CHASING THE LIMITS OF COMPUTING EFFICIENCY IN TERRESTRIAL MODELING: MEMORY BANDWIDTH

Michal Belda, Ludek Kucera

Charles University, Faculty of Mathematics and Physics, Prague, Czech Republic

The corresponding author: Ludek Kucera, ludek@kam.mff.cuni.cz

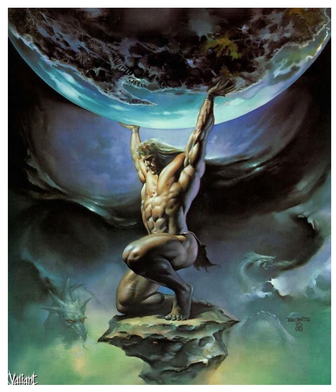

Type	Institution	Peak [PFlop/s]	Top500
Cray XC50	Japan Meteorological Agency	9.1	28
Cray XC50	Japan Meteorological Agency	9.1	29
Cray XC40	UK Meteorological Office	8.1	23
SGI ICE XA	NCAR	5.3	36
Cray XC40	ECMWF	4.2	42
Cray XC40	ECMWF	4.2	43
Cray XC40	Indian Inst. Tropical Meteorology	4.0	45
Bull DLC720	Deutsches Klimarechenzentrum	4.0	55
Cray XC40	UK Meteorological Office	3.0	66
Cray XC40	UK Meteorological Office	3.0	67

Meteorology uses very powerful computers

Terrestrial Modelling: How many PFlop per second can you get from your computer

HPCG (High Performance Conjugate Gradients)

HPCG models heat diffusion using Conjugate Gradients Method

A standard HPC benchmark

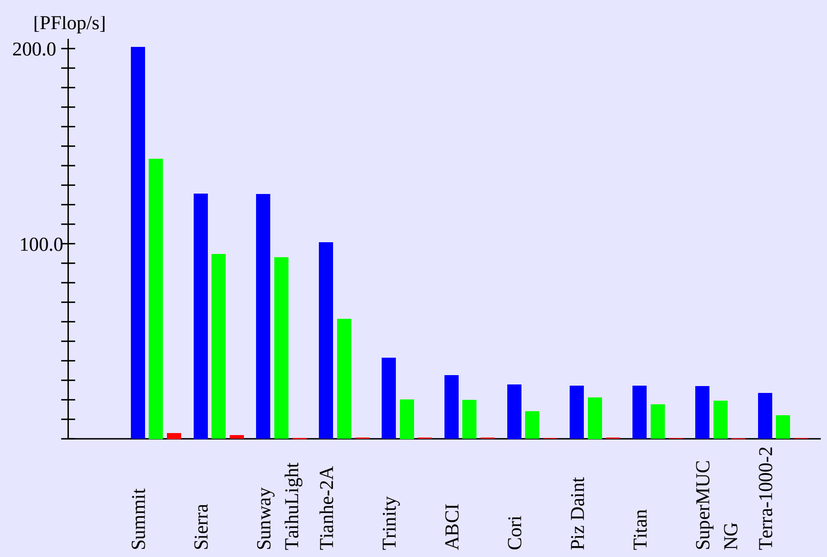

Computer	Peak [PFlop/s]	LINPACK [PFlop/s]	HPCG [PFlop/s]	HPCG/Peak [%]
Summit	200.8	143.5	2.926	1.5
Sierra	125.7	94.6	1.796	1.4
Sunway TL	125.4	93.0	0.481	0.4
Tianhe-2A	100.7	61.4	0.580	0.6
Trinity	41.5	20.2	0.546	1.3
ABCI	32.6	19.9	0.509	1.6
Cori	27.9	14.0	0.355	1.3
Piz Daint	27.2	21.2	0.497	1.8
Titan	27.1	17.6	0.322	1.2
SuperMUC	26.9	19.5	0.208	0.8

Peak, LINPACK, and HPCG computing power of the Top10 supercomputers

Flop/Byte ratio of the HPCG is about 0.25

The main part of the HPCG is SpMV product

2 arithmetic operations

(addition and multiplication)

per each 8 byte matrix element

Processor	Memory Bandwidth [GB/s]	Enough Data for DP HPCG [PFlop/s]	Peak Performance [PFlop/s]	Efficiency Bound [%]
NVIDIA Volta-100	900	0.25*900=225	7800	2.88
Volta-100/NVLink	300	0.25*300= 75	7800	0.96
Intel Xeon Phi "KNL"	480+120	0.25*600=150	3000	5.00
KNL/DRAM	120	0.25*120= 30	3000	1.00

Theoretical Efficiency of the Top 2 Processors

WRF (Weather Research and Forecast)

3 principal subroutines have been studied

representing together ~40 % of the running time

Left Plot: how SP Flop/Byte ratio depends on the cache size

plus the corresponding SP performance of Summit (400 SP PFlop/s)

