
International Conference
Russian Supercomputing Days
September 25-26, 2017, Moscow

International Conference

Russian Supercomputing Days

Supported by the Russian Foundation for Basic Research

Platinum Sponsor – Educational Partner:

Platinum Sponsors:

Golden Sponsor:

Silver Sponsors:

In Cooperation with:

Conference Operator: Top Level Meetings

September 25-26, 2017, Moscow

Russian Academy of Sciences: President Elections
(September, 25th-29th)

Е.Н.Каблов

Г.Я.Красников

Р.И.Нигматулин

В.Я.Панченко

А.М.Сергеев

World Supercomputing News
(IDC-Projected Exascale Systems Details)

U.S.

- Sustained ES: 2023
- Peak ES: 2021
- Vendors: U.S.
- Processors: U.S.
- Initiatives: NSCI/ECP
- Cost: $300-500M per system, plus heavy R&D investments

EU

- Sustained ES: 2023-24
- Peak ES: 2021
- Vendors: U.S., Europe
- Processors: U.S., ARM
- Initiatives: PRACE, ETP4HPC
- Cost: $300-$350 per system, plus heavy R&D investments

China

- Sustained ES: 2023
- Peak ES: 2020
- Vendors: Chinese
- Processors: Chinese (plus U.S.?)
- 13th 5-Year Plan
- Cost: $350-500M per system, plus heavy R&D

Japan

- Sustained ES: 2023-24
- Peak ES: Not planned
- Vendors: Japanese
- Processors: Japanese
- Cost: $600-850M, this includes both 1 system and the R&D costs…will also do many smaller size systems

"Hyperion (IDC) Paints a Bullish Picture of HPC Future", © Hyperion, 2017

World Supercomputing News
(IDC-Projected Exascale Investment Levels)

U.S.

- $1 to $2 billion a year in R&D (including NRE)
- Investments by both governments & vendors
- Plans are to purchase multiple exascale systems

EU

- About 5 billion euros in total
- Investments in multiple exascale and pre-exascale systems
- Investments mostly by country governments with a little from the EU

China

- Over $1billion a year in R&D
- Investments by both governments & vendors
- Plans are to purchases multiple exascale systems each year
- Already investing in 3 pre-exascale systems by 2017/18

Japan

- Planned investment of just over $1billion* (over 5 years) for both the R&D and purchase of 1 exascale system
- To be followed by a number of smaller systems ~$100M to $150M each
- Creating a new processor and a new software environment

"Hyperion (IDC) Paints a Bullish Picture of HPC Future", © Hyperion, 2017

World Supercomputing News

China Upgrading Milky Way 2 Supercomputer to 95 Petaflops

September 19, 2017 by Rich Brueckner Leave a Comment

We have some breaking news from the IHPC Forum in Guangzhou today. Researchers in China are busy upgrading the MilkyWay 2 (Tianhe-2) system to nearly 95 Petaflops (peak). This should nearly double the performance of the system, which is currently ranked at #2 on TOP500 with 33.86 Petaflops on the Linpack benchmark. The upgraded system, dubbed Tianhe -2A, should be completed in the coming months.

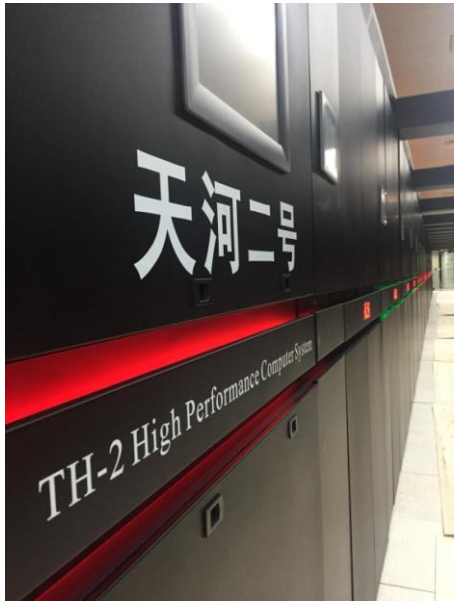

Components	TianHe-2	TianHe-2A
Nodes and performance	16,000 nodes with Intel CPUs + KNC	17,792 nodes with Intel CPUs + Matrix-2000
	54.9 Pflop/s	94.97 Pflop/s
Interconnection	10 Gbps, 1.57 us	14 Gbps, 1 us
Memory	1.4 PB	3.4 PB
Storage	12.4 PB, 512 GB/s	19 PB, 1 TB/s (upgrading, maybe larger)
Energy efficiency	17.8 MW, 1.9 Gflop/s per Watt	16.9MW, >5 Gflop/s per Watt (predicted)
Heterogeneous software	MPSS for Intel KNC	OpenMP/OpenCL for Matrix-2000

Commission on Supercomputing Technologies
(Ministry of Telecom and Mass Communications)

В Минкомсвязи России обсудили перспективы развития суперкомпьютерных и распределенных вычислений

Москва, 24 августа 2017 года. — В Министерстве связи и массовых коммуникаций Российской Федерации состоялось первое заседание Межведомственной комиссии по суперкомпьютерным и грид-технологиям при президиуме Совета при Президенте Российской Федерации по модернизации экономики и инновационному развитию. Представители министерств и ведомств обсудили перспективы развития соответствующей отрасли в России, ее актуальные проблемы и пути их решения. Мероприятие прошло под председательством главы Минкомсвязи России Николая Никифорова.

