Tutorial:
Parallel programming technologies on hybrid architectures

HybriLIT Team

Laboratory of Information Technologies

Joint Institute for Nuclear Research

The Helmholtz International Summer School

"**Lattice QCD, Hadron Structure and Hadronic Matter**"

1, September 2014

HybriLIT Team for heterogeneous computations:

Gheorghe Adam	Evgeny Aleksandrov
Alexander Ayriyan	Dmitriy Belyakov
Dmitry Golub	Margarit Kirakosayn
Mikhail Matveev	Dmitry Podgainy
Alexey Prikhodko	Tatiana Sapozhnikova
Alexej I. Streltsov	Oksana Streltsova
Tatiana Strizh	Shushanik Torosyan
Lucia Valova	Martin Vala
Tatiana Zaikina	Elena Zemlyanaya
Peter Zrelov	Maxim Zuev

Goal: Efficient parallelization of complex numerical problems in computational physics

- OpenMP, MPI technologies
- CUDA technology
- OpenCL technology

September 2014

- MPI+CUDA hybrid technology;
- research on comparative analysis of the efficiency of using GPU, multi-core CPU and Intel Xeon Phi coprocessors.

TOP500 List – June 2014

Rank	Site	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	National Super Computer Center in Guangzhou China	**Tianhe-2 (MilkyWay-2)** - TH-IVB-FEP Cluster, Intel Xeon E5-2692 12C 2.200GHz, TH Express-2, Intel Xeon Phi 31S1P NUDT	3120000	33862.7	54902.4	17808
2	DOE/SC/Oak Ridge National Laboratory United States	**Titan** - Cray XK7 , Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA K20x Cray Inc.	560640	17590.0	27112.5	8209
3	DOE/NNSA/LLNL United States	**Sequoia** - BlueGene/Q, Power BQC 16C 1.60 GHz, Custom IBM	1572864	17173.2	20132.7	7890
4	RIKEN Advanced Institute for Computational Science (AICS) Japan	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect Fujitsu	705024	10510.0	11280.4	12660
5	DOE/SC/Argonne National Laboratory United States	**Mira** - BlueGene/Q, Power BQC 16C 1.60GHz, Custom IBM	786432	8586.6	10066.3	3945
	• • •					
41	Barcelona Supercomputing Center Spain	**MareNostrum** - iDataPlex DX360M4, Xeon E5-2670 8C 2.600GHz, Infiniband FDR IBM	48896	925.1	1017.0	1016
42	Moscow State University - Research Computing Center Russia	**Lomonosov** - T-Platforms T-Blade2/1.1, Xeon X5570/X5670/E5630 2.93/2.53 GHz, Nvidia 2070 GPU, PowerXCell 8i Infiniband QDR T-Platforms	78660	901.9	1700.2	2800
43	Rensselaer Polytechnic Institute United States	**AMOS** - BlueGene/Q, Power BQC 16C 1.6GHz, Custom Interconnect IBM	81920	894.4	1048.6	411

