

INTRODUCTION TO HPC MAELSTROM BOOTCAMP JÜLICH

27 September 2022 | Dr. Andreas Herten | Forschungszentrum Jülich, MAELSTROM



Member of the Helmholtz Association

Outline

Introduction Hardware **Comparison to PC** HPC vs. PC HPC **HPC System Overview Historical Machines** JUWELS JUWELS Cluster JUWELS Booster **GPUs**

Software 1: Core Utilization 2: Parallelization 3: Distribution Enablement Conclusion





What is **HPC**?

High Performance Computing is computing with a powerful machine using the available resources efficiently.

My interpretation

• What kind of CPU does your computer have?

CPU generation, clock speed rate, number of cores, vector length



- What kind of CPU does your computer have? CPU generation, clock speed rate, number of cores, vector length
- How much memory does your computer have? Amount of memory, type, links (GB and GB/s)





- What kind of CPU does your computer have?
 CPU generation, clock speed rate, number of cores, vector length
- How much memory does your computer have? Amount of memory, type, links (GB and GB/s)
- What kind of GPU do you have?

GPU generation, number of cores, watt intake (TDP)





- What kind of CPU does your computer have?
 CPU generation, clock speed rate, number of cores, vector length
- How much memory does your computer have? Amount of memory, type, links (GB and GB/s)
- What kind of GPU do you have?

GPU generation, number of cores, watt intake (TDP)

How fast is your network ?

Throughput, latency





Slide 4124

Amount of memory, type, links (GB <u>and</u> GB/s)What kind of GPU do you have?

How much memory does your computer have?

What kind of CPU does your computer have?

GPU generation, number of cores, watt intake (TDP)

CPU generation, clock speed rate, number of cores, vector length

How fast is your network ?

Throughput, latency





















HPC Node

64 cores, 2.2 GHz, $2 \times$ multi-threading, large caches, advanced instructions, 220 W TDP

Specs based on JURECA DC





y 1 TB size, DDR4, 3200 MHz rate, 190 GiB/s bandwidth



HPC Node

64 cores, 2.2 GHz, $2 \times$ multi-threading, large caches, advanced instructions, 220 W TDP







1 TB size, DDR4, 3200 MHz rate, 190 GiB/s bandwidth



HPC Node

64 cores, 2.2 GHz, $2 \times$ multi-threading, large caches, advanced instructions, 220 W TDP



108 *cores*, 1400 MHz, 2048 bit vector size, double-precision support, 400 W TDP, no graphics

Specs based on JURECA DC



1 TB size, DDR4, 3200 MHz rate, 190 GiB/s bandwidth



HPC Node

64 cores, 2.2 GHz, $2 \times$ multi-threading, large caches, advanced instructions, 220 W TDP

108 cores, 1400 MHz, 2048 bit vector size, double-precision support, 400 W TDP, no graphics 40 GB (80 GB also available), HBM2, 1.555 GB/s bandwidth

Specs based on JURECA DC







- Usually, 2 CPUs sockets, each with 64 cores; use mostly as one CPU with one memory
- 4 distinct GPUs , connected with each other (600 GB/s)
- 4 network connections, each 200 Gbit/s in each direction (*InfiniBand HDR-200*)

HPC Node



Specs based on JURECA DC, JUWELS



Specs based on JURECA DC, JUWELS



Specs based on JURECA DC, JUWELS

	┤┋┋╝┋┋╝┝	▏▋▋▋▌▋▋▋▎	▏▋▋▋▏▋▋▋▌┝	▏▋▋▋▋▋▋▌	╎┋══┊┋══┊┝	╎══╝┋═╝┝	▏▋▋▋▌▋▋▋	▏▋▋▋▌▋▋▋	
	┤╩╩╝┋╩╝┝					┤┋═╝┋═╝┝	===		
	┥┋┋┋┋┋╞					┤┋┋┇┋┋┋╞			
				aes i					
	┥┋═╝┋═╝┝					┤┋═┋┋┋═┋┝╴			

				╎┈┈╵				



High Performance Computing is computing with a powerful machine using the available resources efficiently.

