

LEVELS OF PARALLELISM AND HIGH PERFORMANCE COMPUTING

IEEE IGARSS 2021 Tutorial on Scalable Machine Learning with High Performance and Cloud Computing

DR. - ING. GABRIELE CAVALLARO
HIGH PRODUCTIVITY DATA PROCESSING RESEARCH GROUP
JÜLICH SUPERCOMPUTING CENTRE
WWW.GABRIELE-CAVALLARO.COM

IT ENVIRONMENTS

Key Priorities


Embedded system



 Power consumption and price


Desktop Computing



 Performance / price ration

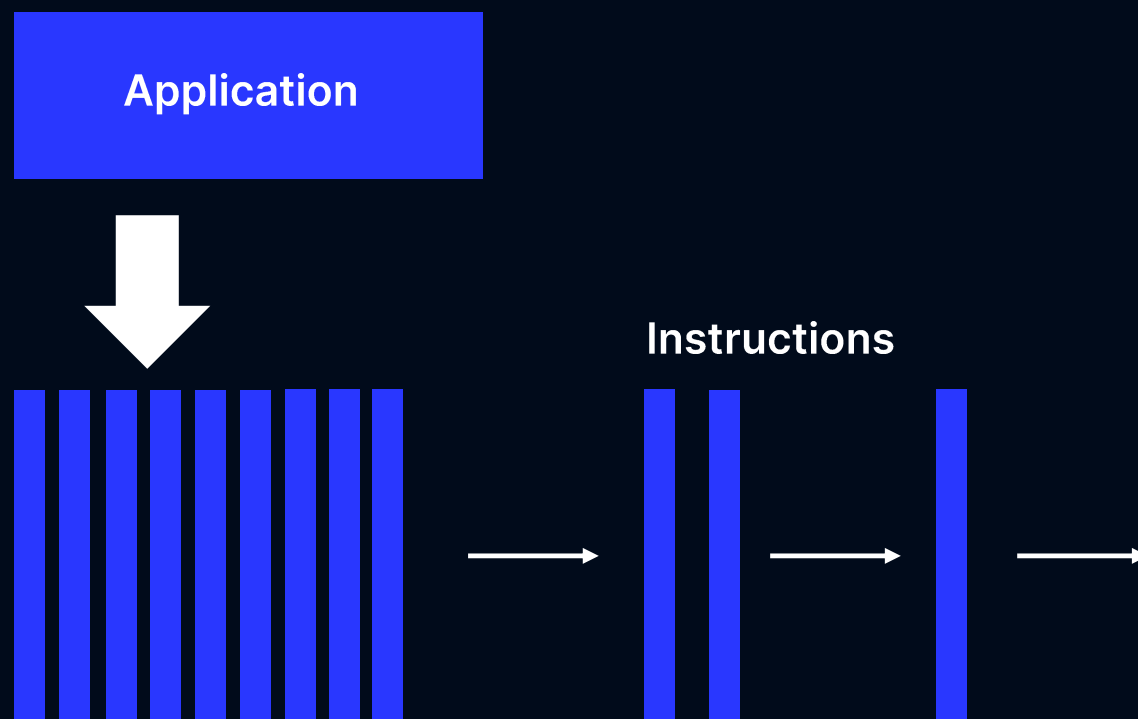
Servers



 Maximum performance and maintainability reliability

HOW TO BECOME FASTER?

Split into instructions



HOW TO BECOME FASTER?

Three Alternative Ways

Work Harder



Clock Speed

Work Smarter



Optimization, Caching


Get Help



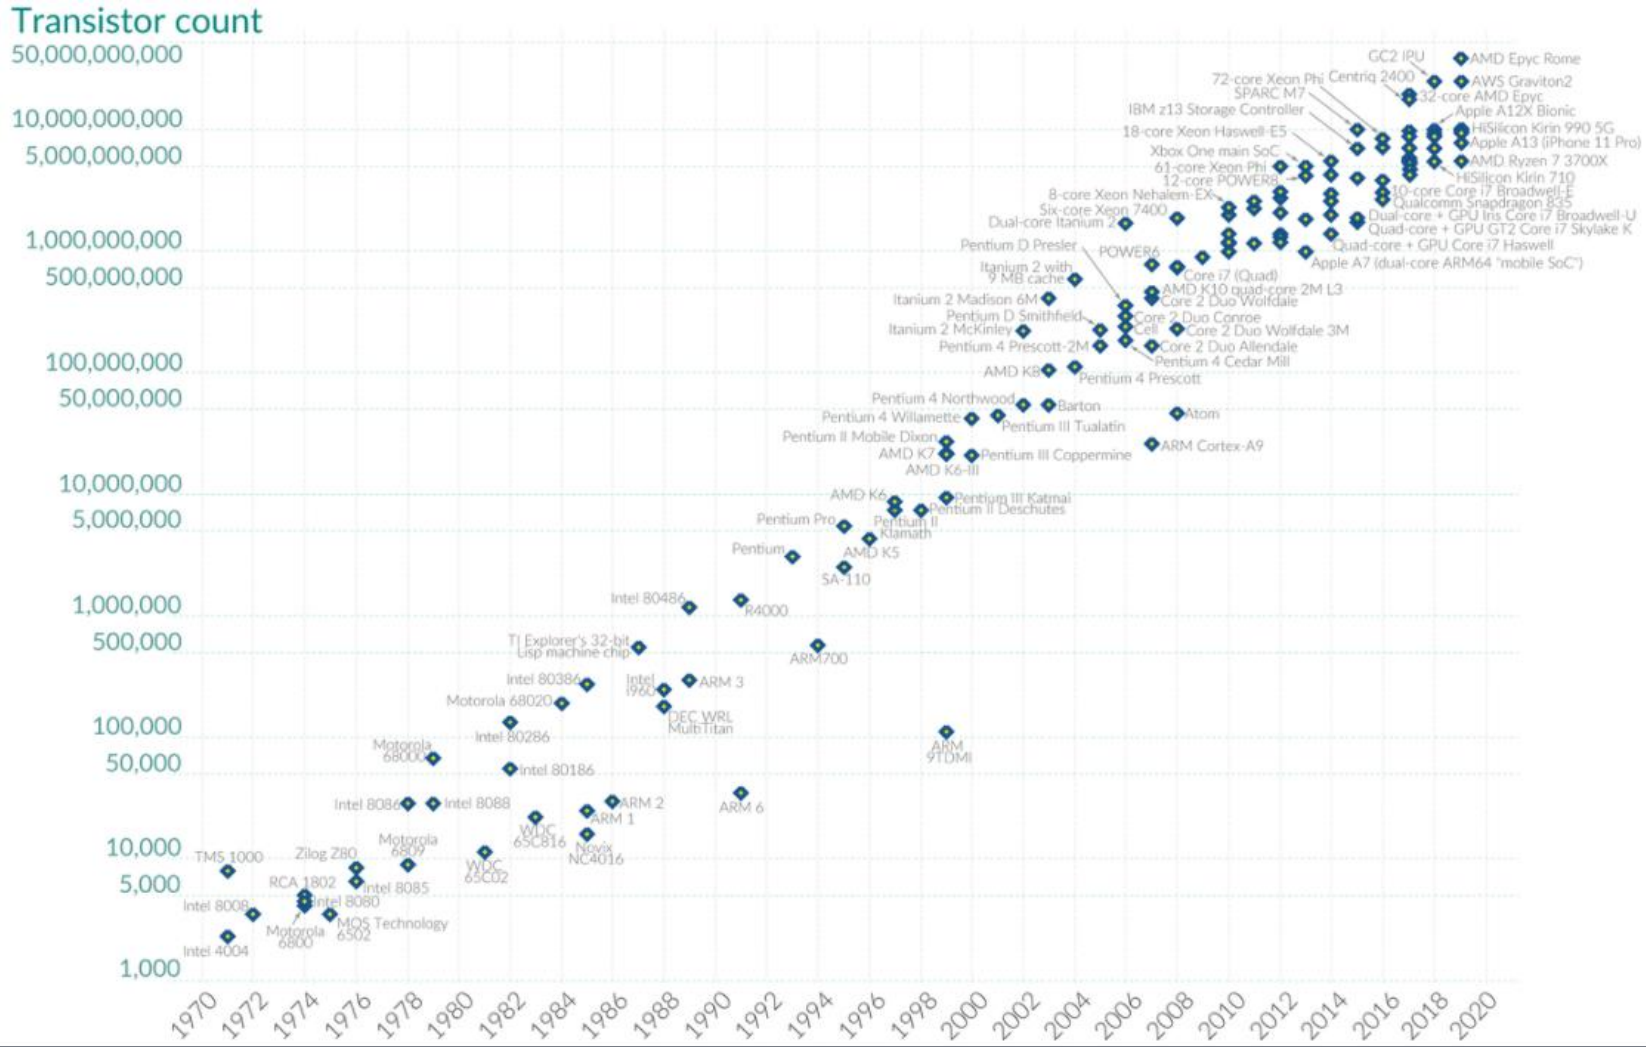
Parallelization

VON NEUMANN ARCHITECTURE

Machine Model




MOORE'S LAW



THE END OF DENNARD SCALING


Why haven't clock speeds increased, even though transistors have continued to shrink?



WILL MOORE'S LAW END?

Paradigm Shift

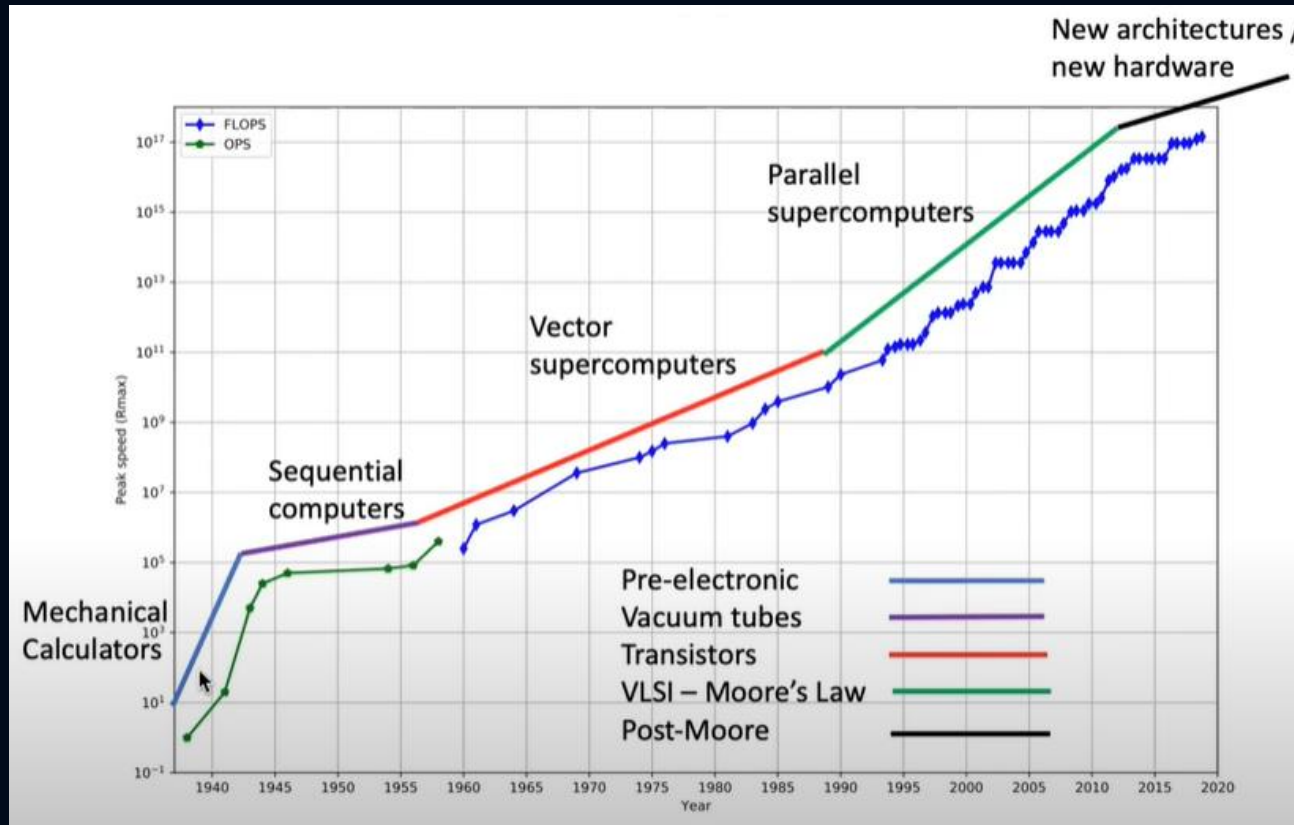
**“There’s no getting around the fact that
we build these things out of atoms”**



Gordon Moore

WHY SUPERCOMPUTER PERFORMANCE KEEP INCREASING?