TOP500 List – June 2014

Performance Share of Accelerators

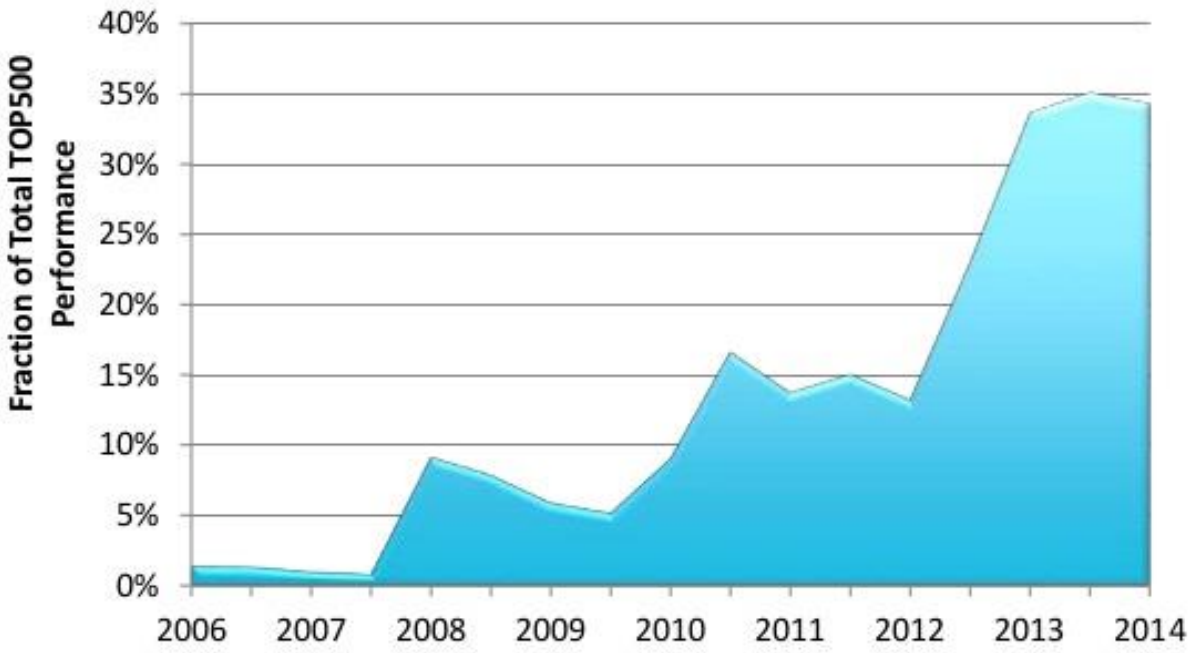

Source:
http://www.top500.org/blog/slides-for-the-43rd-top500-list-now-available/