Top50 of the most powerful Russian supercomputers
(top50.supercomputers.ru)

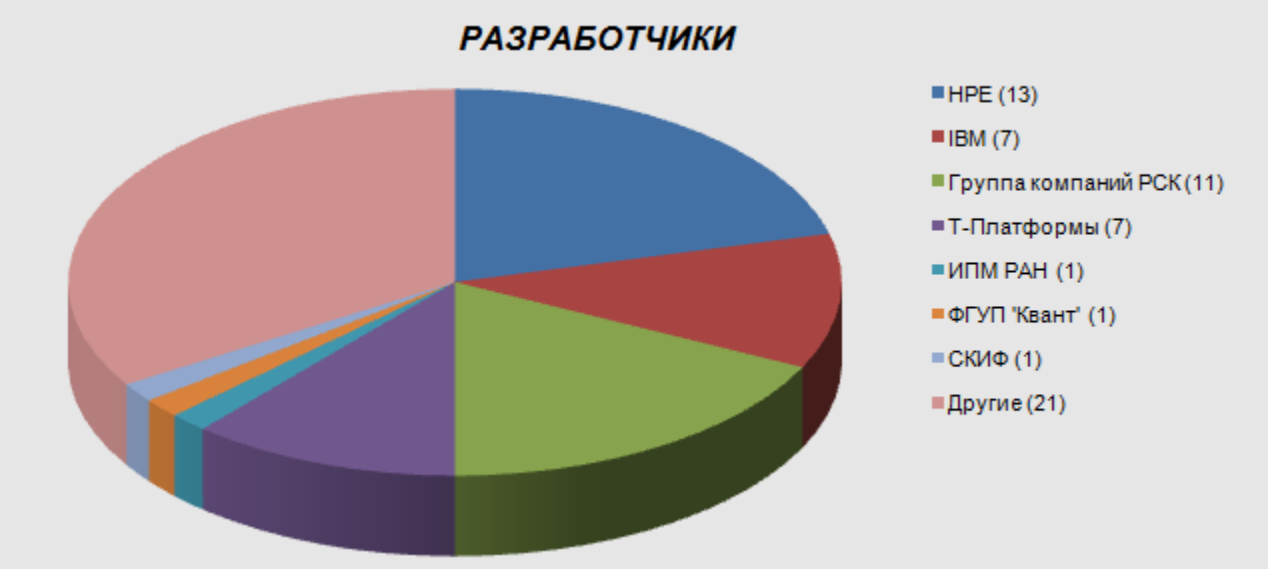

NVIDIA DGX-1 в Top50:

31	Москва ЛНСиГО МФТИ 2017 г.	4/976	узлов: 2 (2xXeon E5-2698v4 [Acc: 8xTesla P100] 2.2 GHz 512 GB RAM) сеть: Infiniband EDR/10 Gigabit Ethernet Copper/Fast Ethernet	56.58	87.68	NVIDIA
32	Москва Сколковский Институт Науки и Технологий 2017 г.	4/976	узлов: 2 (2xXeon E5-2698v4 [Acc: 8xTesla P100] 2.2 GHz 512 GB RAM) сеть: Infiniband EDR/10 Gigabit Ethernet Optical/Fast Ethernet	56.58	87.68	NVIDIA

The most Green Supercomputing Platforms
(June 2017)

