



Heterogeneous cluster HybriLIT: New possibilities for the development of parallel applications

Torosyan Shushanik

on behalf of the

Heterogeneous Computation Team HybriLIT

Laboratory of Information Technologies, JINR

Dubna

July 6-12, 2016

Heterogeneous Computation Team, *HybriLIT*



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



TOP500 List – November 2015

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	3,120,000	33,862.7	54,902.4	17,808
2	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
3	DOE/NNSA/LLNL United States	Sequoia - BlueGene/Q, Power BQC 16C 1.60 GHz, Custom IBM	1,572,864	17,173.2	20,132.7	7,890
35	Moscow State University - Research Computing Center Russia	Lomonosov 2 - T-Platform A-Class Cluster, Xeon E5-2697v3 14C 2.6GHz, Infiniband FDR, Nvidia K40m	37,120	1,849.0	2,575.9	

TOP500 List – November 2015

A total of **104** systems on the list are using accelerator/co-processor technology, up from **90** on June 2015.


- Sixty-six (**66**) of these use NVIDIA chips,
- three (**3**) use ATI Radeon,

- there are now **27** systems with **Intel Xeon Phi** technology.

- Four (**4**) systems use a combination of Nvidia and Intel Xeon Phi accelerators/co-processors.

- four (**4**) – other coprocessors

PERFORMANCE OF ACCELERATORS



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

What we see: modern Supercomputers are hybrid with heterogeneous nodes

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

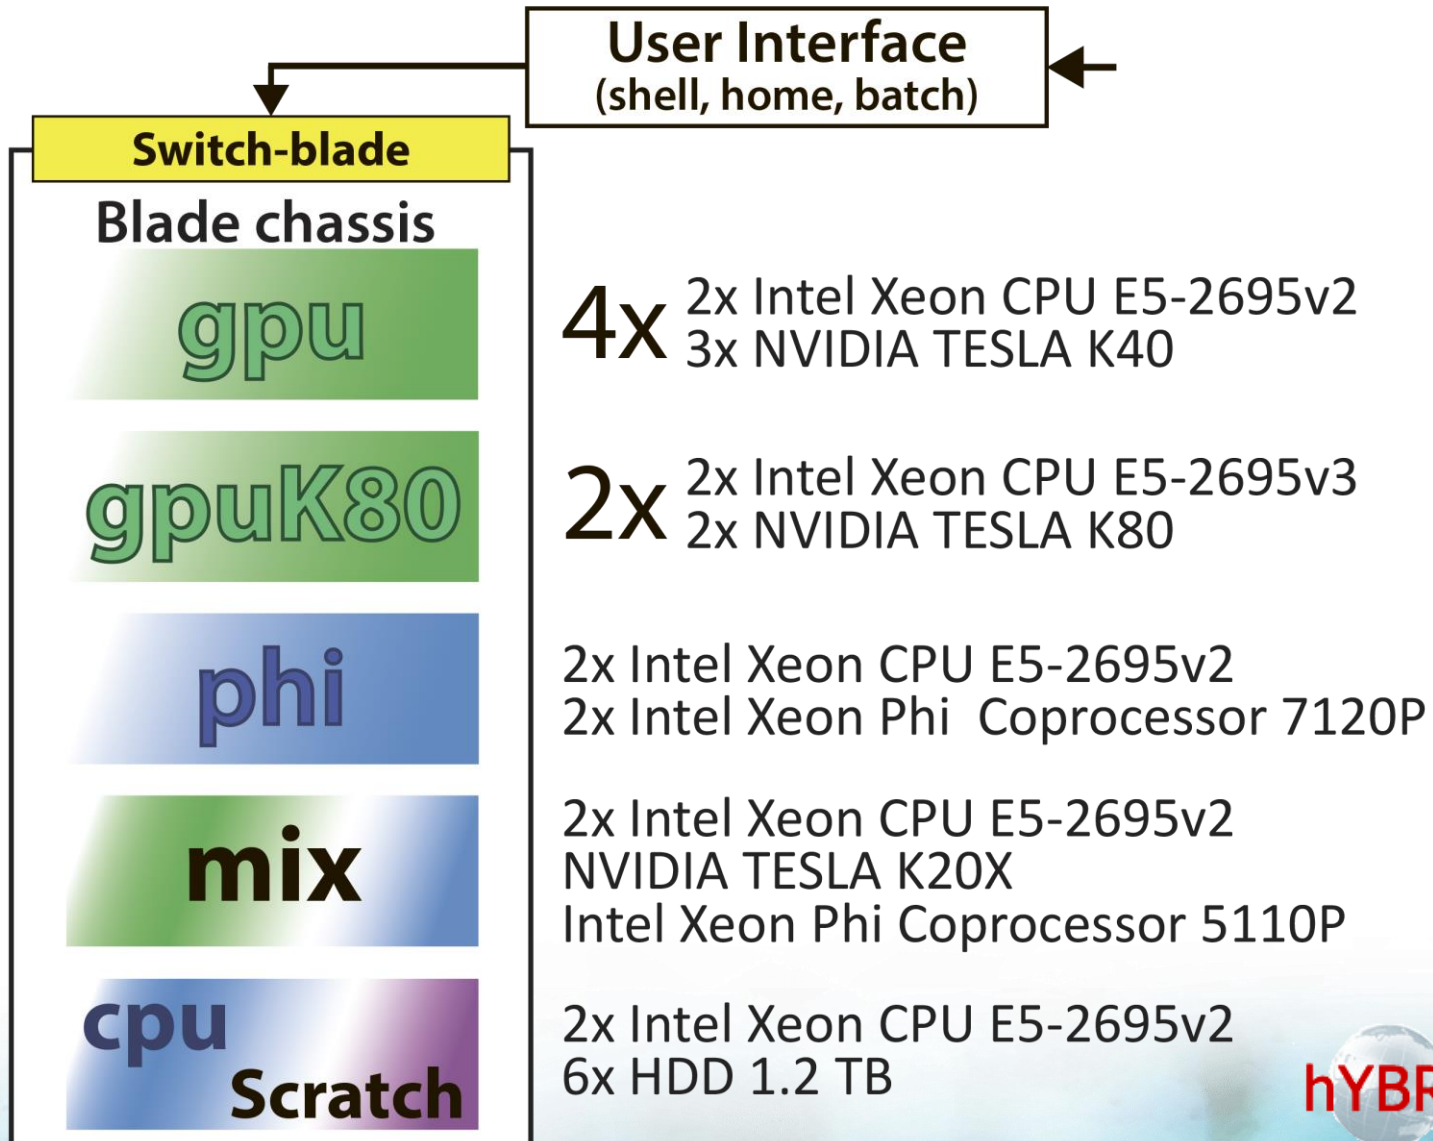
- Multicore/many-core CPUs with shared memory
- Multiple Coprocessors

- 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**



HybriLIT: heterogeneous computation cluster



HybriLIT: heterogeneous computation cluster

Computing resources:

14x CPU Intel Xeon	224 cores
14x GPU K20, K40 & K80	57216 cores
3x Intel Xeon Phi 5110P & 7210P	182 cores

RAM **1920** GB

EOS storage **55.2** TB

Ethernet

InfiniBand **40** GB/s

Power consumption **8.29** kW

Energy efficiency **4.56** GFlops/W

Peak power demand **10** kW

Peak performance for floating point computations:

Single precision **105.7** TFlops

Double precision **37.8** TFlops

NVIDIA Tesla K40 "Atlas" GPU Accelerator



NVIDIA Tesla K80 GPU Accelerator



Specifications

Tesla K40


2880 CUDA GPU cores
Memory 12 GB GDDR5
**Peak precision floating
point performance**
4.29 TFLOPS single-precision
1.43 TFLOPS double-precision