My interpretation

<u>High Performance Computing is</u> <u>computing with a powerful machine</u> using the available resources efficiently.

My interpretation.

Powerful Machines

- Now Powerful nodes (large CPUs, accelerating GPUs, much memory)
 - Many nodes (well-connected through high-speed interconnect)
 - $\rightarrow\,$ Beefed-up versions of commodity computers, with slight specializations; many



Powerful Machines

- Now Powerful nodes (large CPUs, accelerating GPUs, much memory)
 - Many nodes (well-connected through high-speed interconnect)
 - $\rightarrow\,$ Beefed-up versions of commodity computers, with slight specializations; many
- Past First computers: Supercomputers! Mainframe machines: Large installations with most powerful hardware at the time
 - PC era: Even then, specialized computers, like vector machines, or many low-speed CPUs (well-connected)
 - Recent history: x86, then PowerPC, then GPU accelerators, then specialized Arm CPUs





- CDC 6600 supercomputer
- Around 1965
- First supercomputer
- 3 MFLOP/s
- See Wikipedia for more
- Picture by Control Data Corporation



Member of the Helmholtz Association

27 September 2022

Slide 11|24



- CDC 6600 supercomputer
- Around 1965
- First supercomputer
- 3 MFLOP/s
- See Wikipedia for more
- Picture by Control Data Corporation



Member of the Helmholtz Association

27 September 2022

Slide 11|24



HPC performance measured in FLOP/s.

 Floating-point (like 3.14) operations per second



Member of the Helmholtz Association

27 September 2022

Slide 11/24



HPC performance measured in $\ensuremath{\mathsf{FLOP}}\xspace/s.$

- Floating-point (like 3.14) operations per second
- Example: Processor with 2 GHz; 10 cores; per core: 2 multiplications and 2 additions (*FMA*) per cycle



Member of the Helmholtz Association

Slide 11|24


HPC performance measured in FLOP/s.

- Floating-point (like 3.14) operations per second
- Example: Processor with 2 GHz; 10 cores; per core: 2 multiplications and 2 additions (*FMA*) per cycle

 2×10^9 1/s 1/core * 10 core* * (2 + 2) floating-point operation



Member of the Helmholtz Association



HPC performance measured in FLOP/s.

- Floating-point (like 3.14) operations per second
- Example: Processor with 2 GHz; 10 cores; per core: 2 multiplications and 2 additions (*FMA*) per cycle

 $2 \times 10^{9} \text{ 1/s 1/core * 10 core*}$ * (2 + 2) floating-point operation $= 2 * 10^{9} * 10 * 4 \text{ fl-op/s}$



Member of the Helmholtz Association

27 September 2022



HPC performance measured in FLOP/s.

- Floating-point (like 3.14) operations per second
- Example: Processor with 2 GHz; 10 cores; per core: 2 multiplications and 2 additions (*FMA*) per cycle _____

 2×10^9 1/s 1/core * 10 core* * (2 + 2) floating-point operation =2 * 10⁹ * 10 * 4 fl-op/s =80 * 10⁹ FLOP/s =80 GFLOP/s



Member of the Helmholtz Association

27 September 2022



- Cray-1 supercomputer
- Around 1978
- Very successful
- 160 MFLOP/s
- Probably pictured at NERSC



Member of the Helmholtz Association

27 September 2022



- Intel XP/S 140 supercomputer
- Around 1994
- 3680 Intel i860 RISC processors; large-scale parallel system
- 143 GFLOP/s
- Picture by top500.org





- JUGENE supercomputer
- 2008
- 294 912 PowerPC 450 cores; energy-efficient
- 800 TFLOP/s
- Picture by top500.org





- Summit supercomputer
- 2018
- 27 000 GPUs hosted by POWER9 CPUs; first #1 GPU
 - supercomputer
- 200 PFLOP/s
- Picture by Oak Ridge National Lab





- Fugaku supercomputer
- **2020**
- 7 630 848 Arm A64FX cores; #1 supercomputer
- 537 PFLOP/s
- Picture by RIKEN





JUWELS Cluster – Jülich's Scalable System

- 2500 nodes with Intel Xeon CPUs (2 \times 24 cores)
- 46 + 10 nodes with 4 NVIDIA Tesla V100 cards (16 GB memory)
- 10.4 (CPU) + 1.6 (GPU) PFLOP/s peak performance (Top500: #86)







JUWELS Booster – Scaling Higher!