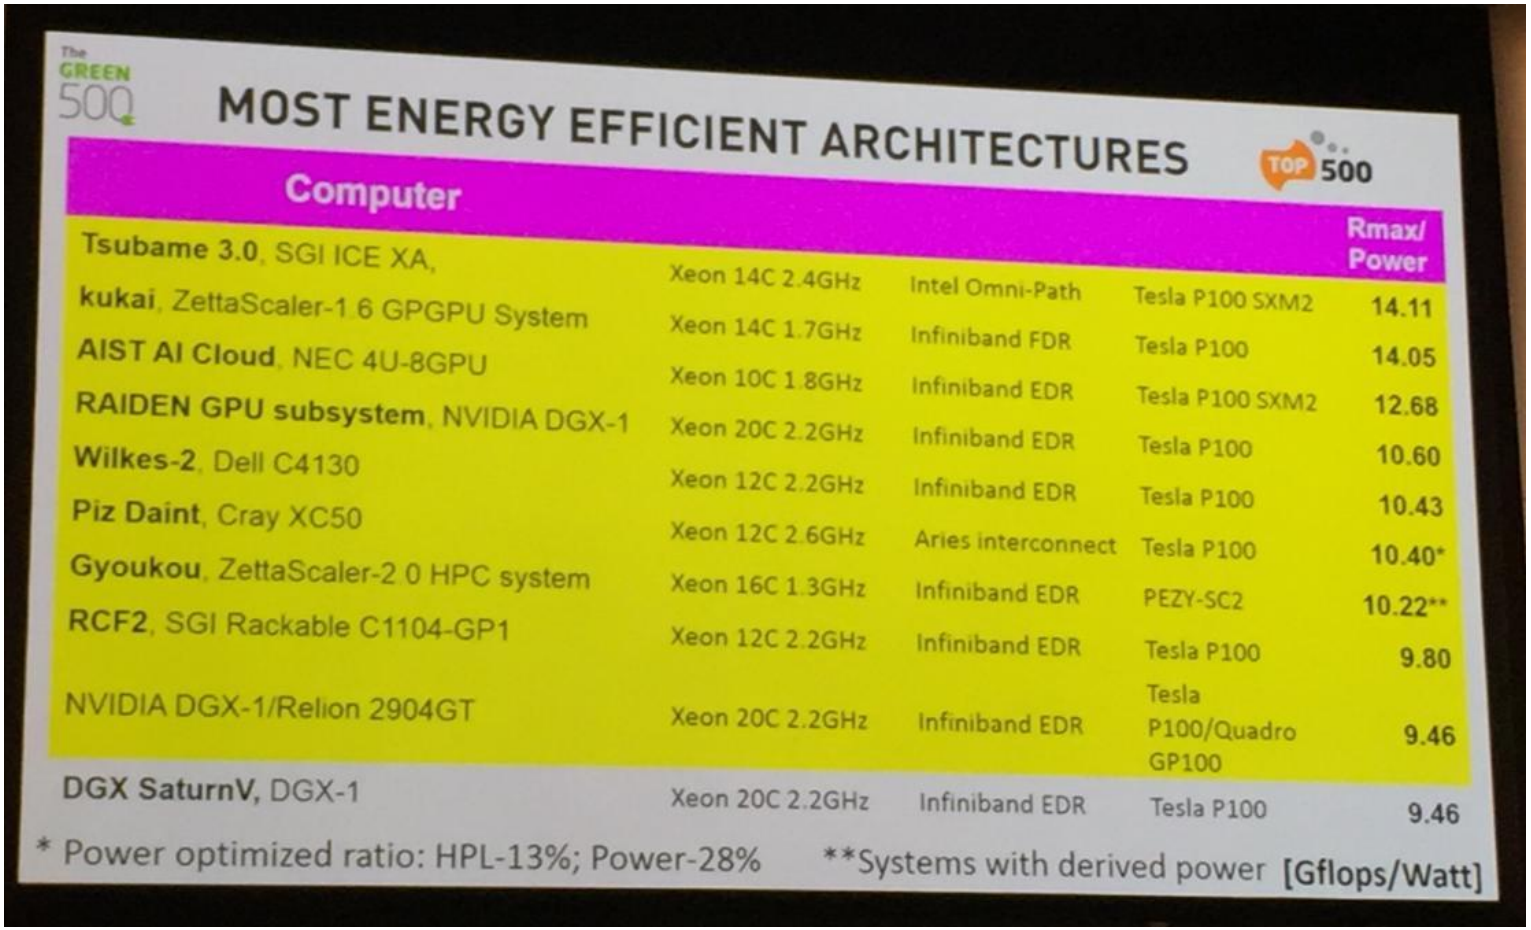

The GREEN 500

MOST ENERGY EFFICIENT ARCHITECTURES

TOP 500

Computer				Rmax/ Power
Tsubame 3.0, SGI ICE XA,	Xeon 14C 2.4GHz	Intel Omni-Path	Tesla P100 SXM2	14.11
kukai, ZettaScaler-1.6 GPGPU System	Xeon 14C 1.7GHz	Infiniband FDR	Tesla P100	14.05
AIST AI Cloud, NEC 4U-8GPU	Xeon 10C 1.8GHz	Infiniband EDR	Tesla P100 SXM2	12.68
RAIDEN GPU subsystem, NVIDIA DGX-1	Xeon 20C 2.2GHz	Infiniband EDR	Tesla P100	10.60
Wilkes-2, Dell C4130	Xeon 12C 2.2GHz	Infiniband EDR	Tesla P100	10.43
Piz Daint, Cray XC50	Xeon 12C 2.6GHz	Aries interconnect	Tesla P100	10.40*
Gyoukou, ZettaScaler-2.0 HPC system	Xeon 16C 1.3GHz	Infiniband EDR	PEZY-SC2	10.22**
RCF2, SGI Rackable C1104-GP1	Xeon 12C 2.2GHz	Infiniband EDR	Tesla P100	9.80
NVIDIA DGX-1/Relion 2904GT	Xeon 20C 2.2GHz	Infiniband EDR	Tesla P100/Quadro GP100	9.46
DGX SaturnV, DGX-1	Xeon 20C 2.2GHz	Infiniband EDR	Tesla P100	9.46

* Power optimized ratio: HPL-13%; Power-28% **Systems with derived power [Gflops/Watt]

[Gflops/Watt]

Courtesy of E.Strohmaier, LBNL, USA


CREDIT SUISSE
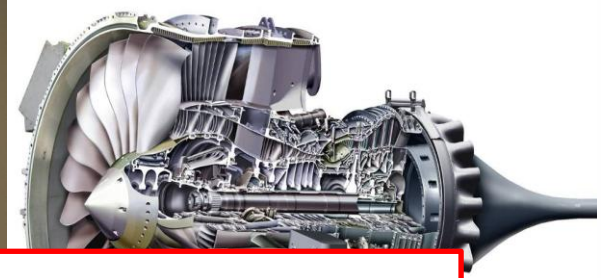
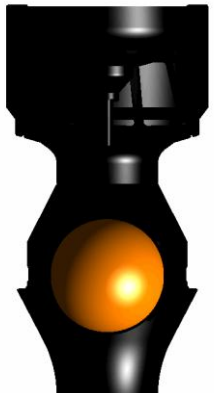
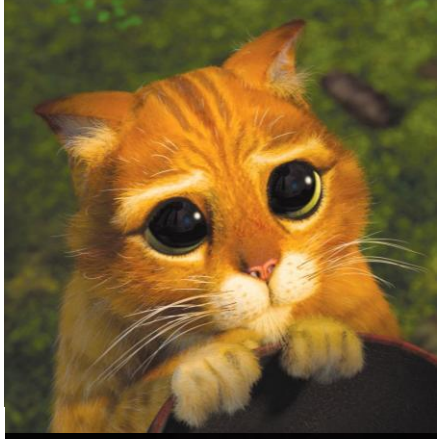

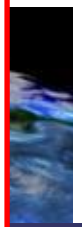
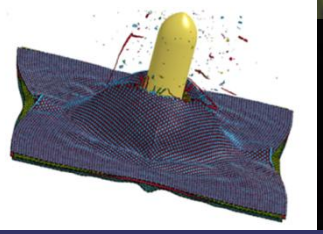

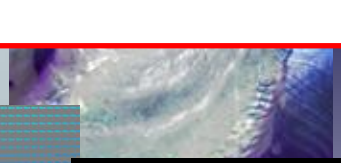
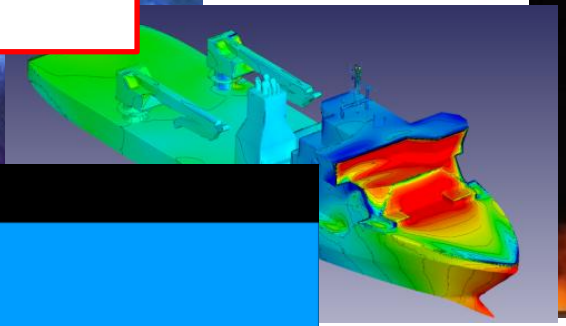
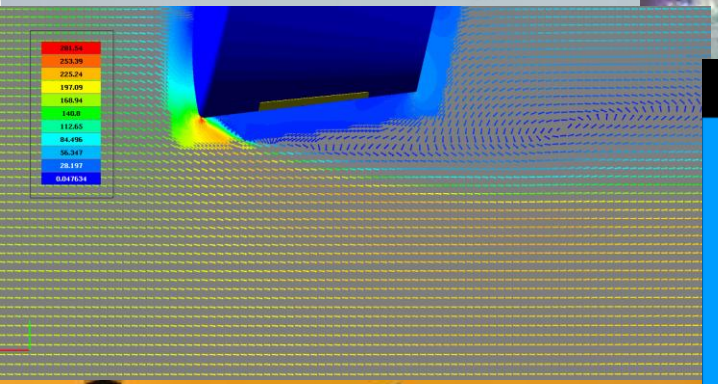
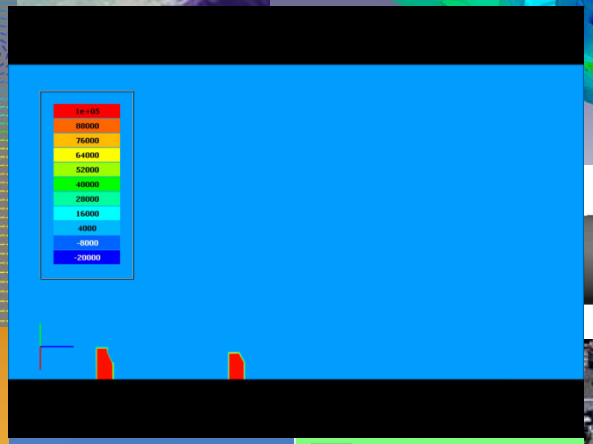
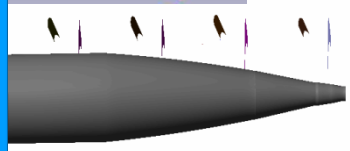


