Computational Science...

...after von Neumann...

Exaflop/s?

Petaflop/s

Teraflop/s

Gigaflop/s

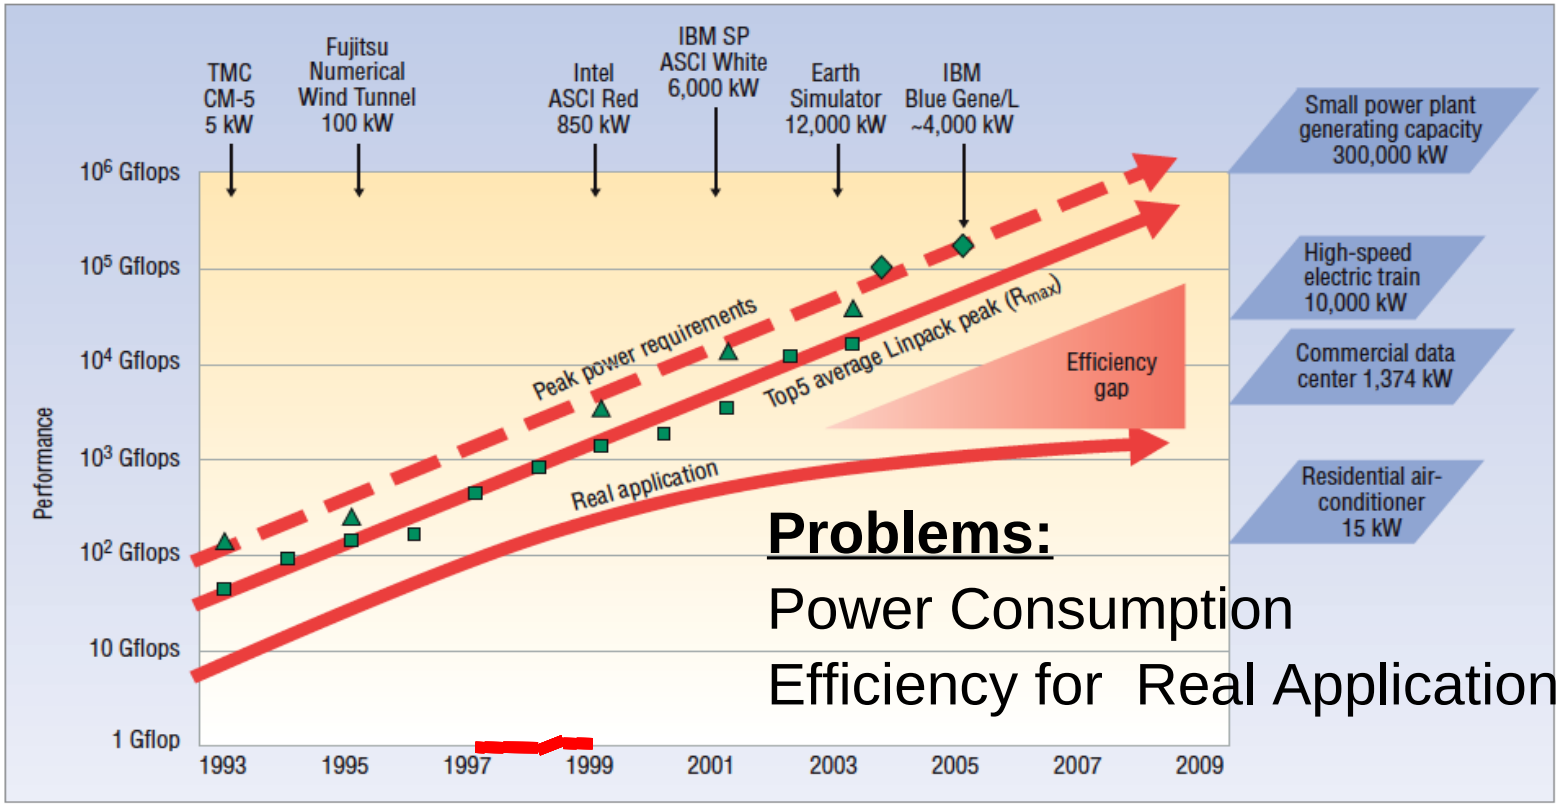

Figure 1. Rising power requirements. Peak power consumption of the top supercomputers has steadily increased over the past 15 years.