Parallel computing is going mainstream




Rob Schreiber, High Performance Computing: Beyond Moore's Law
<https://www.youtube.com/watch?v=LOf57fdxln4&t=258s>

THE EASY TIMES HAVE GONE

Responsibility for better performance is on the software developers

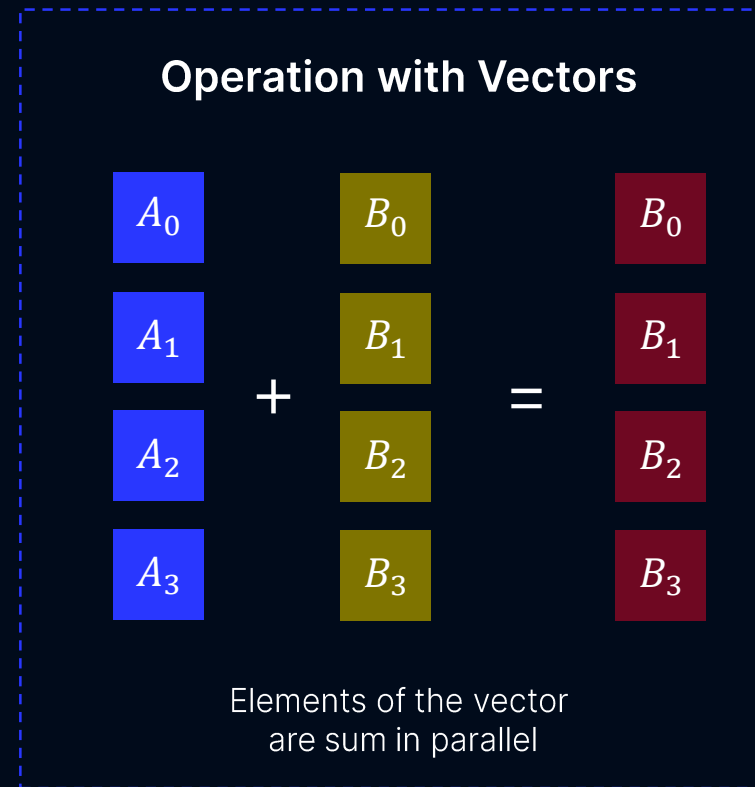
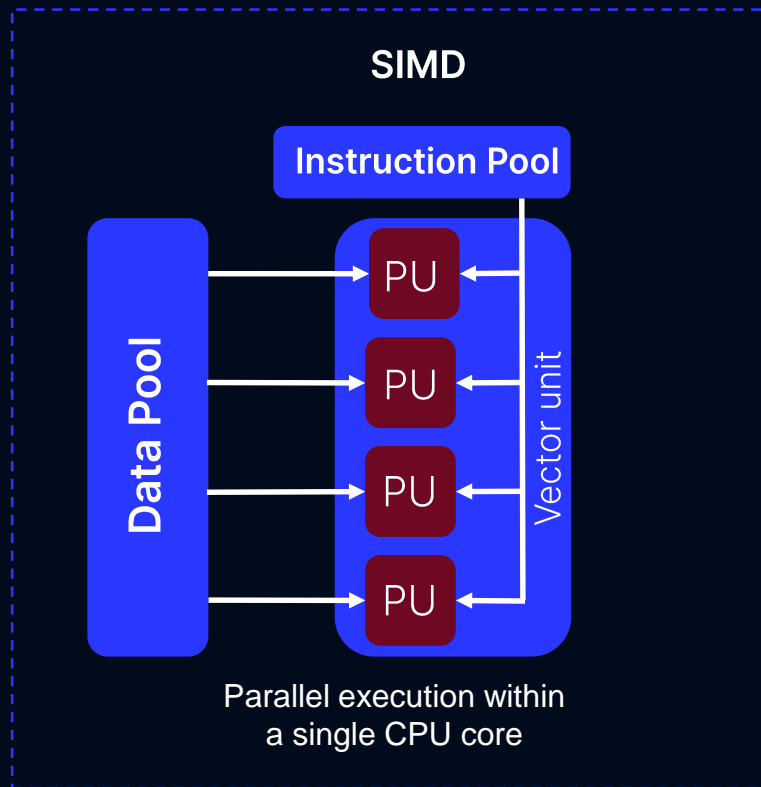


PARALLEL COMPUTING



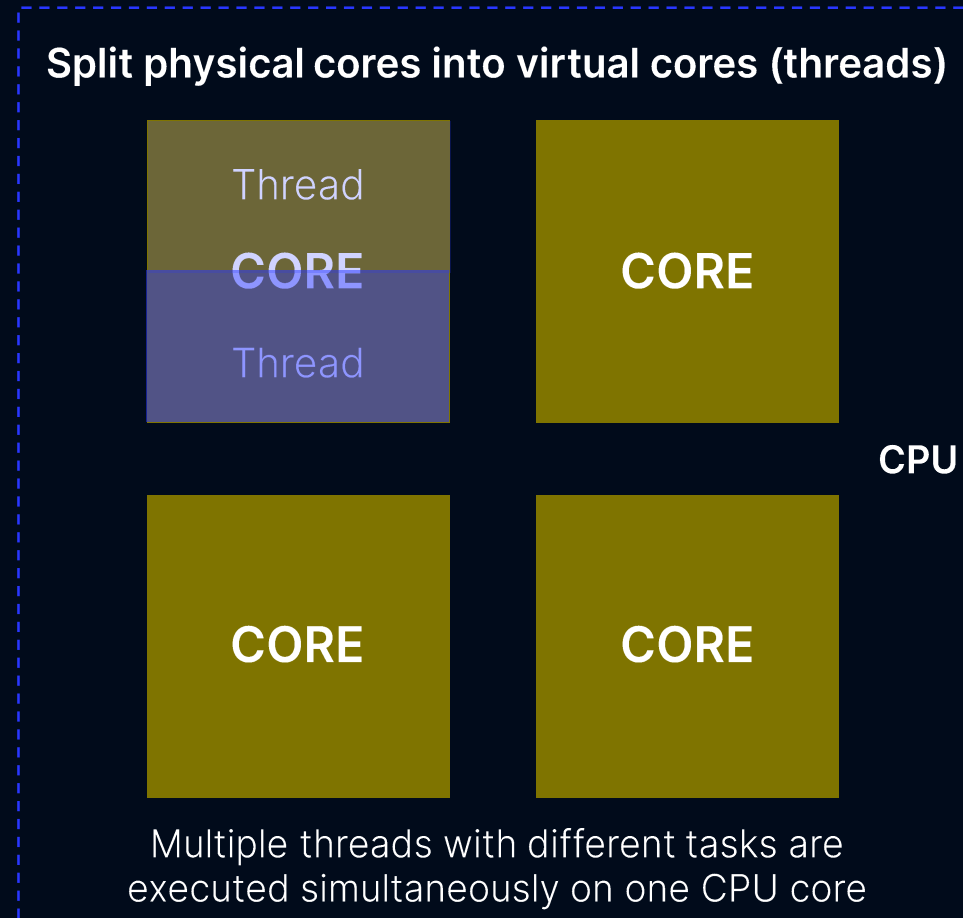
SINGLE INSTRUCTION MULTIPLE DATA (SIMD)

In-core parallelism



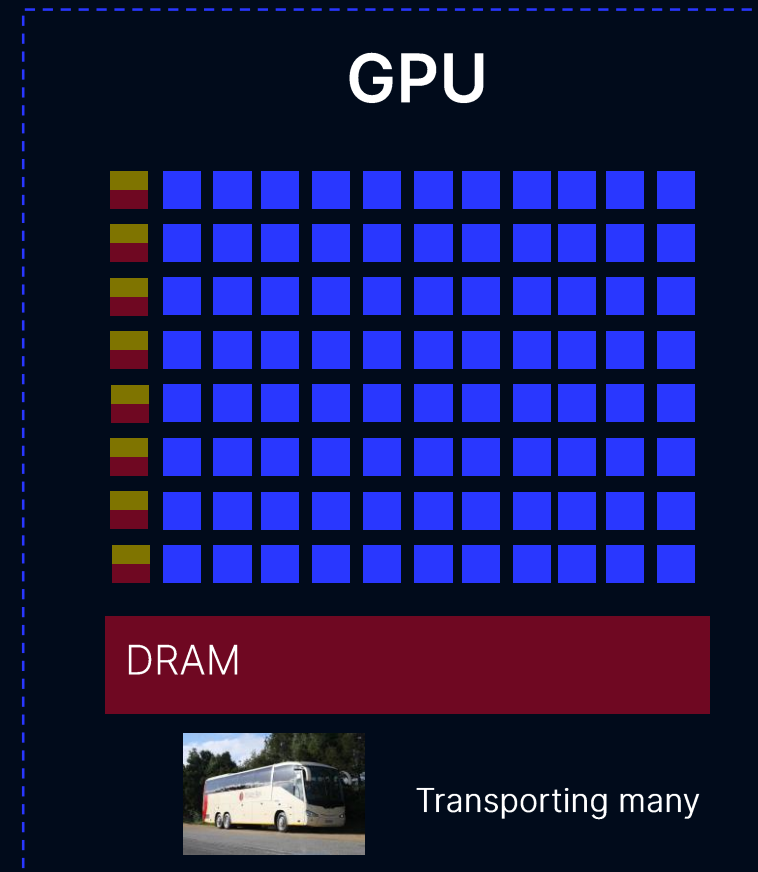
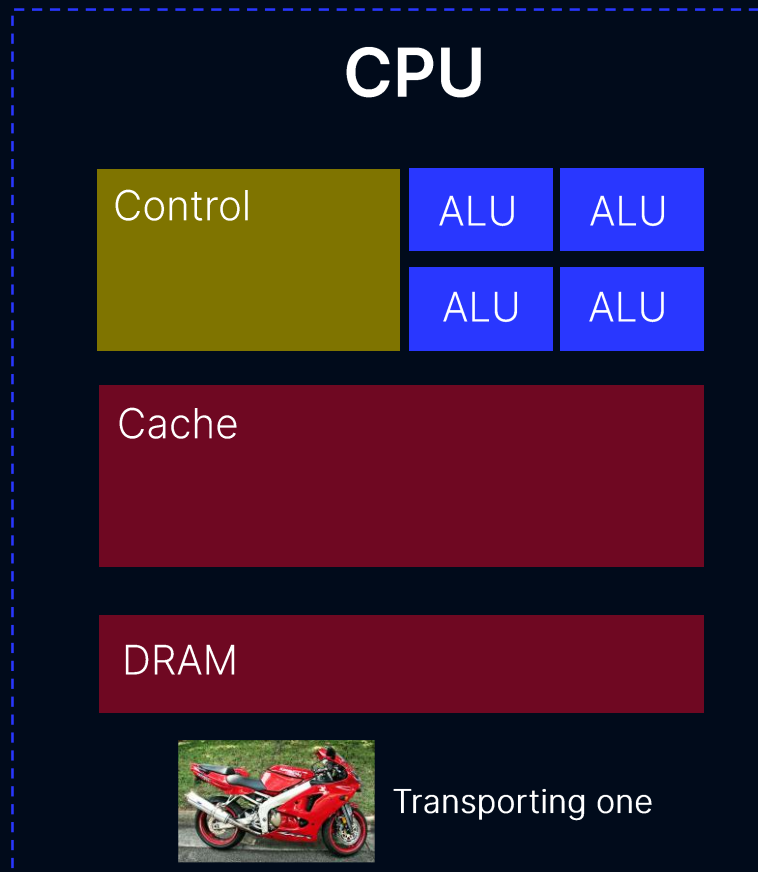
SIMULTANEOUS MULTITHREADING (SMT)

