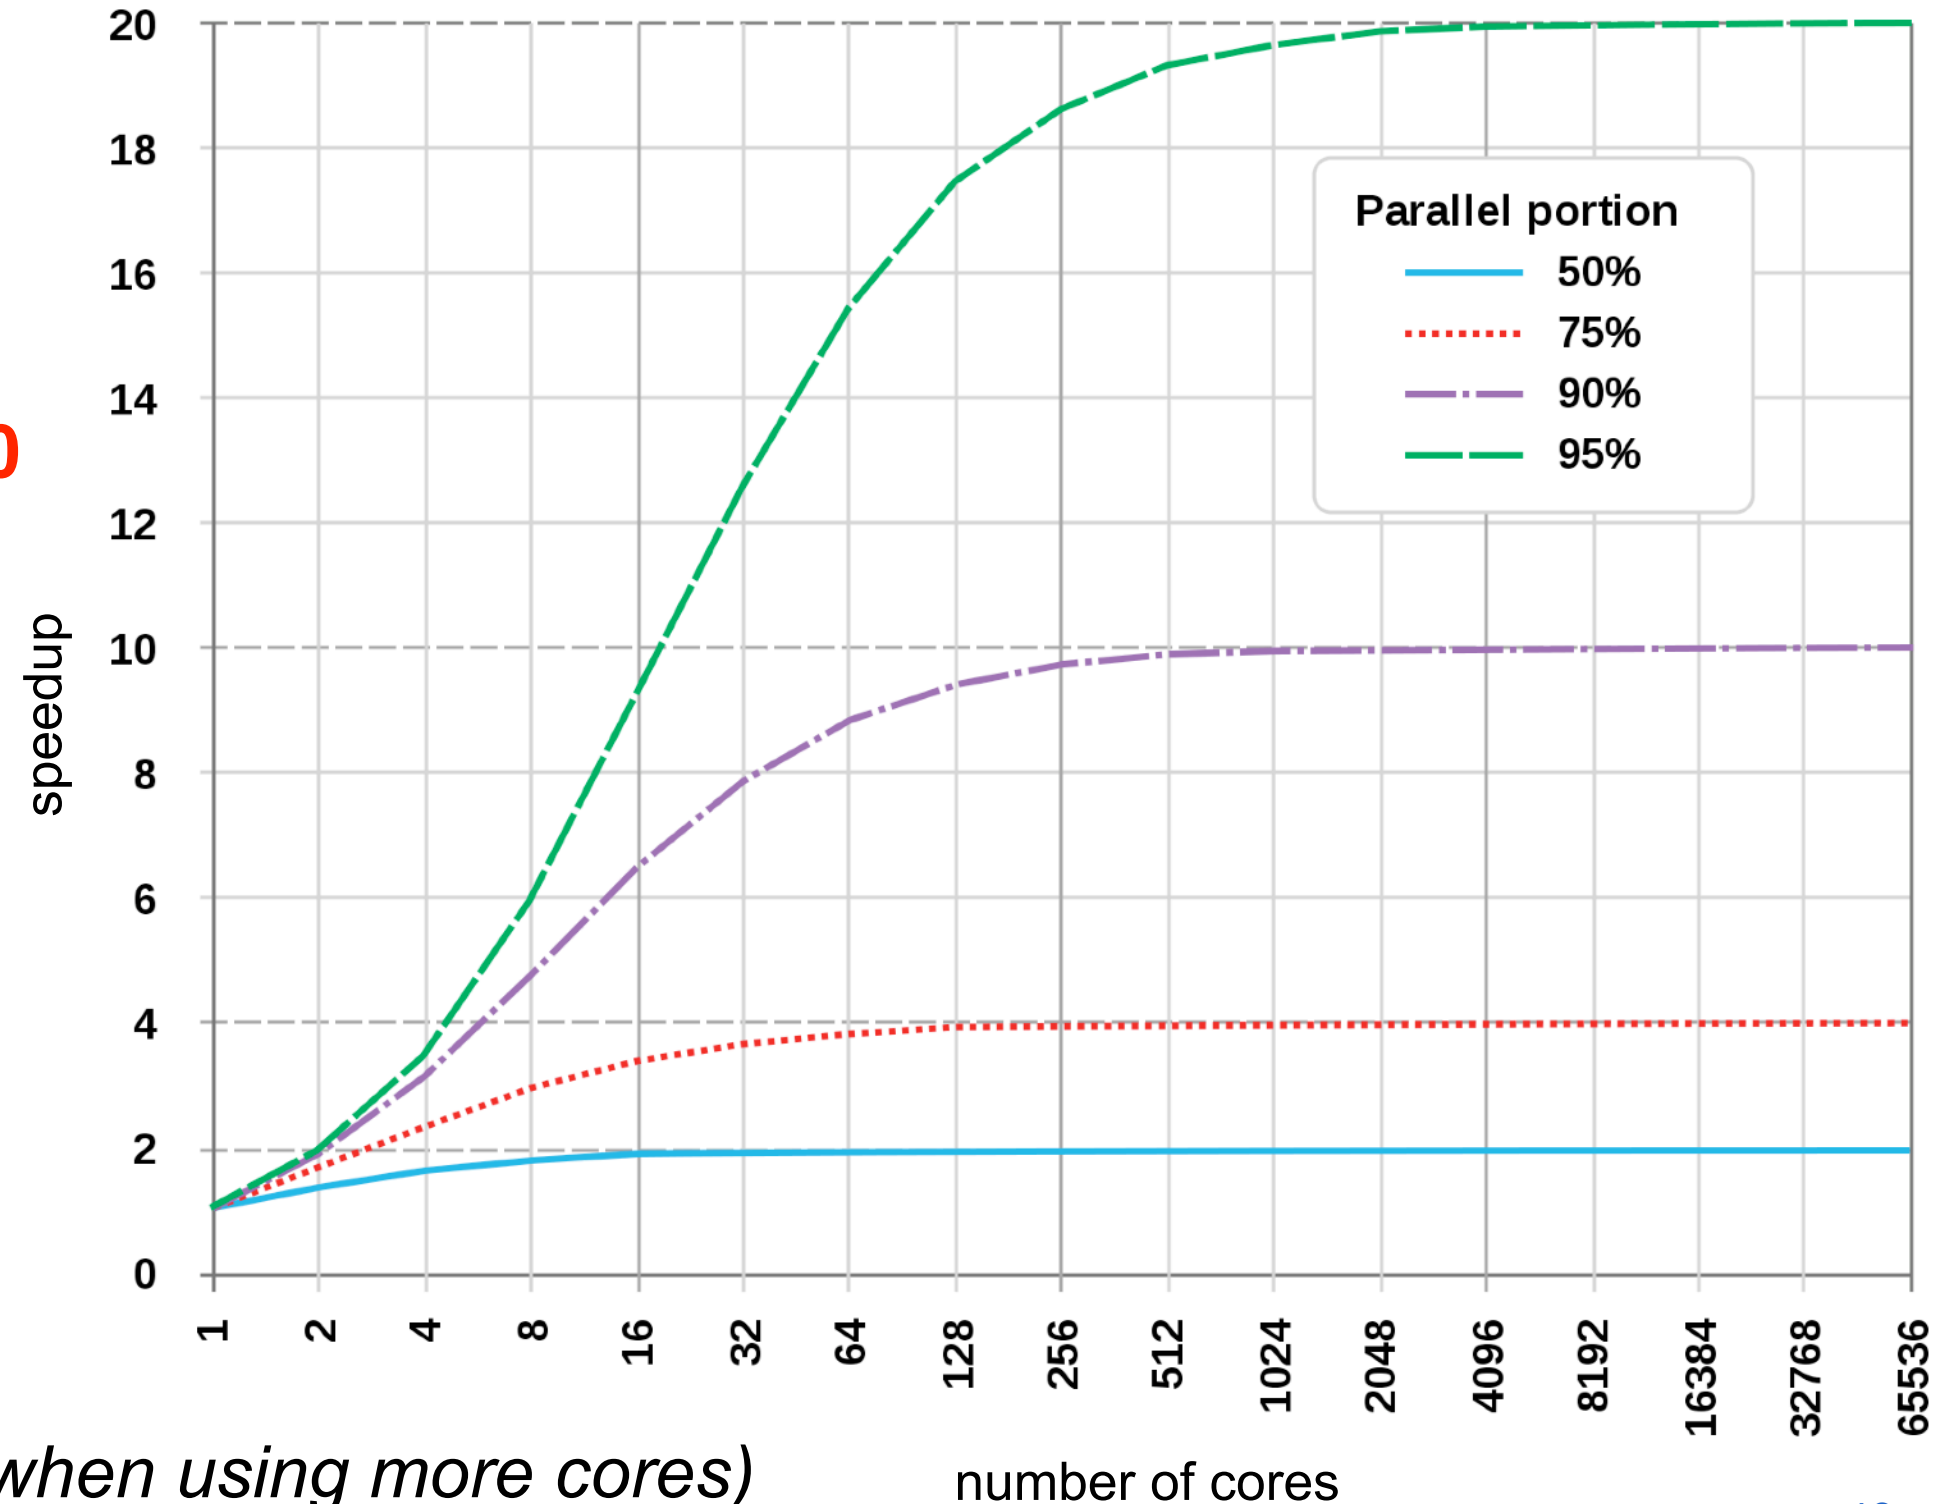
c : number of cores
 p : parallel portion ($[0, 1]$)

Examples:

- $p = 0.5$, 2 cores \rightarrow speedup = 1.33
- $p = 0.5$, 64 cores \rightarrow speedup = 2
- **$p = 0.99$, 1 million cores \rightarrow speedup = 100**

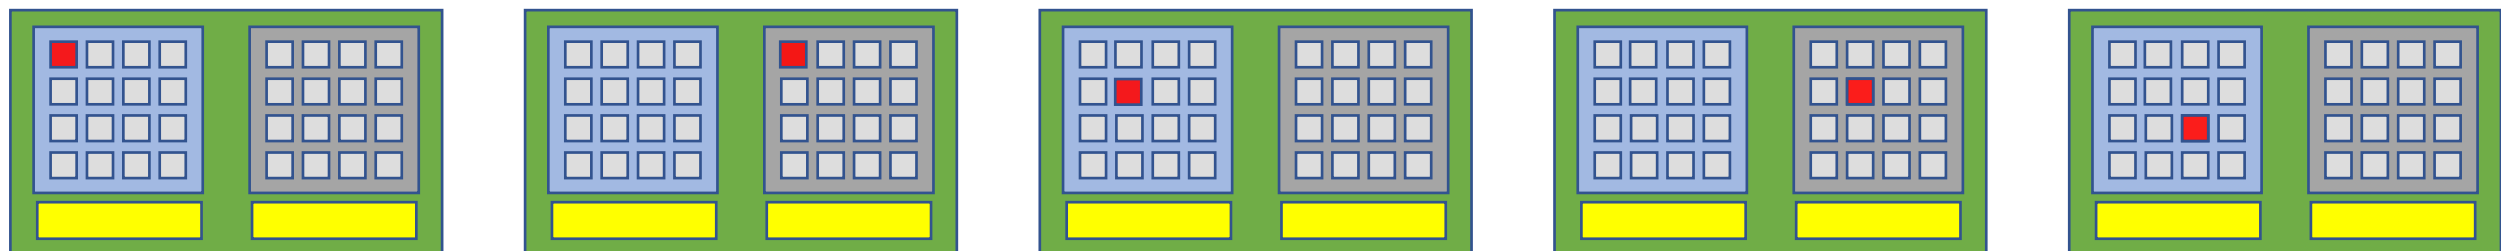
Assumptions:

- Fixed workload (strong scaling*)
- Parallel portion is *perfectly* scalable
- No bottlenecks, overhead, etc.