Thanks to Horst Simon, LBNL/NERSC for this diagram.

SPECIAL HARDWARE

CPUs

Central Processing Units

General Purpose oriented

1-12 Cores

Up to 4 pipes per core using Vector Units

Fully Programmable, many languages available

Very well studied

Max. 125W per processor

GPUs

Graphic Processing Units

Graphics oriented

16-512 Cores

Massively Parallel Architecture, specialized instructions for parallel processing

Fully programmable, but limited languages

Algorithms not fully explored

Max. 400W per card

FPGAs

Field Programmable Gate Arrays

Custom designs, best for processing streaming data

Programmable Logic, Architecture is custom-built for the required application

Requires extensive knowledge to program, development time is longer than CPUs and GPUs

Application interface is custom built on each case

Max. 60W per FPGA

ASICs

Application Specific Integrated Circuits

Fully custom designs, built for a specific application

Not flexible, cannot be changed once it is built

Development is even more specialized than FPGAs

Power consumption varies with the application, usually best performance per Watt

Slide: Guillermo Marcus

HITS

HARDWARE

GRAPE-6 Gravity/Coulomb Part

- G6 Chip: 0.25μ 2MGate ASIC, 6 Pipelines
- at 90MHz, 31Gflops/chip
- 48Tflops full system (March 2002)
- Plan up to 72Tflops full system (in 2002)
- Installed in Cambridge, Marseille, Drexel, Amsterdam, New York (AMNH), Mitaka (NAO), Tokyo, etc..