Microphysics Driver - WSM3

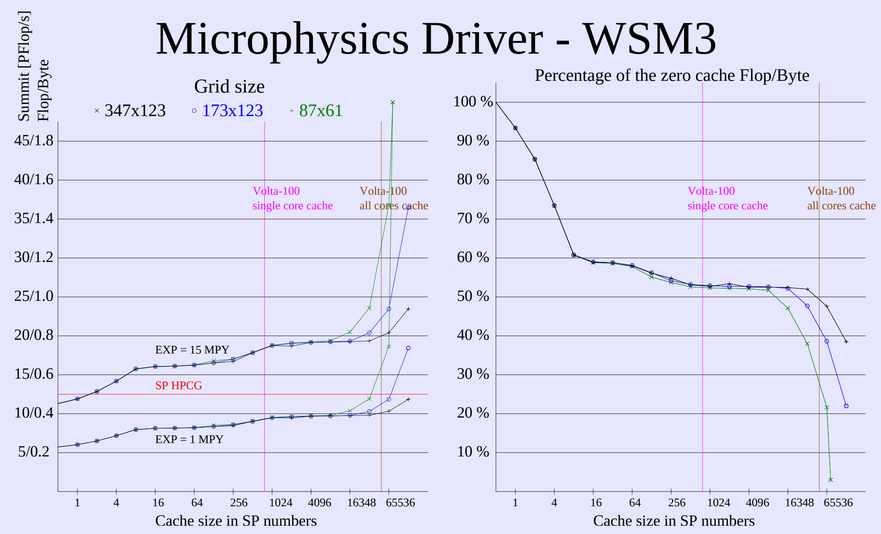

Radiation Driver - RRTM

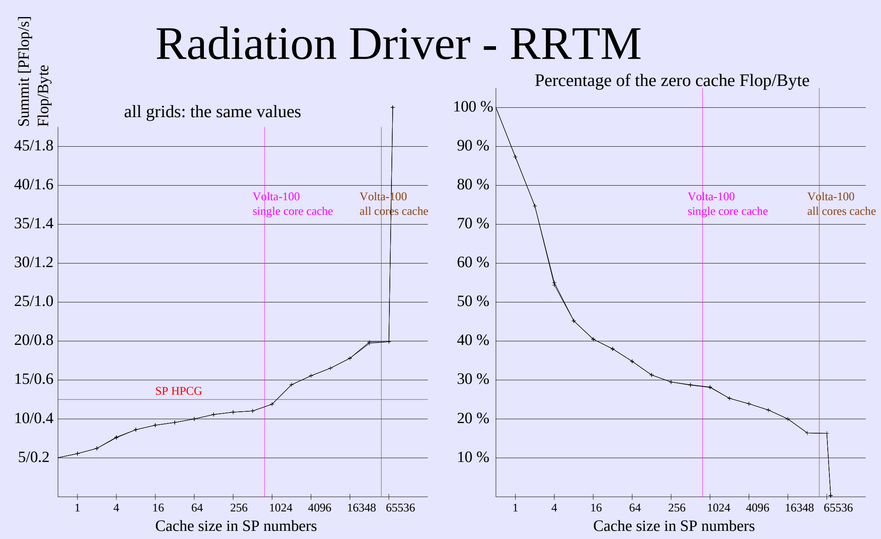

AdvectScalar

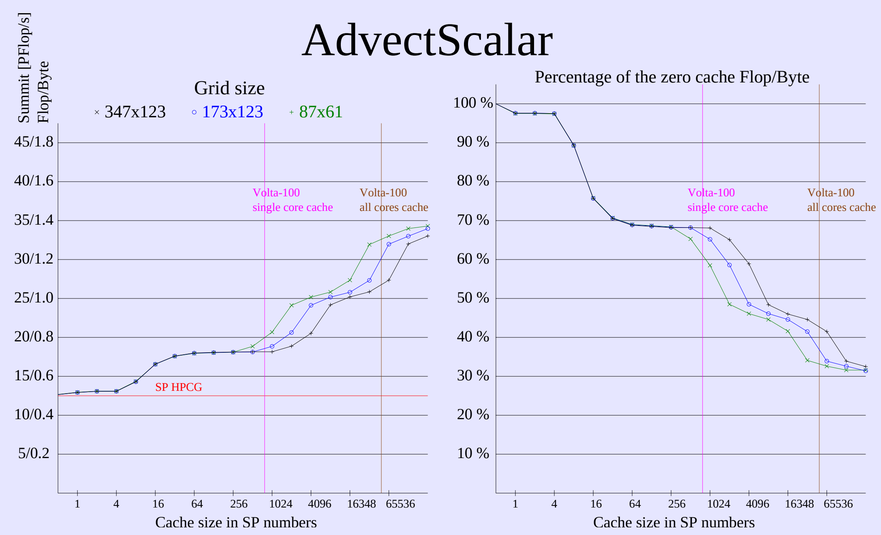

Please ask the representatives of the European Processor Initiative to take memory bandwidth into account