EMBARASSINGLY PARALLEL TASKS

- What if you need to run a lot of very small (single-core) calculations?
- a.k.a. High-Throughput Computing (HTC)
- Only uses part of available resources (no need for fast interconnect or shared storage)
- Same software can be started multiple times within a single job
- Each instance is running independent of the others (no communication between them)
- Tools: GNU parallel, dask, ...



PROGRAMMING LANGUAGES IN SCIENTIFIC COMPUTING

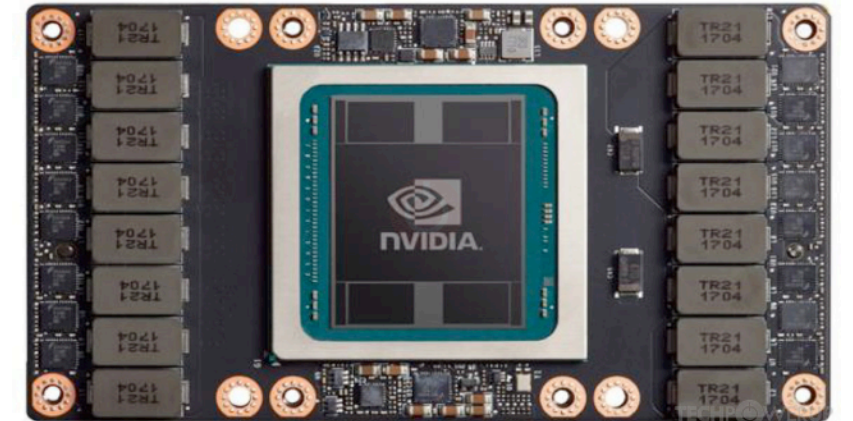


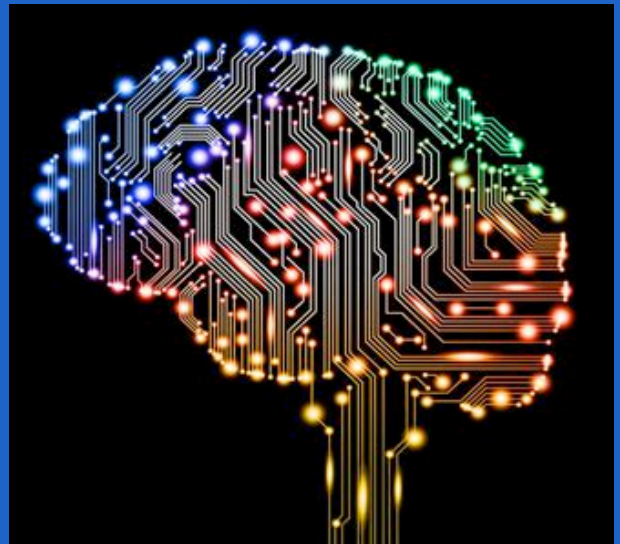
(back in [top 20 of programming languages!](#))



ACCELERATED COMPUTING

- Some type of calculations can benefit from special-purpose hardware
- Prime example of these *accelerators* are GPUs
- Originally created for detailed video game rendering
- Now also GPGPU: General Purpose Graphics Processing Unit
- Programming paradigm: CUDA, OpenCL, OpenACC, ...
- Difficult to use efficiently: separate GPU memory, lots of tiny cores, ...
- Potentially spectacular speedup: ~10-100x (but not for everything...)





artificial intelligence



cloud computing

TRENDS IN HPC

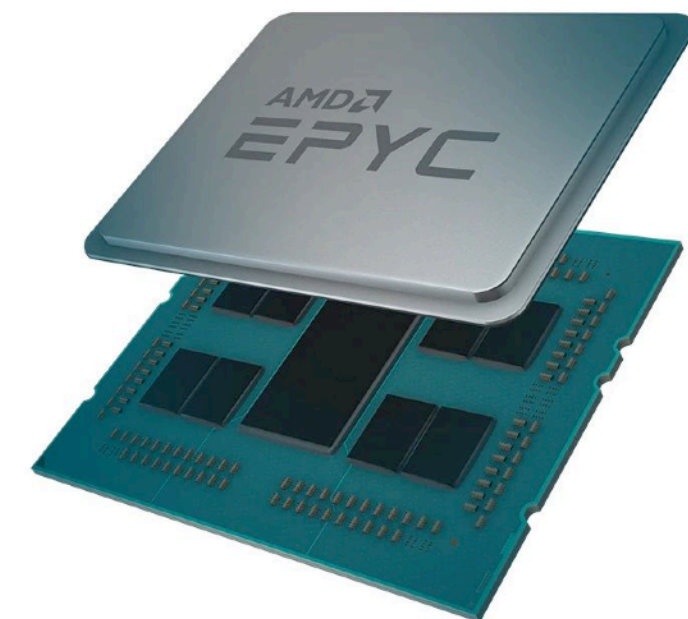
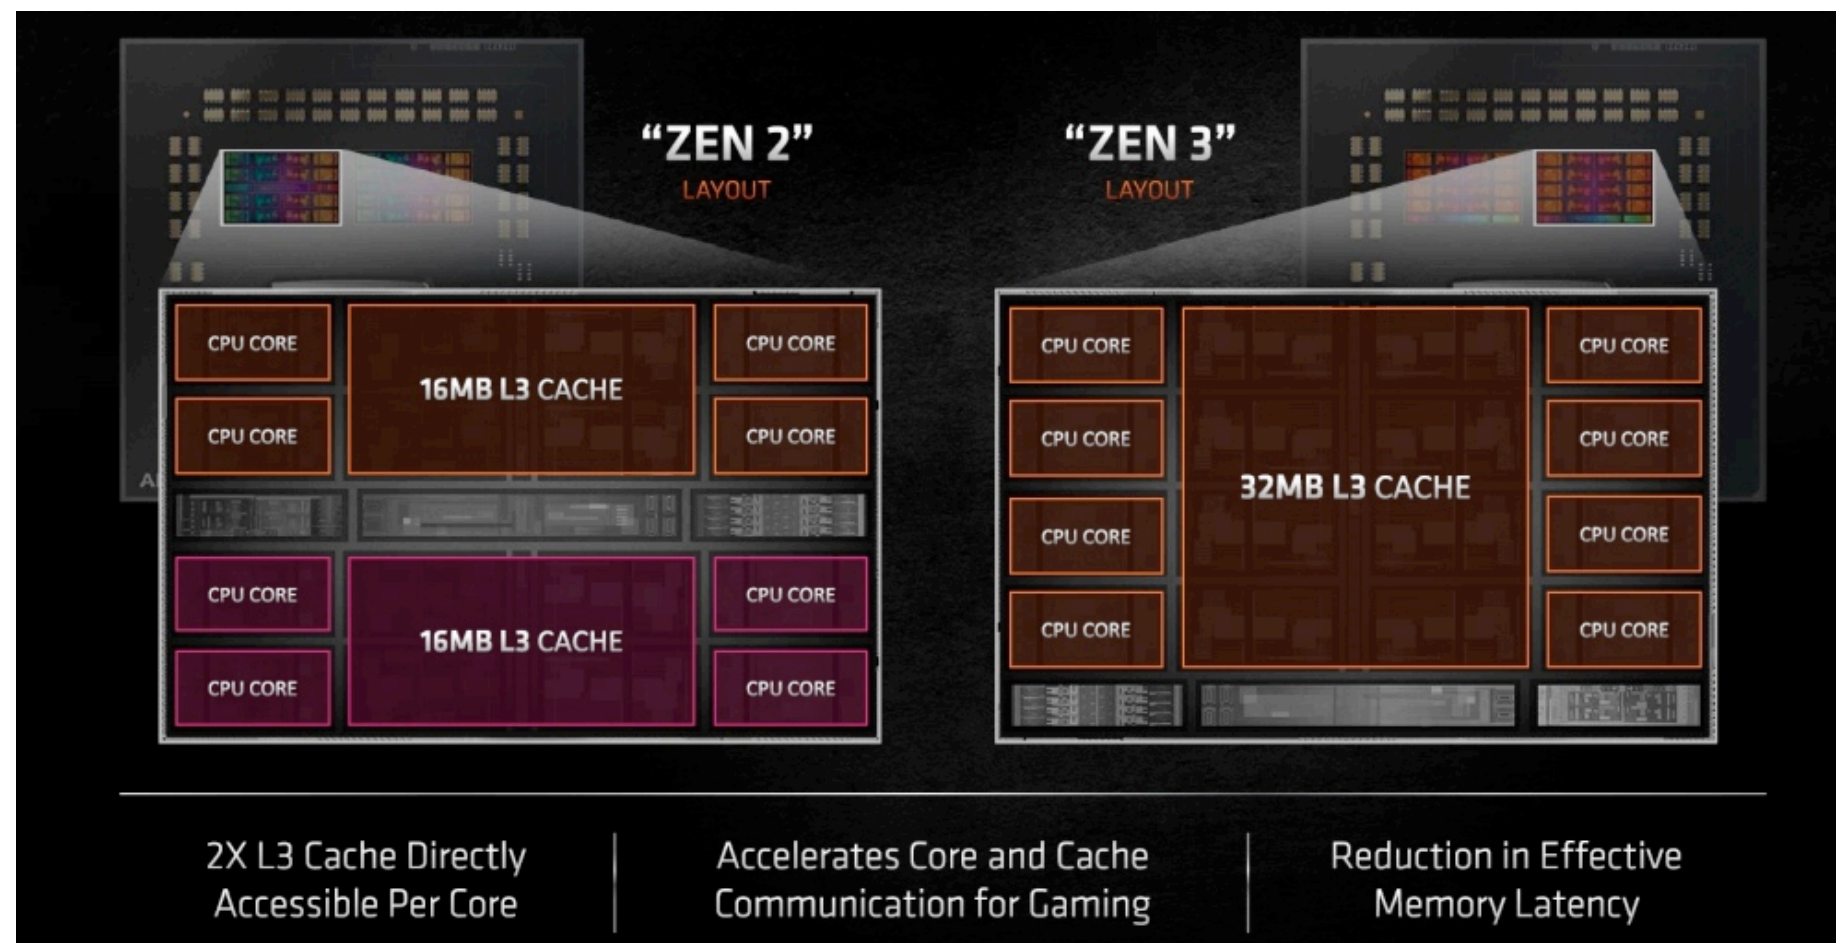
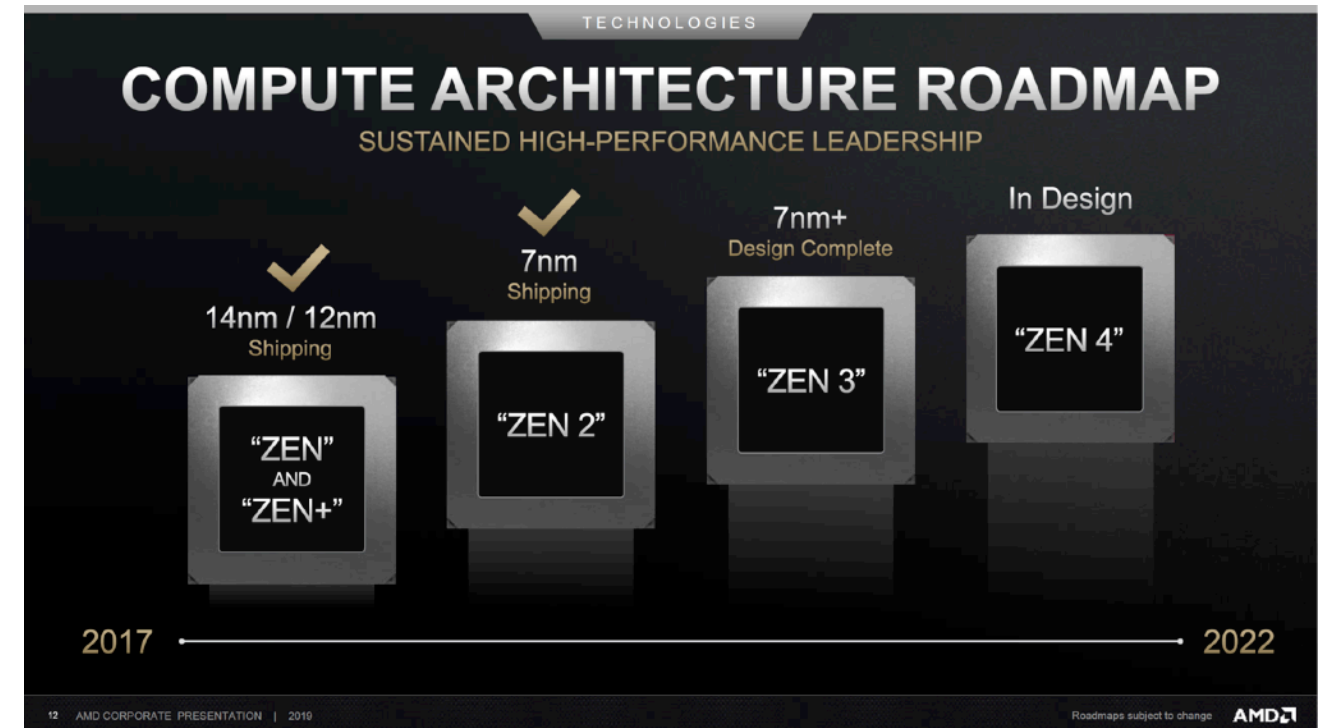
RETURN OF THE PROCESSOR WARS!

