



Образовательная программа по технологиям НРС на базе гетерогенного кластера HybriLIT (ЛИТ ОИЯИ)

В.В. Кореньков, Д.В. Подгайный, О.И. Стрельцова

Лаборатория информационных технологий

II Научно-практическая конференция "Природа, общество, человек"
Круглый стол «Подготовка ИТ-специалистов для наукоёмких проектов ОИЯИ»

*Дубна
30 ноября, 2016*



Heterogeneous Computation Team, *HybriLIT*

TOP500 List – NOVEMBER 2016 (48 edition)

Rank	Site	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	National Supercomputing Center in Wuxi China	Sunway TaihuLight - Sunway MPP, Sunway SW26010 260C NRCPC	10,649,600	93,014.6	125,435.9	15,371
2	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	3,120,000	33,862.7	54,902.4	17,808
3	DOE/SC/Oak Ridge National Laboratory United States	Titan - Cray XK7 , Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA K20x Cray Inc.	560,640	17,590.0	27,112.5	8,209
• • •						
52	Moscow State University - Research Computing Center Russia	Lomonosov 2 - T-Platform A-Class Cluster, Xeon E5-2697v3 14C 2.6GHz, Infiniband FDR, NVIDIA K40m T-Platforms	42,688	2,102.0	2,962.3	1,079

TOP500 List – November 2016 (48th edition)

A total of **86** systems on the list are using accelerator/co-processor technology.

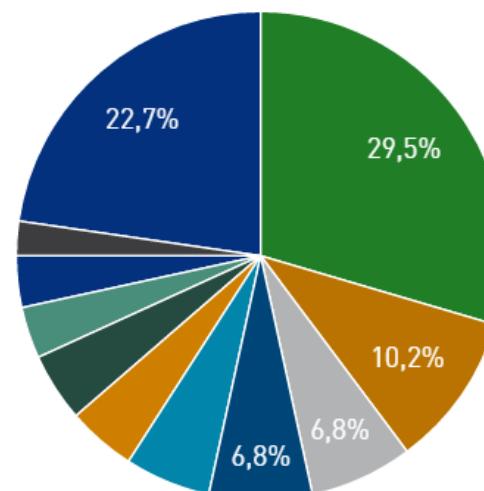
- Sixty (**60**) of these use NVIDIA chips,
- one (**1**) use ATI Radeon, **1** use PEZY technology

- Three (**3**) systems use a combination of NVIDIA and Intel Xeon Phi accelerators/co-processors.

- there are now **21** systems with **Intel Xeon Phi** technology.

10 Systems now use Xeon Phi as the main processing unit.

Accelerator/Co-Processor System Share



- NVIDIA Tesla K40
- NVIDIA Tesla K80
- Intel Xeon Phi 5110P
- NVIDIA Tesla K20x
- Intel Xeon Phi 5120D
- Intel Xeon Phi 7120P
- NVIDIA 2050
- NVIDIA Tesla K20
- NVIDIA Tesla K20m
- NVIDIA 2090
- Others

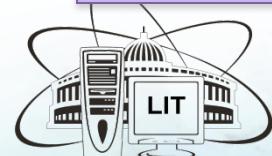
Source:

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

What we see: modern Supercomputers are hybrid with heterogeneous nodes

- Multicore/many-core CPUs with shared memory
- Multiple GPUs
- Multicore/many-core CPUs
- GPUs
- Coprocessors

The practical decisions concerning the implementation of a heterogeneous cluster **HybriLIT** followed from the world wide trends in high performance computing



Heterogeneous Computation Team, *HybriLIT*



HybriLIT: heterogeneous computation cluster

Summary of current version

252 CPU-cores, 77184 GPU-cores,
182 PHI-cores; 2.4 TB RAM; 55.2 TB HDD

Peak performance:

with single precision **142 TFlops**;
with double precision **50 TFlops**

Dell PowerEdge

- 2x Intel Xeon CPU E5-2695v3; - 4x NVIDIA TESLA K80

Supermicro SuperBlade Chassis

-2x Intel Xeon
CPU E5-2695v3
-2x NVIDIA TESLA K80

- 2x Intel Xeon
CPU E5-2695v2
3x NVIDIA TESLA K40

- 2x Intel Xeon
CPU E5-2695v2
- 6x HDD 1.2 TB

- 2x Intel Xeon
CPU E5-2695v2
- NVIDIA TESLA K20X
- Intel Xeon Phi
Coprocessor 5110P

