



Курс: Цифровые технологии в научных исследованиях

ЛЕКЦИЯ 9

LESS MOOR MORE BRAIN

10 декабря 2020

Less Moore, More Brain

Меньше Мура, Больше Мозга

Prof. Dr. Thomas Ludwig
German Climate Computing Center (DKRZ)

BIG DATA ANALYTICS X ARTIFICIAL INTELLIGENCE X MODELING & SIMULATION

RUNNING ON ONE MACHINE IN MISSION-CRITICAL WORKFLOWS

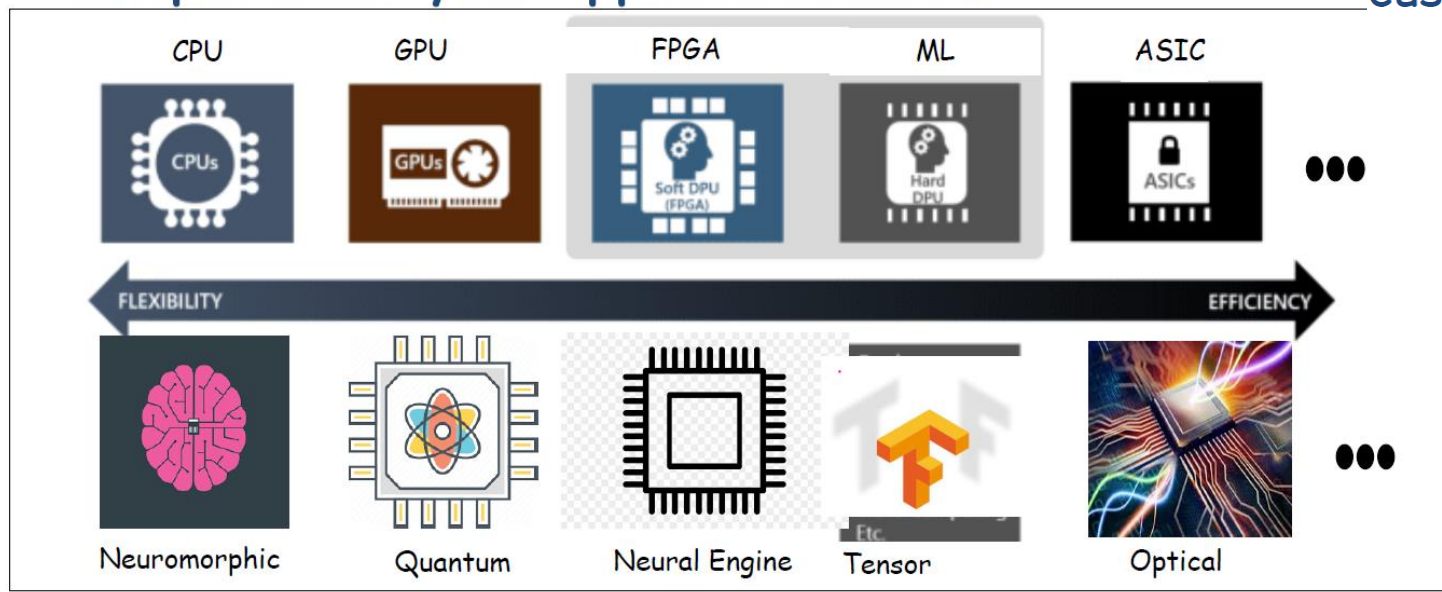
EXASCALE ERA

Mike Woodacre
HPE Fellow/VP, CTO for HPC & MCS

Future HPC Systems Will be Customized...

◆ You will be able to dial up what you need in your computer for your application mix ...

◆ HPC will have extreme heterogeneity and build custom systems for each important application.