TOP500 List – June 2014

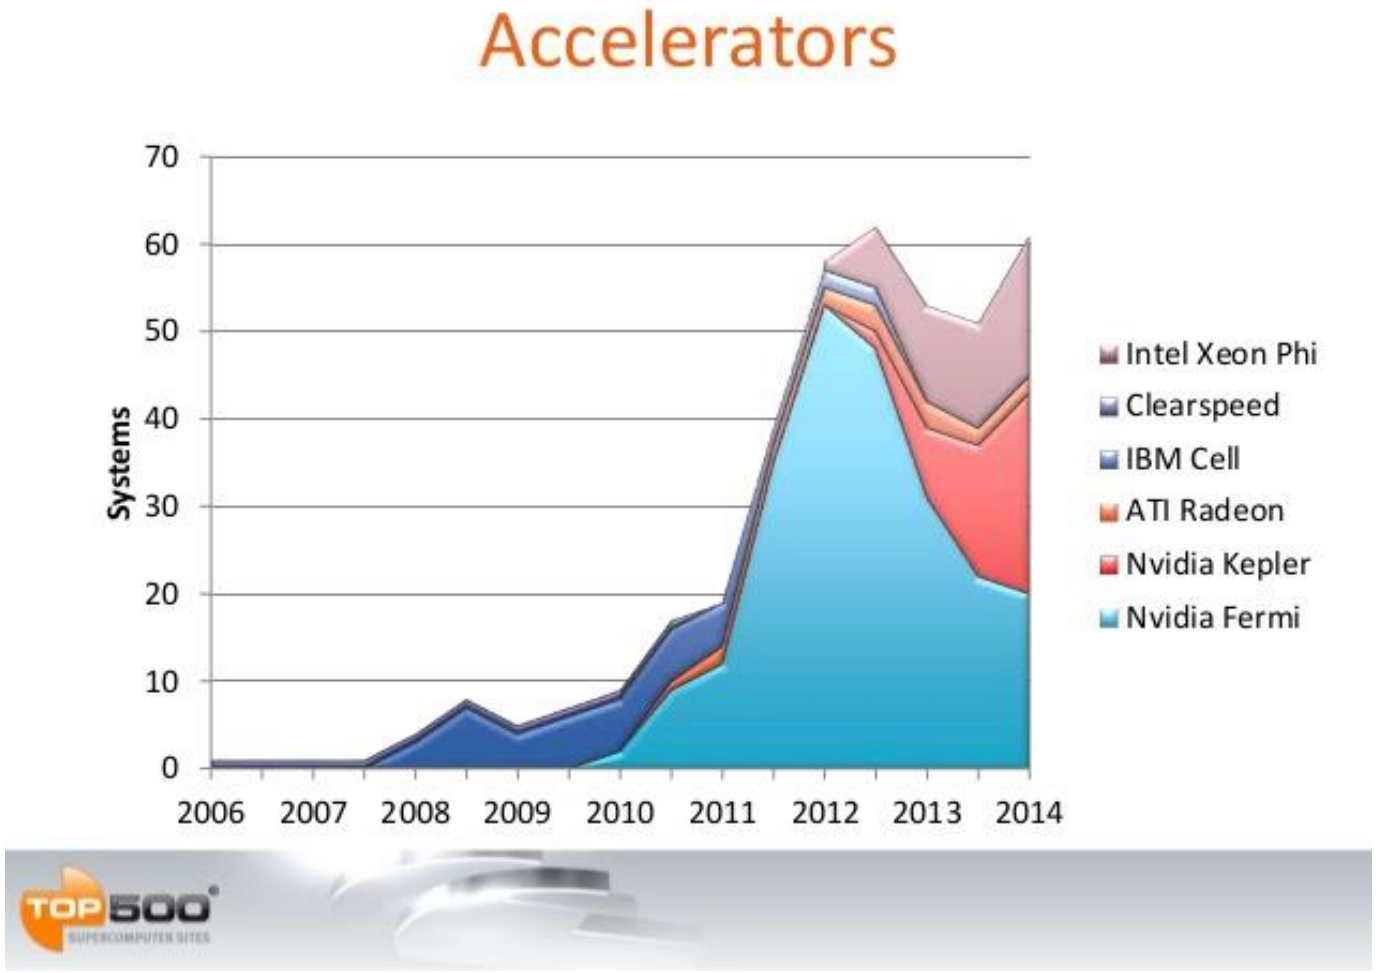

Source:
http://www.top500.org/blog/slides-for-the-43rd-top500-list-now-available/

«Lomonosov» Supercomputer , MSU

МГУ
НИВЦ

>**5000** computation nodes
Intel Xeon X5670/X5570/E5630, PowerXCell 8i
~36 Gb DRAM
2 x **nVidia Tesla** X2070 6 Gb GDDR5 (448 CUDA-cores)
InfiniBand QDR

LIT

NVIDIA Tesla K40 “Atlas” GPU Accelerator

- **Custom languages such as CUDA and OpenCL**
- **Specifications**
  - **2880** CUDA GPU cores
  - **Peak precision floating point performance**

    4.29 TFLOPS single-precision

    1.43 TFLOPS double-precision
  - **memory**

    12 GB GDDR5

    Memory bandwidth up to **288** GB/s

Supports Dynamic Parallelism and HyperQ features

«Tornado SUSU» Supercomputer,
South Ural State University, Russia

«Tornado SUSU» supercomputer took the 157 place in 43-th issue of TOP500 rating (June 2014).

480 computing units (compact and powerful computing blade-modules)
960 processors Intel Xeon X5680
(Gulftown, 6 cores with frequency 3.33 GHz)
384 coprocessors Intel Xeon Phi SE10X (61 cores with frequency 1.1 GHz)

Intel® Xeon Phi™ Coprocessor

Intel Many Integrated Core Architecture (Intel **MIC**) is a multiprocessor computer architecture developed by Intel.

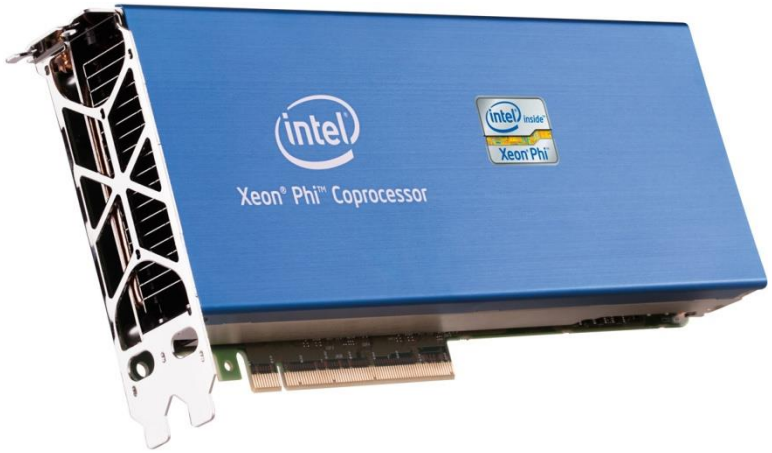

At the end of 2012, Intel launched the first generation of the Intel Xeon Phi product family.

Intel Xeon Phi 7120P

Clock Speed	**1.24 GHz**
L2 Cache	**30.5 MB**
TDP	**300 W**
Cores	**61**
More threads	**244**

The core is capable of supporting **4 threads** in hardware.