-2x Intel Xeon
CPU E5-2695v2
- 2x Intel Xeon Phi
Coprocessor 7120P



The heterogeneous computing cluster may serve

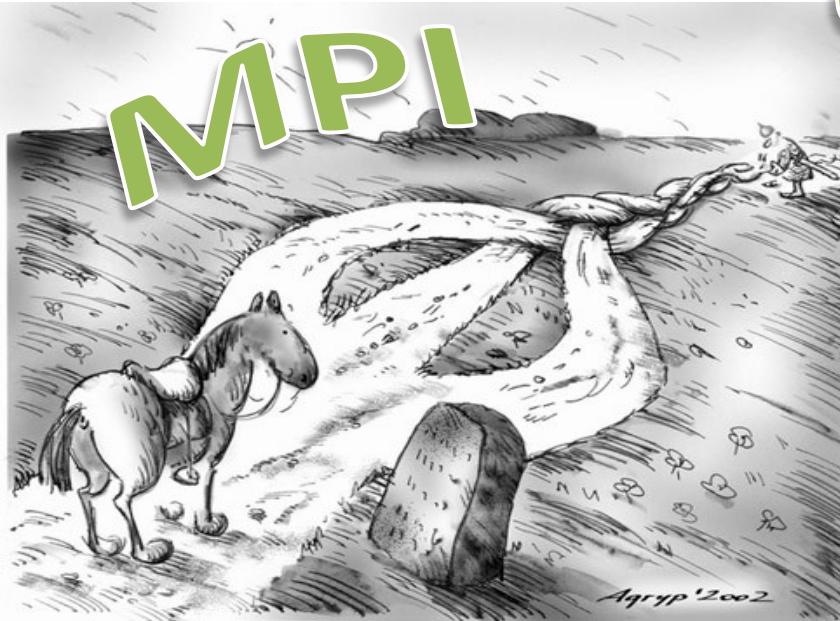
- **to create** own software for investigations demanding resource-intensive computations
- **to use** already developed software products and applied mathematical libraries for calculations on hybrid architectures
- **to develop** parallel algorithms for experimental data processing and analysis using programming paradigms for specialized computing systems consisting of graphic accelerators and co-processors