Jack Dongarra
University of Tennessee
Oak Ridge National Laboratory
University of Manchester



ПОЛИТЕХ

ТЕКУЩИЙ TOP500

Rmax from LINPACK MPP



June 2020: The TOP 10 Systems (43% of the Total Performance of Top500)

$Ax=b$, dense problem

Rank	Site	Computer	Country	Cores	Rmax [Pflops]	% of Peak	Power [MW]	GFlops/Watt
1	RIKEN Center for Computational Science	Fugaku, ARM A64FX (48C, 2.2 GHz), Tofu D Interconnect	 Japan	7,299,072	415.	81	28.3	14.7
2	DOE / OS Oak Ridge Nat Lab	Summit, IBM Power 9 (22C, 3.0 GHz), Nvidia GV100 (80C), Mellanox EDR	 USA	2,397,824	149.	74	10.1	14.7
3	DOE / NNSA L Livermore Nat Lab	Sierra, IBM Power 9 (22C, 3.1 GHz), Nvidia GV100 (80C), Mellanox EDR	 USA	1,572,480	94.6	75	7.44	12.7
4	National Super Computer Center in Wuxi	Sunway TaihuLight, SW26010 (260C) + Custom	 China	10,649,000	93.0	74	15.4	6.05
5	National Super Computer Center in Guangzhou	Tianhe-2A NUDT, Xeon (12C) + MATRIX-2000 + Custom	 China	4,981,760	61.4	61	18.5	3.32
6	Eni S.p.A	HPC5, Dell EMC PowerEdge C4140, Xeon (24C, 2.1 GHz) + Nvidia V100 (80C), Mellanox HDR	 Italy	669,760	35.5	69	2.25	15.8
7	NVIDIA Corporation	Selene, Nvidia DGX AMD (64C, 2.25 GHz) + Nvidia A100 (108C), Mellanox HDR	 USA	277,760	27.6	80	1.34	20.6
8	Texas Advanced Computing Center / U of Texas	Frontera, Dell C6420, Xeon Platinum, 8280 28C 2.7 GHz, Mellanox HDR	 USA	448,448	23.5	61		
9	CINECA	Marconi-100, IBM Power System AC922, P9 (16C, 3 GHz) + Nvidia V100 (80C), Mellanox EDR	 Italy	347,776	21.6	74	1.98	10.9
10	Swiss CSCS	Piz Daint, Cray XC50, Xeon (12C) + Nvidia P100 (56C) + Custom	 Swiss	387,872	21.2	78	2.38	8.90

TOP500 List in June 2020

Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	Supercomputer Fugaku - Supercomputer Fugaku, A64FX 48C 2.2GHz, Tofu interconnect D, Fujitsu RIKEN Center for Computational Science Japan	7,299,072	415,530.0	513,854.7	28,335
		↑ 3.0x	↑ 2.8x	↑ 2.5x	↑ 2.8x
2	Summit - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM DOE/SC/Oak Ridge National Laboratory United States	2,414,592	148,600.0	200,794.9	10,096

Rmax from LINPACK MPP

$Ax=b$, dense problem

А что при другой нагрузке ?

hpcg-benchmark.org

HPCG Results; The Other Benchmark

- High Performance Conjugate Gradients (HPCG).
- Solves $Ax=b$, A large, sparse, b known, x computed.

ДЕГРАДАЦИЯ ПРОИЗВОДИТЕЛЬНОСТИ ПРИ ДРУГОЙ НАГРУЗКЕ

HPCG Benchmark June 2020



Rank	Site	Computer	Cores	HPL Rmax (Pflop/s)	TOP500 Rank	HPCG (Pflop/s)	Fraction of Peak
1	RIKEN Center for Computational Science Japan	Fugaku , Fujitsu A64FX, Tofu	7,299,072	415.53	1	13.4	2.5%
2	DOE/SC/ORNL USA	Summit , AC922, IBM POWER9 22C 3.7GHz, Dual-rail Mellanox FDR, NVIDIA Volta V100, IBM	2,414,592	143.50	2	2.926	1.5%
3	DOE/NNSA/LLNL USA	Sierra , S922LC, IBM POWER9 20C 3.1 GHz, Mellanox EDR, NVIDIA Volta V100, IBM	1,572,480	94.64	3	1.796	1.4%
4	Eni S.p.A. Italy	HPC5 , PowerEdge, C4140, Xeon Gold 6252 24C 2.1 GHz, Mellanox HDR, NVIDIA Volta V100	669,760	35.45	6	0.860	2.4%
5	DOE/NNSA/LANL/SNL USA	Trinity , Cray XC40, Intel Xeon E5-2698 v3 16C 2.3GHz, Aries, Cray	979,072	20.16	11	0.546	1.3%
6	NVIDIA USA	Selene , DGX SuperPOD, AMD EPYC 7742 64C 2.25 GHz, Mellanox HDR, NVIDIA Ampere A100	277,760	27.58	7	0.5093	1.8%
7	Natl. Inst. Adv. Industrial Sci. and Tech. (AIST) Japan	ABCI , PRIMERGY CX2570M4, Intel Xeon Gold 6148 20C 2.4GHz, Infiniband EDR, NVIDIA Tesla V100, Fujitsu	391,680	16.86	12	0.5089	1.7%
8	Swiss National Supercomputing Centre (CSCS) Switzerland	Piz Daint , Cray XC50, Intel Xeon E5-2690v3 12C 2.6GHz, Cray Aries, NVIDIA Tesla P100 16GB, Cray	387,872	19.88	10	0.497	1.8%
9	National Supercomputing Center in Wuxi China	Sunway TaihuLight , Sunway MPP, SW26010 260C 1.45GHz, Sunway, NRCPC	10,649,600	93.01	4	0.481	0.4%
10	Korea Institute of Science and Technology Information Republic of Korea	Nurion , CS500, Intel Xeon Phi 7250 68C 563584C 1.4GHz, Intel Omni-Path, Intel Xeon Phi 7250, Cray	570,020	13.93	18	0.391	1.5%