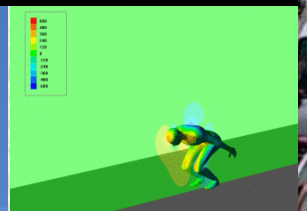


World Supercomputing News, Trends, Perspectives
(Deep Learning / Artificial Intelligence deluge…)

DL: from prediction to active control!

DL/AI: from double and single precision to half-precision
(16 bits: exponent (5 bits), fraction (10 bits), sign (1 bit)):
AI-precision, AI-flops, AI-Eflops (130 Pflops in Japan…).

AI: the notion should be used carefully, imitation Intelligence…
Intelligence achieved through artificial means!

AI/DL haven't shown exascale computing demand as some scientific projects do.

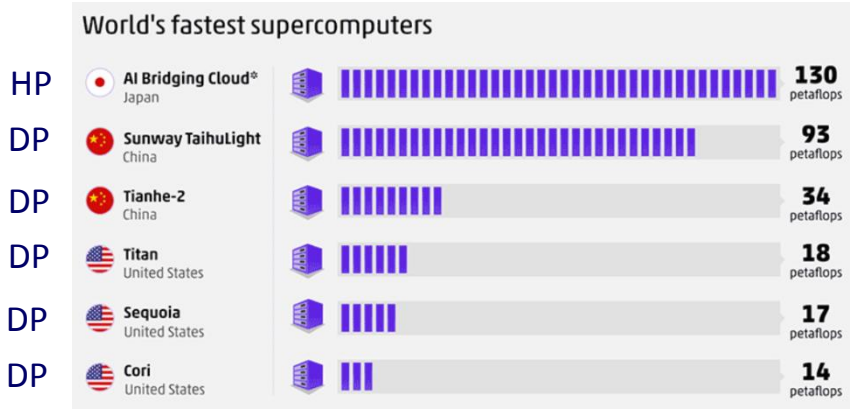

(edition.cnn.com/2017/06/13/tech/…)

World Supercomputing News, Trends, Perspectives
(huge diversity of computer devices)