Heterogeneous Computation Team, *HybriLIT*



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



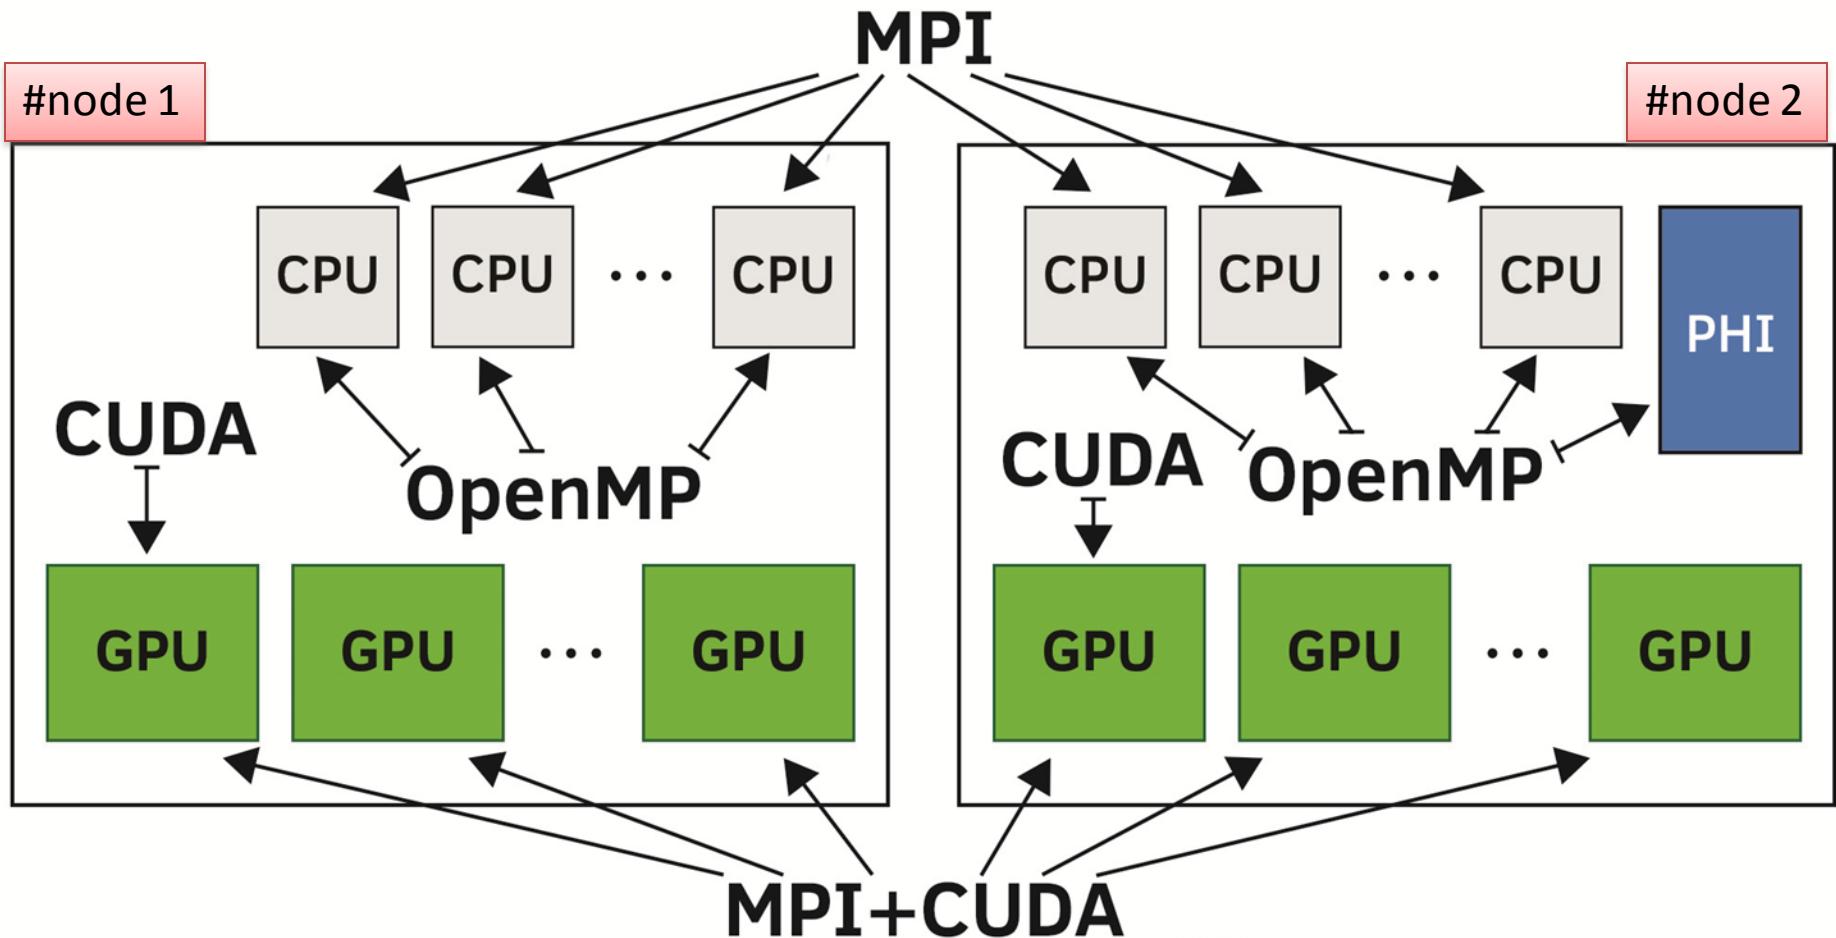
It is not easy to follow modern trends.
Modification of the existing codes or developments of new ones ?



