

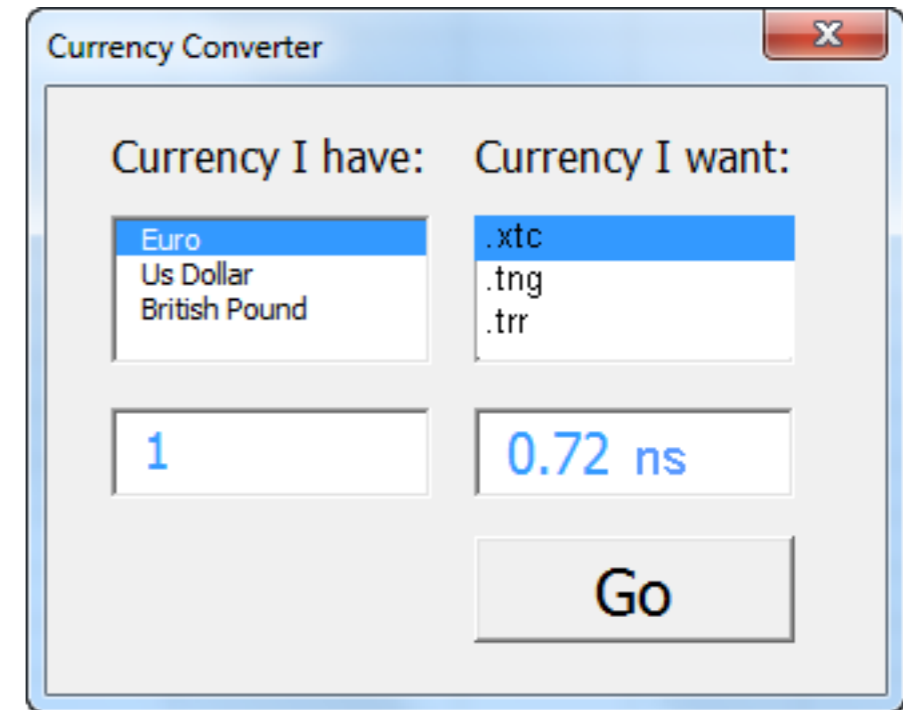
Carsten Kutzner
Theoretical & Computational Biophysics
MPI for biophysical Chemistry

SAVING MONEY WITH GRMACS

Cost-efficient MD simulations 2018

MOTIVATION

- from a fixed hardware budget ...
- ... produce as much MD trajectory as possible!
- use tax payer's money responsibly
- therefore, measure MD performance and get hardware prices in an ongoing effort



OUTLINE

- **2014:**

SOFTWARE NEWS AND UPDATES

WWW.C-CHEM.ORG

Journal of
**COMPUTATIONAL
CHEMISTRY**

Best Bang for Your Buck: GPU Nodes for GROMACS Biomolecular Simulations

Carsten Kutzner,^{*[a]} Szilárd Páll,^[b] Martin Fechner,^[a] Ansgar Esztermann,^[a]
Bert L. de Groot,^[a] and Helmut Grubmüller^[a]

- **RECAP:** what were our conclusions in 2014/15?
- **UPDATE:** hardware & software developments + their impact

WHAT IS THE 'OPTIMAL' HARDWARE TO BUY?

WHAT DO YOU WANT?

general-purpose cluster for all kinds of applications

- ▶ ~~large RAM~~ GROMACS uses 250 MB - 1 GB of CPU RAM per process
- ▶ nodes connected by a ~~high-performance network~~
- ▶ double-prec. ~~GPU performance~~
- ▶ ~~large GPU memory~~ even a 2M atom system requires only 225 MB GPU RAM

WHAT CAN WE SPARE?

maximize cost-efficiency by specialization

- ▶ GROMACS only



max. sampling, many separate simulations



WE OPTIMIZE FOR HIGH THROUGHPUT



~~single long trajectories~~



run these @ national HPC centers

FIND OPTIMAL HARDWARE FOR GROMACS!