In-core parallelism



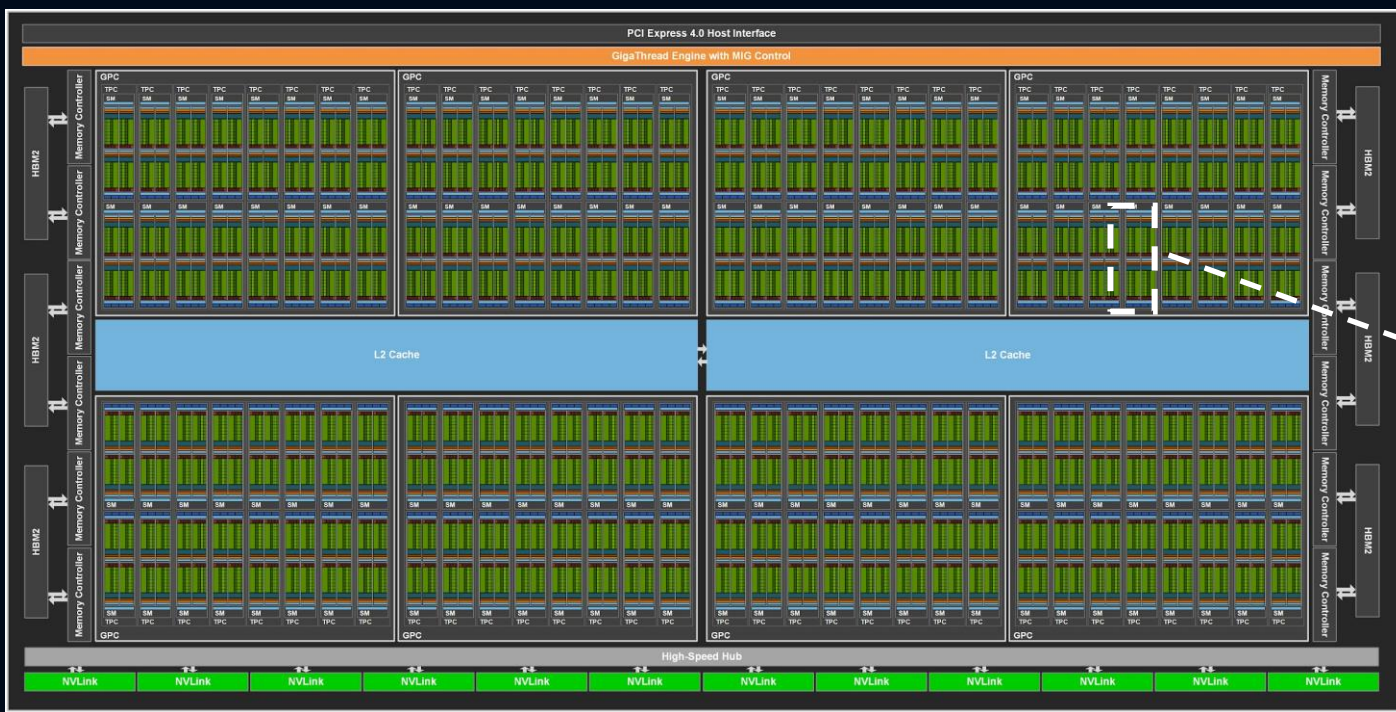
GRAPHICS PROCESSING UNIT (GPU) VS CPU

Made of many simple Cores



NVIDIA AMPERE GPU ARCHITECTURE

Streaming Multiprocessors (SMs)



NVIDIA Ampere Architecture
(128 SMs)

Levels of Parallelism and HPC

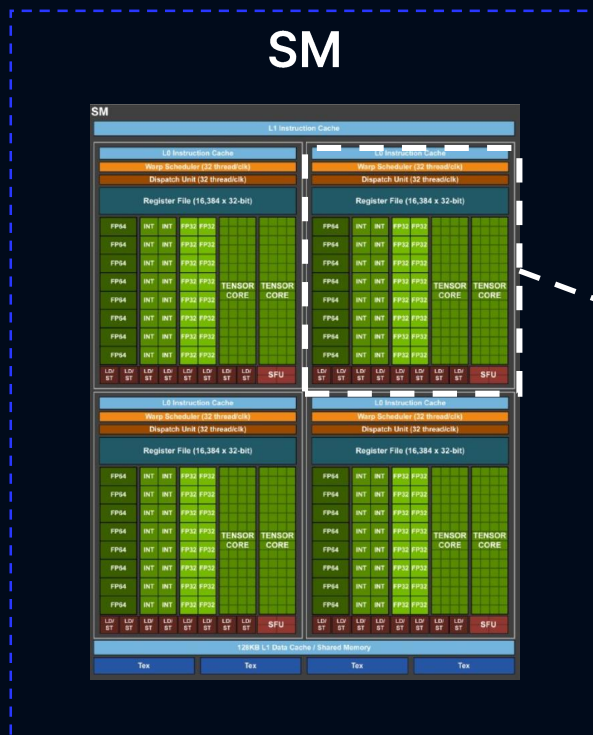
Ronny Krashinsky, et al., NVIDIA Ampere Architecture In-Depth
<https://developer.nvidia.com/blog/nvidia-ampere-architecture-in-depth/>

SM \cong CPU core



NVIDIA AMPERE GPU ARCHITECTURE

Streaming Processors (SPs)



SP (CUDA core)

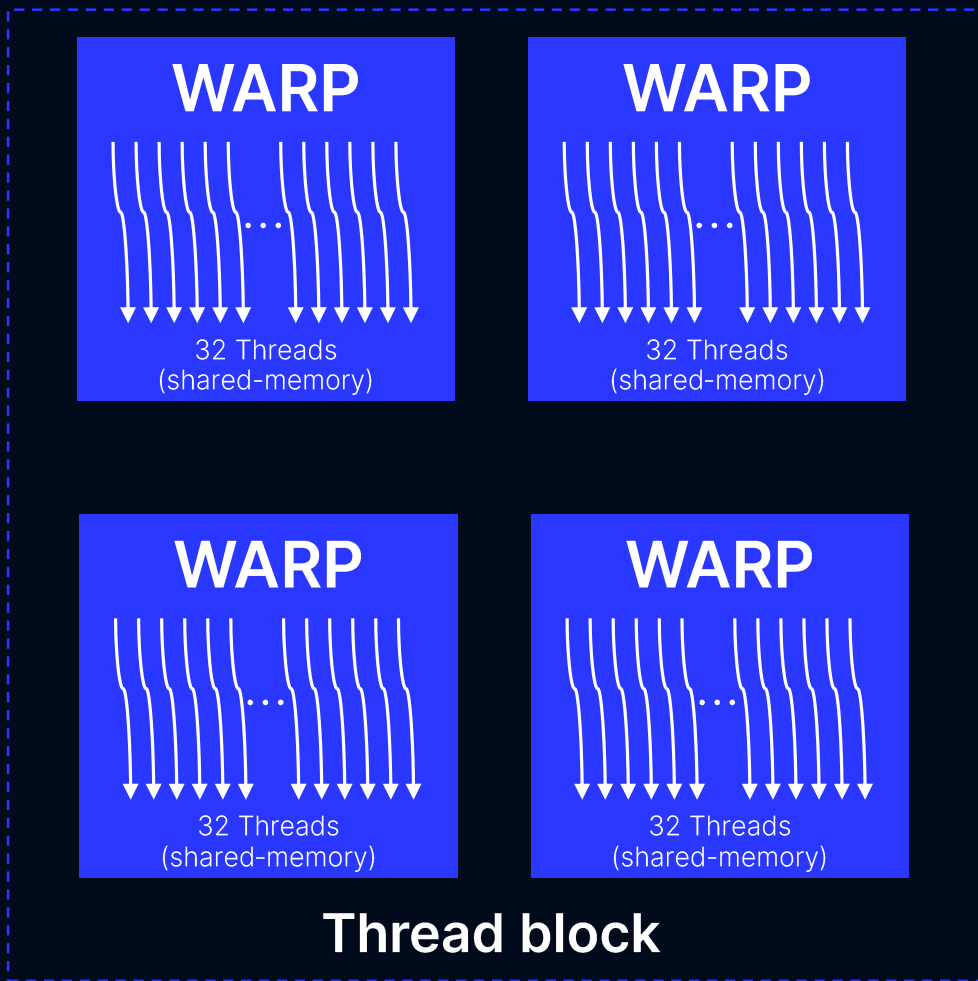


Compute elements:
16 32-bit integer point units
16 32-bit floating point units
8 64-bit floating point units

Ronny Krashinsky, et al., NVIDIA Ampere Architecture In-Depth
<https://developer.nvidia.com/blog/nvidia-ampere-architecture-in-depth/>

SINGLE INSTRUCTION MULTIPLE THREADS (SIMT)

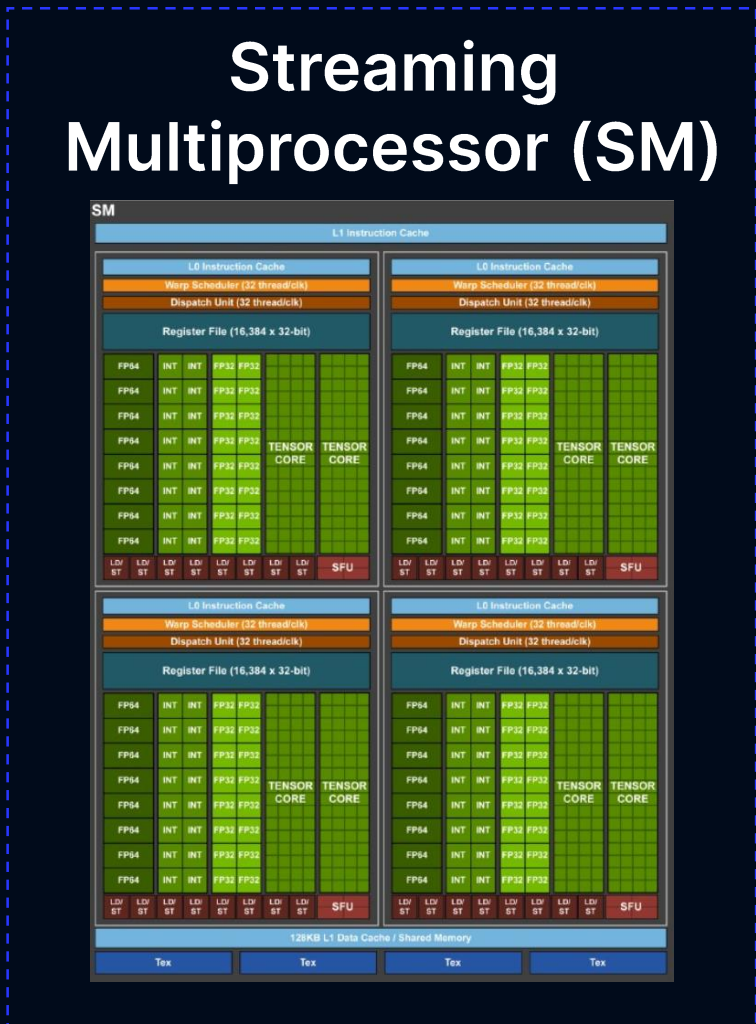
In-processor parallelism - SIMT = SIMD + SMT