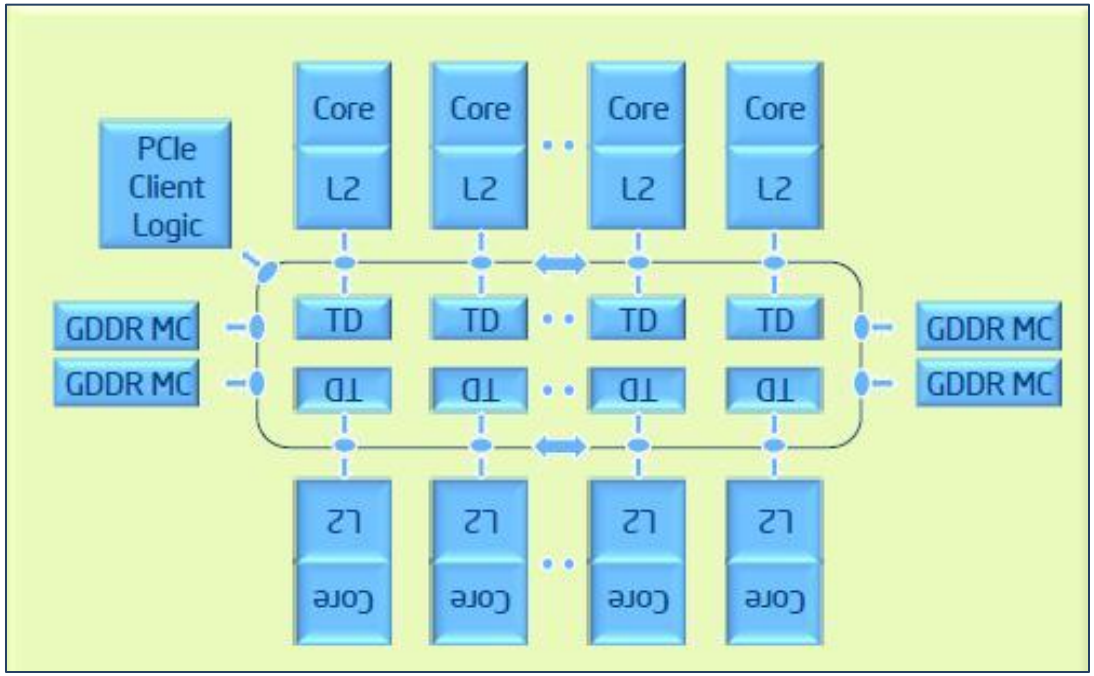

HybriLIT: heterogeneous computation cluster

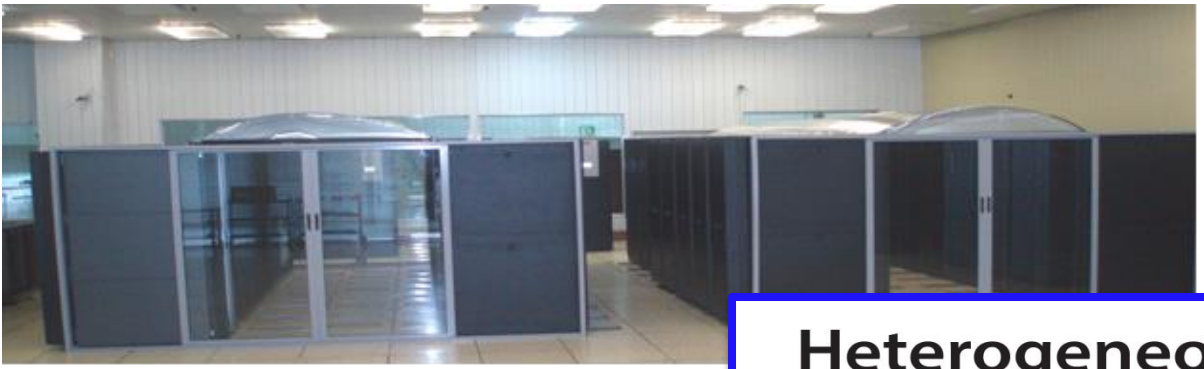

CICC comprises
2582 Cores
Disk storage capacity
1800 TB

August, 2014

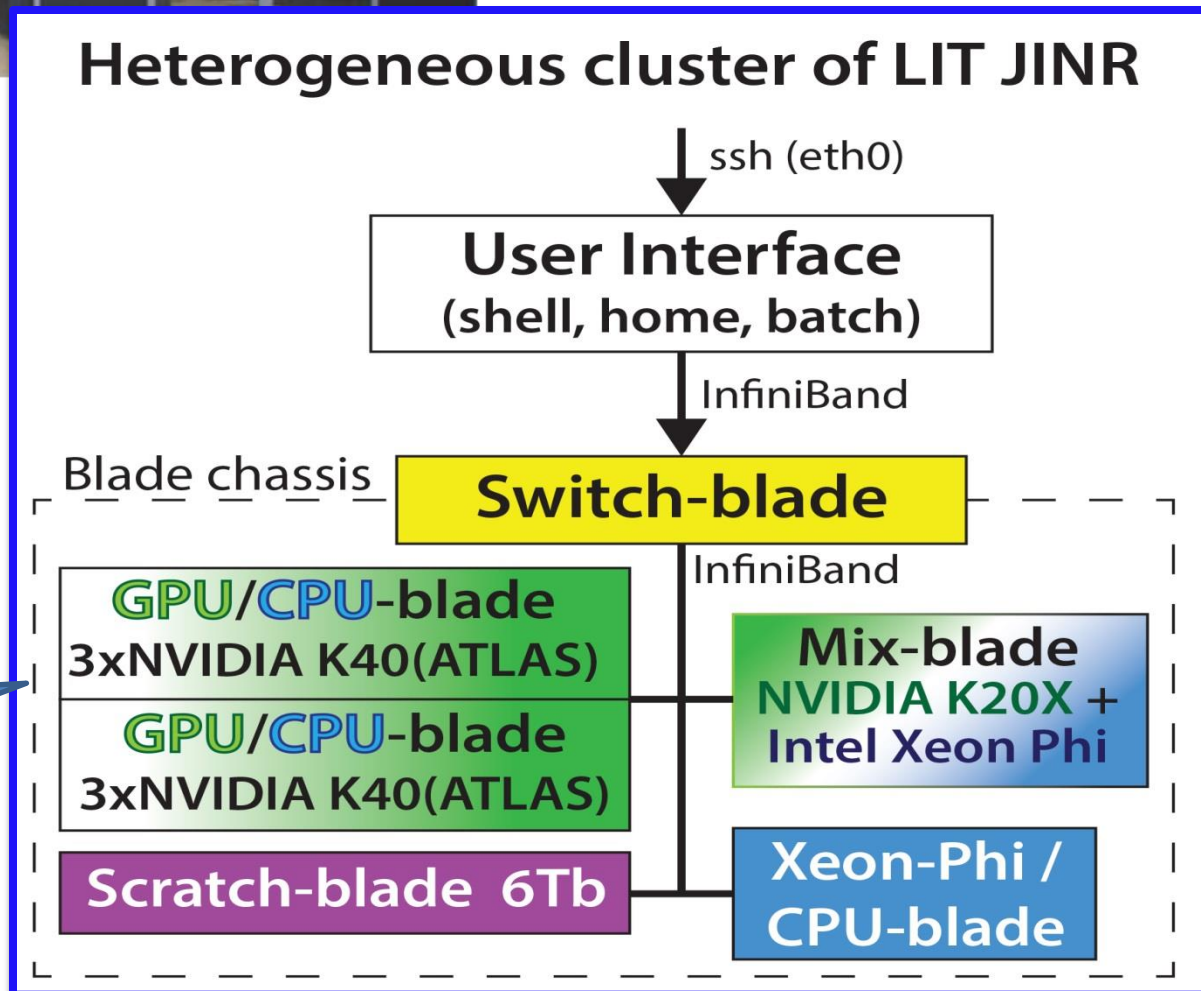

HybriLIT: heterogeneous computation cluster