Tesla K80

Number of GPUs: **2x Kepler GK210**
4992 (2496 per GPU) CUDA GPU cores
Memory 24 GB (12 GB per GPU)
**Peak precision floating point
performance**
8.74 TFLOPS single-precision (with GPU Boost)
2.91 TFLOPS double-precision (with GPU Boost)



Intel® Xeon Phi™ Coprocessor



Intel Xeon Phi 7120P

Clock Speed **1.24 GHz**

L2 Cache **30.5 MB**

Cores **61**

More threads **244**

Peak precision floating point performance

2.42 TFLOPS single-precision

1.21 TFLOPS double-precision



Heterogeneous Computation Team, *HybriLIT*



HybriLIT

Software and
Information
Environment

**OS: Scientific
Linux 6.7
Nano RAMFS
(bootloader Linux)**

**SLURM
(workload manager)**

**NFS
(file system)**

**EOS
(file system)**

**CernVM-FS
(Virtual Software Appliance)**

MODULES

System Level

**Software for parallel
computing:**
OpenMPI 1.6.5, 1.8.1;
CUDA 5.5, 6.0, 6.6, 7.0,
7.5;
GNU 4.4.7, 4.8.4, 4.9.3
Intel Parallel Studio XE
2016;
PGI 15.3

**FreeIPA
(identity manager
solution)**

HybriLIT web-site

<http://hybrilit.jinr.ru/>

Indico:

<http://indico-hybrilit.jinr.ru>

GitLab:

<https://gitlab-hybrilit.jinr.ru>

HybriLIT user support:

<https://pm.jinr.ru/projects/hybrilit-user-support>

Monitoring:

<https://stat-hlit.jinr.ru/>

User level

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

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

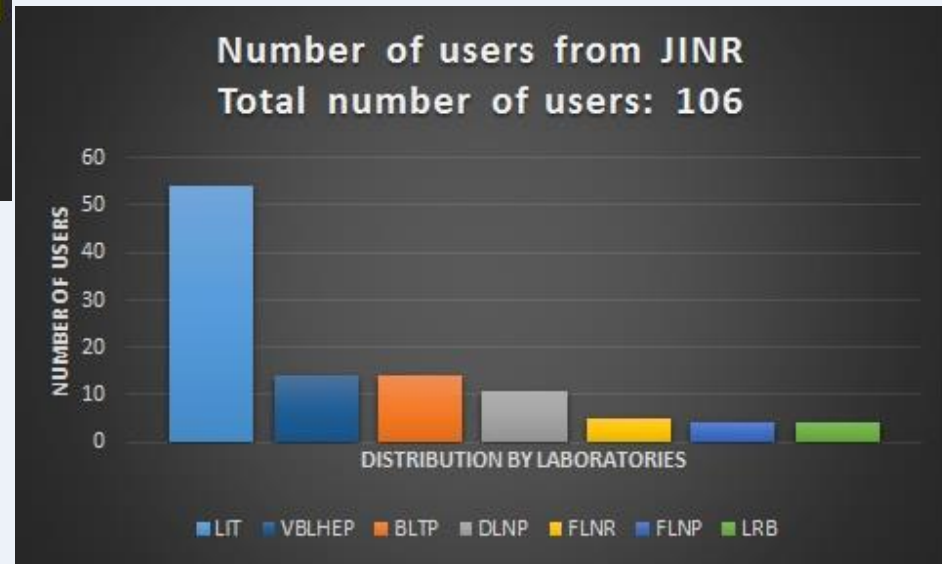


Fig.2. Number of participants from JINR




Conclusions

- ✓ Heterogeneous cluster HybriLIT at LIT JINR is being actively developed by the heterogeneous computation team.
- ✓ The developed software and information environment allows users of the cluster carry out resource-demanding computations, develop parallel applications on the newest computation architectures for solution of a variety of tasks.



Thank you for attention!



Heterogeneous Computation Team, *HybriLIT*