HETEROGENEOUS COMPUTATIONS TEAM, *HybriLIT*



Parallel technologies: levels of parallelism



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

Intel: compilers, tools for developing, debugging and profiling parallel applications, mathematical library

Intel Inspector

Memory and Threading
Checking

Intel VTune

Amplifier
Performance Profiler

Intel Advisor

Vectorization
Optimization and
Thread Prototyping

Intel Math Kernel Library

Optimized Routines for
Science, Engineering,
and Financial

Intel Data Analytics Acceleration Library

Optimized for Data
Analytics & Machine
Learning

Intel® Integrated Performance Primitives

Image, Signal, and
Compression Routines

Intel Math Kernel Library

Optimized Routines for
Science, Engineering,
and Financial

Intel Threading Building Blocks

Task-Based Parallel C++
Template Library

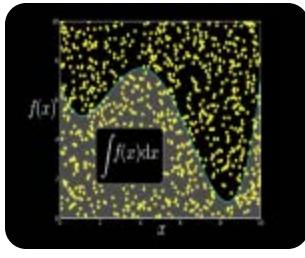
Source: <https://software.intel.com/en-us/articles/intel-parallel-studio-xe-release-notes>



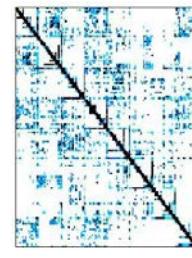
Some GPU-accelerated Libraries



NVIDIA cuBLAS



NVIDIA cuRAND



NVIDIA cuSPARSE



NVIDIA NPP



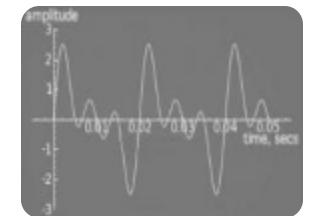
Vector Signal
Image Processing



GPU Accelerated
Linear Algebra



Matrix Algebra
on GPU and
Multicore



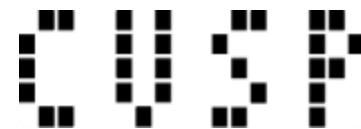
NVIDIA cuFFT



ROGUE WAVE
SOFTWARE
IMSL Library



ArrayFire Matrix
Computations



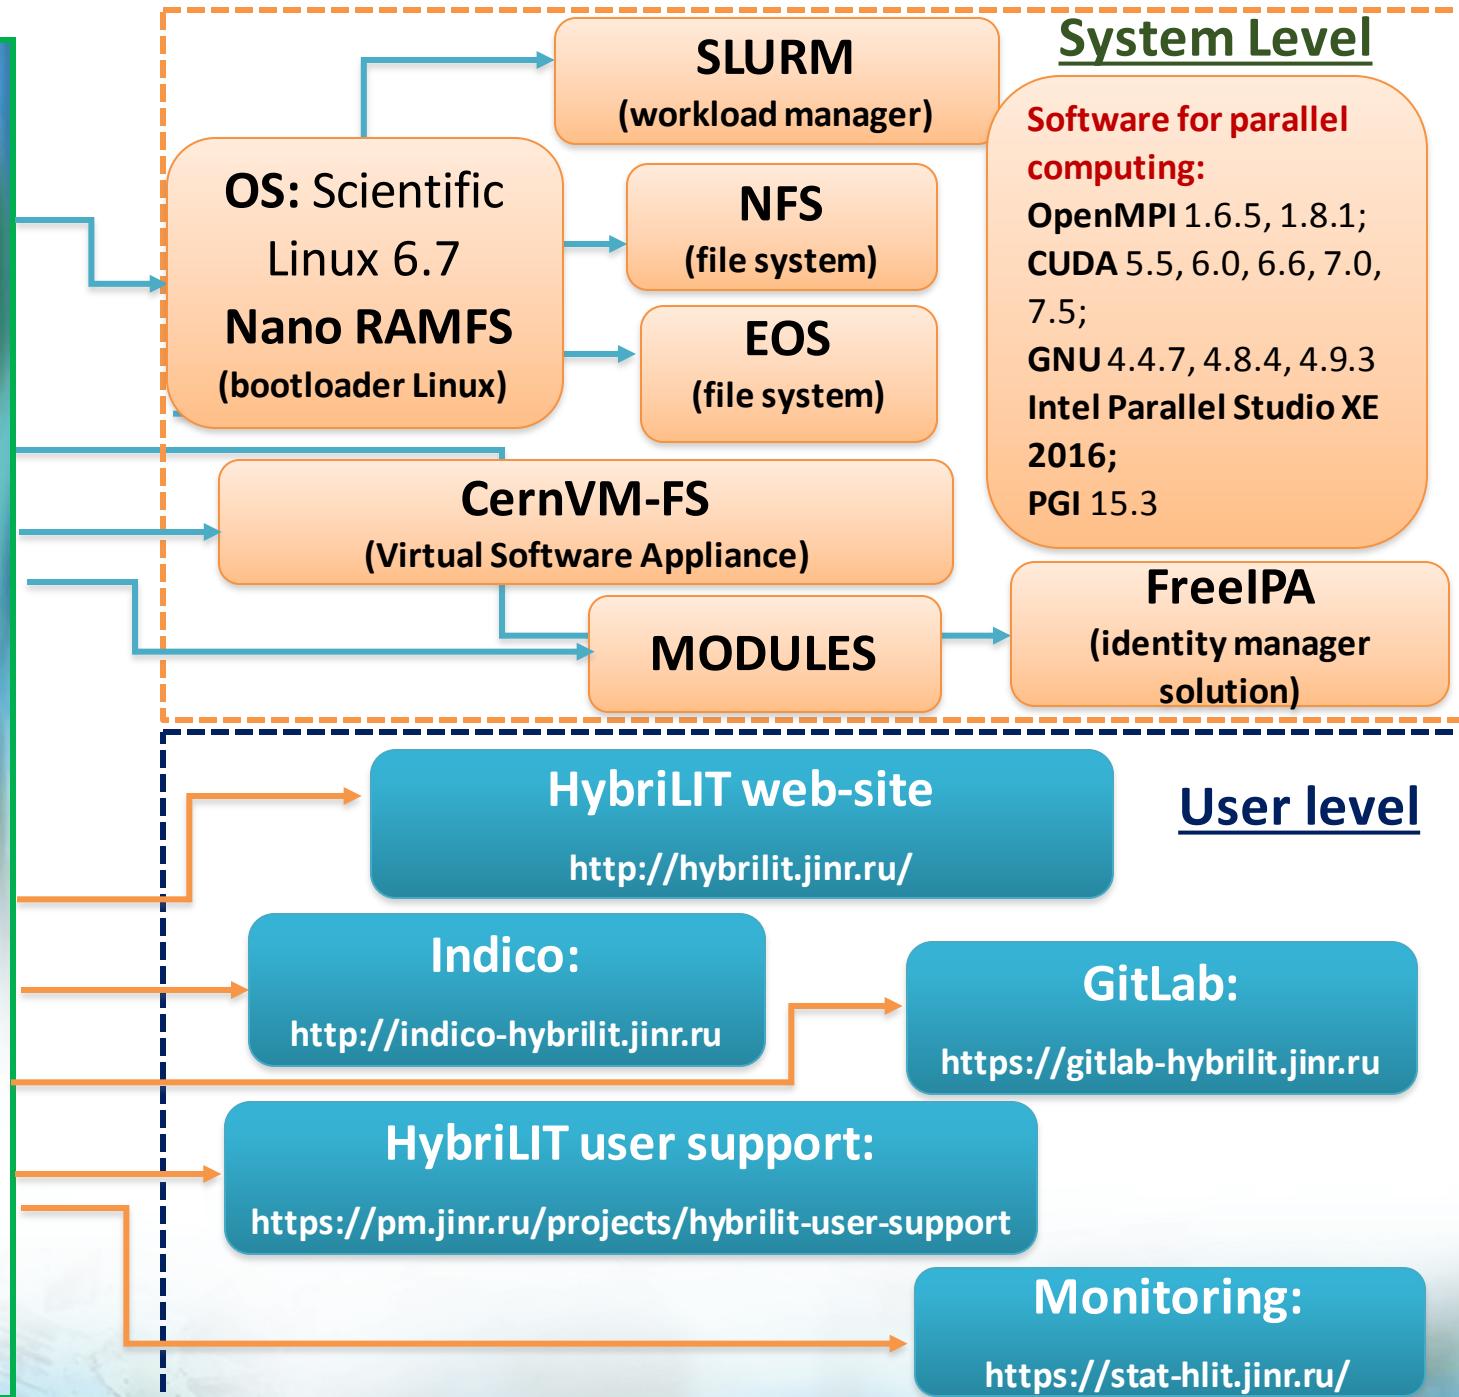
Sparse Linear
Algebra



C++ STL
Features for
CUDA

HybriLIT