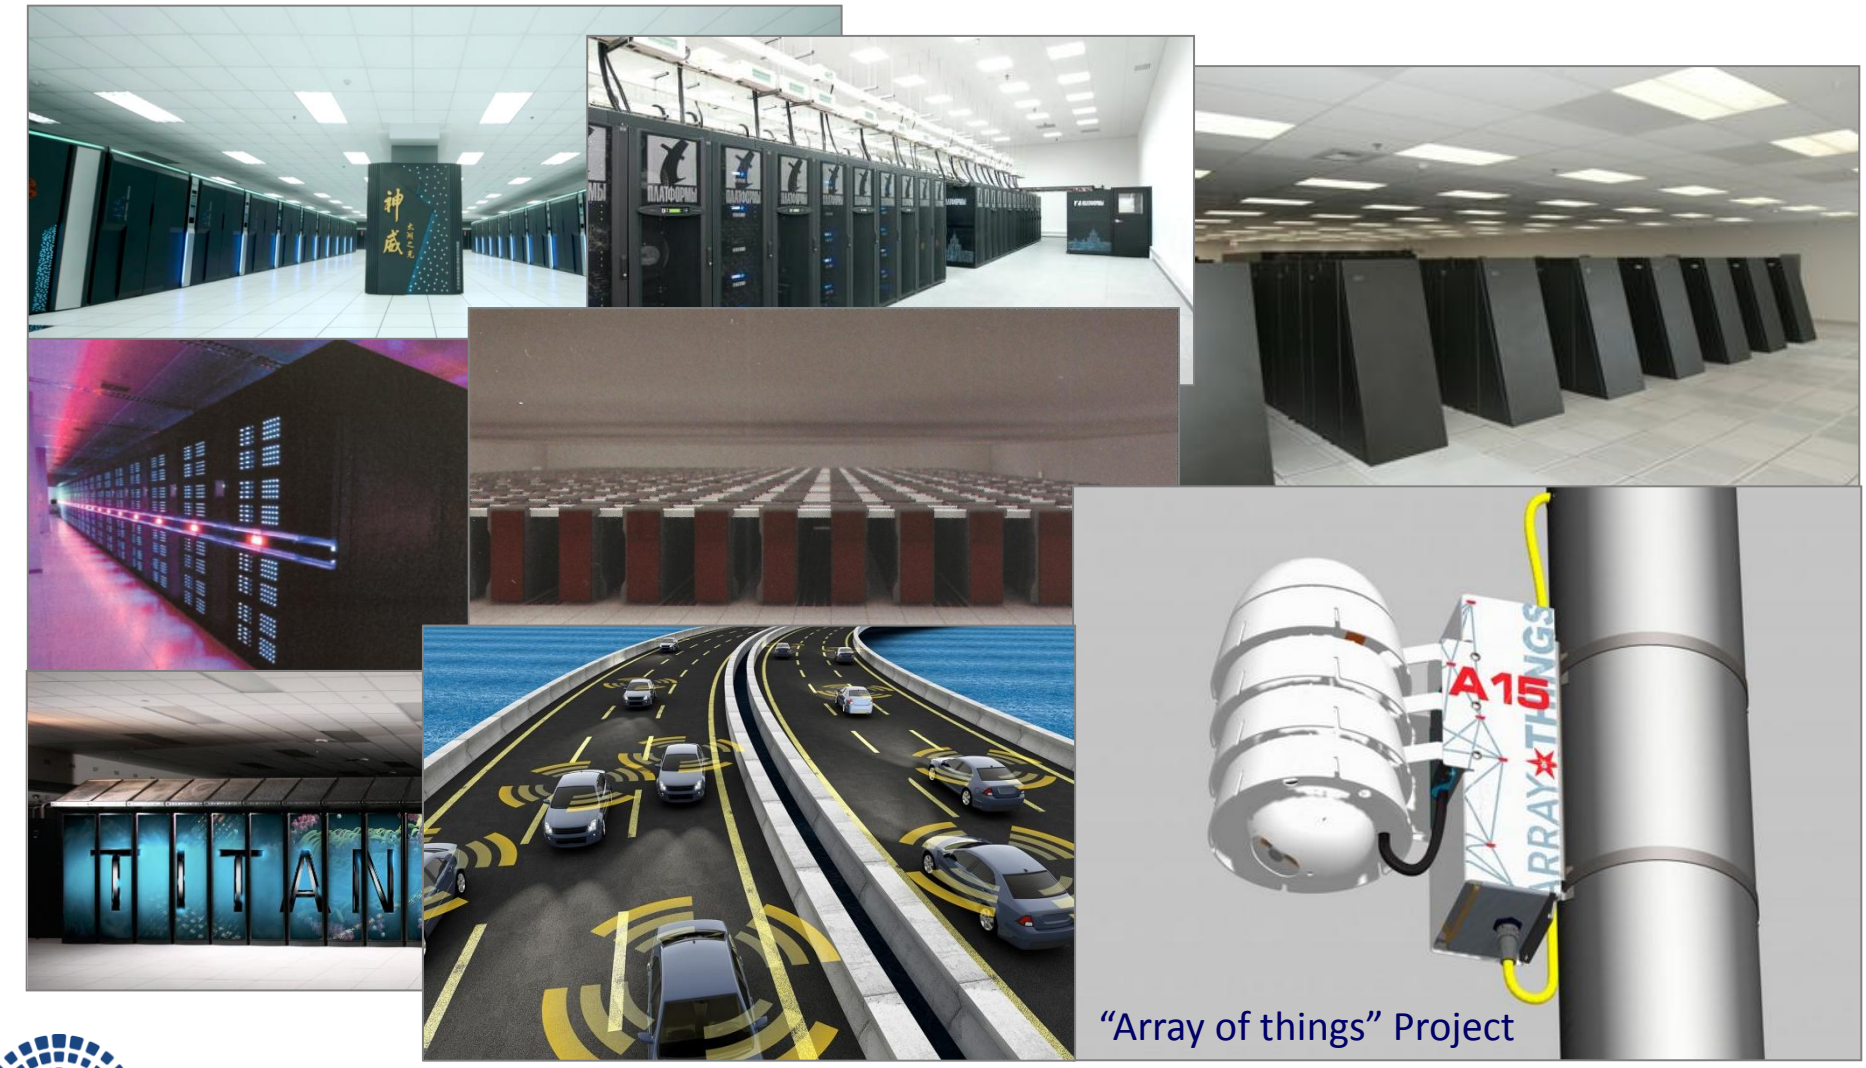

“Array of things” Project

Square Kilometre Array: A Leading Big Data Challenge for 2020

Antennas

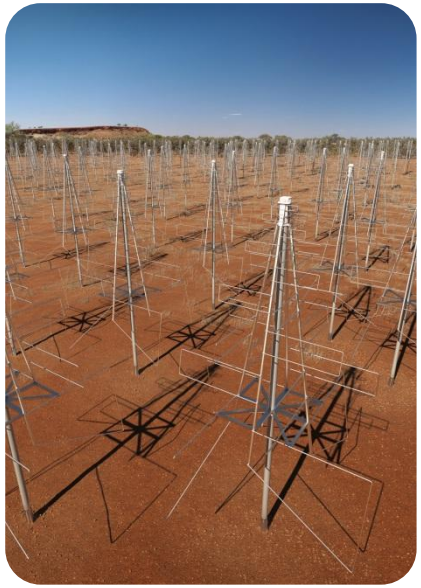

Digital Signal Processing (DSP)

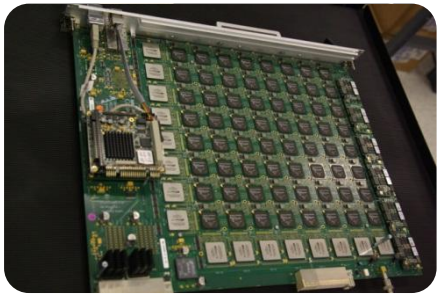

Transfer antennas to DSP
2020: 20,000 PBytes/day
2028: 200,000 PBytes/day

Over 10's to 1000's kms

Imaging – HPC problem
2020: 100 PBytes/day
2028: 10,000 PBytes/day

Over 10's to 1000's kms

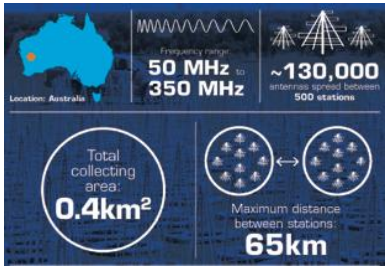

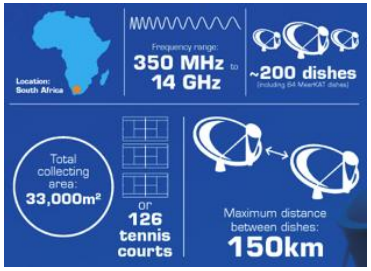

High Performance Computing Facility (HPC)

HPC Processing
2020: 300 Pflop/s
2028: 30 Eflop/s

Russian Supercomputing Days

Courtesy of J.Dongarra, UT, ORNL, USA

World Supercomputing News, Trends, Perspectives
(complexity of computer devices)

Floating point operations per cycle per core

- Most of the recent computers have FMA (Fused multiple add): (i.e. x ←x + y*z in one cycle)
- Intel Xeon earlier models and AMD Opteron have SSE2
  - 2 flops/cycle DP & 4 flops/cycle SP
- Intel Xeon Nehalem (’09) & Westmere (’10) have SSE4
  - 4 flops/cycle DP & 8 flops/cycle SP
- Intel Xeon Sandy Bridge(’11) & Ivy Bridge (’12) have AVX
  - 8 flops/cycle DP & 16 flops/cycle SP
- Intel Xeon Haswell (’13) & (Broadwell (’14)) AVX2
  - 16 flops/cycle DP & 32 flops/cycle SP
- Intel Xeon Skylake (server) AVX 512
  - 32 flops/cycle DP & 64 flops/cycle SP
  - Knight’s Landing

We are here 

World Supercomputing News, Trends, Perspectives
(CPU access latencies)

In 167 cycles can do 2672 DP Flops!