- = 936 nodes with AMD EPYC Rome CPUs (2 \times 24 cores)
- Each with 4 NVIDIA A100 Ampere GPUs (each: FP64TC: 19.5 FP64: 9.7 TFLOP/s, 40 GB memory)
- InfiniBand DragonFly+ HDR-200 network; $4 \times 200 \text{ Gbit/s per node}$







Top500 List Nov 2021:

- #1 Europe
- #8 World
- #4* Top/Green500



JUWELS Booster – Scaling Higher!

- = 936 nodes with AMD EPYC Rome CPUs (2 \times 24 cores)
- Each with 4 NVIDIA A100 Ampere GPUs (each: FP64TC: 19.5 FP64: 9.7 TFLOP/s, 40 GB memory)
- InfiniBand DragonFly+ HDR-200 network; $4 \times 200 \text{ Gbit/s per node}$



 Current fastest supercomputer: Frontier at Oak Ridge (USA) with 38 000 AMD MI250X GPUs – 1.102 EFLOP/s; also most energy-efficient!





- Current fastest supercomputer: Frontier at Oak Ridge (USA) with 38 000 AMD MI250X GPUs – 1.102 EFLOP/s; also most energy-efficient!
- 2022: Aurora at Argonne with > 60 000 Intel Ponte Vecchio GPUs - > 2 EFLOP/s
- 2023: El Capitan at Lawrence Livermore with AMD MI300 GPUs -> 2 EFLOP/s





Slide 14|24

- Current fastest supercomputer: Frontier at Oak Ridge (USA) with 38 000 AMD MI250X GPUs – 1.102 EFLOP/s; also most energy-efficient!
- 2022: Aurora at Argonne with > 60 000 Intel Ponte Vecchio GPUs - > 2 EFLOP/s
- 2023: El Capitan at Lawrence Livermore with AMD MI300 GPUs -> 2 EFLOP/s
- 2023: JUPITER at JSC -> 1 EFLOP/s! GPUs, but details TBD







- Current fastest supercomputer: Frontier at Oak Ridge (USA) with 38 000 AMD MI250X GPUs – 1.102 EFLOP/s; also most energy-efficient!
- 2022: Aurora at Argonne with > 60 000 Intel Ponte Vecchio GPUs - > 2 EFLOP/s
- 2023: El Capitan at Lawrence Livermore with AMD MI300 GPUs -> 2 EFLOP/s
- 2023: JUPITER at JSC -> 1 EFLOP/s!
 GPUs, but details TBD







- GPUs: Exascale Enablers
- Processors efficient at applying same (/similar) instruction on large set of data (image)
- Over last 15 years, extended from rendering to variable computing
- Not good for every task, but **great for some**, which happen to be computing with large amounts of similar data



- GPUs: Exascale Enablers
- Processors efficient at applying same (/similar) instruction on large set of data (image)
- Over last 15 years, extended from rendering to variable computing
- Not good for every task, but **great for some**, which happen to be computing with large amounts of similar data
- Programming model: SIMT, SIMD ⊗ SMT (vectors ⊗ threads)
- JUWELS Booster thread 100 % occupancy: 3744 GPUs \times 108 SMs \times 2048 threads/SM = 828 112 896 threads





- GPUs: Exascale Enablers
- Processors efficient at applying same (/similar) instruction on large set of data (image)
- Over last 15 years, extended from rendering to variable computing
- Not good for every task, but **great for some**, which happen to be computing with large amounts of similar data
- Programming model: SIMT, SIMD ⊗ SMT (vectors ⊗ threads)
- JUWELS Booster thread 100 % occupancy: 3744 GPUs \times 108 SMs \times 2048 threads/SM = 828 112 896 threads
- Important vendors: First NVIDIA, then AMD; soon also Intel