FROM EDGE TO CLOUD: АРХИТЕКТУРА РАСПРЕДЕЛЕННОЙ ГЕТЕРОГЕННОЙ СРЕДЫ СКЦ ПОЛИТЕХНИЧЕСКИЙ (ПРОДВИГАЕМ НАЧИНАЯ С 2018)

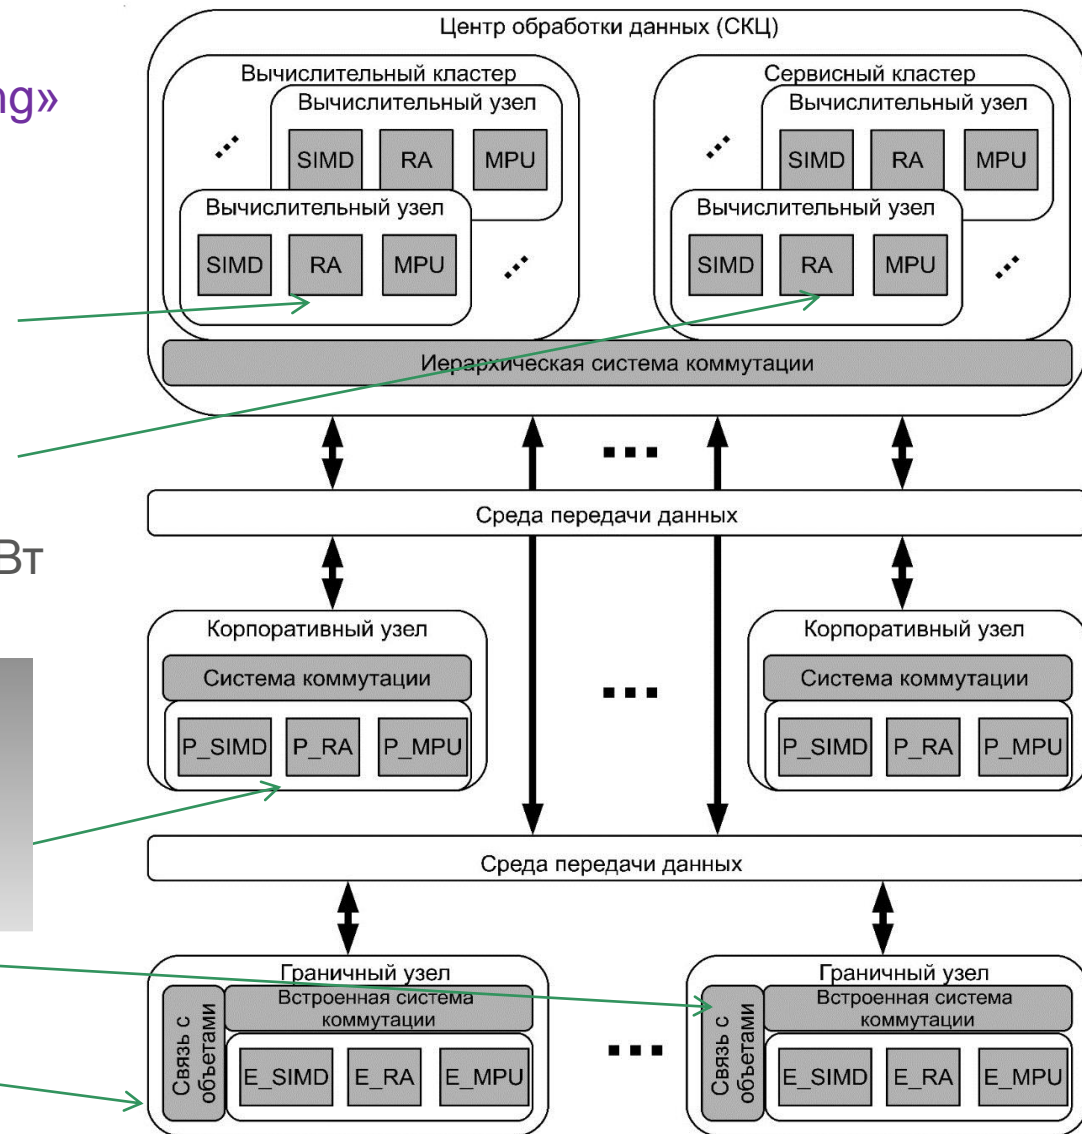
Уровень «объяснения» и «моделирования/simulation/modeling»
>4 Гфлопс/Вт