The WARPS are executed simultaneously by a SM

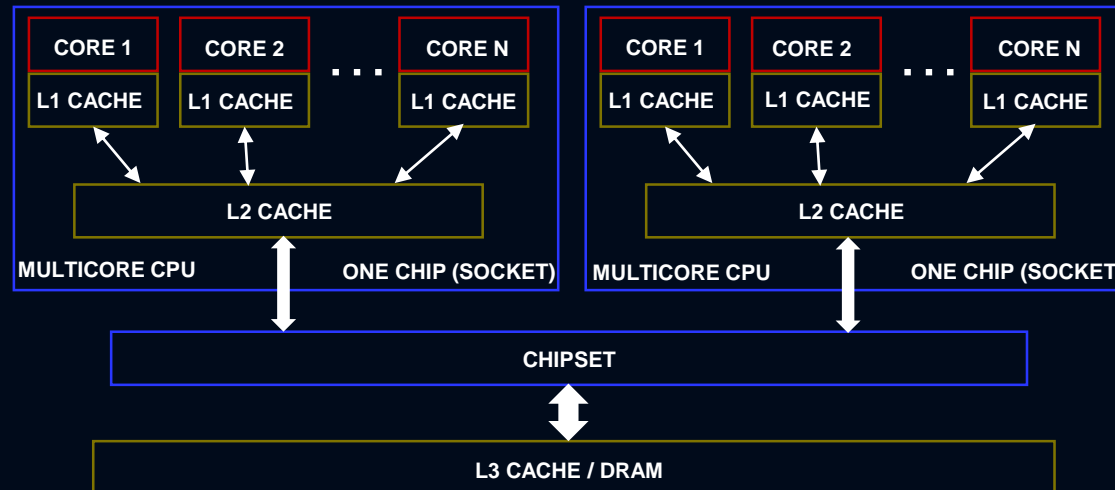


The threads of a WARP execute the same instruction (as SIMD)