New Jersey, Indiana, Heidelberg

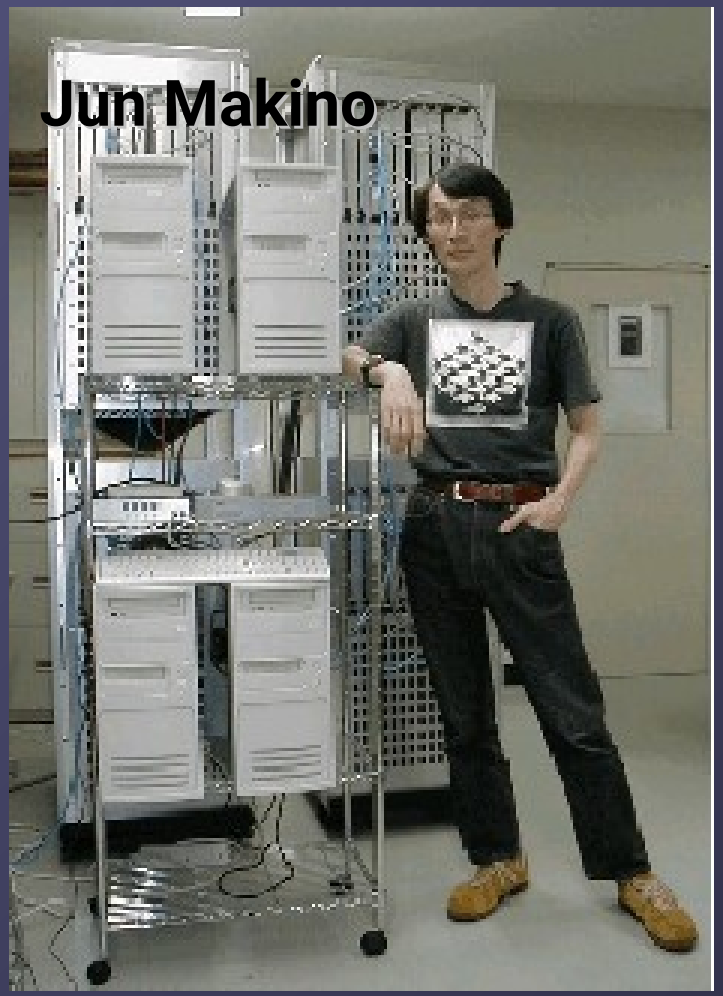

VolkswagenStiftung

kepler1 Cluster 2006

4 Tflops (32 micro-GRAPE6)
Dual Port Infiniband
4 MPRACE-1 reconfigurable
(soon: 32 MPRACE-2)

GRAPE + MPRACE
= GRACE

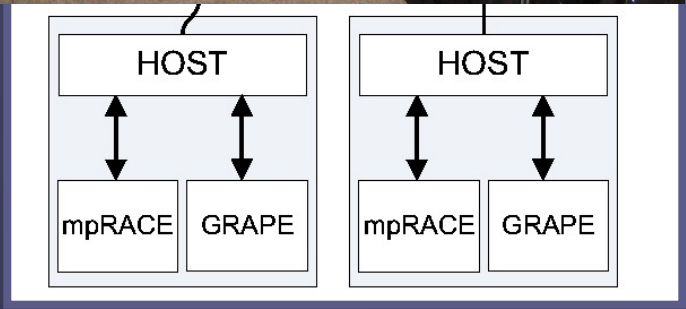

GPU: NAOC laohu cluster Beijing, China

Heidelberg

Kepler GPU cluster

VolkswagenStiftung

Kepler GPU cluster

12 nodes = 12 x 16 = 192 CPU cores (@ 2 GHz)

12 x 64 GB = 768 GB RAM CPU memory

12 GPUs K20m = 12 x 2496 ~ 30k GPU threads

12 x 4.8 GB ~ 57 GB GPU device memory

4 x Xilinx Virtex-6 FPGA (ML 605)

since beg. 2013 operated.

Since 2019 three more nodes:
1x Quadro P6000
1x RTX 2080 Ti
1x A5000

Our own φGRAPE/GPU N-body code

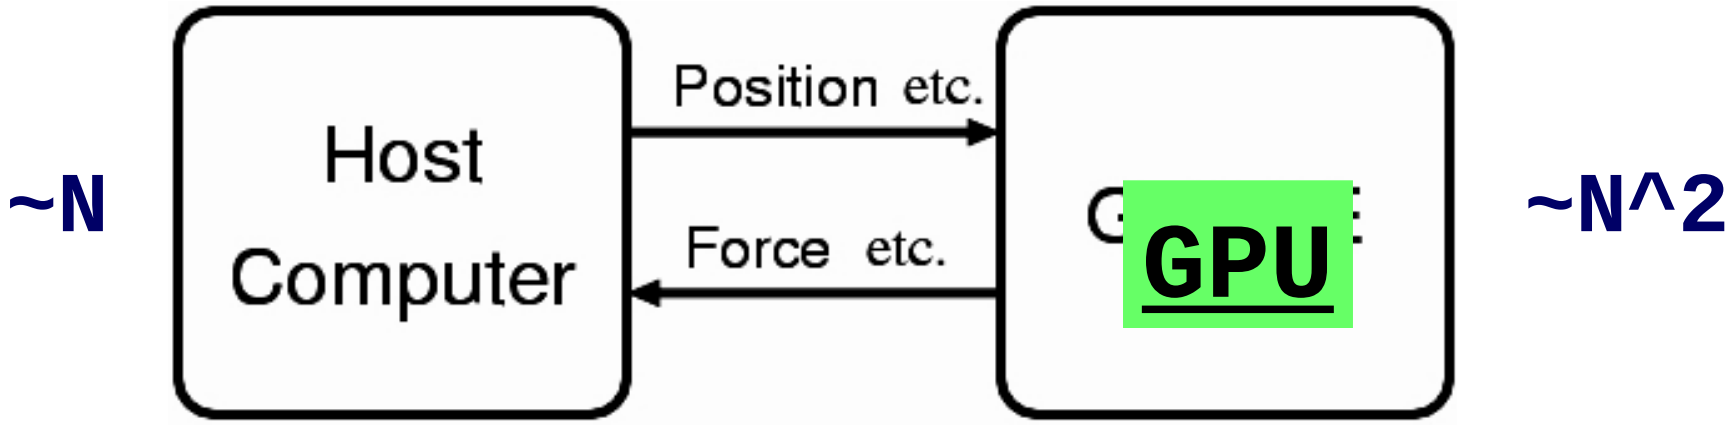

$$\vec{a}_i = \sum_{j=1;j\neq i}^{N} \vec{f}_{ij} \qquad \vec{f}_{ij} = -\frac{G \cdot m_j}{\left(r_{ij}^2 + \varepsilon^2\right)^{3/2}} \vec{r}_{ij}$$