- In the last decade, Intel was king (w.r.t. microprocessors in servers)
- AMD is back in the game since ~2017 (AMD EPYC processor generation)
- Interest in ARM CPUs is booming (see Fugaku, Apple M1, NVIDIA Grace)
- IBM POWER is still relevant too (POWER9, POWER10 coming soon)
- In the (near?) future: "open source" high-performance RISC-V processors?



AMD EPYC PROCESSORS (ZEN* MICROARCHITECTURE)

- AMD is competitive again with Intel
- Smaller manufacturing process (7nm, Intel is still at 14nm...)
- Quite different microarchitecture ("chiplet" design), same instruction set architecture (x86_64)



INTEREST IN ARM CPUS IS PICKING UP SPEED

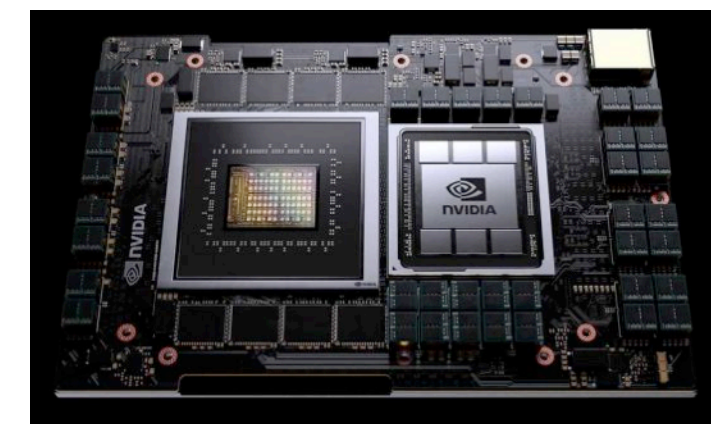
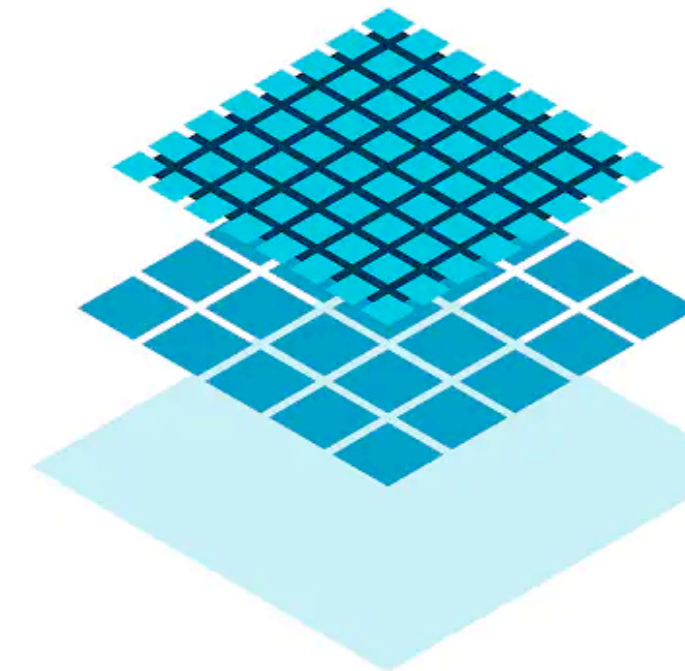
- ARM is a different processor Instruction Set Architecture (ISA)
- ARM processors are the most common in the world
 - > 1 billion units per year, used in smartphones, tablets, wearables, ...
- License from Arm Ltd. is required to use ARM CPU designs
- ARM CPUs are now also used in recent large supercomputers:
Fugaku (Japan), Isambard (UK), Astra (US), ...
- NVIDIA has announced its first ARM-based server CPU "Grace"
- NVIDIA is trying to buy ARM (deal is not approved yet by regulators)



Implementation

Microarchitecture

Architecture



RISC-V OPEN STANDARD ISA

- RISC-V is the 5th generation of RISC (Reduced Instruction Set Computer) ISA
- **Open source licensed:** free to use by anyone (no license fees)
- Several companies are designing RISC-V processors (some general purpose)
- Long term goal of [European Processor Initiative \(EPI\)](#) project is to design and produce a RISC-V processor for supercomputers & co, manufactured in Europe
- Barcelona Supercomputing Centre is has announced the [eProcessor](#) project, to build an open source out-of-order processor + European full stack ecosystem based on RISC-V
- Will we see a (European?) RISC-V based supercomputer soon?
- Lots of work to be done on software side: libraries, applications, etc.



RISE OF CLOUD COMPUTING



- Several companies provide "cloud computing" services

- "Rent" computing time & co, pay only for what you use



- Easy (?) to **create your own (throwaway) supercomputer!**

- Including shared storage, GPUs, fast interconnect, ...

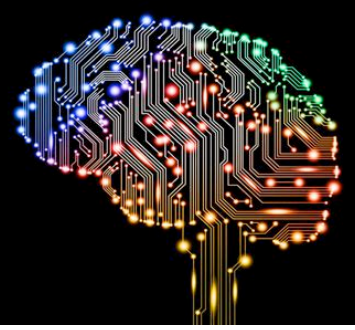


- Very flexible, but also expensive (or is it?)