- GPUs: Exascale Enabler
- Processors efficient at a on large set of data (ima
- Over last 15 years, exten computing
- Not good for every task, be computing with large
- Programming model: S
- JUWELS Booster thread SMs × 2048 threads/SM
- Important vendors: Firs

SМ								
			L1 Instru	ction Cac	he			
L0 Instruction Cache					LO	Instruction C	ache	
Warp Scheduler (32 thread/clk)				Warp Scheduler (32 thread/clk)				
Dispatch Unit (32 thread/clk)				Dispatch Unit (32 thread/clk)				
	Regist	er File (16,38	4 x 32-bit)		Registe	r File (16,38-	4 x 32-bit)	
INT32 INT32 INT32 INT32 INT32 INT32 INT32 INT32 INT32 INT32 INT32 INT32 INT32 INT32 INT32 INT32	FP32 FP32 FP32 FP32	FP64 FP64 FP64 FP64 FP64 FP64 FP64 FP64	TENSOR CORE 4 ¹⁹ GENERATION	INT32 INT32 INT32 INT32 INT32 INT32 INT32 INT32 INT32 INT32 INT32 INT32 INT32 INT32 INT32 INT32	F932 F932 F932 F932	FP64 FP64 FP64 FP64 FP64 FP64 FP64 FP64	TENSOR (4 th GENER)	
LD/ ST	LOY LOY LU ST ST S	Y LOY LOY ST ST ST	ST ST SFU	LUY ST	LOV LOV LOV ST ST ST	LDY LDY ST ST	LOV LOV ST ST	

	L0 Instruction C	ache			L0 Instruct	ion Ci	ache	
Warp Scheduler (32 thread/clk)				Warp Scheduler (32 thread/clk)				
Dispatch Unit (32 thread/clk)				Dispatch Unit (32 thread/clk)				
Register File (16,384 x 32-bit)				Register File (16,384 x 32-bit)				
NT32 FP32 NT32 FP32	FP32 FP64 FP32 FP64		INT32 INT32	FP32 FP FP32 FP	32 FP6 32 FP6	4		
NT32 FP32 NT32 FP32	FP32 FP64 FP32 FP64		INT32	FP32 FP FP32 FP	32 FP6 32 FP6	-		
NT32 FP32 NT32 FP32	FP32 FP64 FP32 FP64		INT32 INT32	FP32 FP FP32 FP	32 FP6 32 FP6	4	TENSOR COP	
NT32 FP32 NT32 FP32	FP32 FP64 FP32 FP64	TENSOR CORE	INT32 INT32	FP32 FP FP32 FP	32 FP6 32 FP6	4		SOR CORE
NT32 FP32 NT32 FP32	FP32 FP64 FP32 FP64	4" GENERATION	INT32 INT32	FP32 FP FP32 FP	32 FP6 32 FP6	4	4 th GENERATIO	
NT32 FP32 NT32 FP32	FP32 FP64 FP32 FP64		INT32 INT32	FP32 FP FP32 FP	32 FP6 32 FP6	4		
NT32 FP32 NT32 FP32	FP32 FP64 FP32 FP64		INT32 INT32	FP32 FP FP32 FP	32 FP6 32 FP6	4 4		
NT32 FP32 NT32 FP32	FP32 FP64 FP32 FP64		INT32 INT32	FP32 FP FP32 FP	32 FP6 32 FP6	4		
LDY LDY LD ST ST ST	LDV LDV LDV ST ST ST	ST ST SFU	LD/ ST	LDV LDV ST ST	LDV LDV ST ST		LD/ L ST I	SFU



Tex



Pictures by NVIDIA [8]

GPUs

GPUs: Exascale Enablers

- Processors eff^{: -:} on large set of
- Over last 15 ye computing
- Not good for e be computing
- Programming
- JUWELS Boos





- SMs imes 2048 threads/SM = 828 112 896 threads
- Important vendors: First NVIDIA, then AMD; soon also Intel



Slide 15124

Picture by AMD [9]

- GPUs: Exascal
- Processors eff on large set of
- Over last 15 ye computing
- Not good for e be computing
- Programming
- JUWELS Boos SMs × 2048 th







<u>High Performance Computing is</u> <u>computing with a powerful machine</u> using the available resources efficiently.