No softening for NBODY6++GPU

Nbody6++ Structure

Slide of: Long Wang

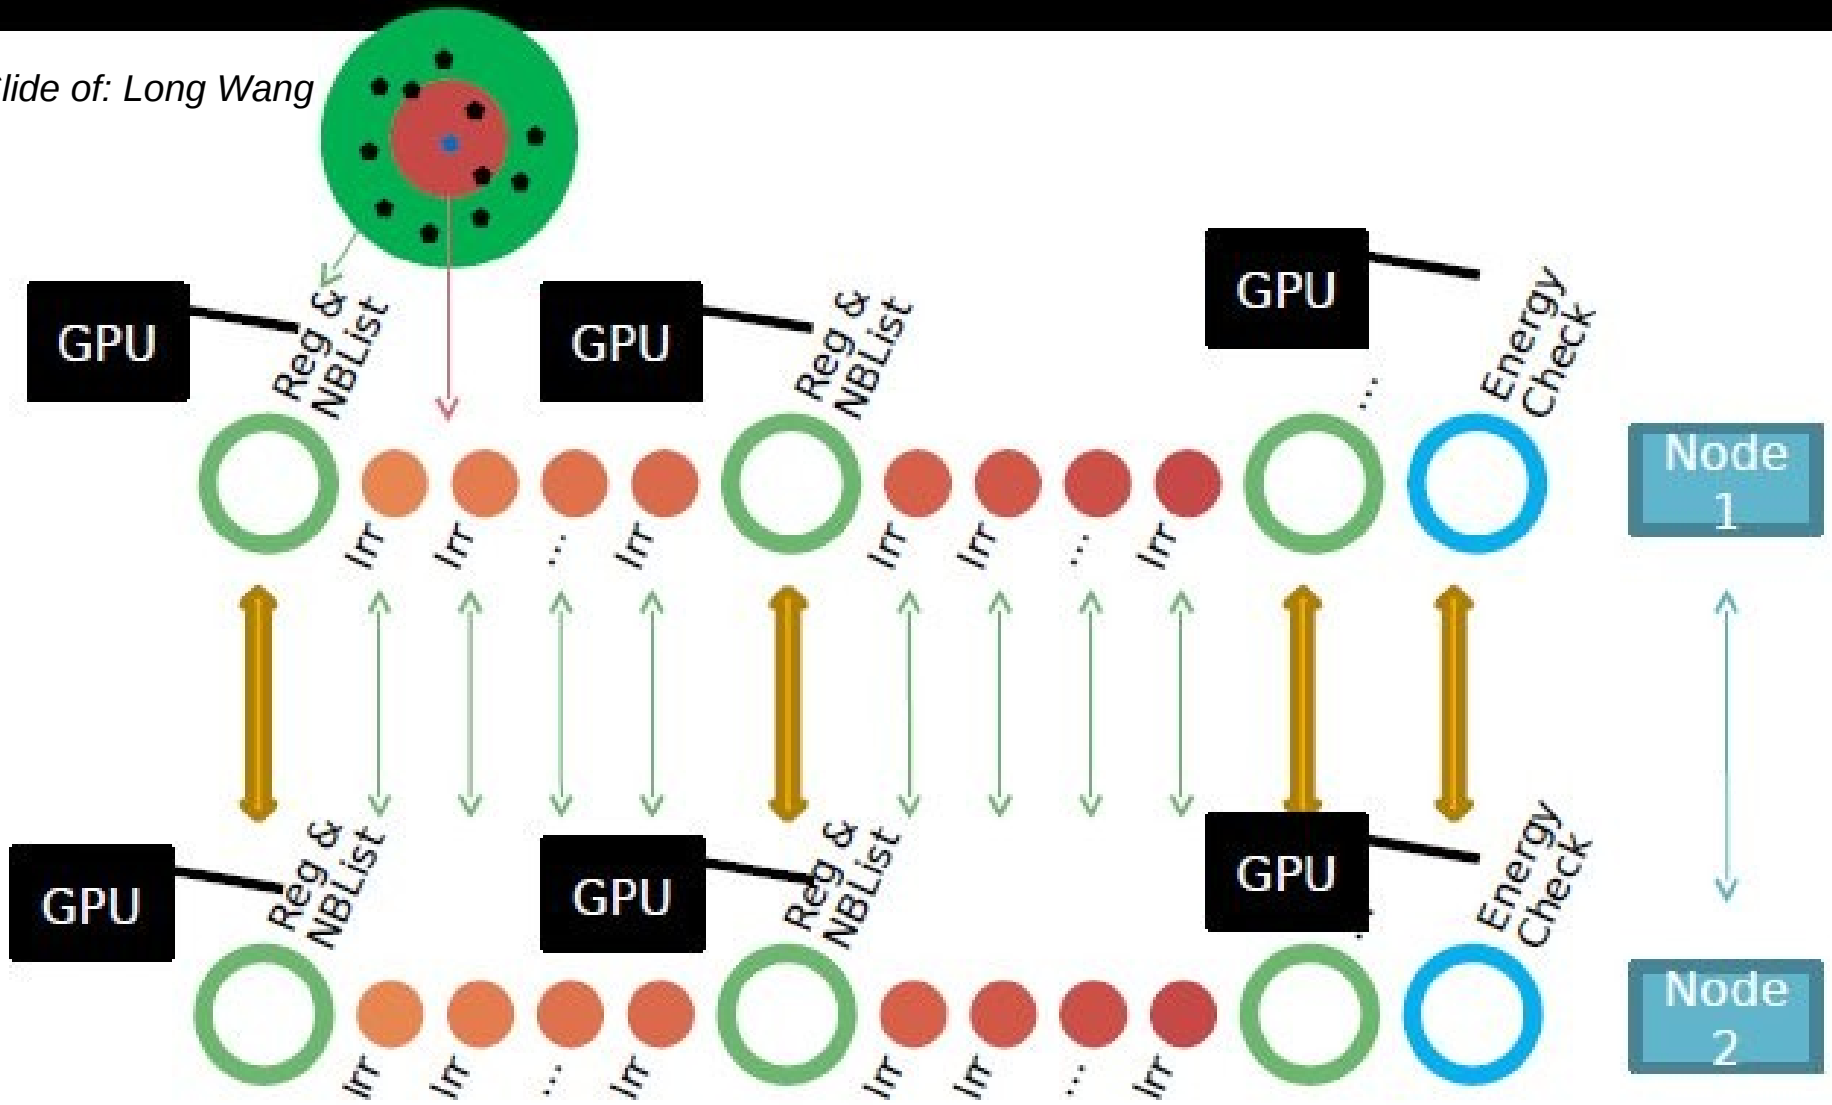

Parallel code on cluster

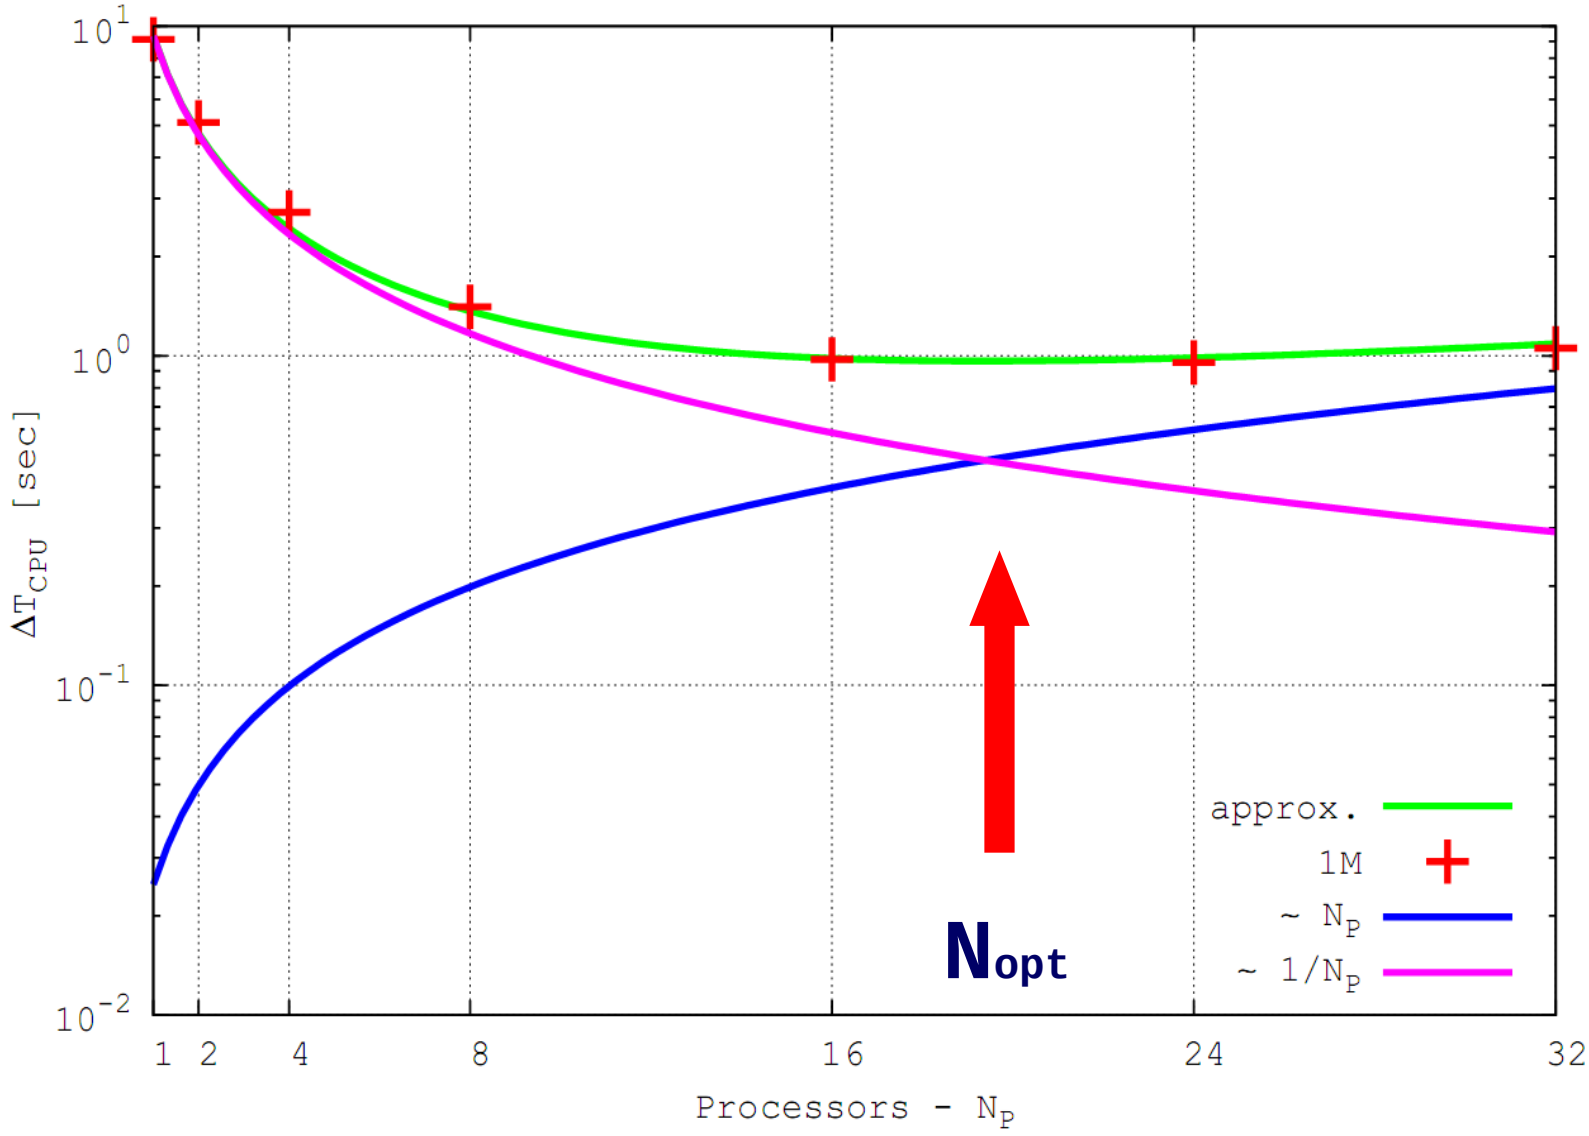