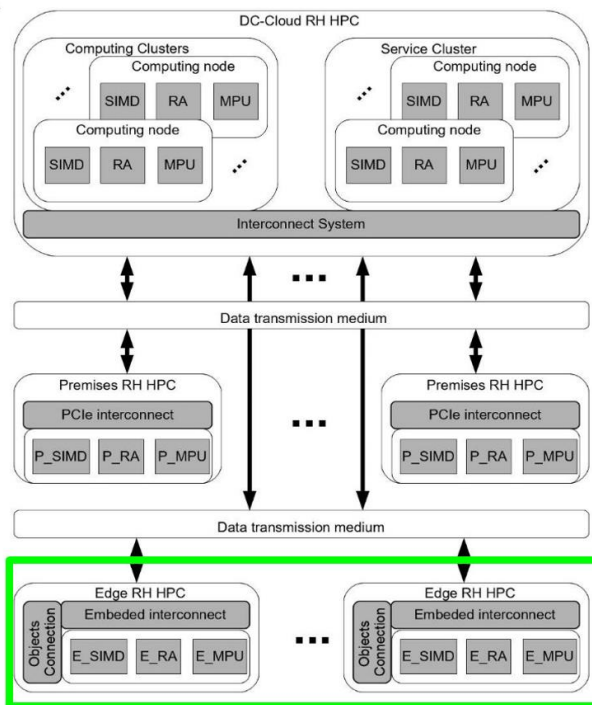
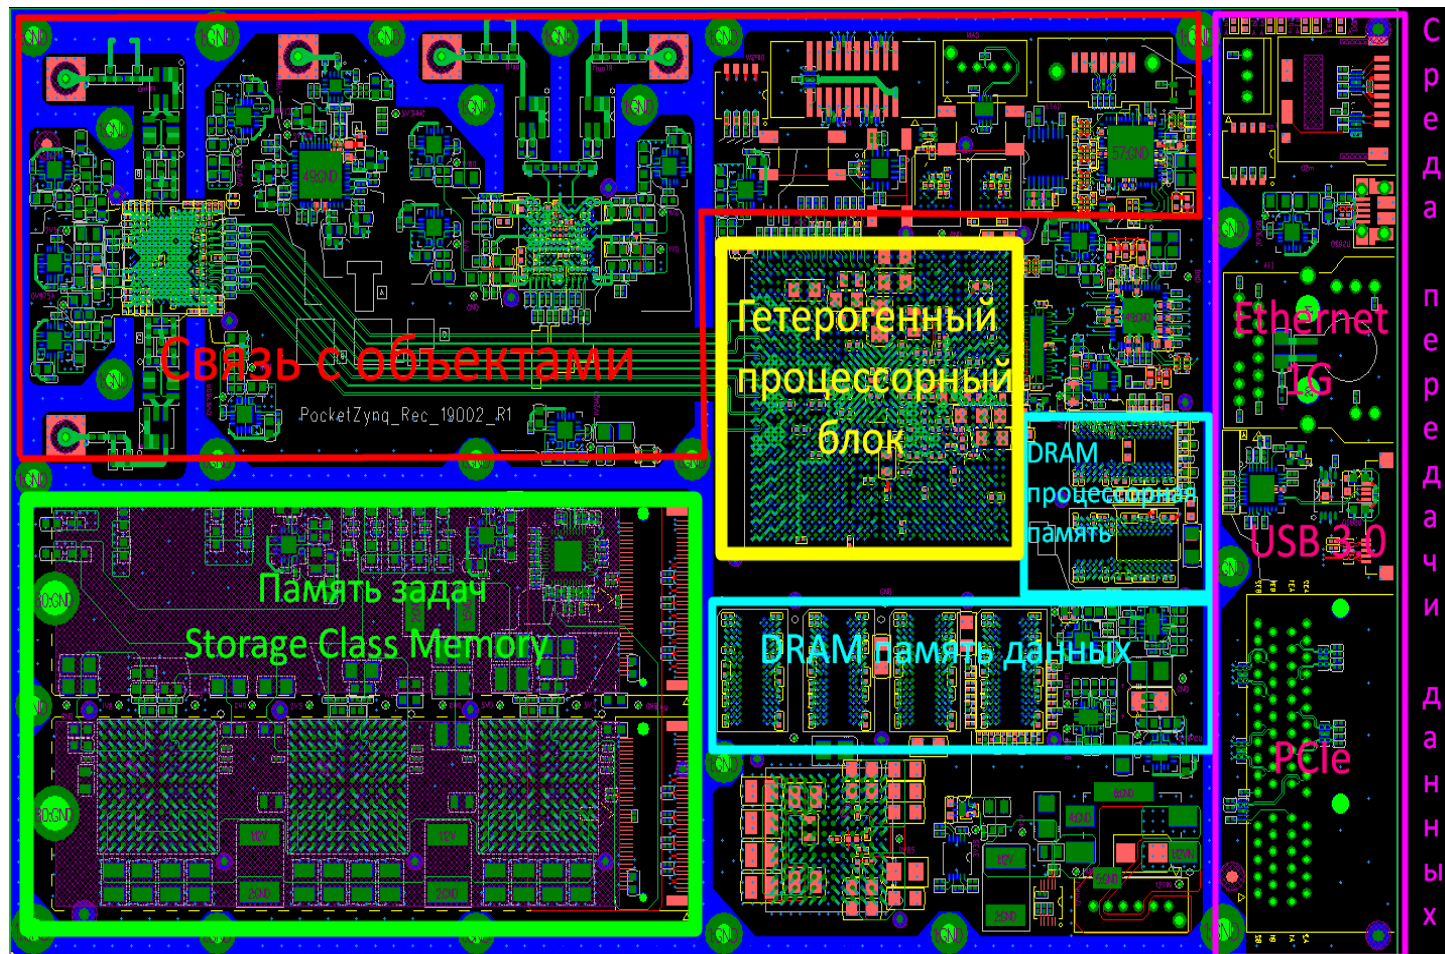
Уровень «агрегации» и «машинного обучения» >10 Гфлопс/Вт



Уровень доступа и предобработки «больших данных» >20
Гфлопс/Вт



Интерфейсы взаимодействия с экзо интеллектуальными системами

This block contains several images of circuit boards with highlighted components and labels:

- Связь с объектами** (Connection with objects): A red-bordered image of a circuit board with a central processor area.
- Гетерогенный процессорный блок** (Heterogeneous processor block): A yellow-bordered image of a processor chip.
- DRAM процессорная память** (DRAM processor memory): A blue-bordered image of a DRAM chip.
- DRAM память данных** (DRAM data memory): A cyan-bordered image of a DRAM chip.
- Пам'ять задач Storage Class Memory** (Task memory Storage Class Memory): A green-bordered image of a memory array.
- Ethernet 1G**: A purple-bordered image of an Ethernet network interface.
- USB 3.0**: A pink-bordered image of a USB 3.0 port.
- PCIe**: A pink-bordered image of a PCIe interface.

On the right side of the collage, there is a vertical text label: **Среда передачи данных** (Data transmission environment).

Interface Functions