Software and Information Environment



41 tutorials and lectures on parallel programming technologies have been held on the basis of the heterogeneous computing cluster HybriLIT

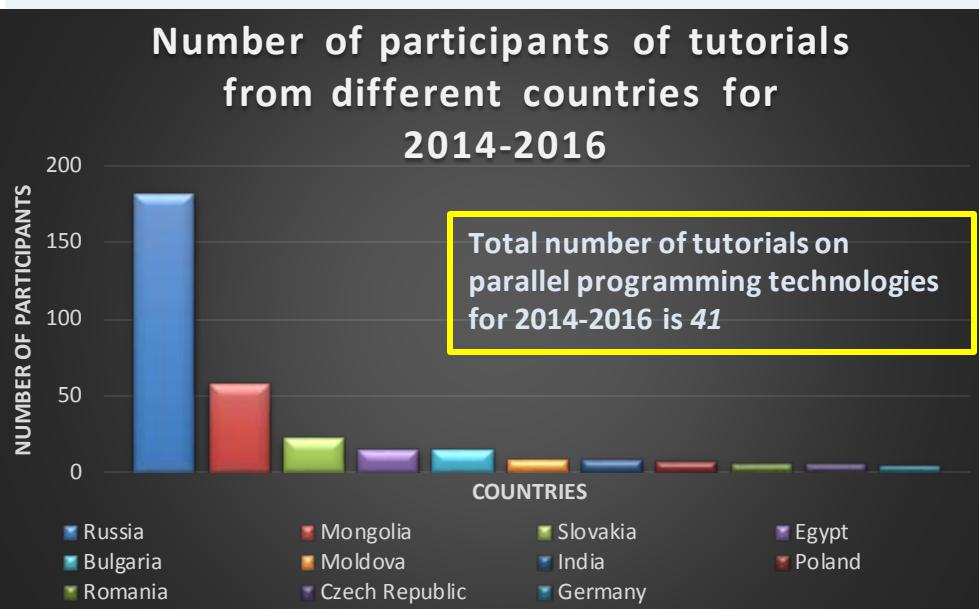


Fig 1. Number of participants from different countries for 2014-2016 yy.

Participants of the tutorials are students and scientists from Russia, Mongolia, Slovakia, Egypt, Bulgaria, India, etc.

The total number of participants comprises over 300 people from different universities and scientific centers

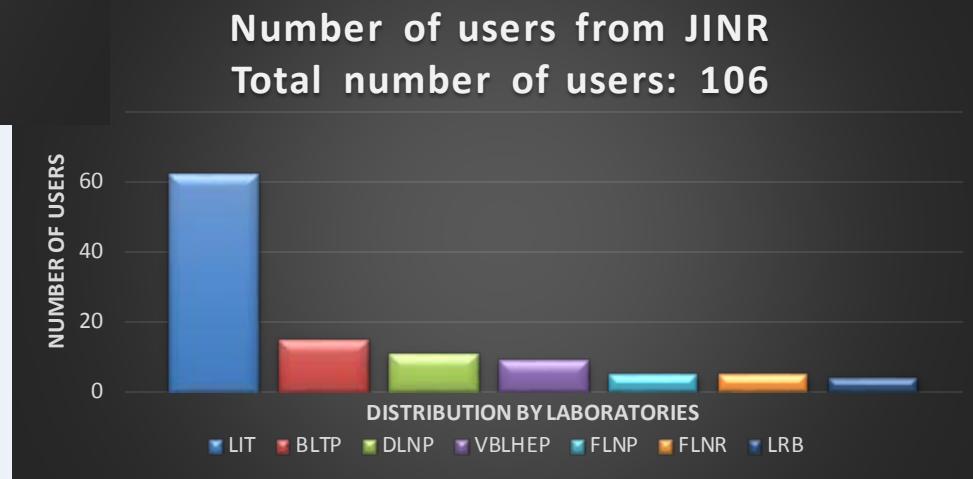


Fig.2. Number of participants from JINR

Dubna University at HybriLIT

Summer school of students from the Dubna University: “Parallel Programming Technologies” - **13 students**