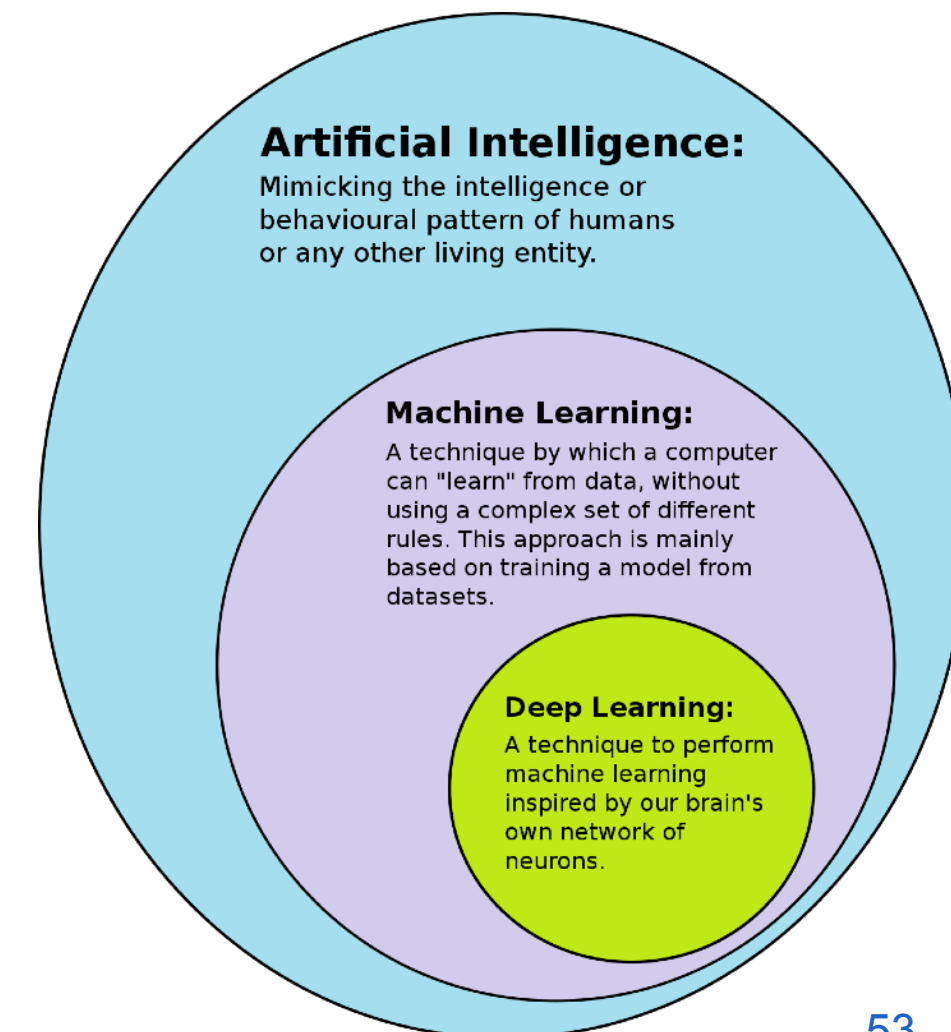
- "There is no cloud, just somebody else's computer."



ARTIFICIAL INTELLIGENCE: A NEW HOPE



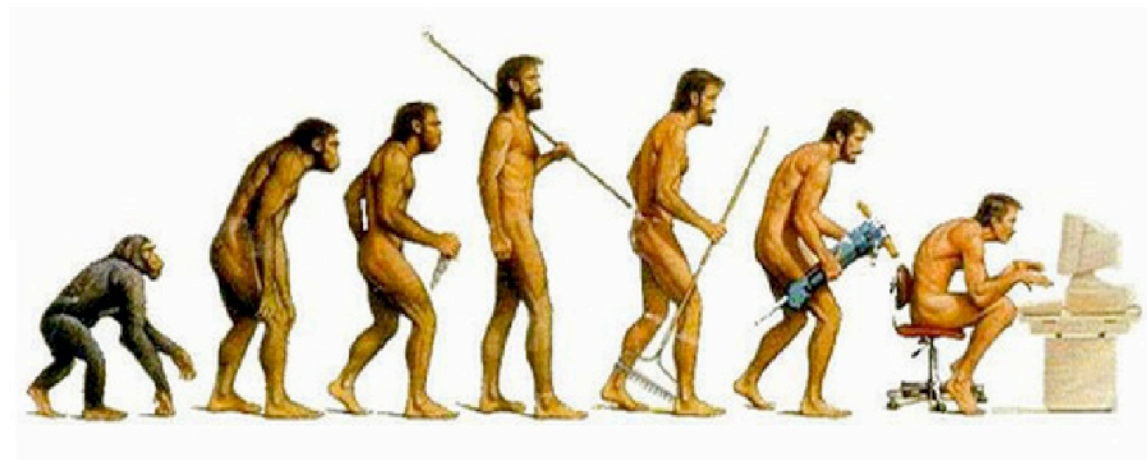
- Increased interest and adoption of machine learning techniques
- Fueled by capabilities of high-end computers (incl. supercomputers, GPUs)
- Typically requires large amounts of data ("big data")
- Lots of applications: digital imaging, self-driving cars, ...
- Deep learning "big bang" in 2009 fed by GPU boom

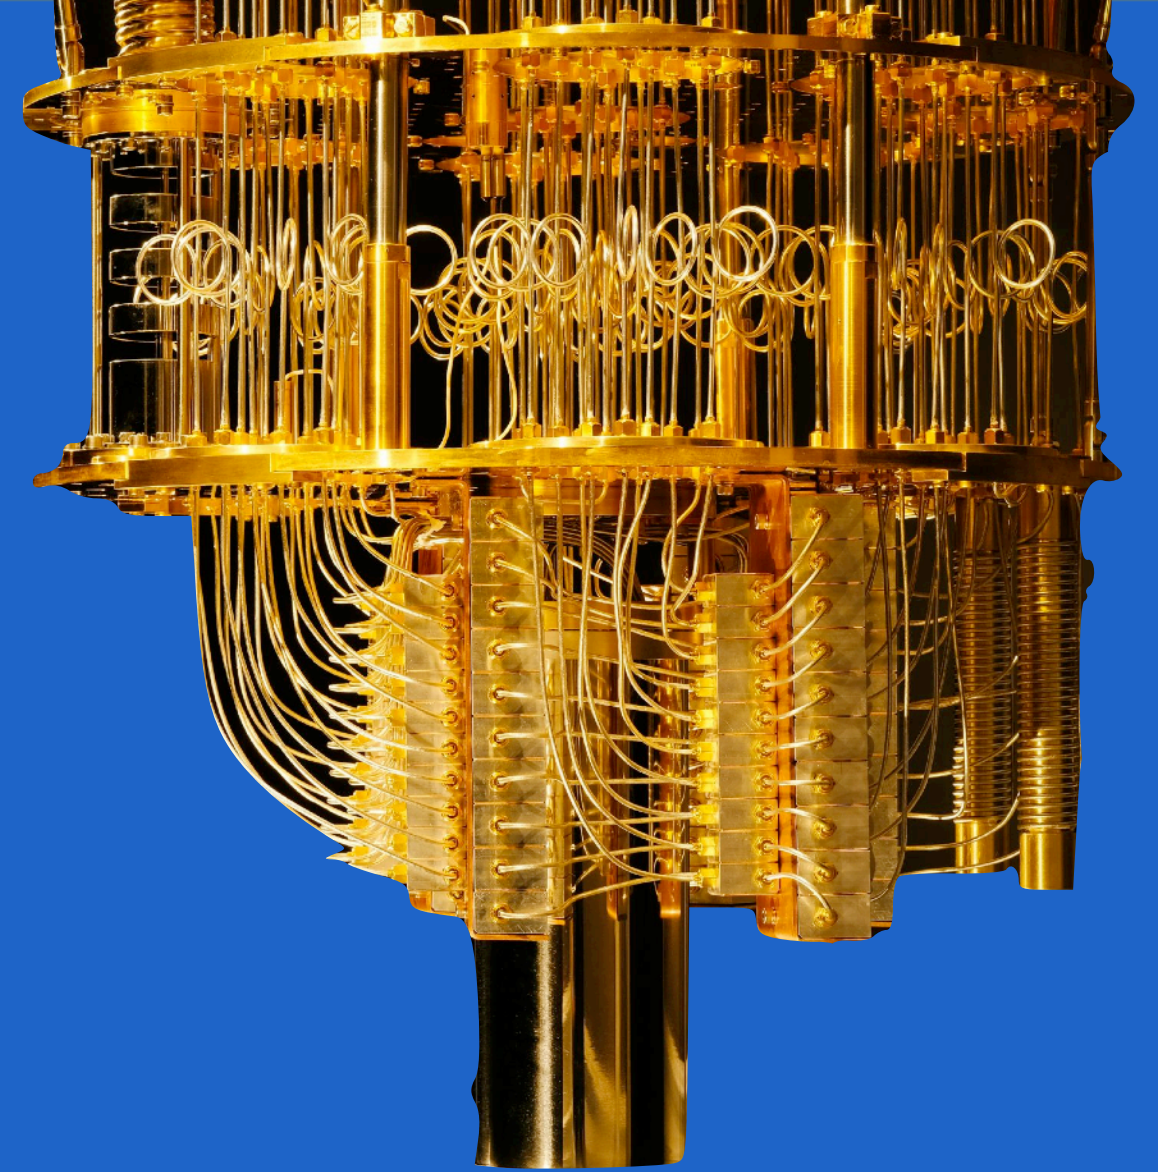


POLL

What will be the next big (r)evolution in scientific computing?

- Exascale supercomputers
- End of Intel/AMD processors as most prominent CPUs, switch to Arm
- High-performance RISC-V microprocessors
- HPC in the cloud
- Quantum computing
- Something else





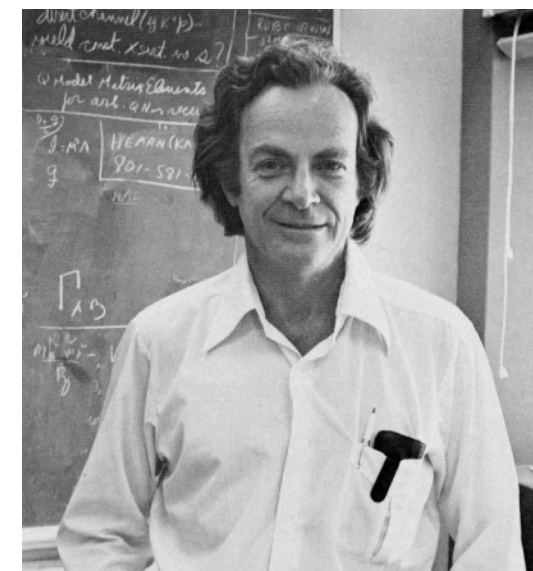
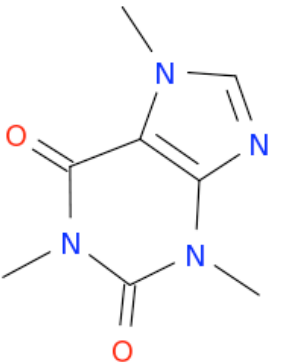
QUANTUM COMPUTING (IN A NUTSHELL)

DISCLAIMER

- I am *not* an expert in quantum computing.
- I have been trying to follow progress in quantum computing for a couple of years.
- Statements I make may be oversimplifying things too much (or may even be incorrect).
- I am *totally* skipping the math behind all this (complex numbers, probability, etc.)
- Some of this may be outdated already.
- Google was my friend when puzzling this together...
- Please don't invest money based on what I say (unless you don't need it).