My interpretation.

High Performance Computing is computing with a powerful machine <u>using the available resources efficiently</u>.

My interpretation.

Resource Utilization



Exploit all capabilities of processing entity (core)



Resource Utilization



Exploit all capabilities of processing entity (core)

Parallelize to all processing entities of node



Resource Utilization



Exploit all capabilities of processing entity (core)

Parallelize to all processing entities of node

Distribute to all nodes



- Modern CPUs: many advanced instructions, high clock rate, large caches, high memory bandwidth
- Use via tailored algorithms, specific functions (*intrinsics*), modern compilers, optimized libraries



- Modern CPUs: many advanced instructions, high clock rate, large caches, high memory bandwidth
- Use via tailored algorithms, specific functions (*intrinsics*), modern compilers, optimized libraries
- Example: Vectorization/SIMD





- Modern CPUs: many advanced instructions, high clock rate, large caches, high memory bandwidth
- Use via tailored algorithms, specific functions (*intrinsics*), modern compilers, optimized libraries
- Example: Vectorization/SIMD



- \times 4 multiplications
- + 4 additions
- = 4 assignments
- \rightarrow 8 instructions



- Modern CPUs: many advanced instructions, high clock rate, large caches, high memory bandwidth
- Use via tailored algorithms, specific functions (*intrinsics*), modern compilers, optimized libraries
- Example: Vectorization/SIMD

A ₀		B_0		C_0		D ₀
A_1	×	B_1	+	C_1	_	D_1
A ₂	B ₂	'	<i>C</i> ₂		<i>D</i> ₂	
A ₃		<i>B</i> ₃		<i>C</i> ₃		<i>D</i> ₃

- \times 4 multiplications + 4 additions SIMD 4 assignments
- \rightarrow 8 instructions

- \times 1 multiplication
- 1 addition
- 1 assignment
- \rightarrow 2 instructions



- Modern CPUs: many advanced instructions, high clock rate, large caches, high memory bandwidth
- Use via tailored algorithms, specific functions (*intrinsics*), modern compilers, optimized libraries

SIMD

Example: Vectorization/SIMD



- × 4 multiplications
 - 4 additions
 - 4 assignments
- \rightarrow 8 instructions

- × 1 multiplication + 1 addition = 1 assignment
 - \rightarrow 2 instructions





- Modern CPUs: many advanced instructions, high clock rate, large caches, high memory bandwidth
- Use via tailored algorithms, specific functions (*intrinsics*), modern compilers, optimized libraries
- Example: Vectorization/SIMD



1 multiplication **CPU Instruction:** \times 4 multiplications VADDPD 4 additions 1 addition 4 assignments 1 assignment 2 instructions mm256 add pd(); 8 instructions MASIMD multiplication addition assignment 1 instruction



- Modern CPUs: many advanced instructions, high clock rate, large caches, high memory bandwidth
- Use via tailored algorithms, specific functions (*intrinsics*), modern compilers, optimized libraries
- Example: Vectorization/SIMD



1 multiplication **CPU Instruction:** \times 4 multiplications VADDPD 4 additions 1 addition 4 assignments 1 assignment mm256 add pd(); \rightarrow 8 instructions MASIMO 2 instructions **CPU Instruction:** multiplication VFMADD132PD addition L assignment mm256 fmadd pd; 1 instruction



- Modern CPUs: many advanced instructions, high clock rate, large caches, high memory bandwidth
- Use via tailored algorithms, specific functions (*intrinsics*), modern compilers, optimized libraries
- Example: Vectorization/SIMD



 $\begin{array}{c} \times & 4 \text{ multiplications} \\ + & 4 \text{ additions} \\ = & 4 \text{ assignments} \\ \rightarrow & 8 \text{ instructions} \end{array} \xrightarrow{\text{SIMD}} \begin{array}{c} \times & 1 \text{ multiplication} \\ + & 1 \text{ addition} \\ = & 1 \text{ assignment} \\ \rightarrow & 2 \text{ instructions} \end{array} \xrightarrow{\text{Compiler!}} \begin{array}{c} \text{Compiler!} \end{array}$



- Modern CPUs: many advanced instructions, high clock rate, large caches, high memory bandwidth
- Use via tailored algorithms, specific functions (*intrinsics*), modern compilers, optimized libraries
- Example: Vectorization/SIMD



1 multiplication \times 4 multiplications 4 additions 1 addition Improve 4 assignments 1 assignment throughput! 2 instructions 8 instructions MASIMD multiplication Improve addition throughput L assignment 1 instruction more!