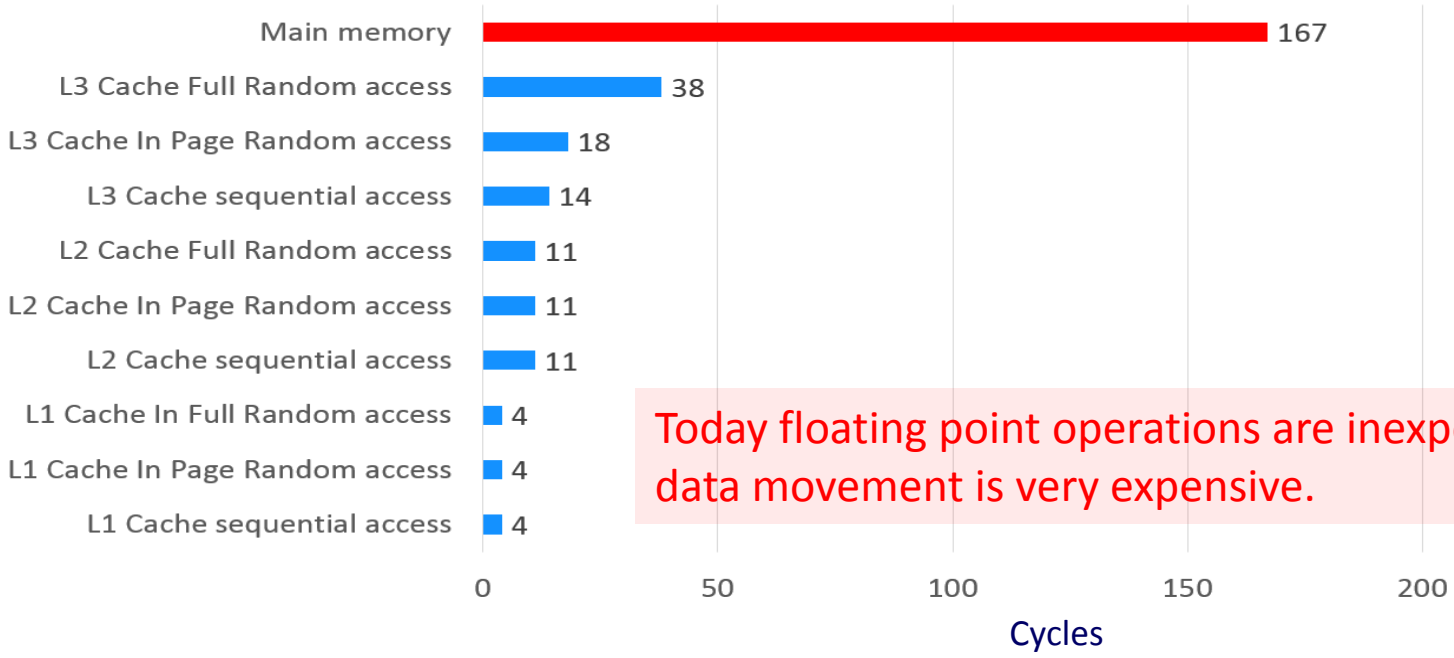

Today floating point operations are inexpensive, data movement is very expensive.

Yutaka Ishikawa about Post-K supercomputer:
"Easy of use is one of our Key Performance Indicators".

World Supercomputing News, Trends, Perspectives
(huge diversity of computer devices)

How to protect legacy software assets ?

“Array of things” Project

Файл Правка Вид Журнал Закладки Инструменты Справка

Open Encyclopedia of Parall

https://algowiki-project.org/en/Open_Encyclopedia_of_Parallel_Algorithmic_Feature 67% Поиск

AlgoWiki

Page Discussion Read View source View history Search

Open Encyclopedia of Parallel Algorithmic Features

(Redirected from Open Encyclopedia of Algorithms' Properties)

Main page
Forum
Recent changes

File storage
New files
Upload file

Tools
What links here
Related changes
Special pages
Printable version
Permanent link
Page information

In other languages
Русский

Welcome! Join us!

AlgoWiki is an open encyclopedia of **algorithms' properties** and **features of their implementations** on different hardware and software platforms from mobile to extreme scale, which allows for collaboration with the worldwide computing community on algorithm descriptions.

AlgoWiki provides an exhaustive description of an algorithm. In addition to classical algorithm properties such as serial complexity, AlgoWiki also presents additional information, which together provides a complete description of the algorithm: its parallel complexity, parallel structure, determinacy, data locality, performance and scalability estimates, communication profiles for specific implementations, and many others.

Read more: About project.

Project structure

Algorithm classification — the main section of AlgoWiki which contains descriptions of all algorithms. All algorithms are added to the appropriate category of the classification, and classification is expanded with new sections if necessary.