QUANTUM COMPUTING: HYPE OR REVOLUTION?

- Some problems are just **too big**, even for the largest (future) supercomputers!
 - Would take too long to solve, require too much memory, etc.
- Example: exact simulation of a caffeine molecule
 - ~95 electrons, would require $\sim 10^{48}$ classical bits (1 terabyte = $\sim 2^{40}$ = $\sim 10^{12}$ bytes)
- **Quantum computers** could be the answer...
- We're getting close to "*quantum supremacy*"
(or maybe we are there already...)



Richard Feynman (1918-1988)
"father" of quantum computing

CLASSICAL BITS VS QUBITS

- Classical bits are either 0 or 1 (like flipping a coin: heads or tails)



- **Quantum bits ("qubits")** can be in a state somewhere in between 0 and 1...

- Like a *spinning* coin, in some sense they can be "both 0 and 1 at the same time"



- Measuring (observing) a qubit collapses the state to either 0 or 1 (with a certain probability)

- Geometrical representation of a single qubit: *Bloch sphere*

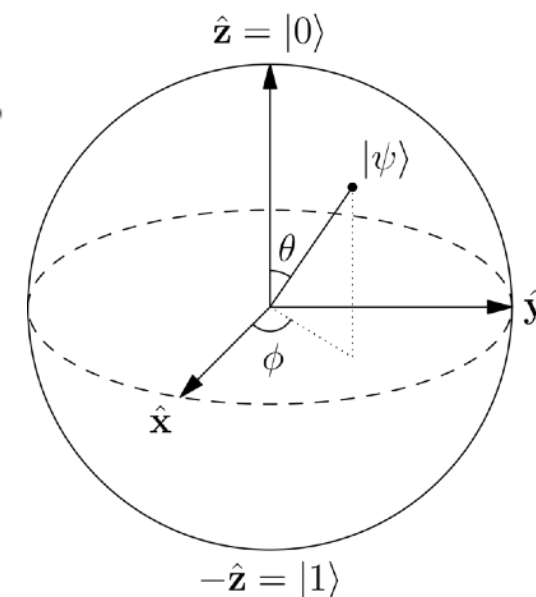
$$|\psi\rangle = \alpha |0\rangle + \beta |1\rangle$$

- North pole : $|0\rangle$ (equivalent with 0 in classical bit)

- South pole: $|1\rangle$ (equivalent with 1 in classical bit)

- Equator: probability of 0.5 of observing 0 or 1 when measuring

- Somewhere else on surface of the sphere: leaning towards either $|0\rangle$ or $|1\rangle$...



LEVERAGING QUANTUM PHYSICS

- **Superposition**

- A qubit can be put in a state between $|0\rangle$ and $|1\rangle$
- A collection of qubits can be manipulated to **favor** the "optimal" state

- **Entanglement**

- Two qubits can be "entangled": measuring one impacts the other
- The distance between the entangled qubits doesn't matter!
- Can be used for *teleportation* of information (faster-than-light communication!)
- Einstein called this "spooky action at a distance"

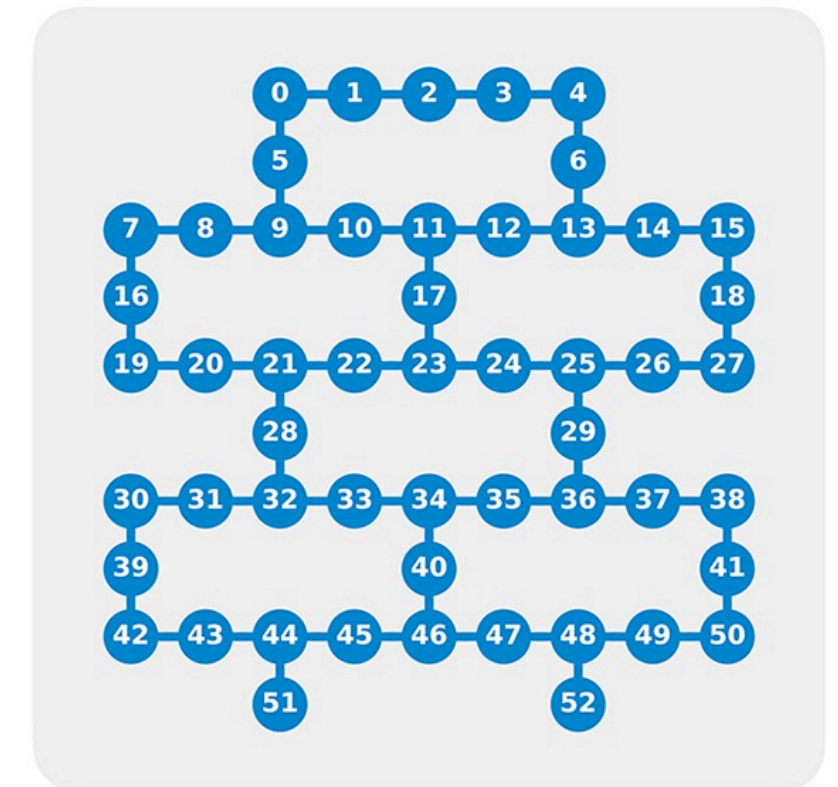
- **Interference**

- The state of a qubit can be changed via constructive or destructive interference

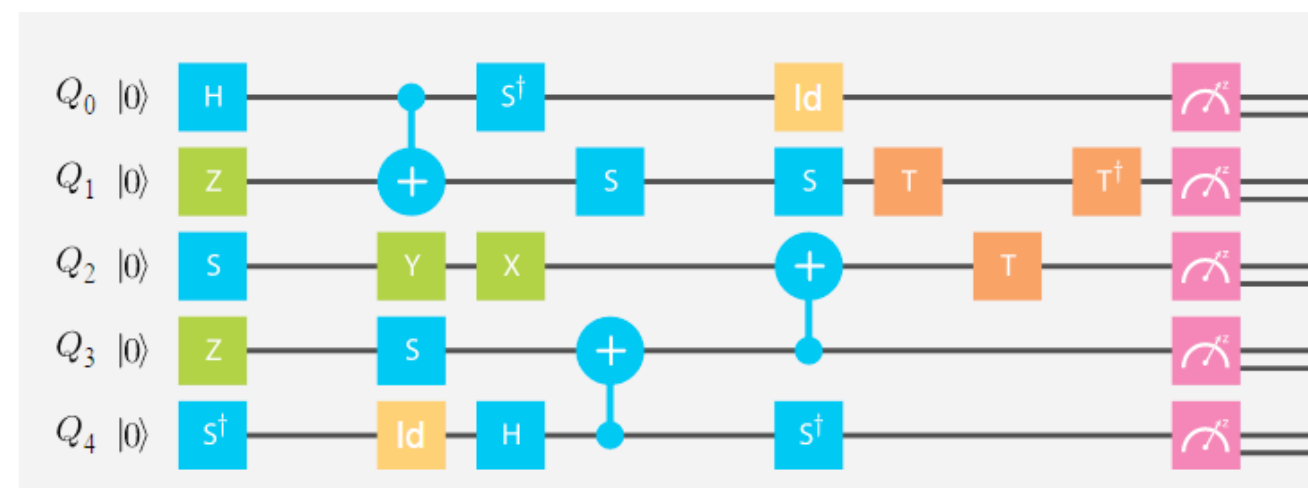
QUBITS AND QUANTUM CIRCUITS

- Combining multiple qubits makes things interesting...
- In some sense, every qubits *doubles* the compute power
 - Assuming each qubit is connected to all others (it won't be)
- Quantum circuits* can be used to manipulate a set of qubits
 - Gate-model quantum computing: a circuit is a series of *gates*
 - Equivalent to a computer program in classical computers

53 Qubit Rochester Device



Operator	Gate(s)
Pauli-X (X)	
Pauli-Y (Y)	
Pauli-Z (Z)	
Hadamard (H)	
Controlled Not (CNOT, CX)	
Controlled Z (CZ)	



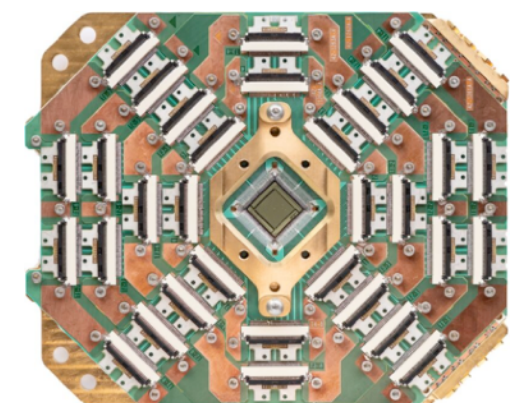
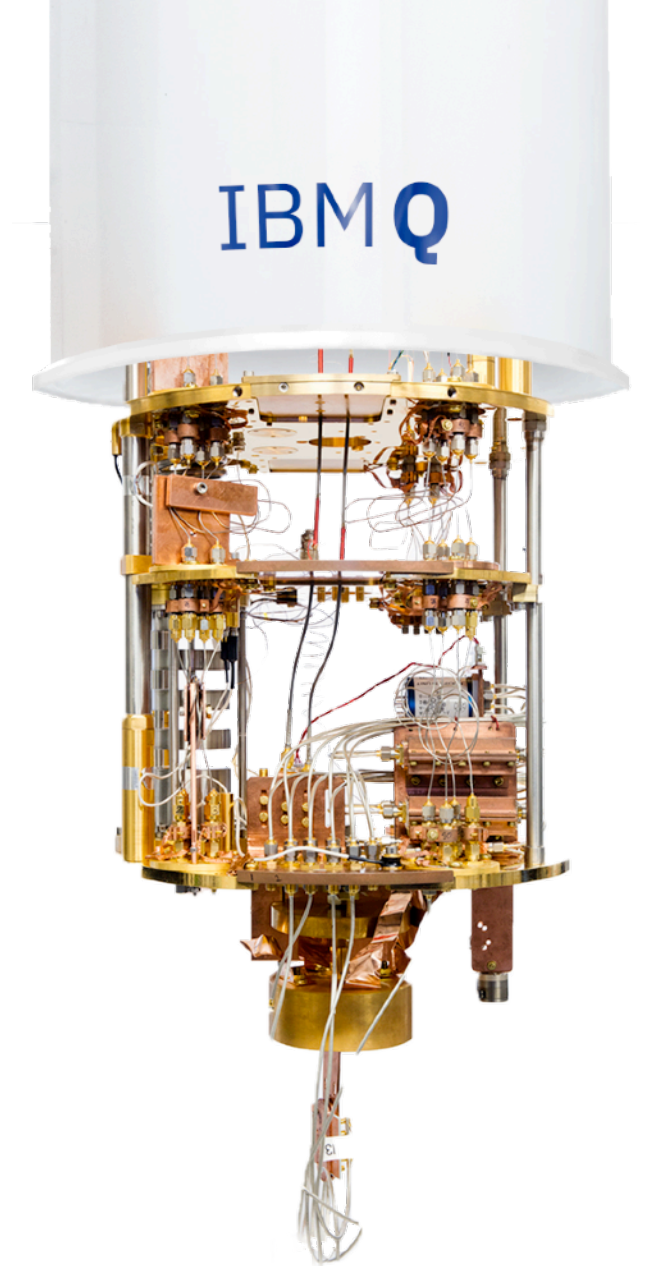
TYPES OF QUANTUM COMPUTERS