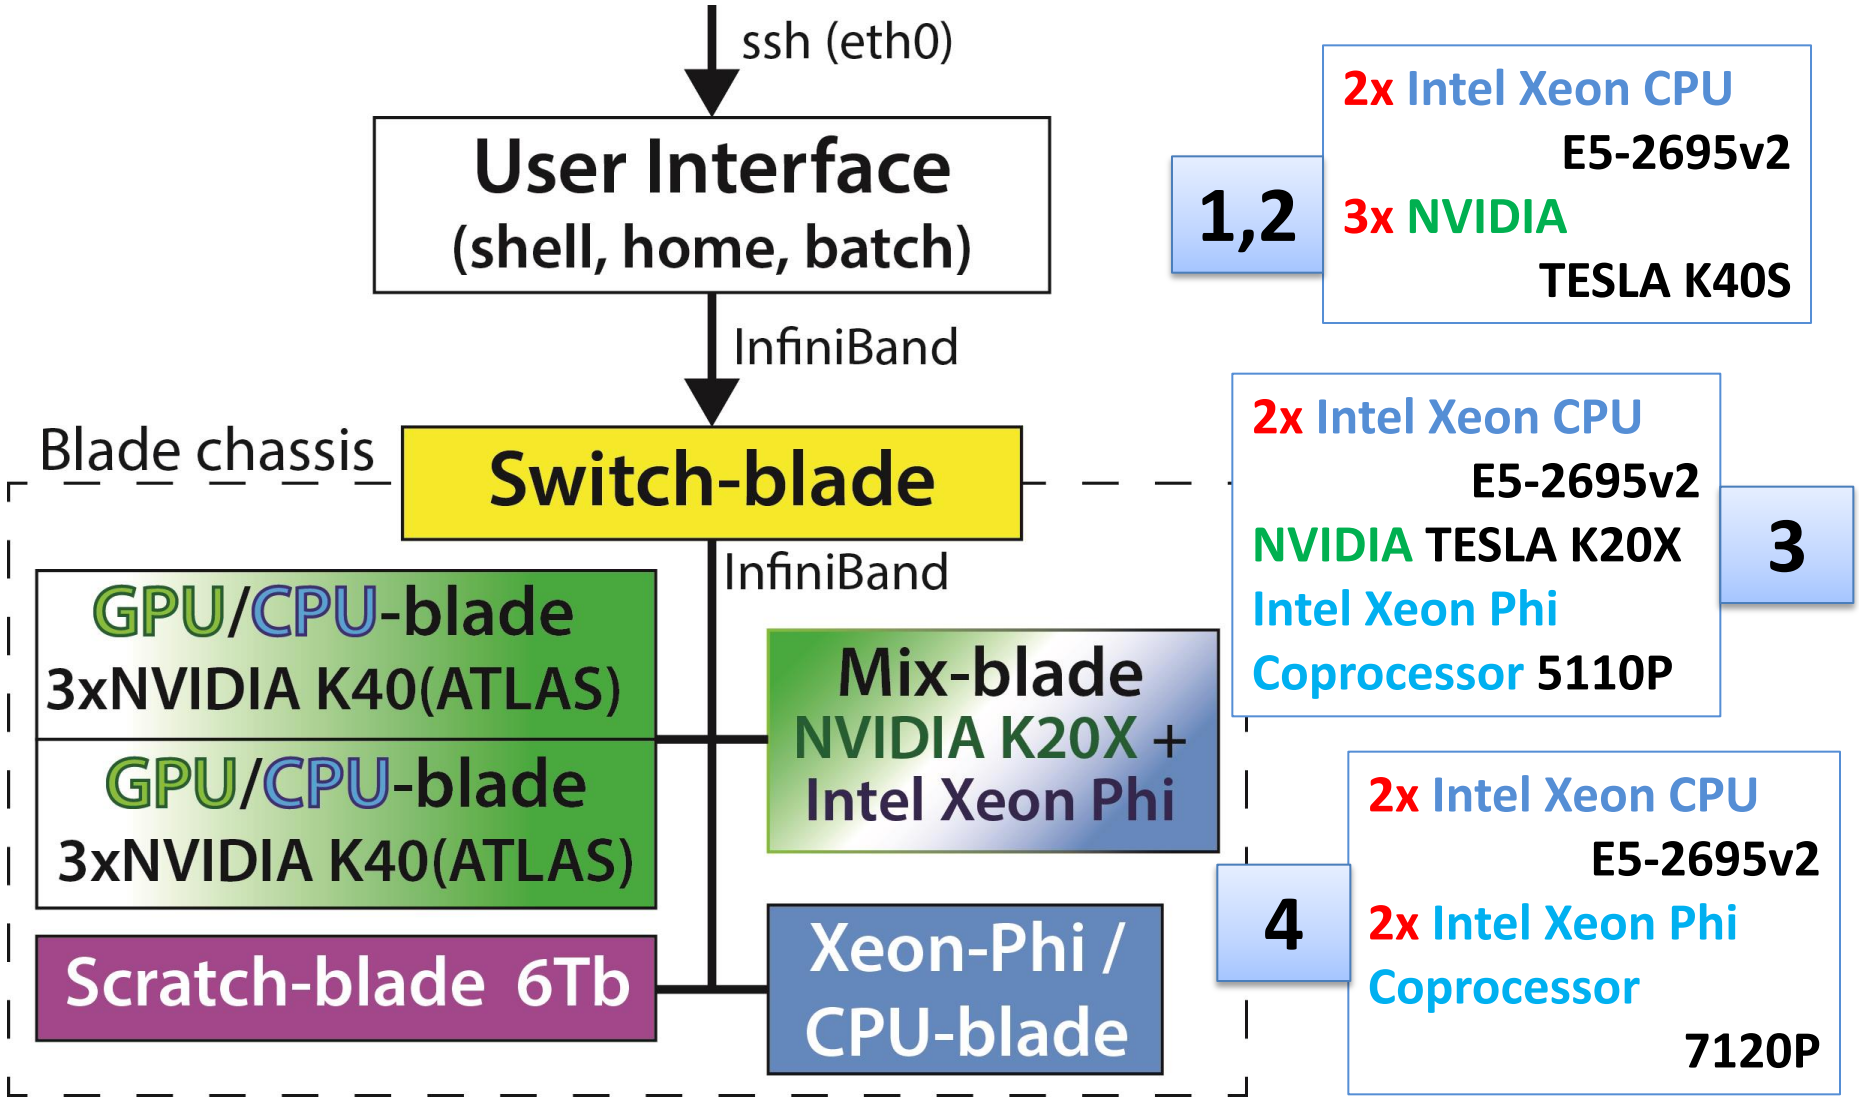

What we see: modern Supercomputers are hybrid with heterogeneous nodes

- Multiple CPU cores with share memory
- Multiple GPU

- Multiple CPU cores with share memory
- Multiple Coprocessor

- Multiple CPU