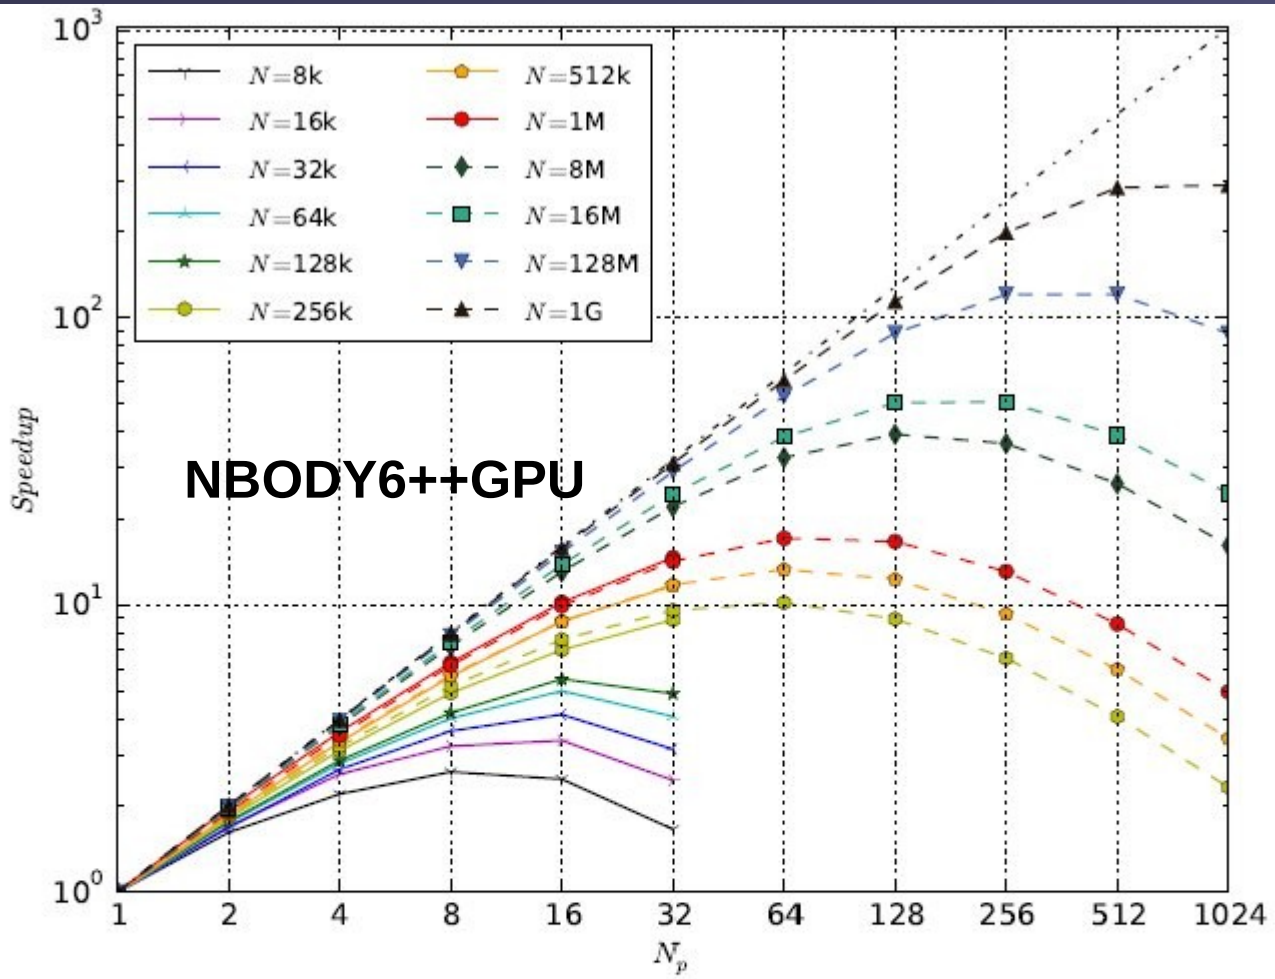

Huang, Berczik, Spurzem, Res. Astron. Astroph. 2016, 16, 11.

Fig. 2 The speed-up (S) of NBODY6++ as a function of particle number (N) and processor number (N_p). Solid points are the measured speed-up ratio between sequential and parallel wall-clock time, dash lines predict the performance of larger scale simulations further. The symbols used in figure have the magnitudes: $1k = 1,024$, $1M = 1k^2$ and $1G = 1k^3$.