- **Circuit-based quantum processors**

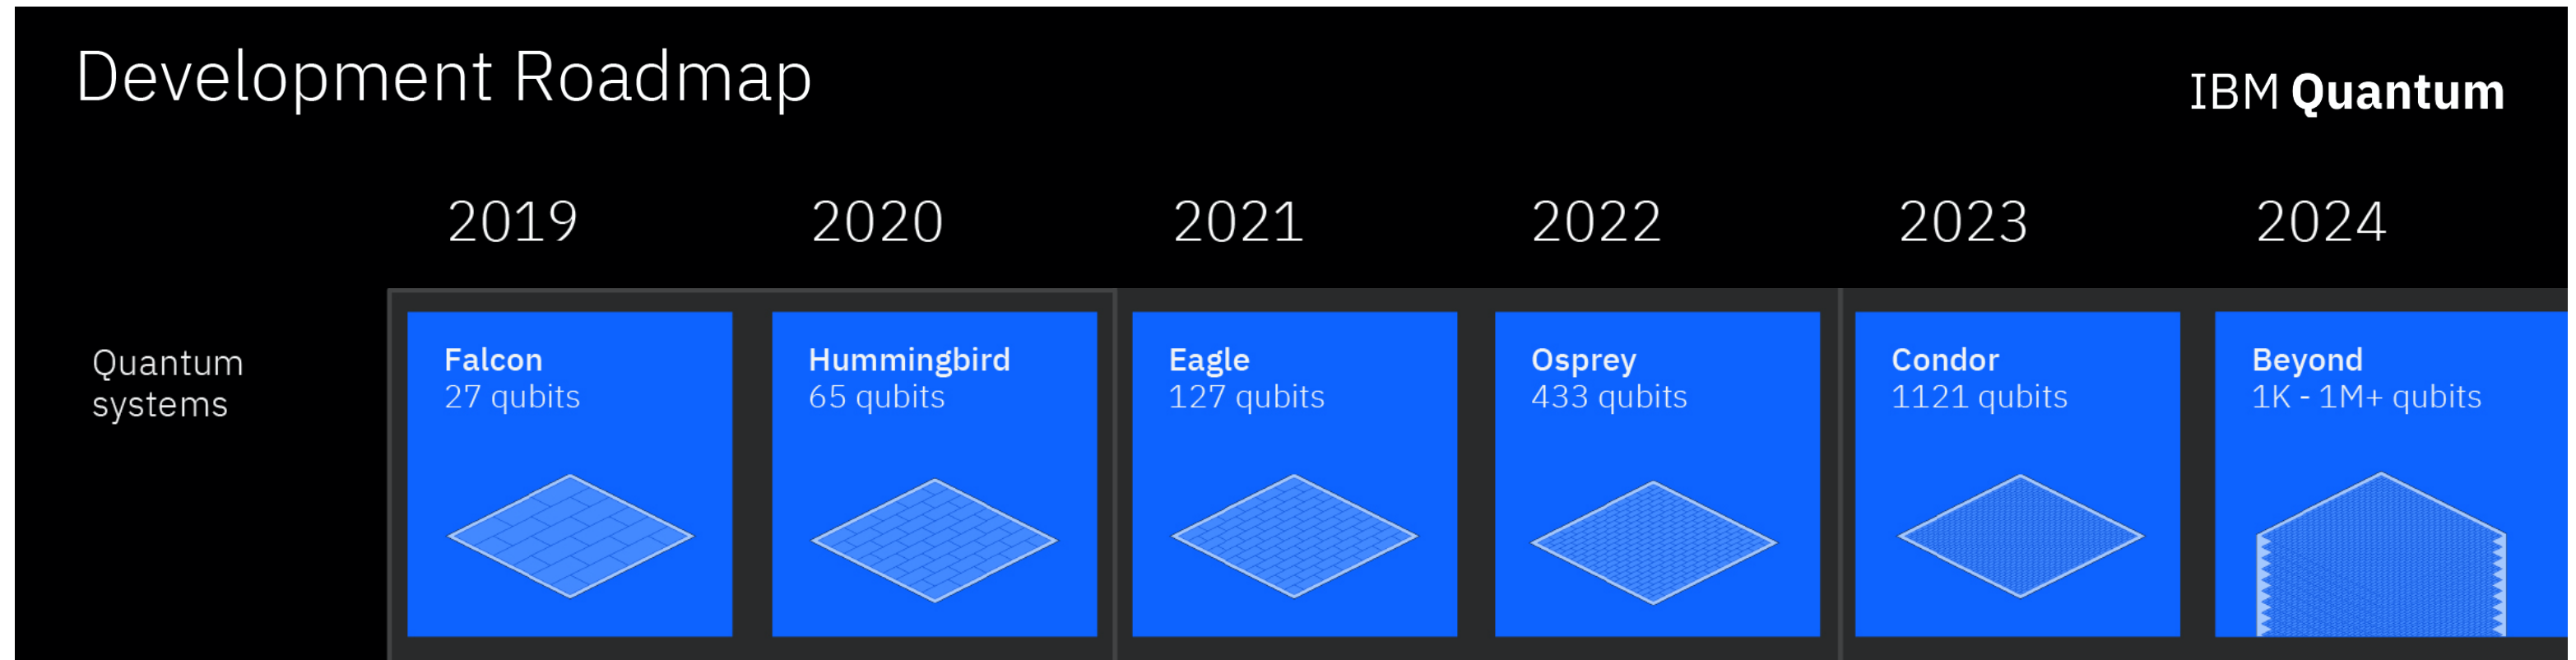
- IBM, Google, Intel, Rigetti, Xanadu Quantum Technologies
- Different technologies: superconducting qubits, photonics, ...
- Currently less than 100 qubits

- **Quantum annealers**

- Can only be used for certain type of problems (optimization)
- D-Wave One (2011): 128 qubits
- D-Wave 2000Q (2017): 2048 qubits
- D-Wave Advantage (2020): 5760 qubits



IBM QUANTUM ROADMAP



- Today: 60-70 qubits
- Soon (this year?): > 100 qubits
- In next 5 years: over 1 million qubits?



APPLICATIONS FOR QUANTUM COMPUTING



Cryptography

- Break popular encryption algorithms (RSA)
- Would be *very* disruptive (quantum computing as a "weapon")
- Post-quantum cryptography (who wins the race?)

Machine learning

- Quantum Machine learning algorithms already exist today
- Artificial Intelligence (AI) more popular than ever
- Could be the way to the "Singularity"

Simulation of quantum physics

- Molecular simulations, drug development, etc.
- Very popular application in supercomputers today
- Some problems are impossible to solve today, even on supercomputers

Search & optimization

- Searching for "best" item in unstructured space
- Can be proved Grover's is the *optimal* way of searching

QUANTUM ALGORITHMS

- **Grover's search algorithm**

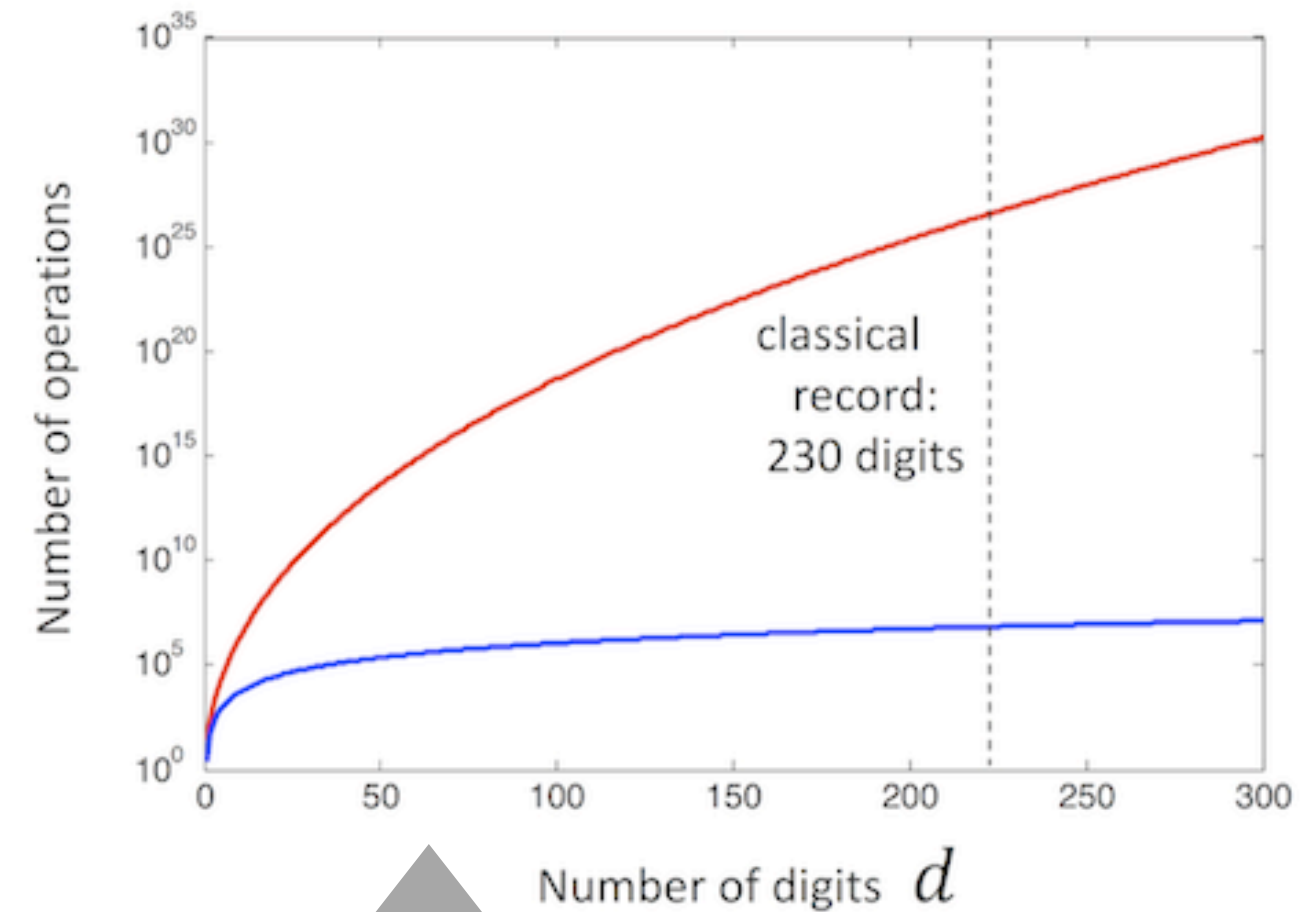
- Unstructured search problem: find best out of N
- On average: $\frac{N}{2}$ classically, \sqrt{N} with Grover's