Today's featured picture

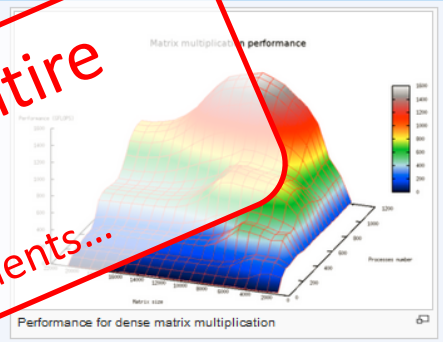

Performance for dense matrix multiplication

Work organization

Description of algorithm properties and structure
Guides to writing sections of the algorithm's description
Glossary
Help with editing

Readiness of articles

Finished articles:

- Single-qubit transform of a state vector
- Two-sided Thomas algorithm, pointwise version
- Poisson equation, solving with DFT
- Thomas algorithm, pointwise version
- Backward substitution
- Dense matrix multiplication
- Dot product
- Horners method
- Givens method
- Cholesky method
- Cholesky decomposition

Featured article

Cholesky decomposition

1 Properties and structure of the algorithm

1.1 General description

The Cholesky decomposition algorithm was first proposed by André-Louis Cholesky (October 15, 1875 - August 31, 1918) at the end of the First World War shortly before he was killed in battle. He was a French military officer and mathematician. The idea of this algorithm was published in 1924 by his fellow officer and, later, was used by Banachiewicz in 1938 [7]. In the Russian mathematical literature, the Cholesky decomposition is also known as the square-root method [1-3] due to the square root operations used in this decomposition and not used in Gaussian elimination.

Properties of the algorithm:

- Sequential complexity: $O(n^3)$
- Height of the parallel form: $O(n)$
- Width of the parallel form: $O(n^2)$
- Amount of input data: $\frac{n(n+1)}{2}$
- Amount of output data: $\frac{n(n+1)}{2}$

Originally, the Cholesky decomposition was used only for dense real symmetric positive definite matrices. At present, the application of this decomposition is much wider. For example, it can also be employed for the case of Hermitian matrices. In order to increase the computing performance, its block versions are often applied.

In the case of sparse matrices, the Cholesky decomposition is also widely used as the main stage of a direct method for solving linear systems. In order to reduce the memory requirements and the profile of the matrix, special reordering strategies are applied to minimize the number of arithmetic operations. A number of reordering strategies are used to identify the independent matrix blocks for parallel computing systems.

1.2 Mathematical description

Input data: a symmetric positive definite matrix A whose elements are denoted by a_{ij}).

Output data: the lower triangular matrix L whose elements are denoted by l_{ij}).

AlgoWiki is a project for the entire computing community !

For scientists, researchers, teachers, students...

Supercomputing Consortium of Russian Universities

Distributed Summer Supercomputing Academy

June, 23rd – July, 1st , 2017

Educational tracks:

- *MPI / OpenMP programming technologies*
- *NVIDIA GPU programming technologies*
- *Design and deployment of supercomputer solutions*
- *Quantum Informatics*
- *OpenFOAM/Salome/Paraview open software*
- *Computational hydrodynamics on supercomputers*

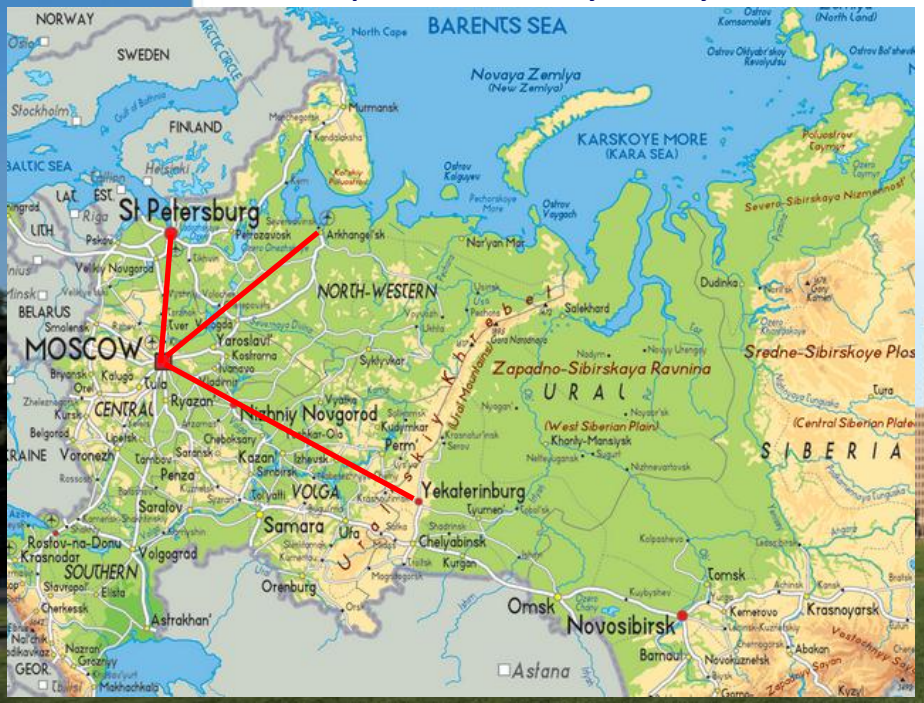

Supported by: IBM, T-Platforms, NVIDIA, RSC

Distributed Summer Supercomputing Academy
(http://Academy.HPC-Russia.ru)