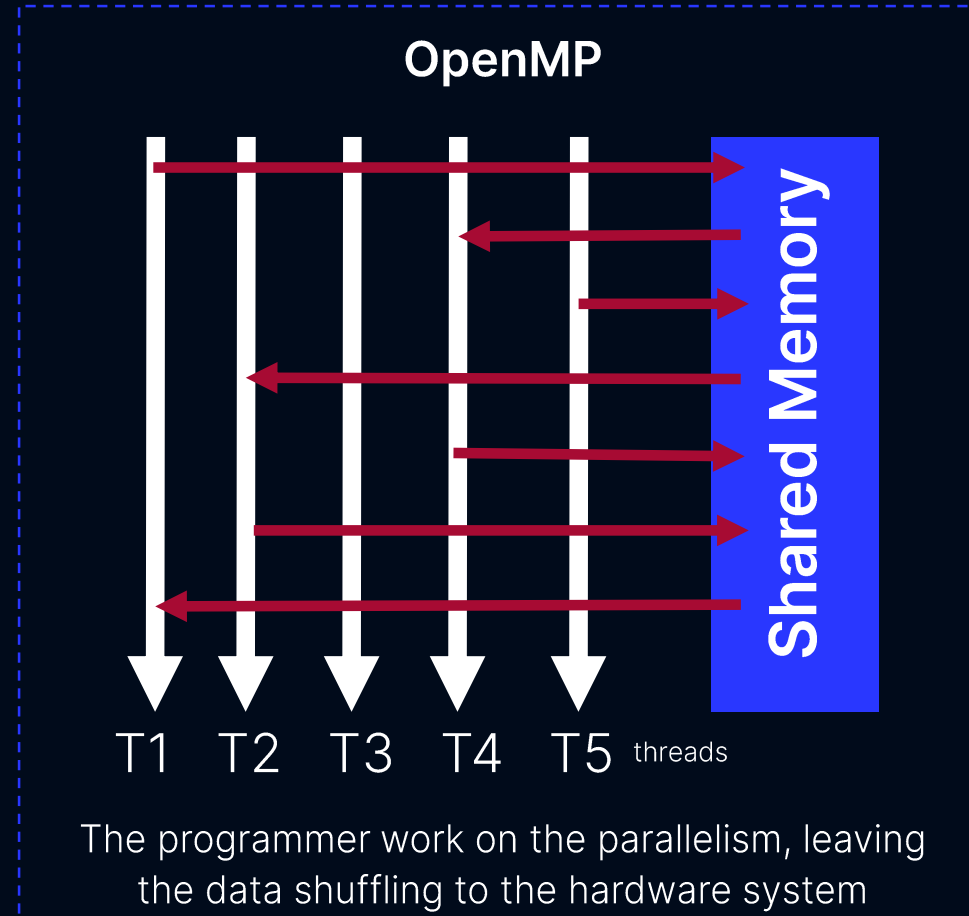
SHARED-MEMORY ARCHITECTURE

Single computer



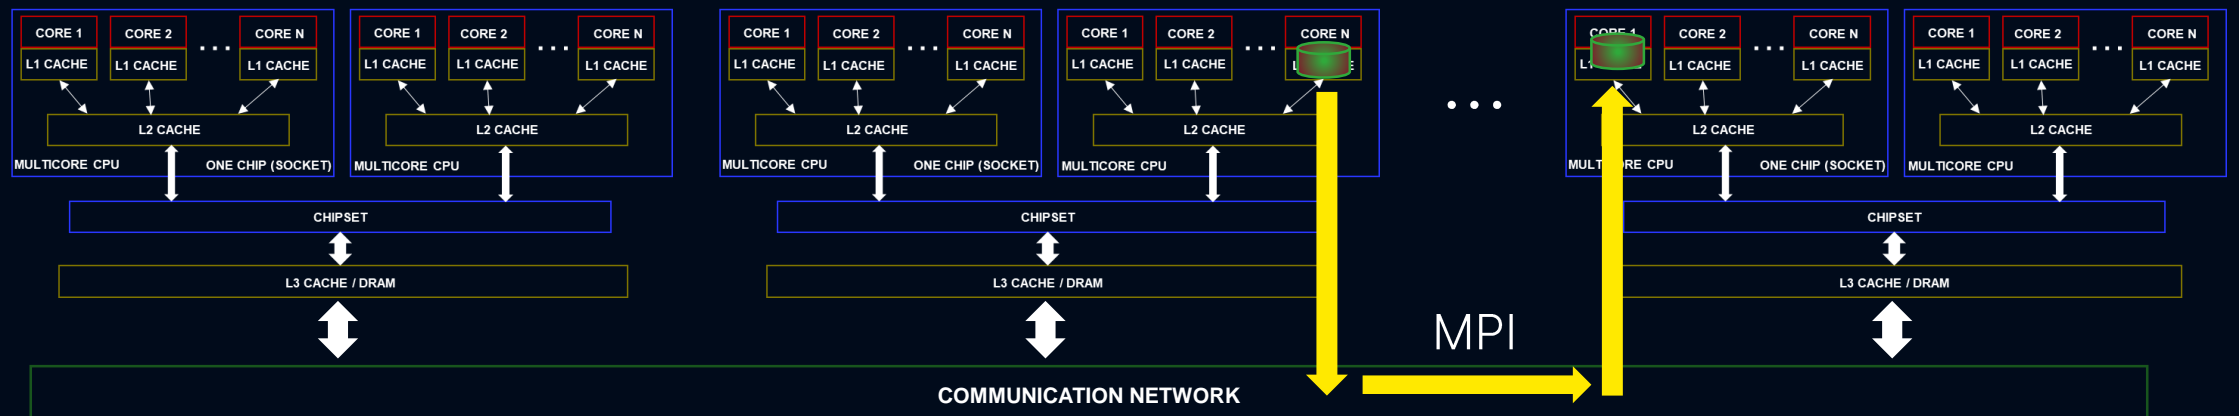
PARALLEL PROGRAMMING MODEL: OPENMP

Single Computer



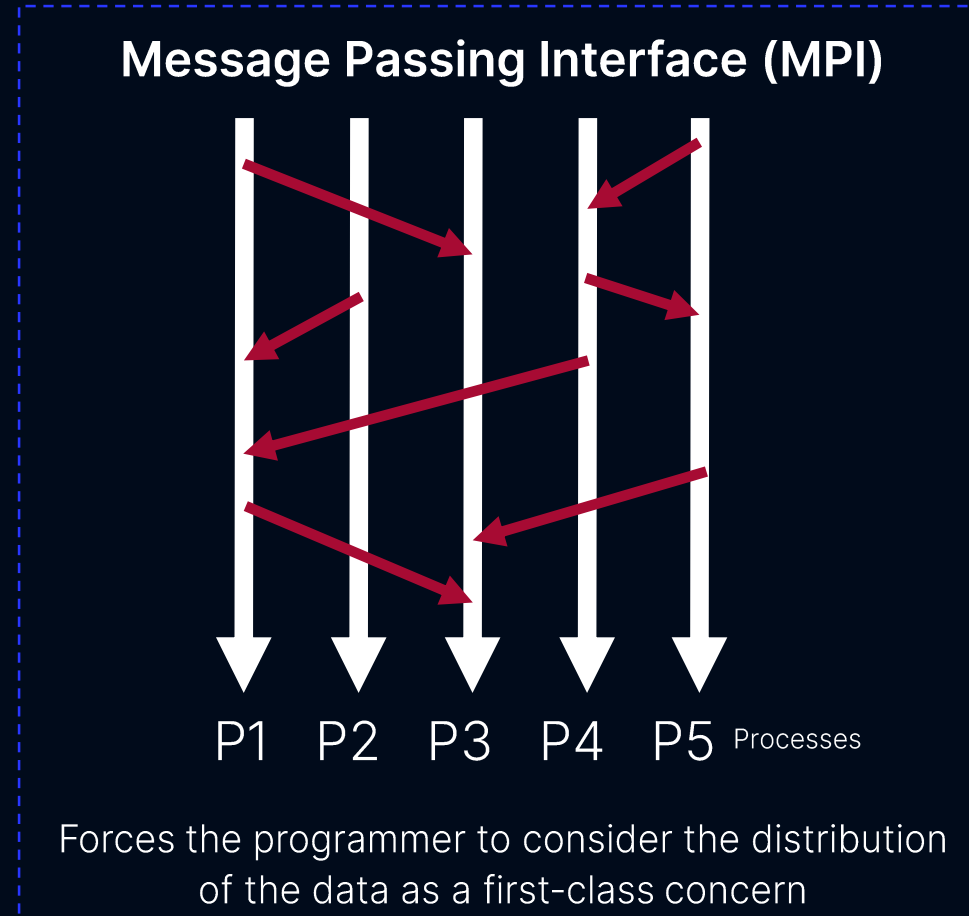
DISTRIBUTED-MEMORY ARCHITECTURE

Multiple Computers



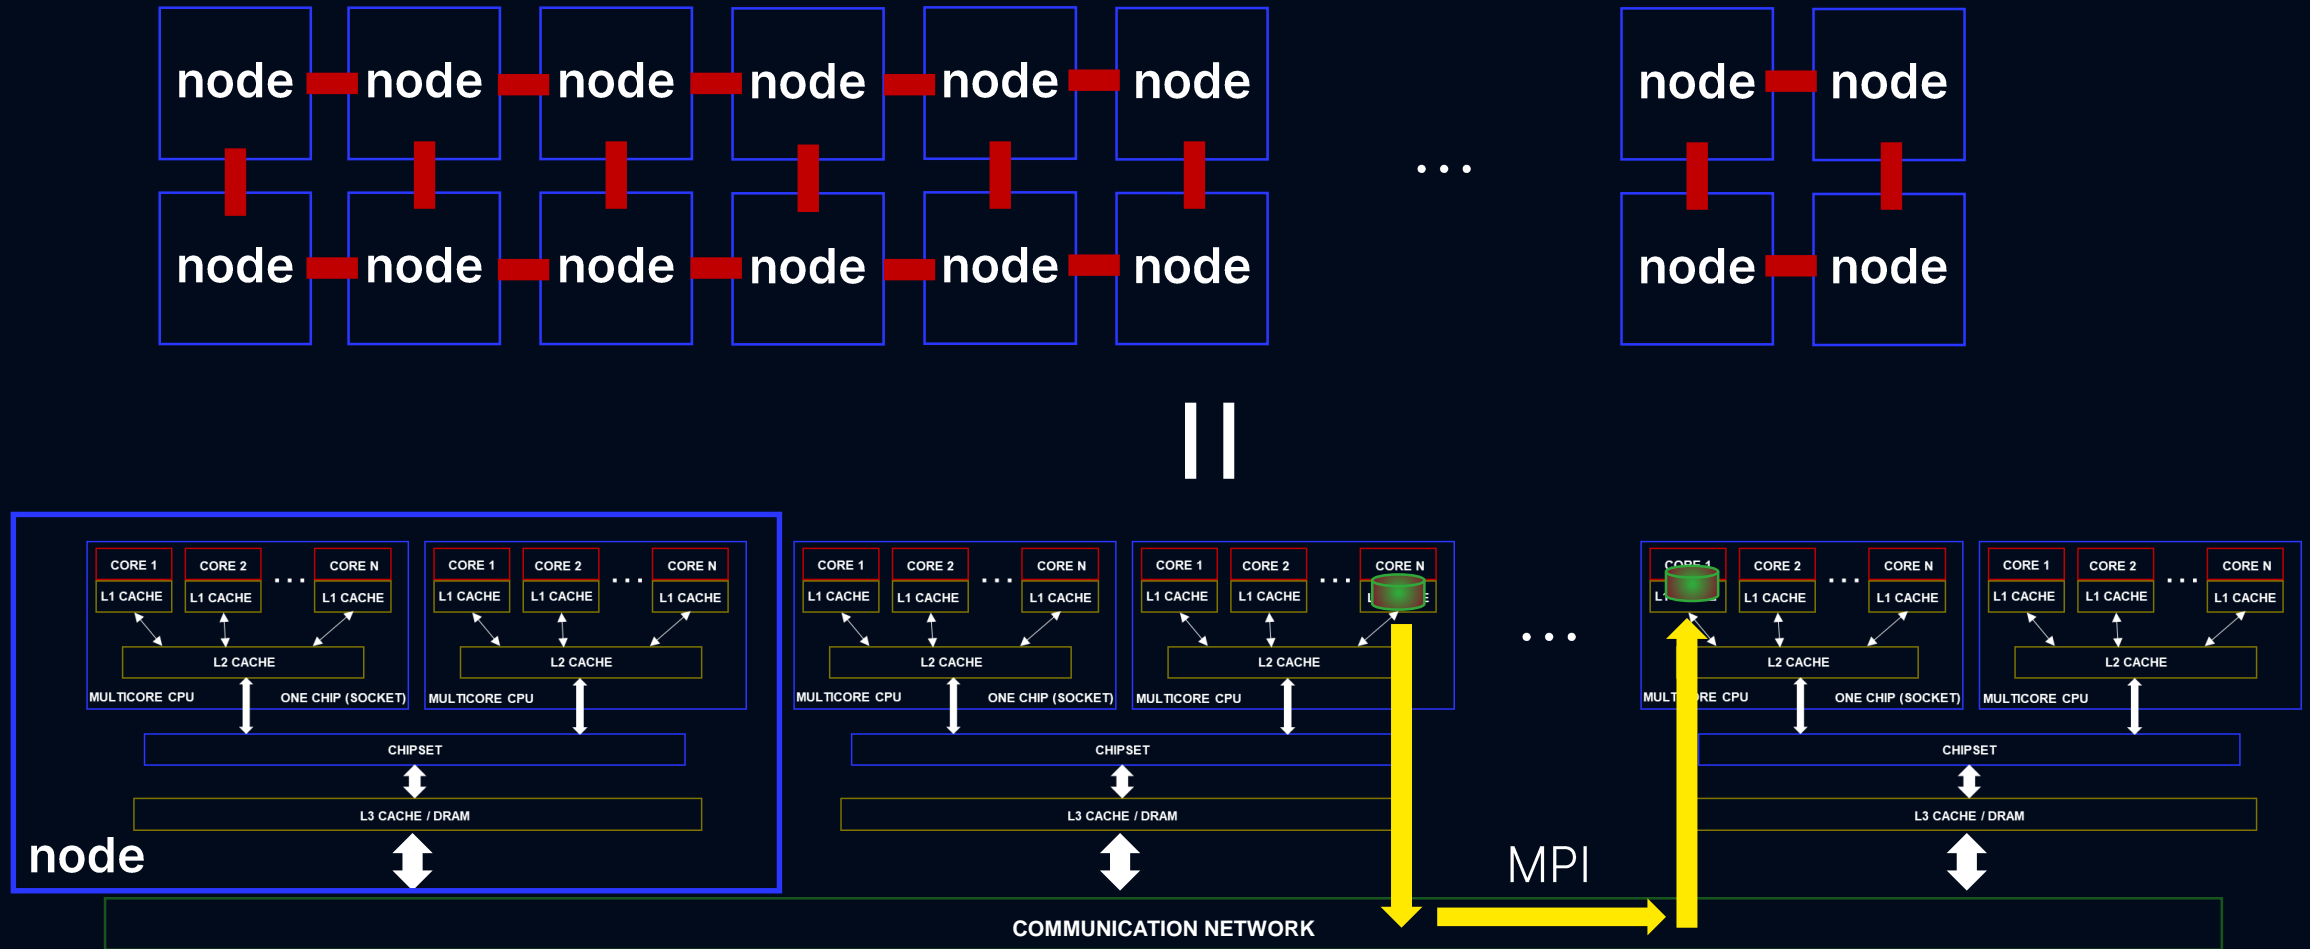
PARALLEL PROGRAMMING MODEL: MPI

Multiple Computers




WHAT IS A SUPERCOMPUTER?

Mixture of shared-memory and distributed-memory architectures



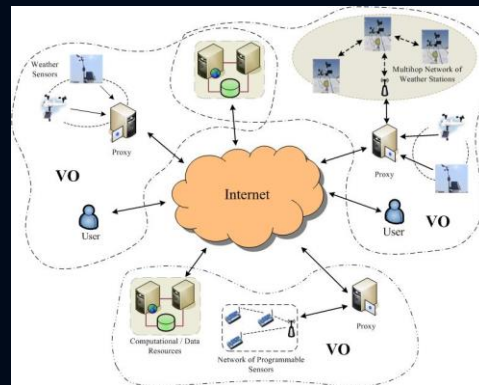
WHAT IS NOT A SUPERCOMPUTER?

Computer cluster



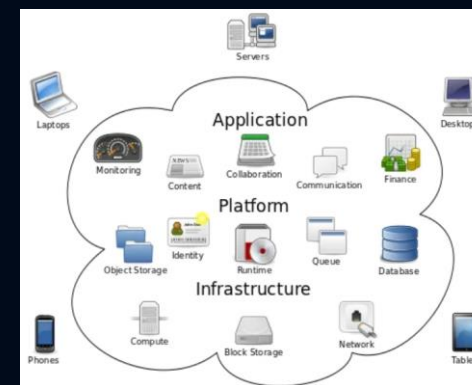
Many computers bound together locally

Grid computing



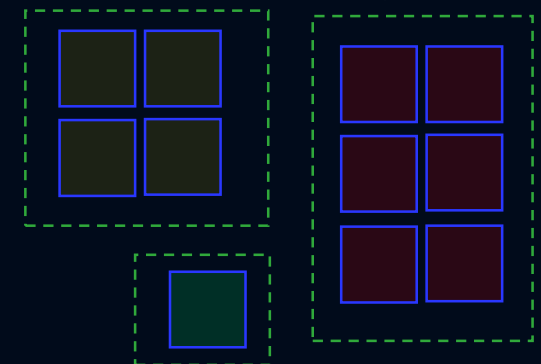
Many clusters working as some kind of supercomputer

Cloud computing



Virtual machines in compute center(s)

High Throughput Computing



Many independent computations


Difference between Grid Computing and Cloud Computing
<http://www.differencebetween.net/>