Analysis/plot by Stepan Nassyr, 2022.

Forschungszentrum

CENTRE

1 Exploit All Capabilities of Processing Entity



From core to cores





- From core to cores
- From CPU cores to GPU cores





- From core to cores
- From CPU cores to GPU cores
- Parallelization: Tasks work on portion of full problem using some local shared memory; *fine-grained* split





- From core to cores
- From CPU cores to GPU cores
- Parallelization: Tasks work on portion of full problem using some local shared memory; *fine-grained* split
- CPU Mostly through operating system capacities
 - OS threads launched on cores
 - Easiest threading interface: OpenMP

#pragma omp parallel for for (int i = 0; i < N; i++) y[i] = x[i] * 3.14 + a[i];</pre>





- From core to cores
- From CPU cores to GPU cores
- Parallelization: Tasks work on portion of full problem using some local shared memory; *fine-grained* split
- CPU Mostly through operating system capacities
 - OS threads launched on cores
 - Easiest threading interface: OpenMP

#pragma omp parallel for for (int i = 0; i < N; i++) y[i] = x[i] * 3.14 + a[i];</pre>

- GPU Through dedicated programming environments
 - Mostly, explicit models

```
int i = threadIdx.x + blockIdx.x * blockDim.x;
y[i] = x[i] * 3.14 + a[i];
```

Also, higher-level models (OpenMP, OpenACC)





Slide 19124



From node to nodes





3 Distribute to All Nodes

- From node to nodes
- Distribution: Tasks work on portion of full problem using distributed memory; *coarse-grained* split





3 Distribute to All Nodes

- From node to nodes
- Distribution: Tasks work on portion of full problem using distributed memory; *coarse-grained* split
- Every task runs on own node with copy of program, defined exchange functions
- High-speed network important! GPUs directly attached to network





3 Distribute to All Nodes

- From node to nodes
- Distribution: Tasks work on portion of full problem using distributed memory; *coarse-grained* split
- Every task runs on own node with copy of program, defined exchange functions
- High-speed network important! GPUs directly attached to network
- Classical programming model: MPI
 MPI_Comm_size(MPI_COMM_WORLD, &size);
 MPI_Comm_rank(MPI_COMM_WORLD, &rank);
 MPI_Send(buffer, 10, MPI_INT, 1, 555, MPI_COMM_WORLD);
 MPI_Allreduce(local_sum, global_sum, 1, MPI_FLOAT, MPI_SUM,
 → MPI_COMM_WORLD);





Compilers **Translate** high-level code to low-level machine code, with general *and* very architecture-specific optimizations



Compilers **Translate** high-level code to low-level machine code, with general *and* very architecture-specific optimizations

Frameworks Offer pre-programmed **function primitives** to build a program upon



ompilers **Translate** high-level code to low-level machine code, with general *and* very architecture-specific optimizations

Frameworks Offer pre-programmed function primitives to build a program upon

Libraries **Back-end**, low-level functions, usually optimized extensively, sometimes by vendors themselves



ompilers **Translate** high-level code to low-level machine code, with general *and* very architecture-specific optimizations

Frameworks Offer pre-programmed function primitives to build a program upon

Libraries **Back-end**, low-level functions, usually optimized extensively, sometimes by vendors themselves

Compilers

CPU GCC, LLVM, Intel, Cray

GPU + NVIDIA CUDA, NVHPC, AMD

> Long history, constantly evolving

Frameworks

MPI OpenMPI, MPICH

Threads pthreads, OpenMP GPU CUDA, HIP, SYCL, pSTL, Kokkos

Libraries

CPU MKL, BLIS, FFTW

GPU cuBLAS, rocBLAS, cuDNN

→ TensorFlow, PyTorch, ELPA



High Performance Computing is <u>computing with a powerful machine</u> using the available resources efficiently.