- GPU
- Coprocessor

Parallel technologies: levels of parallelism

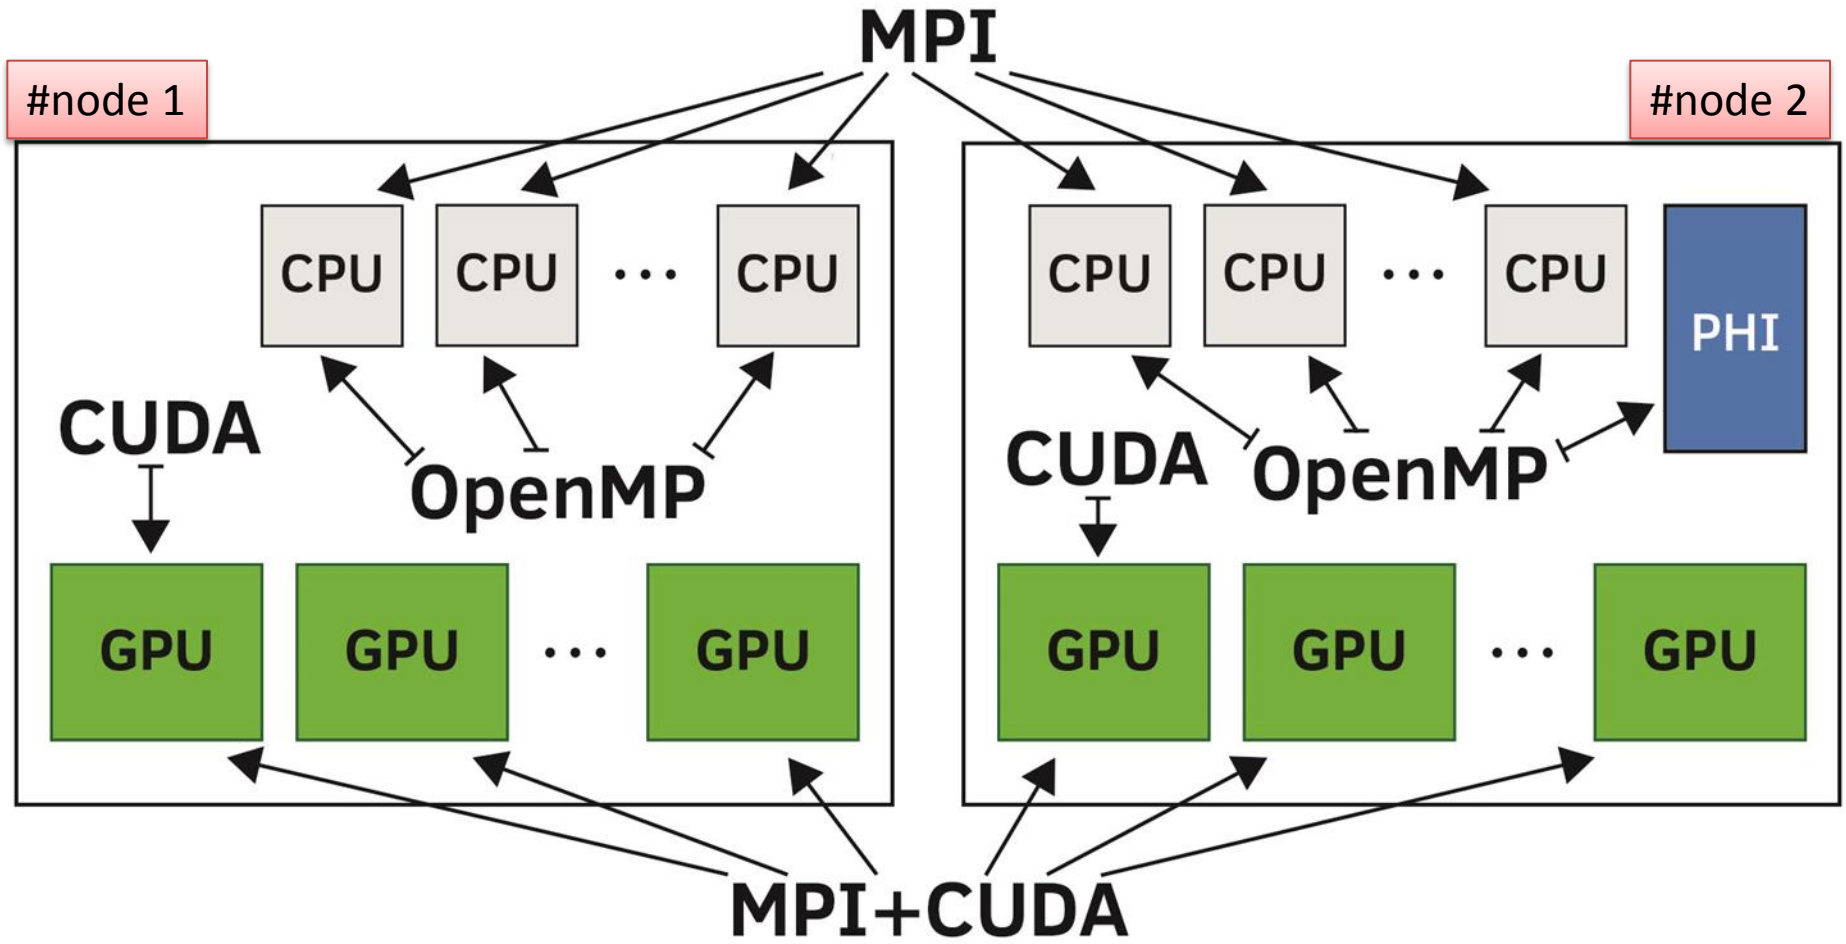

How to control hybrid hardware:
MPI – OpenMP – CUDA - OpenCL ...

In the last decade novel computational facilities and technologies has become available: MPI-OpenMP-CUDA-OpenCL...