- **Shor's algorithm**

- Factoring integer numbers into primes in polynomial time
- Can be used to break RSA encryption (uses 2,048 bit integers, or larger)
- Shor requires $\sim 2N$ *logical* qubits (N: number of bits to represent integer to factor)

- **Quantum phase estimation**

- Central building block for many quantum algorithms

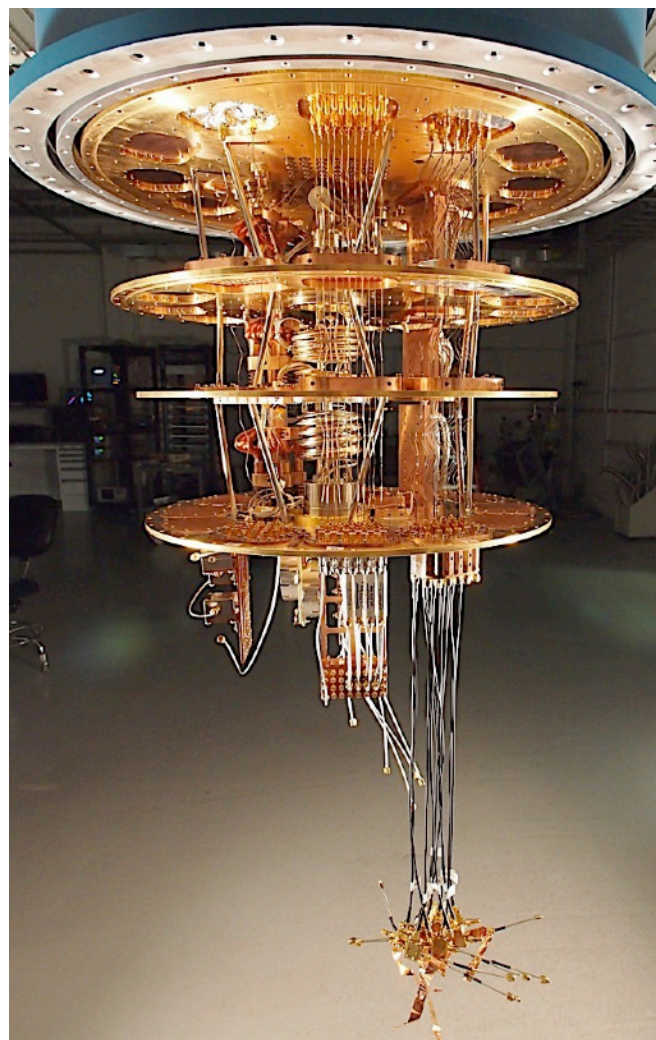


CHALLENGES IN QUANTUM COMPUTING

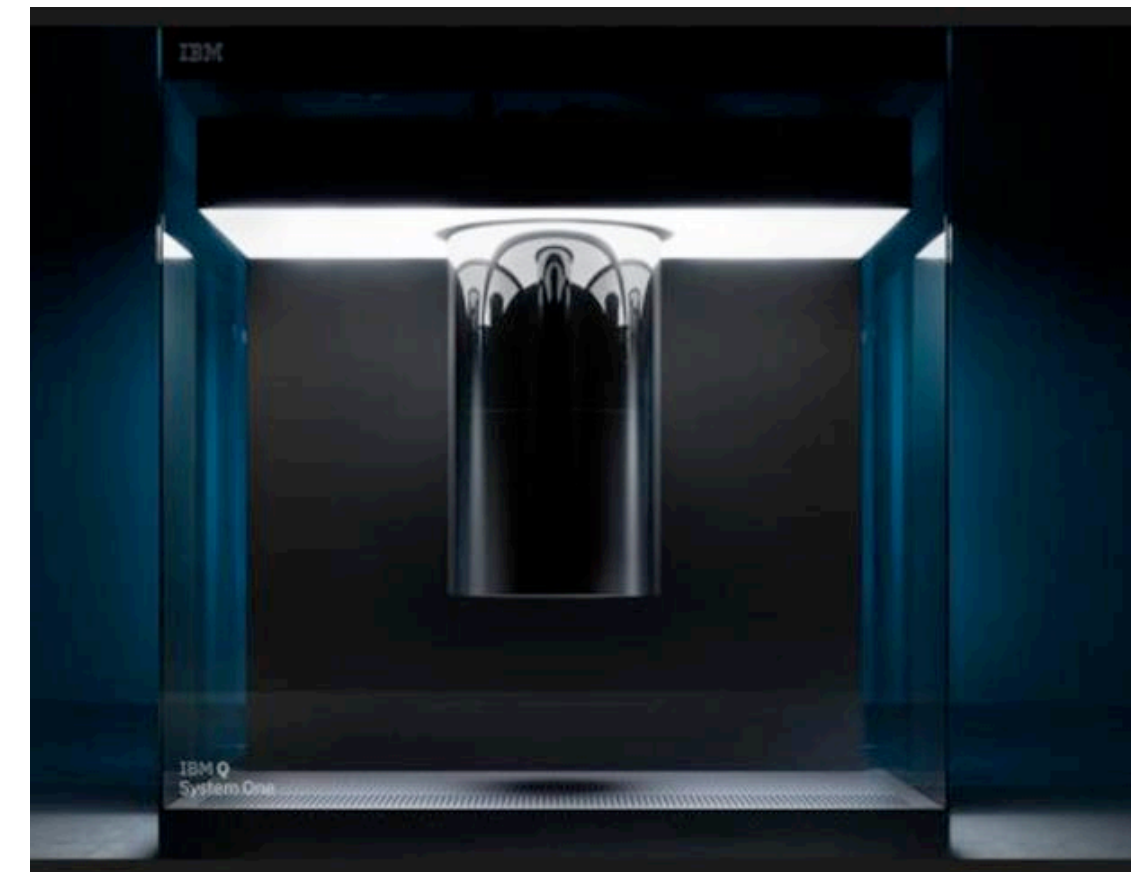
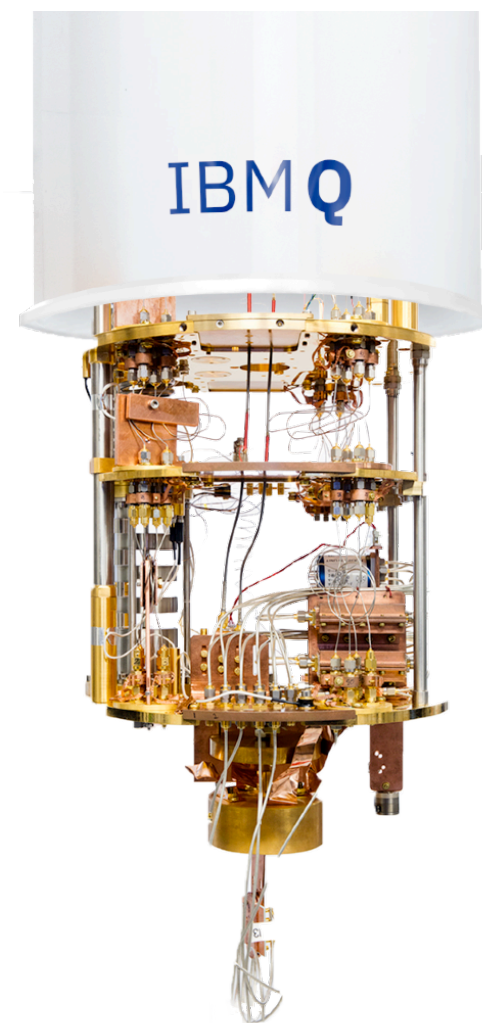
- **Physics**
 - Creating and controlling qubits: superconducting, photonics, ...
 - Keeping out the noise, temperatures close to 0K (colder than outer space!)
- **Decoherence**
 - Qubits "collapse" due to outside influence: noise, vibrations, radiation, ...
 - Quantum "programs" must finish before decoherence happens (< 1 ms currently)
- **Connectivity**
 - Multiple qubits must be as connected as possible so they can be entangled, etc.
 - Quantum "compilers" are used to map quantum circuits on physical qubit layout
- New generation of **programmers** needs to be trained to use quantum computers...

CURRENT QUANTUM COMPUTERS

- Currently in **Noisy Intermediate Scale Quantum (NISQ)** era
- Small number of qubits, very short coherence time (~ 0.1 ms)
- Largest number of superconducting qubits: Google Bristelone (72 qubits)
- Largest number of qubits: Jiuzhang (76 photonic qubits), China



"Naked"
quantum
computers
(Google, IBM)



IBM Q System One
Commercial quantum computer (20 qubits)

- You can play with a **real quantum computer** yourself (or simulate one).
- In the cloud! For free! Using Python code!