<http://indico-hlit.jinr.ru/event/62/>

Architecture Computing Systems,

Practical Training (2016-2017 academic year) : **140 students**

<https://indico-hlit.jinr.ru/category/21/>

Terminal course on parallel programming technologies:

(Faculty of Natural and Engineering Science) : **20** students

<https://indico-hlit.jinr.ru/event/70/>

Elective course on parallel programming technologies :

(first- and second-year students of SAM) : **12** students

<https://indico-hlit.jinr.ru/event/71/>

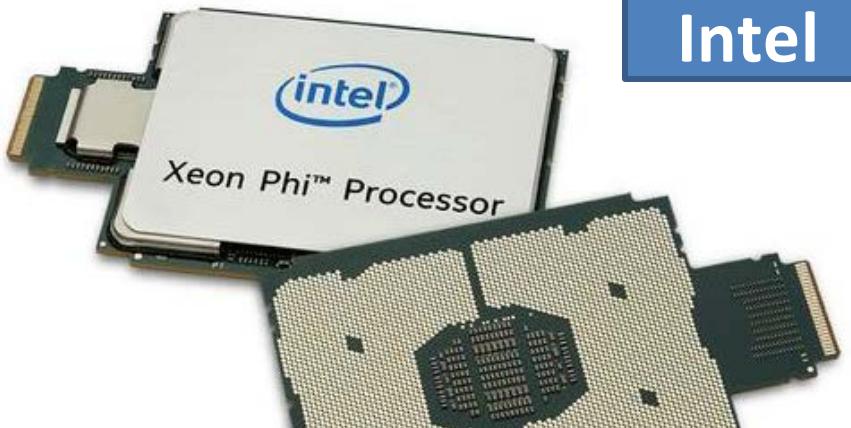
Bachelor's and master's thesis (2014-2016) : **7** students



Heterogeneous Computation Team, *HybriLIT*



New computing architecture and IT-technology



Intel

Intel® Xeon Phi™ Processor 7290F
http://ark.intel.com/products/95831/Intel-Xeon-Phi-Processor-7290F-16GB-1_50-GHz-72-core



NVIDIA

NVIDIA® Tesla® P100 GPU accelerators
<http://www.nvidia.com/object/tesla-p100.html>



ZAO RSC Technologies

<http://www.rscgroup.ru/>



Heterogeneous Computation Team, *HybriLIT*



Future development

Training of specialists in the field of **HPC** in accordance with the needs of JINR laboratories

New educational program specialization



Bio-informatics :
LRB JINR, ISE "Dubna"