TOP500 LIST

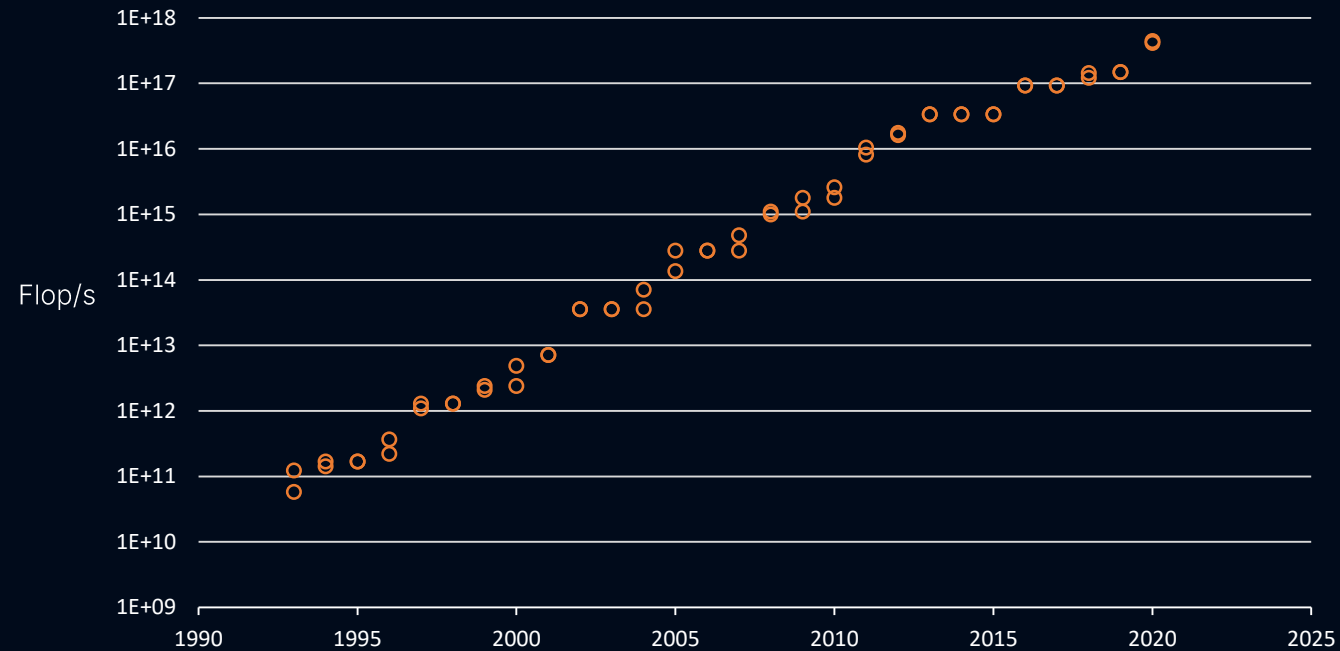
November 2020



<https://www.top500.org/>



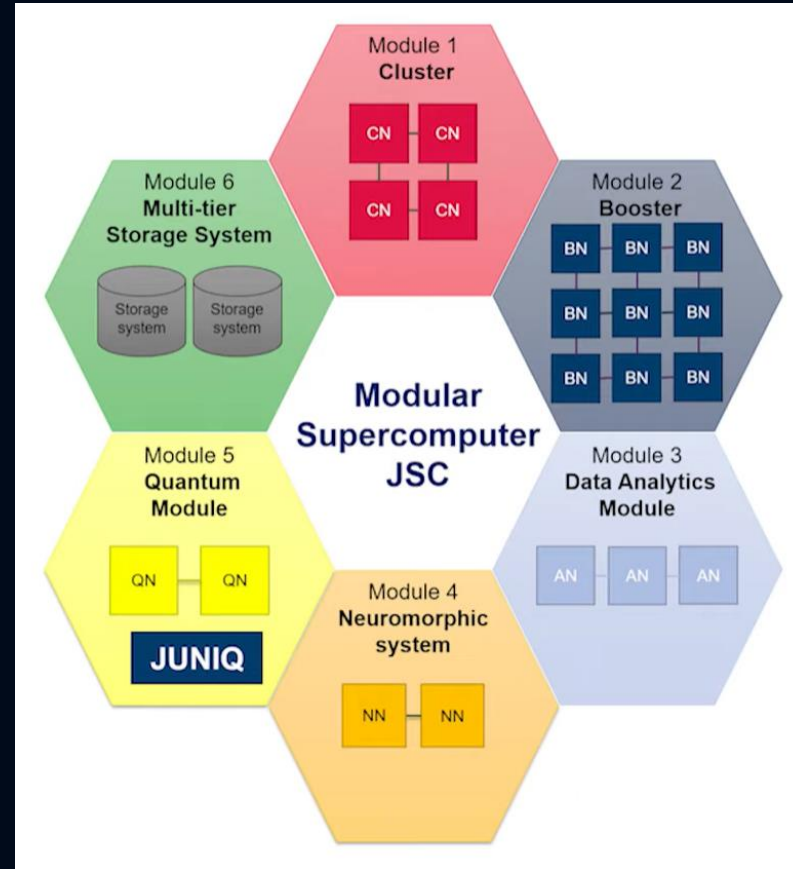
A RACE TOWARD EXASCALE COMPUTING



TOP500, Performance Development
<https://www.top500.org/statistics/perfdevel/>

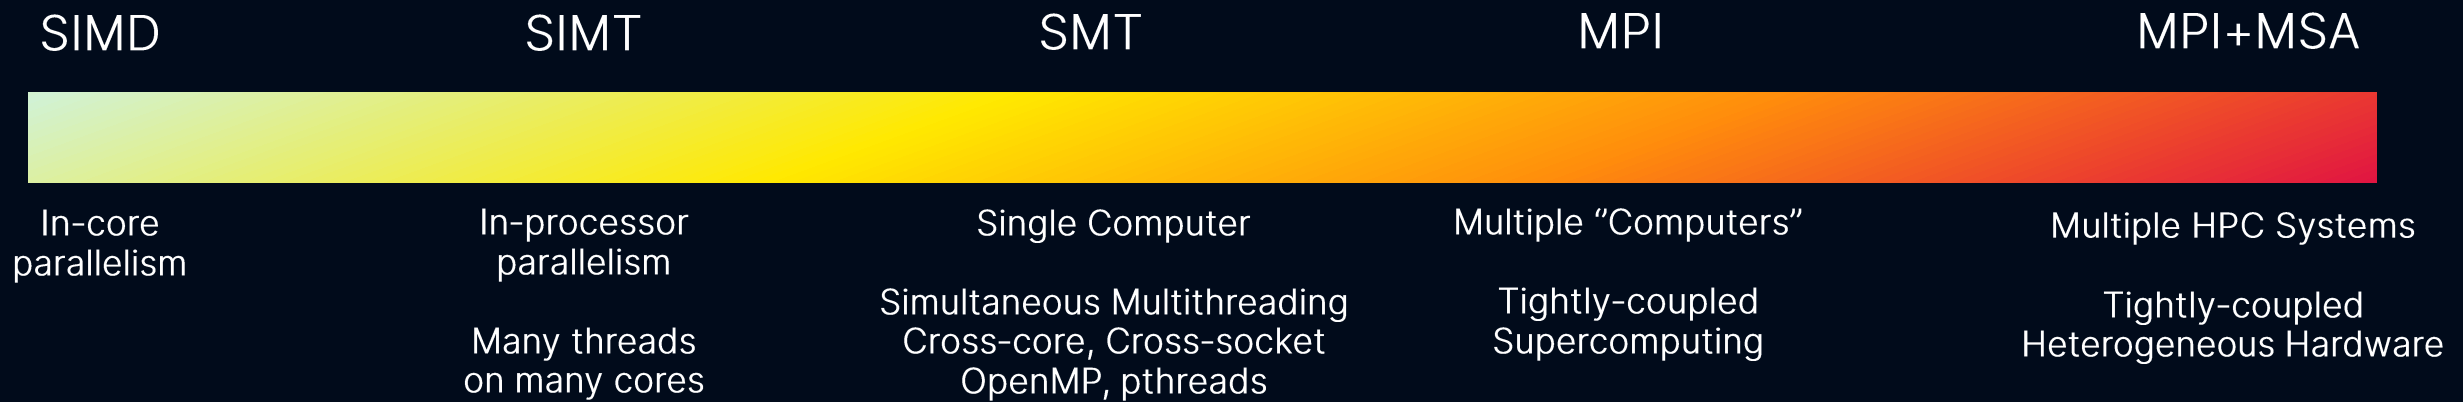
MODULAR SUPERCOMPUTING ARCHITECTURE (MSA)

Heterogeneous HPC clusters (modules) within a single system

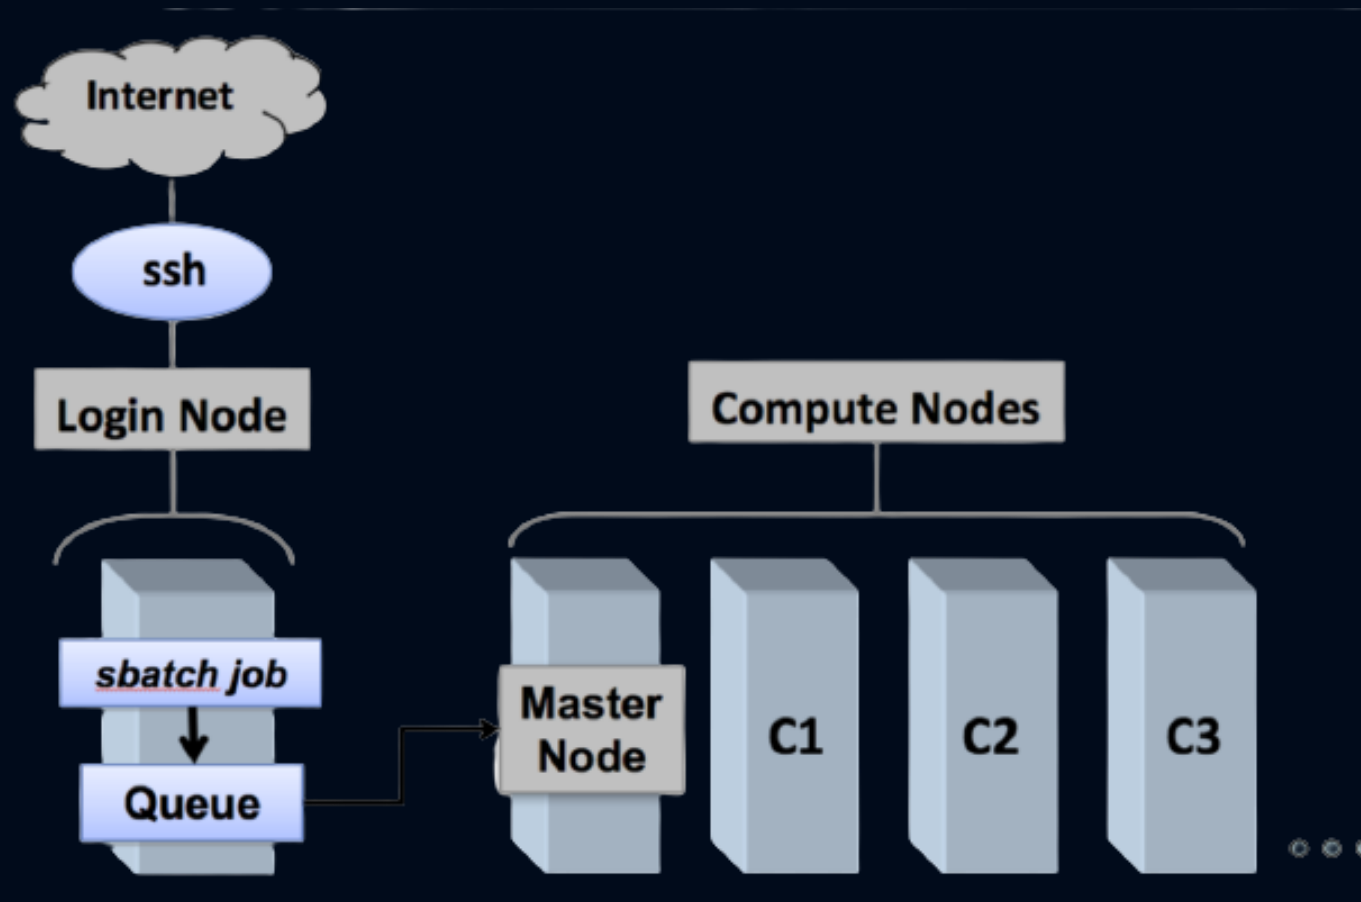


REVIEW ON HARDWARE LEVELS OF PARALLELISM

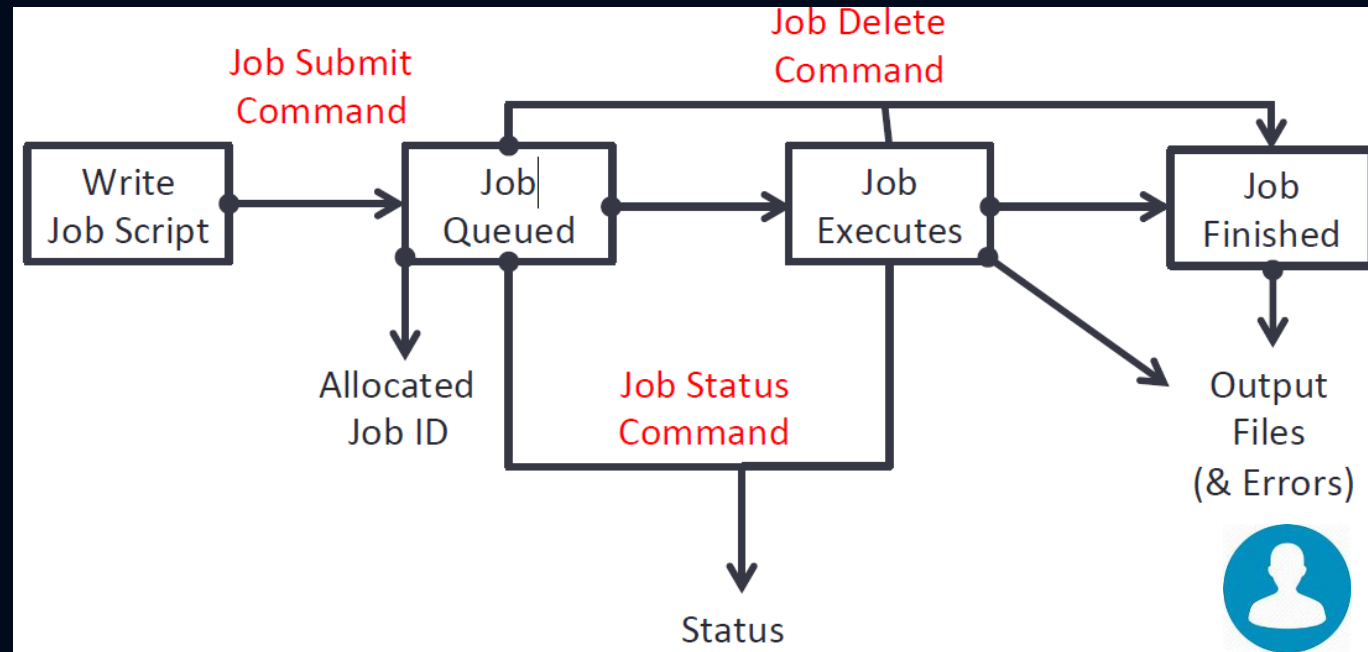
Best performance is achieved with a combination of them!