RIKEN, Kobe, JAPAN

FUGAKU

Nature's Secrets

富岳

Mt. Fuji

The world's fastest Super Computer 2020 /2021

7.6 million cores, 442 Pflop/s

Supercomputer from China: 96/33 Pflop/s Linpack Wuxi/Guangzhou/Tianjin National Supercomputing Center Taihu 10 mill. cores

Tianhe-2 (MilkyWay-2) - TH-IVB-FEP Cluster, Intel Xeon E5-2692 12C 2.200GHz, TH Express-2, Intel Xeon Phi 31S1P

(Xeon Phi replaced by Chinese Matrix-2000)
Total: ~5 million cores

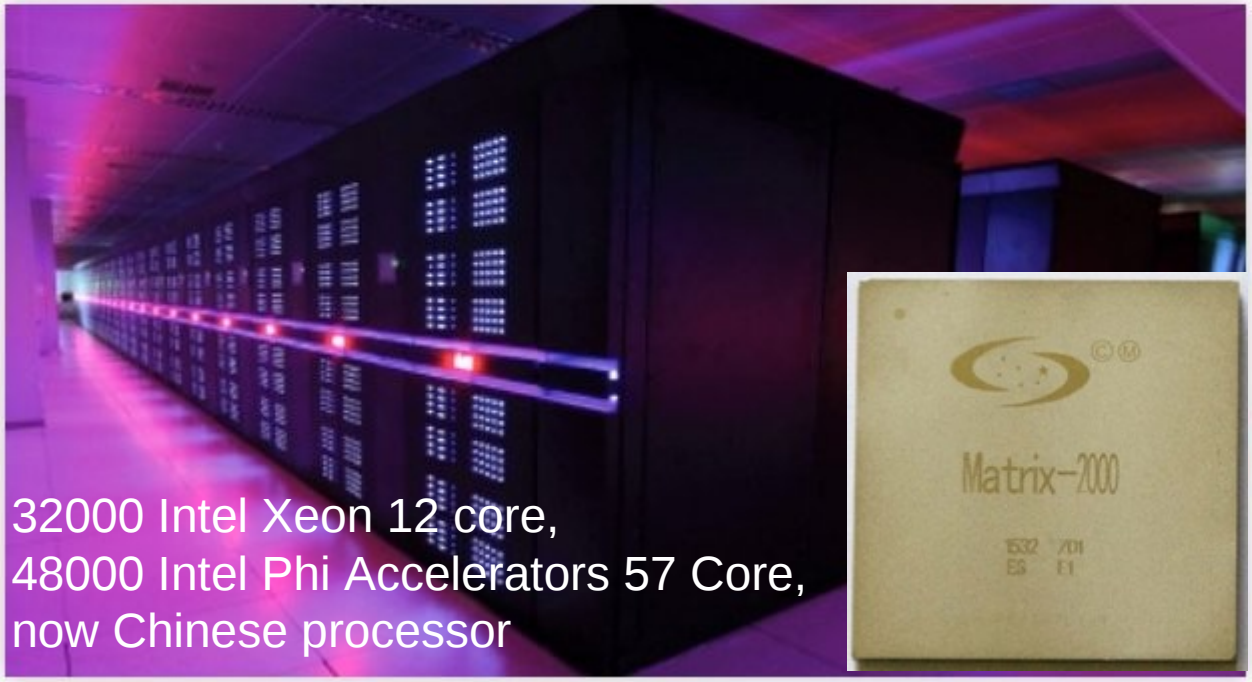

32000 Intel Xeon 12 core,
48000 Intel Phi Accelerators 57 Core,
now Chinese processor

The Matrix-2000 features 128 RISC cores operating at 1.2 GHz achieving 2.46 / 4.92 TFLOPS (DP/SP) with a peak power dissipation of 240 W

LUMI

Supercomputer, Kajaani, Finland

Using only Hydroelectric Power and its Heat used for heating buildings.

No. 3 in top500
No. 7 in green500

2.2 million cores
10.000 AMD GPUs

EuroHPC and LUMI consortium:
Finland, Belgium, Czech Republic, Denmark, Estonia, Iceland, Norway, Poland, Sweden, and Switzerland.

GCS
Gauss Centre for Supercomputing

JUWELS Booster 936 nodes (AMD CPU, 4x Ampere GPU)
~450.000 AMD cores, 25 million NVIDIA Ampere GPU cores
~ 70 Pflop/s SP ~ 44 Pflop/s DP
No. 12 in top500 list, No. 25 in green500 list

Jülich Wizard for European Leadership Science