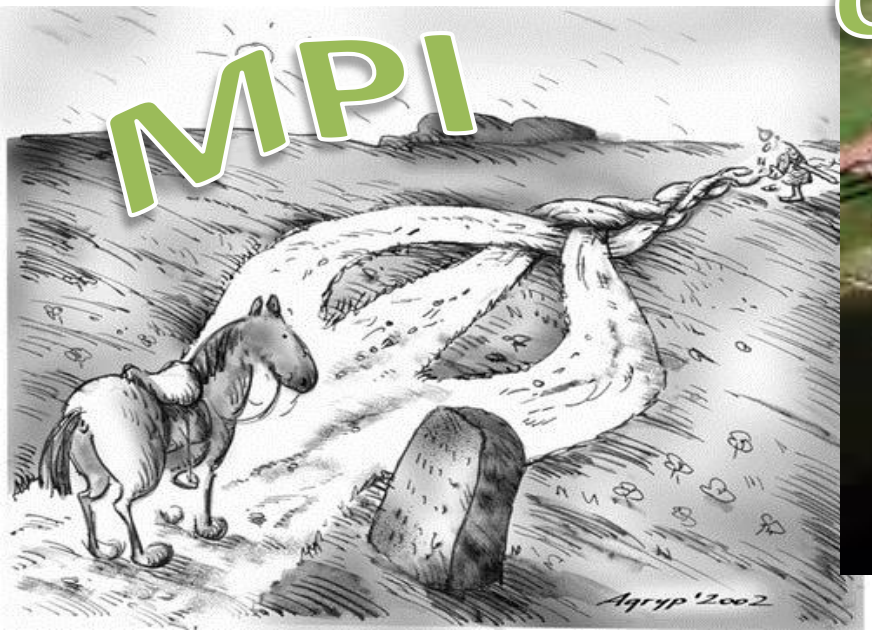

It is not easy to follow modern trends.
Modification of the existing codes or developments of new ones ?

Conclusion

- Modern development of computer technologies (multi-core processors, GPU , coprocessors and other) require the development of new approaches and technologies for parallel programming.
- Effective use of high performance computing systems allow accelerating of researches, engineering development and creation of a specific device.

Thank you for attention!