ANATOMY OF A SUPERCOMPUTER




SUPERCOMPUTER USAGE MODEL



JUSUF

Jülich Support for Fenix



<https://fenix-ri.eu/>

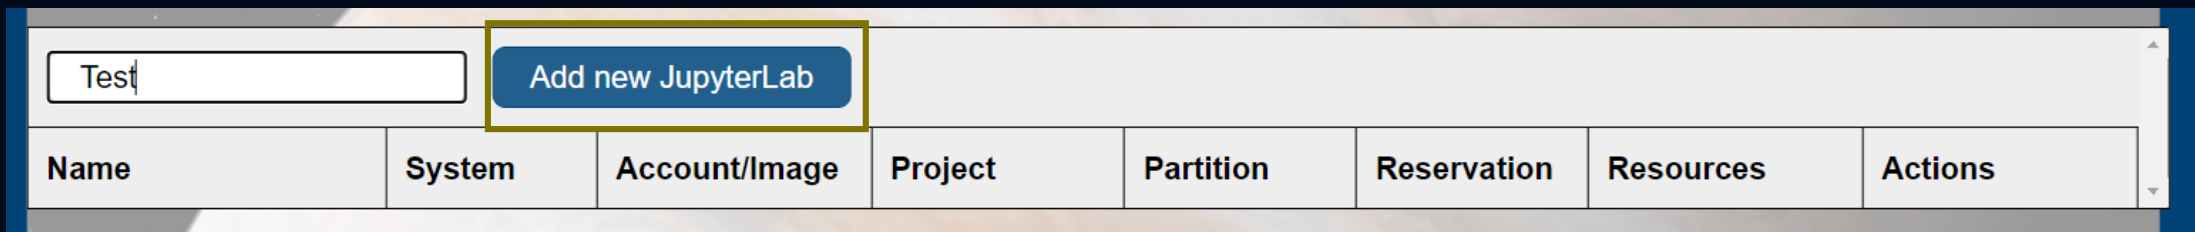
Partition with GPUs

61 NVIDIA V100 GPU with 16 GB Memory

<https://apps.fz-juelich.de/jsc/hps/jusuf/cluster/configuration.html>

PRE-PRACTICAL – STEP 1

Go to <https://jupyter-jsc.fz-juelich.de/>, login and add a new JupyterLab



The screenshot shows a web interface for adding a new JupyterLab instance. At the top, there is a text input field containing the word "Test" and a blue button labeled "Add new JupyterLab". Below this is a table with the following columns: Name, System, Account/Image, Project, Partition, Reservation, Resources, and Actions.

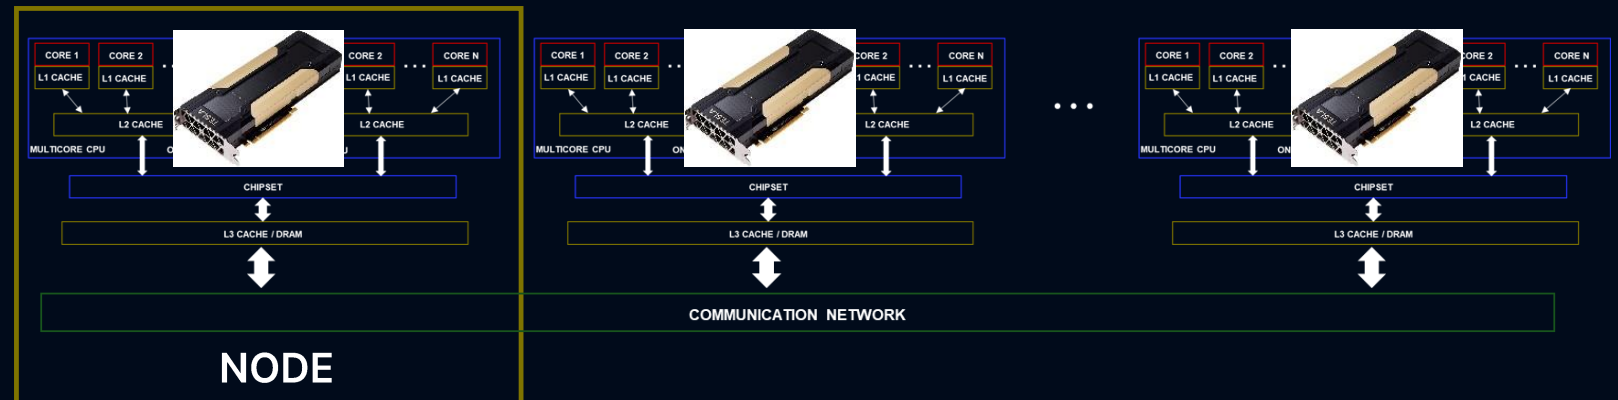
Name	System	Account/Image	Project	Partition	Reservation	Resources	Actions
------	--------	---------------	---------	-----------	-------------	-----------	---------

STEP 2

Select the options and Start

JupyterLab Options

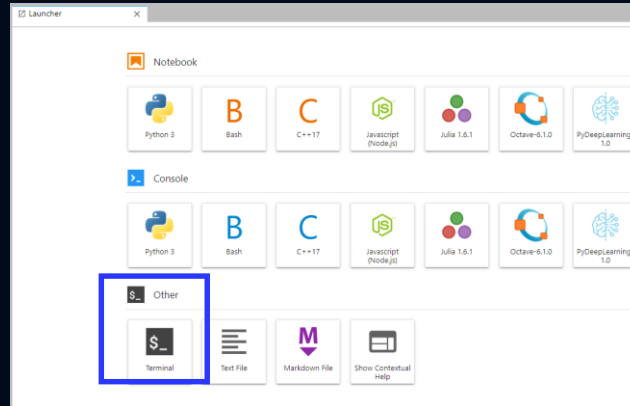
Version	JupyterLab
System	JUSUF
Account	cavallaro1
Project	training2118
Partition	gpu
Nodes [1, 46]	1
Runtime (min) [10, 1440]	30
GPUs [1, 1]	1
Start	



STEP 3

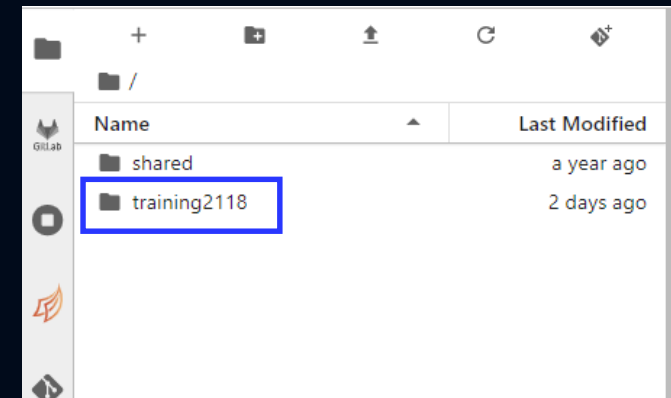
Obtain the tutorial folder

- Open terminal



- Run this command and the tutorial folder training2118 will appear in the navigator on the left

```
cavallaro1@jsfc041:~  
[cavallaro1@jsfc041 ~]$ ln -s /p/project/training2118/
```



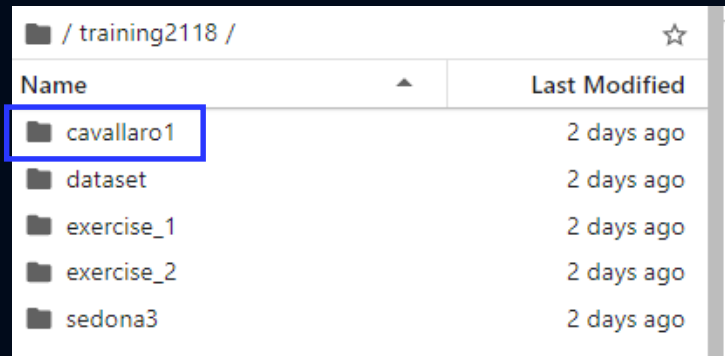
STEP 4

Create your own folder

- Run the following commands in the terminal to navigate to the tutorial folder “training2118” and create your own folder

```
cavallaro1@jsfc041:/p/projec X  
[cavallaro1@jsfc041 ~]$ cd /p/project/training2118/  
[cavallaro1@jsfc041 training2118]$ mkdir $USER
```

- You can check if the folder was created by looking at the navigator on the left



Name	Last Modified
cavallaro1	2 days ago
dataset	2 days ago
exercise_1	2 days ago
exercise_2	2 days ago
sedona3	2 days ago

STEP 5

Download Jupyter notebook

- Go to <https://www.gabriele-cavallaro.com/teaching/tutorial-igarss2021>
- Download the Jupyter notebook "show_resources.ipynb" of Lecture 2

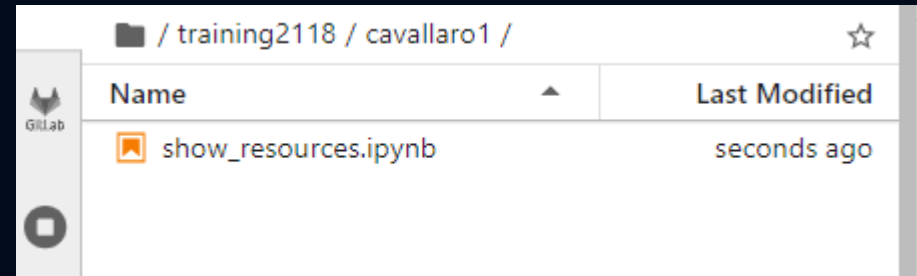
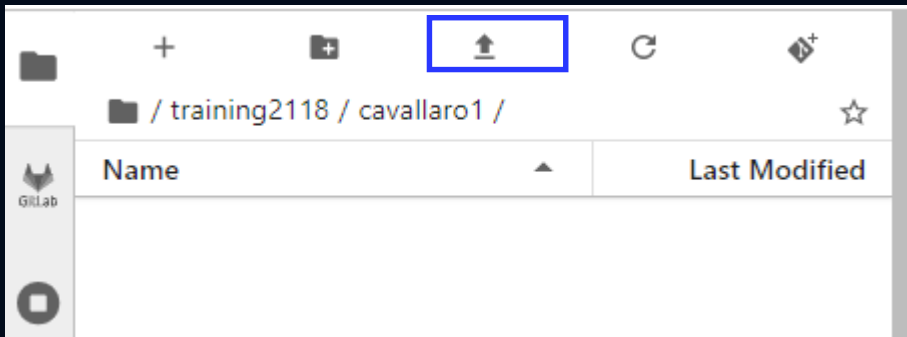
Lecture 2: Levels of Parallelism and High Performance Computing

Get lecture

[Get notebook](#)