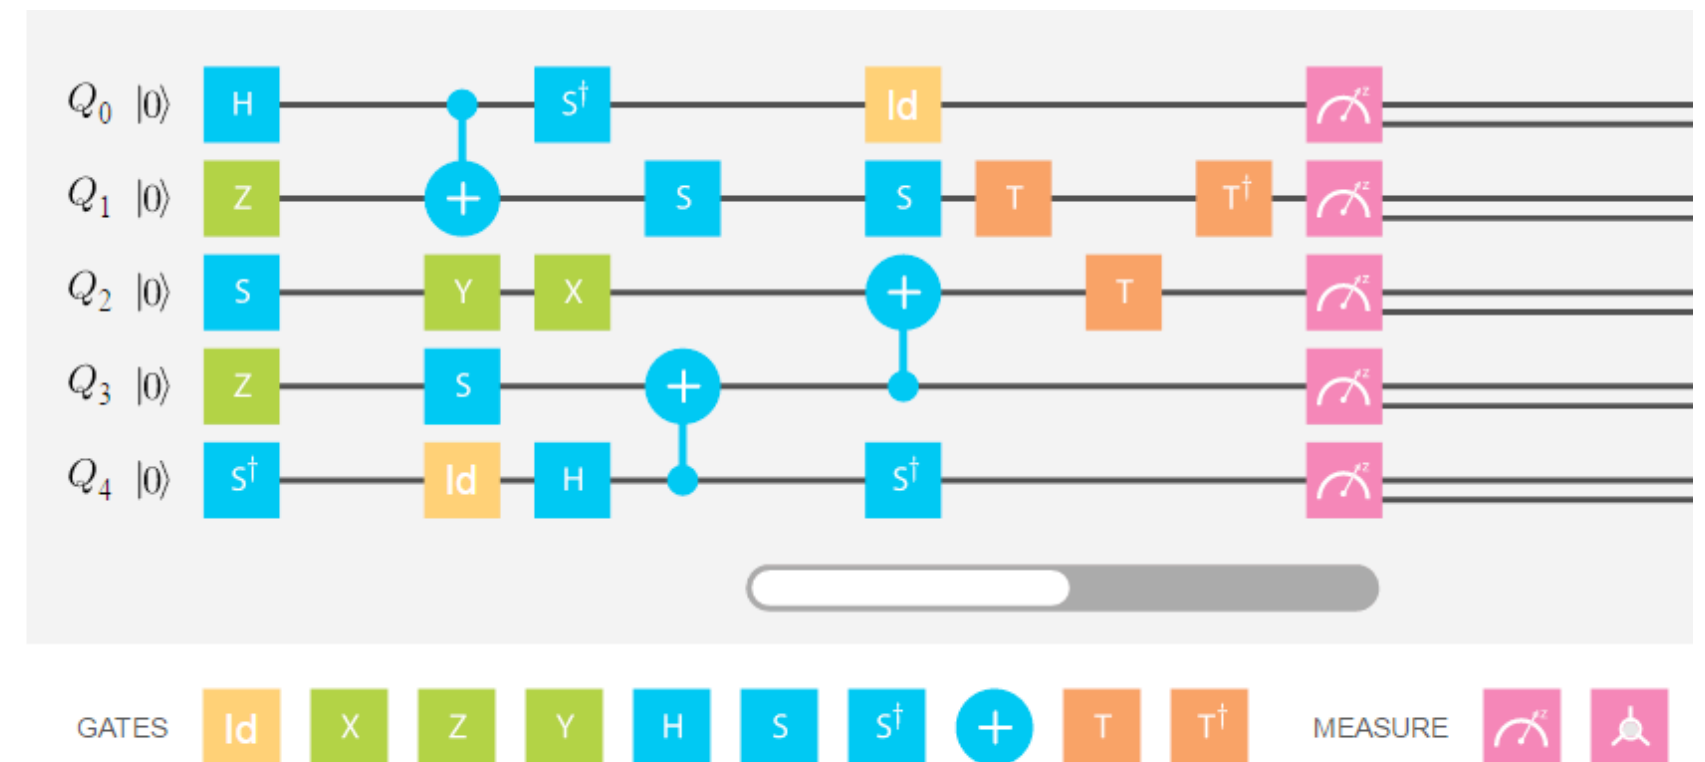
```
from qiskit import QuantumCircuit

# Create a Quantum Circuit (2 qubits, 2 classical bits)
circuit = QuantumCircuit(2, 2)

# Add a H gate on qubit 0
circuit.h(0)

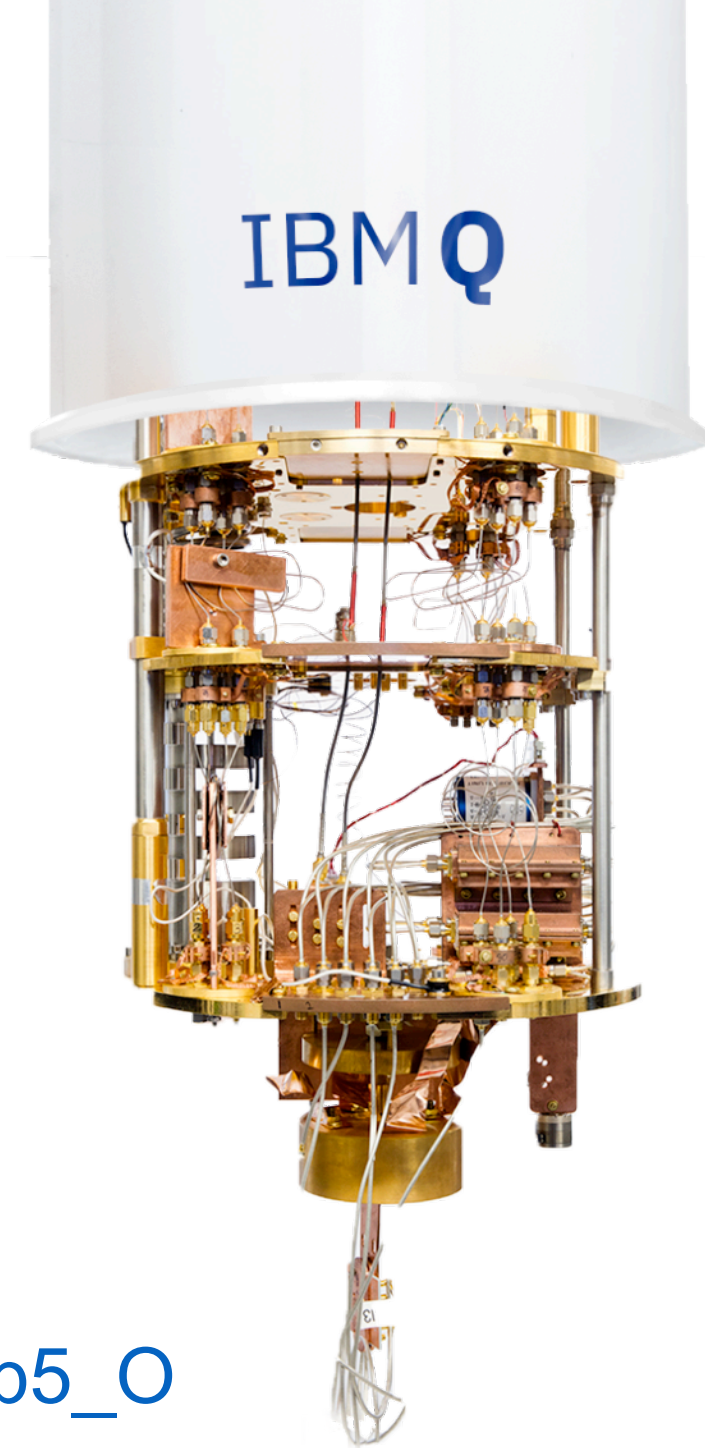
# Add a CX (CNOT) gate on control qubit 0 and target qubit 1
circuit.cx(0, 1)

# Map the quantum measurement to the classical bits
circuit.measure([0,1], [0,1])
```

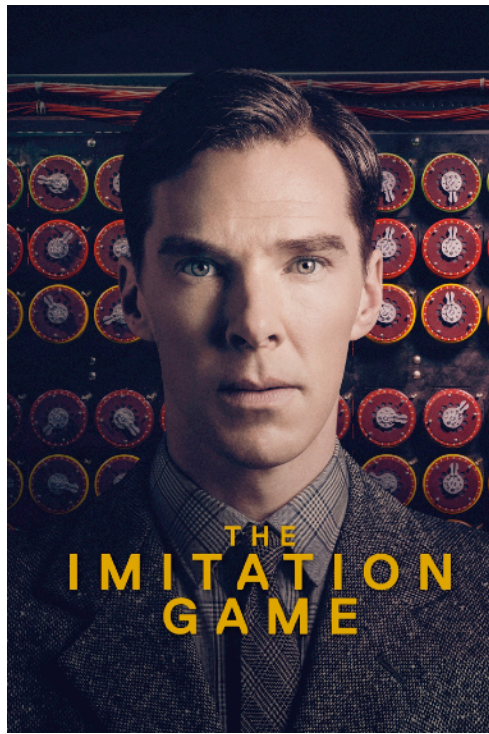


WOULD YOU LIKE TO KNOW MORE?

- <https://www.qosf.org>
- <https://quantum-computing.ibm.com/composer/docs/iqx/guide>
- <https://www.mooc-course.com/subject/quantum-computing>
- https://fosdem.org/2021/schedule/event/cloud_quantum_computing
- https://www.youtube.com/playlist?list=PLqXHJj6bcnY92luCrMt78ThcBOJ5Nb5_O
- <https://www.thequantumdaily.com/2020/05/15/7-quantum-computing-books-for-the-uninitiated>



MOVIES TO WATCH



The Imitation Game (2014)

Biography of Alan Turing, father of theoretical computer science and artificial intelligence, invented the "bombe" to decode Enigma machine of Nazi Germany during World War II.



Hidden Figures (2016)

Story about the African-American female mathematicians who served as "human computers" at NASA in 1961, who taught themselves FORTRAN in order to operate the IBM computer which was set to replace them.

Questions?

Kenneth Hoste

HPC system administrator

GHENT UNIVERSITY
DEPARTMENT OF INFORMATION AND
COMMUNICATION TECHNOLOGY (DICT)



kenneth.hoste@ugent.be - hpc@ugent.be



@kehoste - @hpcugent

<https://www.ugent.be/hpc>