My interpretation.

- HPC is intensive computing with largest machines
- Sometimes like Formula 1, sometimes like a tanker
- Sophisticated hardware is underlying everything, delivering up to 1.1 EFLOP/s
- Advanced software holds everything together and enables science at the frontiers







- HPC is intensive computing with largest machines
- Sometimes like Formula 1, sometimes like a tanker
- Sophisticated hardware is underlying everything, delivering up to 1.1 EFLOP/s
- Advanced software holds everything together and enables science at the frontiers, like
 - Plasma physics simulations
 - Drug discovery
 - Material design
 - Weather and climate modelling
 - Precise Artificial Intelligence







- HPC is intensive computing with largest machines
- Sometimes like Formula 1, sometimes like a tanker
- Sophisticated hardware is underlying everything, delivering up to 1.1 EFLOP/s
- Advanced software holds everything together and enables science at the frontiers, like
 - Plasma physics simulations
 - Drug discovery
 - Material design
 - Weather and climate modelling
 - Precise Artificial Intelligence









- HPC is intensive computing with largest machines
- Sometimes like Formula 1, sometimes like a tanker
- Sophisticated hardware is underlying everything, delivering up to 1.1 EFLOP/s
- Advanced software holds everything together and enables science at the frontiers, like
 - Plasma physics simulations
 - Drug discovery
 - Material design
 - Weather and climate modelling
 - Precise Artificial Intelligence
- We are hiring!
 - go.fzj.de/jsc-jobs







Slide 24124

- HPC is intensive computing with largest machines
- Sometimes like Formula 1, sometimes like a tanker
- Sophisticated hardware is underlying everything, delivering up to 1.1 EFLOP/s
- Advanced software holds everything together and enables science at the frontiers. like Thank you for your attention!
 - Plasma physics simulations
 - Drug discoverv
 - Material design
 - Weather and climate modelling
 - Precise Artificial Intelligence
- We are hiring!
 - go.fzi.de/isc-iobs



a.herten@fz-juelich.de

Appendix

Appendix License References





This slide deck is published under the following license:

CC BY-SA 4.0





References: Images, Graphics I

- [1] Forschungszentrum Jülich GmbH (Ralf-Uwe Limbach). JUWELS Booster.
- [2] Control Data Corporation. Picture: CDC 6600. Computer History Museum. URL: https://www.computerhistory.org/revolution/supercomputers/10/33 (pages 33, 34).
- [3] Sandia National Lab. Picture: Intel XP/S 140. Top500.org. URL: https://www.top500.org/resources/top-systems/intel-xps-140paragon-sandia-national-labs/ (page 41).
- [4] Forschungszentrum Jülich. Picture: JUGENE. JUGENE Press Release. URL: https://www.fzjuelich.de/de/aktuelles/news/pressemitteilungen/2007/index4763_htm (page 42).



References: Images, Graphics II

- [5] OLCF at ORNL. Picture: Summit. Flickr. URL: https://www.flickr.com/photos/olcf/42659222181/ (page 43).
- [6] RIKEN. Picture: Fugaku. Fujitsu.com. URL: https://blog.de.fujitsu.com/data-driven/fugaku-der-aktuellweltweit-leistungsstaerkste-supercomputer/ (page 44).
- [7] OLCF at ORNL. Picture: Frontier. Flickr. URL: https://www.flickr.com/photos/olcf/52117623843/.
- [8] Nvidia Corporation. Pictures: Ampere GPU. Ampere Architecture Whitepaper. URL: http://www.nvidia.com/nvidia-ampere-architecture-whitepaper (pages 55–57).



References: Images, Graphics III

[9] AMD Inc. Pictures: Instinct MI250 GPU. Promotional Material. URL: https://videocardz.net/amd-instinct-mi250 (page 58).