Applied Mathematics
for JINR projects and experiments



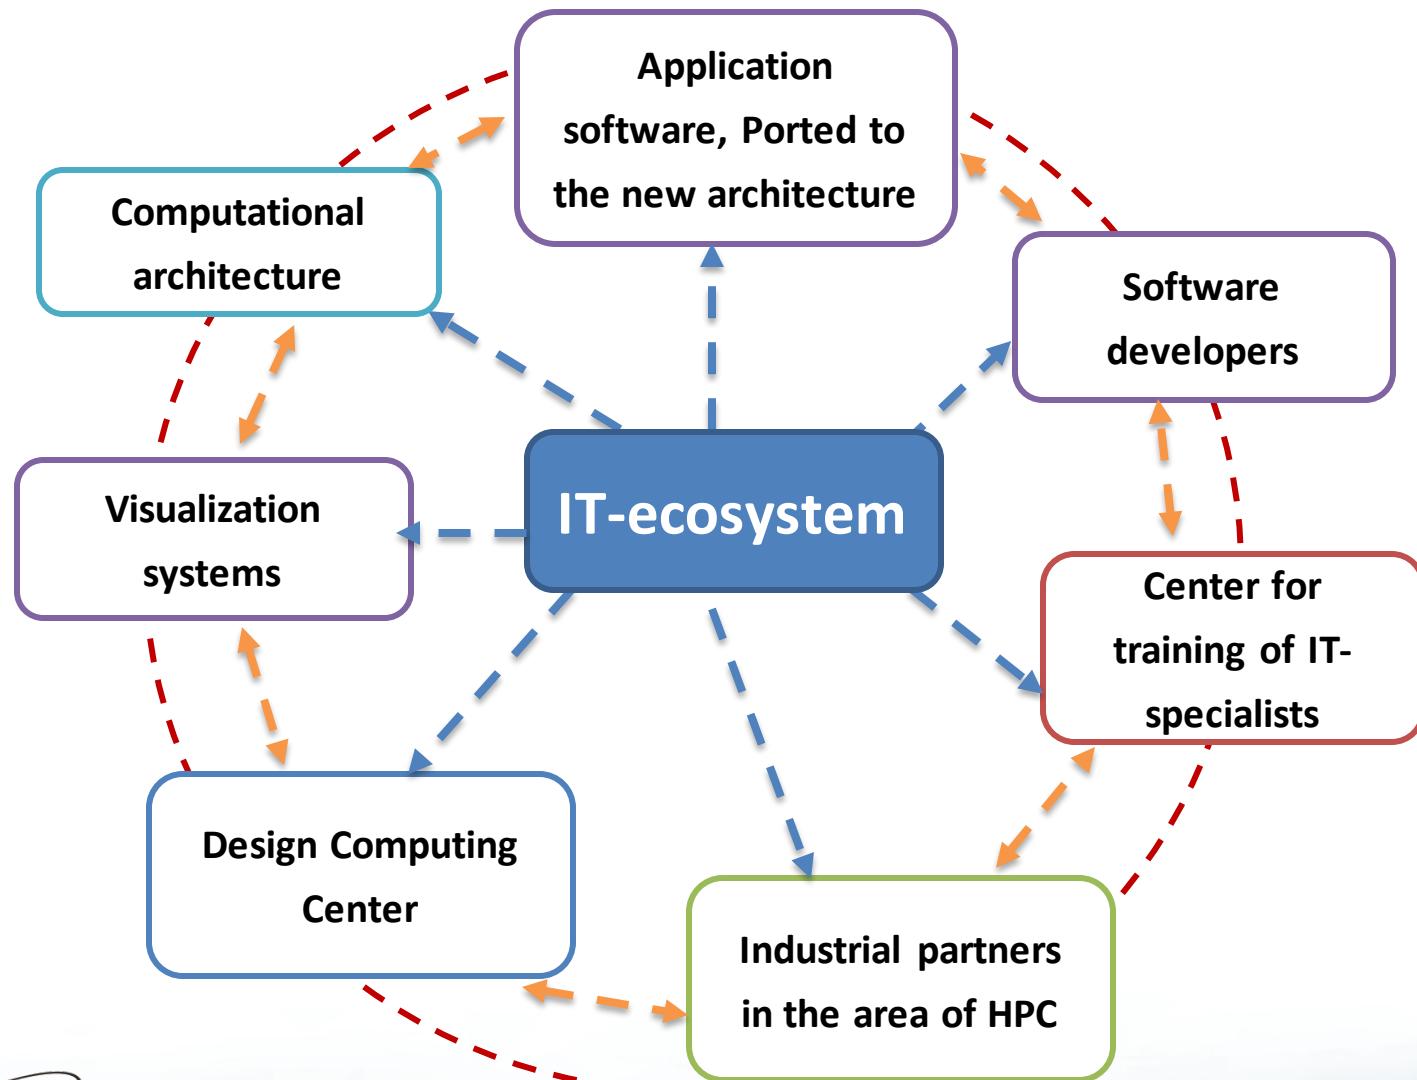
Carrying out tutorials, workshops, lectures made by specialists from the leading companies that develop new architectures and software.



Heterogeneous Computation Team, *HybriLIT*



The proposal on creation of central IT-support mega-projects in high energy physics



Heterogeneous Computation Team, *HybriLIT*



Thank you for attention!



Heterogeneous Computation Team, *HybriLIT*