STEP 6

Upload the Jupyter notebook in your own folder



STEP 7

Run the Jupyter notebook

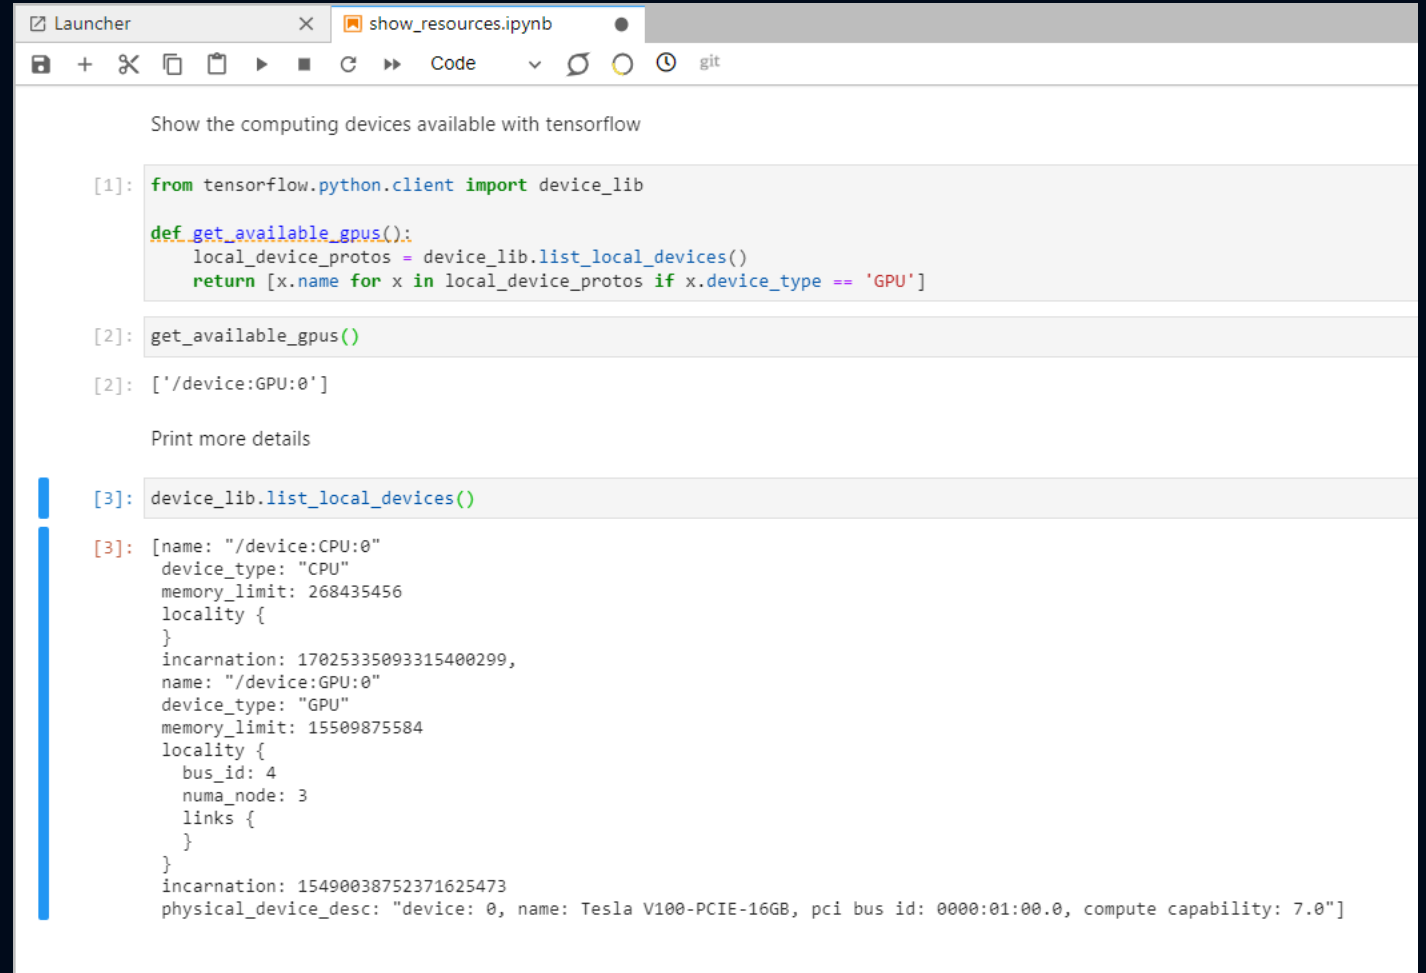
- Select the kernel PyDeepLearning-1.0

Select Kernel

Select kernel for: "show_resources.ipynb"

PyDeepLearning-1.0 ▾

No Kernel Select



The screenshot shows a Jupyter notebook window titled 'show_resources.ipynb'. The code in the notebook is as follows:

```
Show the computing devices available with tensorflow

[1]: from tensorflow.python.client import device_lib

def get_available_gpus():
    local_device_protos = device_lib.list_local_devices()
    return [x.name for x in local_device_protos if x.device_type == 'GPU']

[2]: get_available_gpus()

[2]: ['/device:GPU:0']

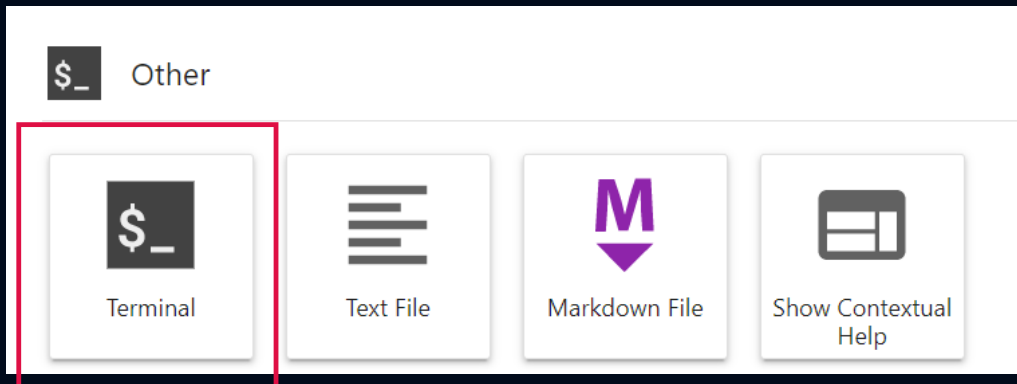
Print more details

[3]: device_lib.list_local_devices()

[3]: [name: "/device:CPU:0"
device_type: "CPU"
memory_limit: 268435456
locality {
}
incarnation: 17025335093315400299,
name: "/device:GPU:0"
device_type: "GPU"
memory_limit: 15509875584
locality {
  bus_id: 4
  numa_node: 3
  links {
  }
}
incarnation: 15490038752371625473
physical_device_desc: "device: 0, name: Tesla V100-PCI-E-16GB, pci bus id: 0000:01:00.0, compute capability: 7.0"]
```

STEP 8

TERMINAL / Shell

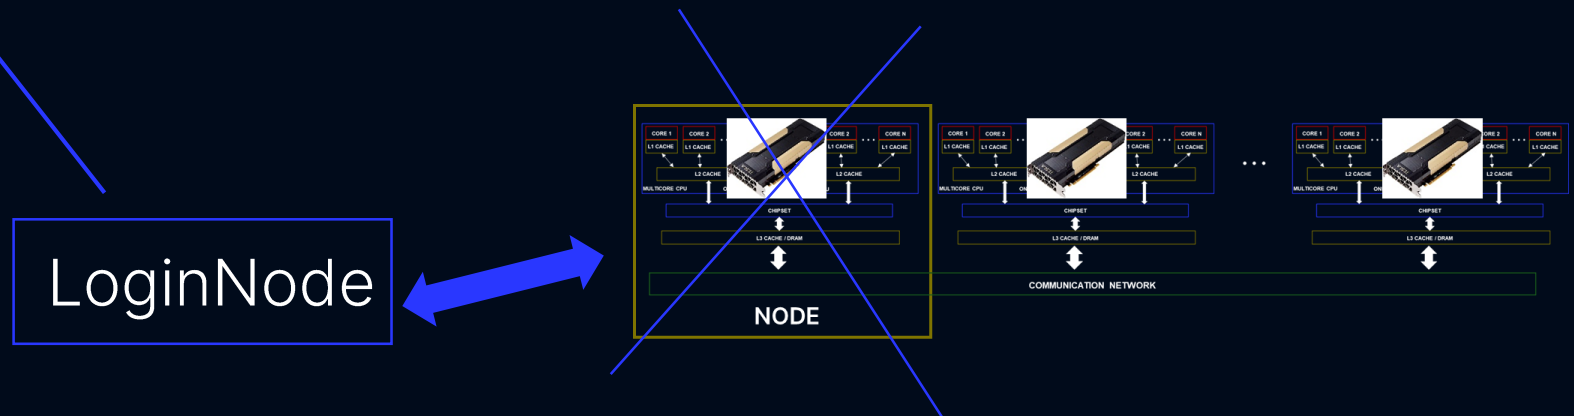

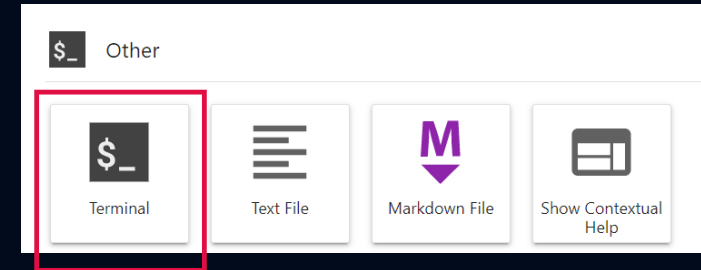
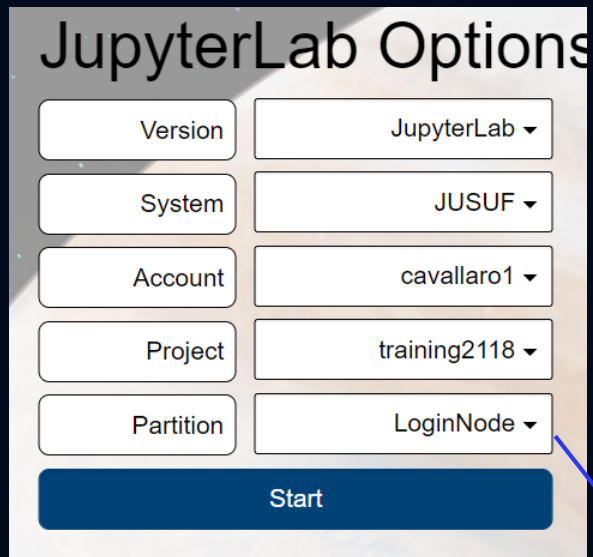


```
cavallaro1@jsfc041:~  
[cavallaro1@jsfc041 ~]$ module unload nvidia-driver/.default  
[cavallaro1@jsfc041 ~]$ srun nvidia-smi  
Fri Jul 9 16:28:37 2021  
+-----+  
| NVIDIA-SMI 460.32.03    Driver Version: 460.32.03    CUDA Version: 11.2    |  
+-----+  
| GPU   Name               Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |  
| Fan  Temp  Perf    Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |  
|                                           MIG M.           |  
+-----+  
|    0   Tesla V100-PCIE...    On          | 00000000:01:00.0 Off  |           0         |  
| N/A   27C    P0     26W / 250W |  0MiB / 16160MiB |    0%      Default  |  
|                                           N/A             |  
+-----+  
+-----+  
| Processes:                                                       GPU Memory |  
| GPU   GI    CI          PID    Type   Process name                               Usage    |  
|      ID    ID                                         |          |  
+-----+  
| No running processes found                                     |          |  
+-----+  
[cavallaro1@jsfc041 ~]$
```

STEP 9

Access on the login node

- Starting point, in this node you do not have GPUs



TOMORROW PRACTICALS

Hands-on - Distributed Deep Learning

- Workflow on the batch system
- Use job scripts to execute algorithms with more nodes (i.e., > 1 GPUs)