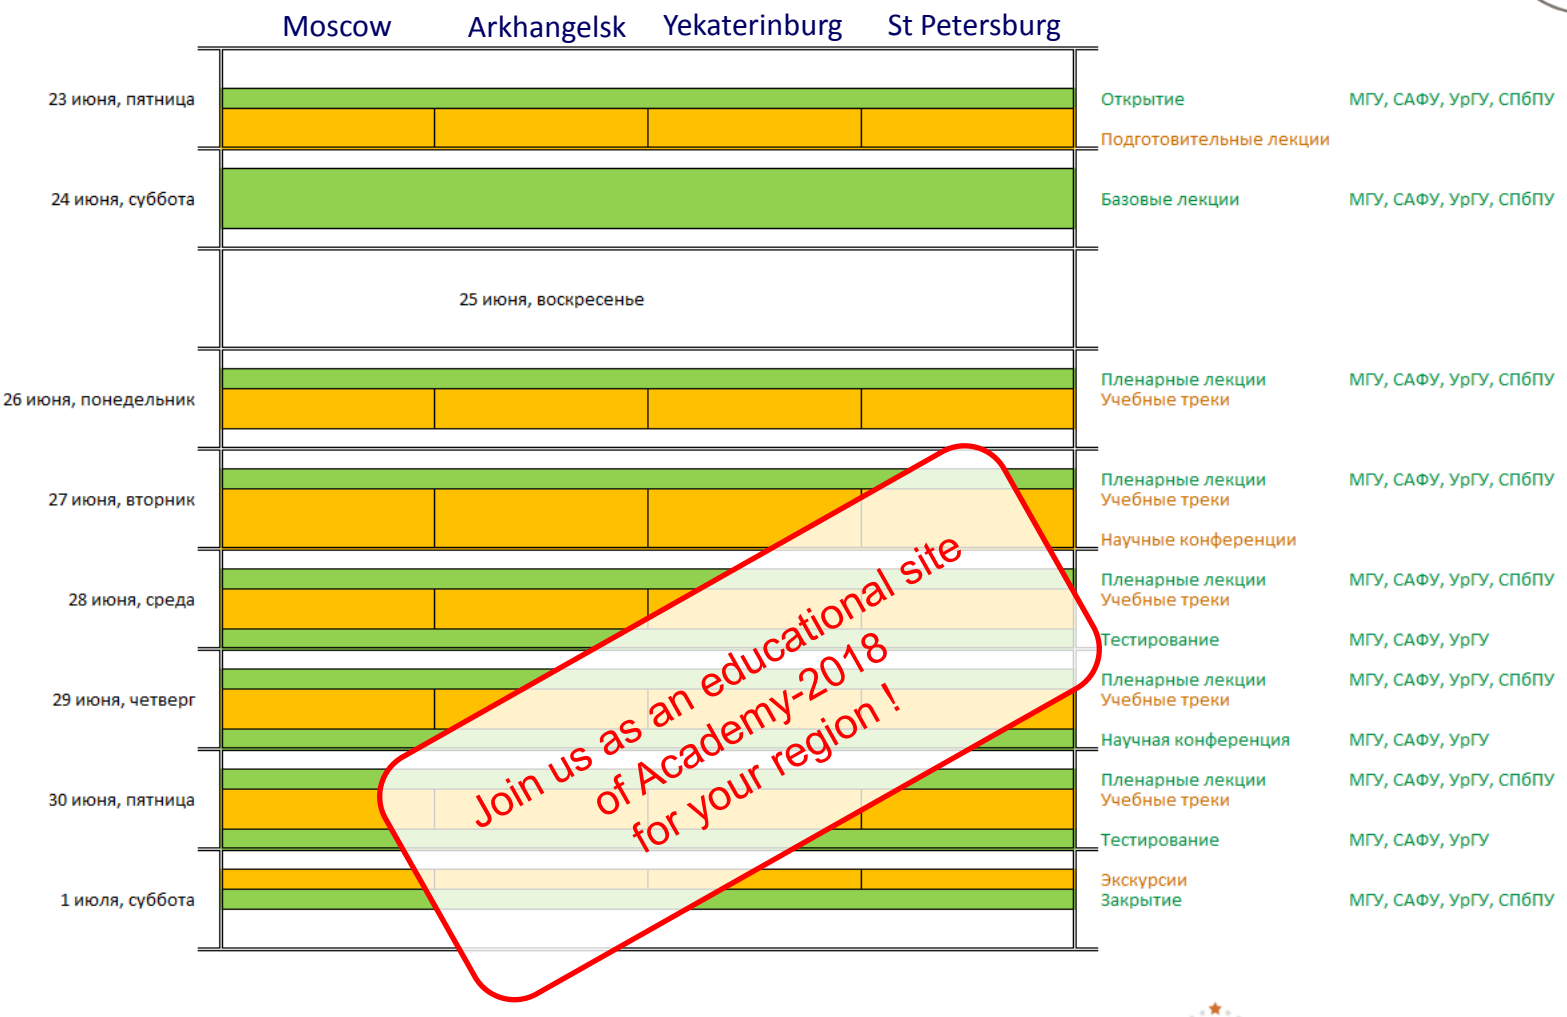

Russian Supercomputing Days

SUPERCOMPUTING CONSORTIUM OF RUSSIAN UNIVERSITIES

“Supercomputing Education” Series: New Books

Russian Supercomputing Days

Russian Supercomputing Days: General Schedule

September, 25th
Monday

September, 26th
Tuesday

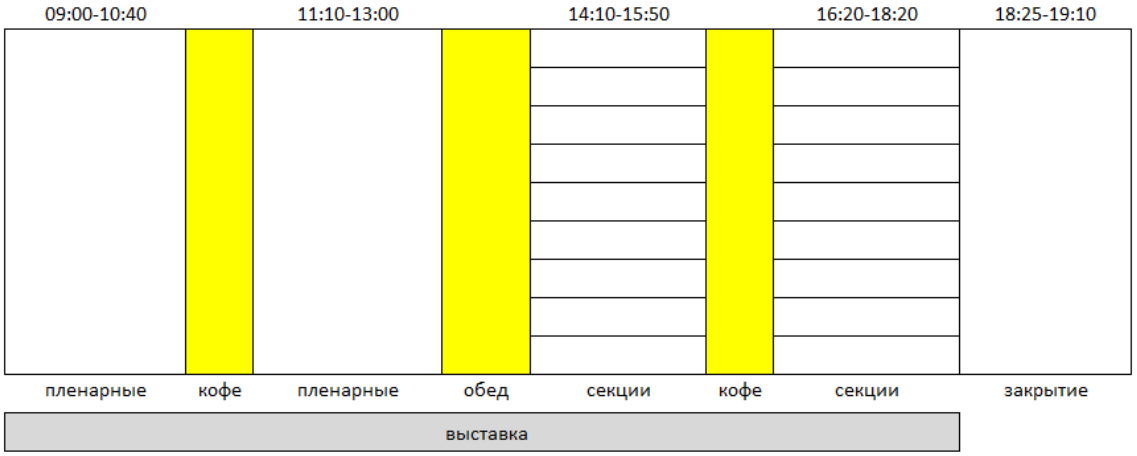

Russian Supercomputing Days

Russian Supercomputing Days: Awards

- Best Research Paper,
- Best Research Paper – Industrial Session,
- Best Research Paper – Young Scientists Session,
- Best Research Poster.

Winners will be announced at the Closing Session:

September, 26th, 18:25


International Conference
Russian Supercomputing Days
September 25-26, 2017, Moscow