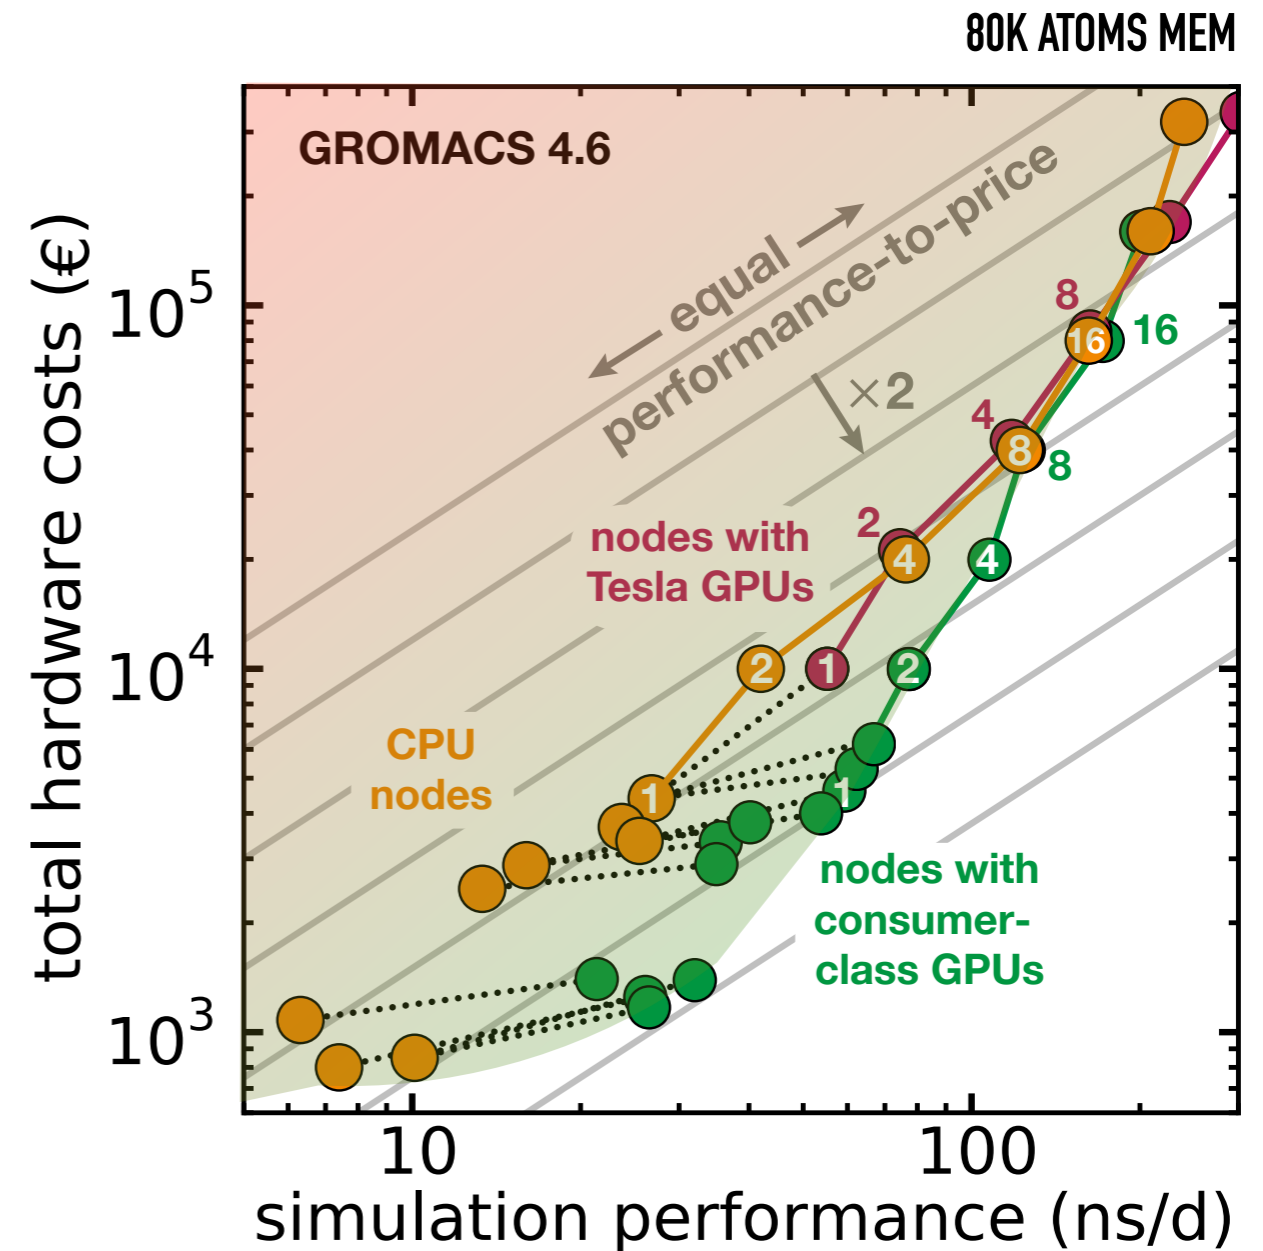
Our criteria:



1. **high performance-to-price (P/P) ratio**
2. **low energy consumption**
3. **reasonably high single-node performance**
4. **low rack space requirements**

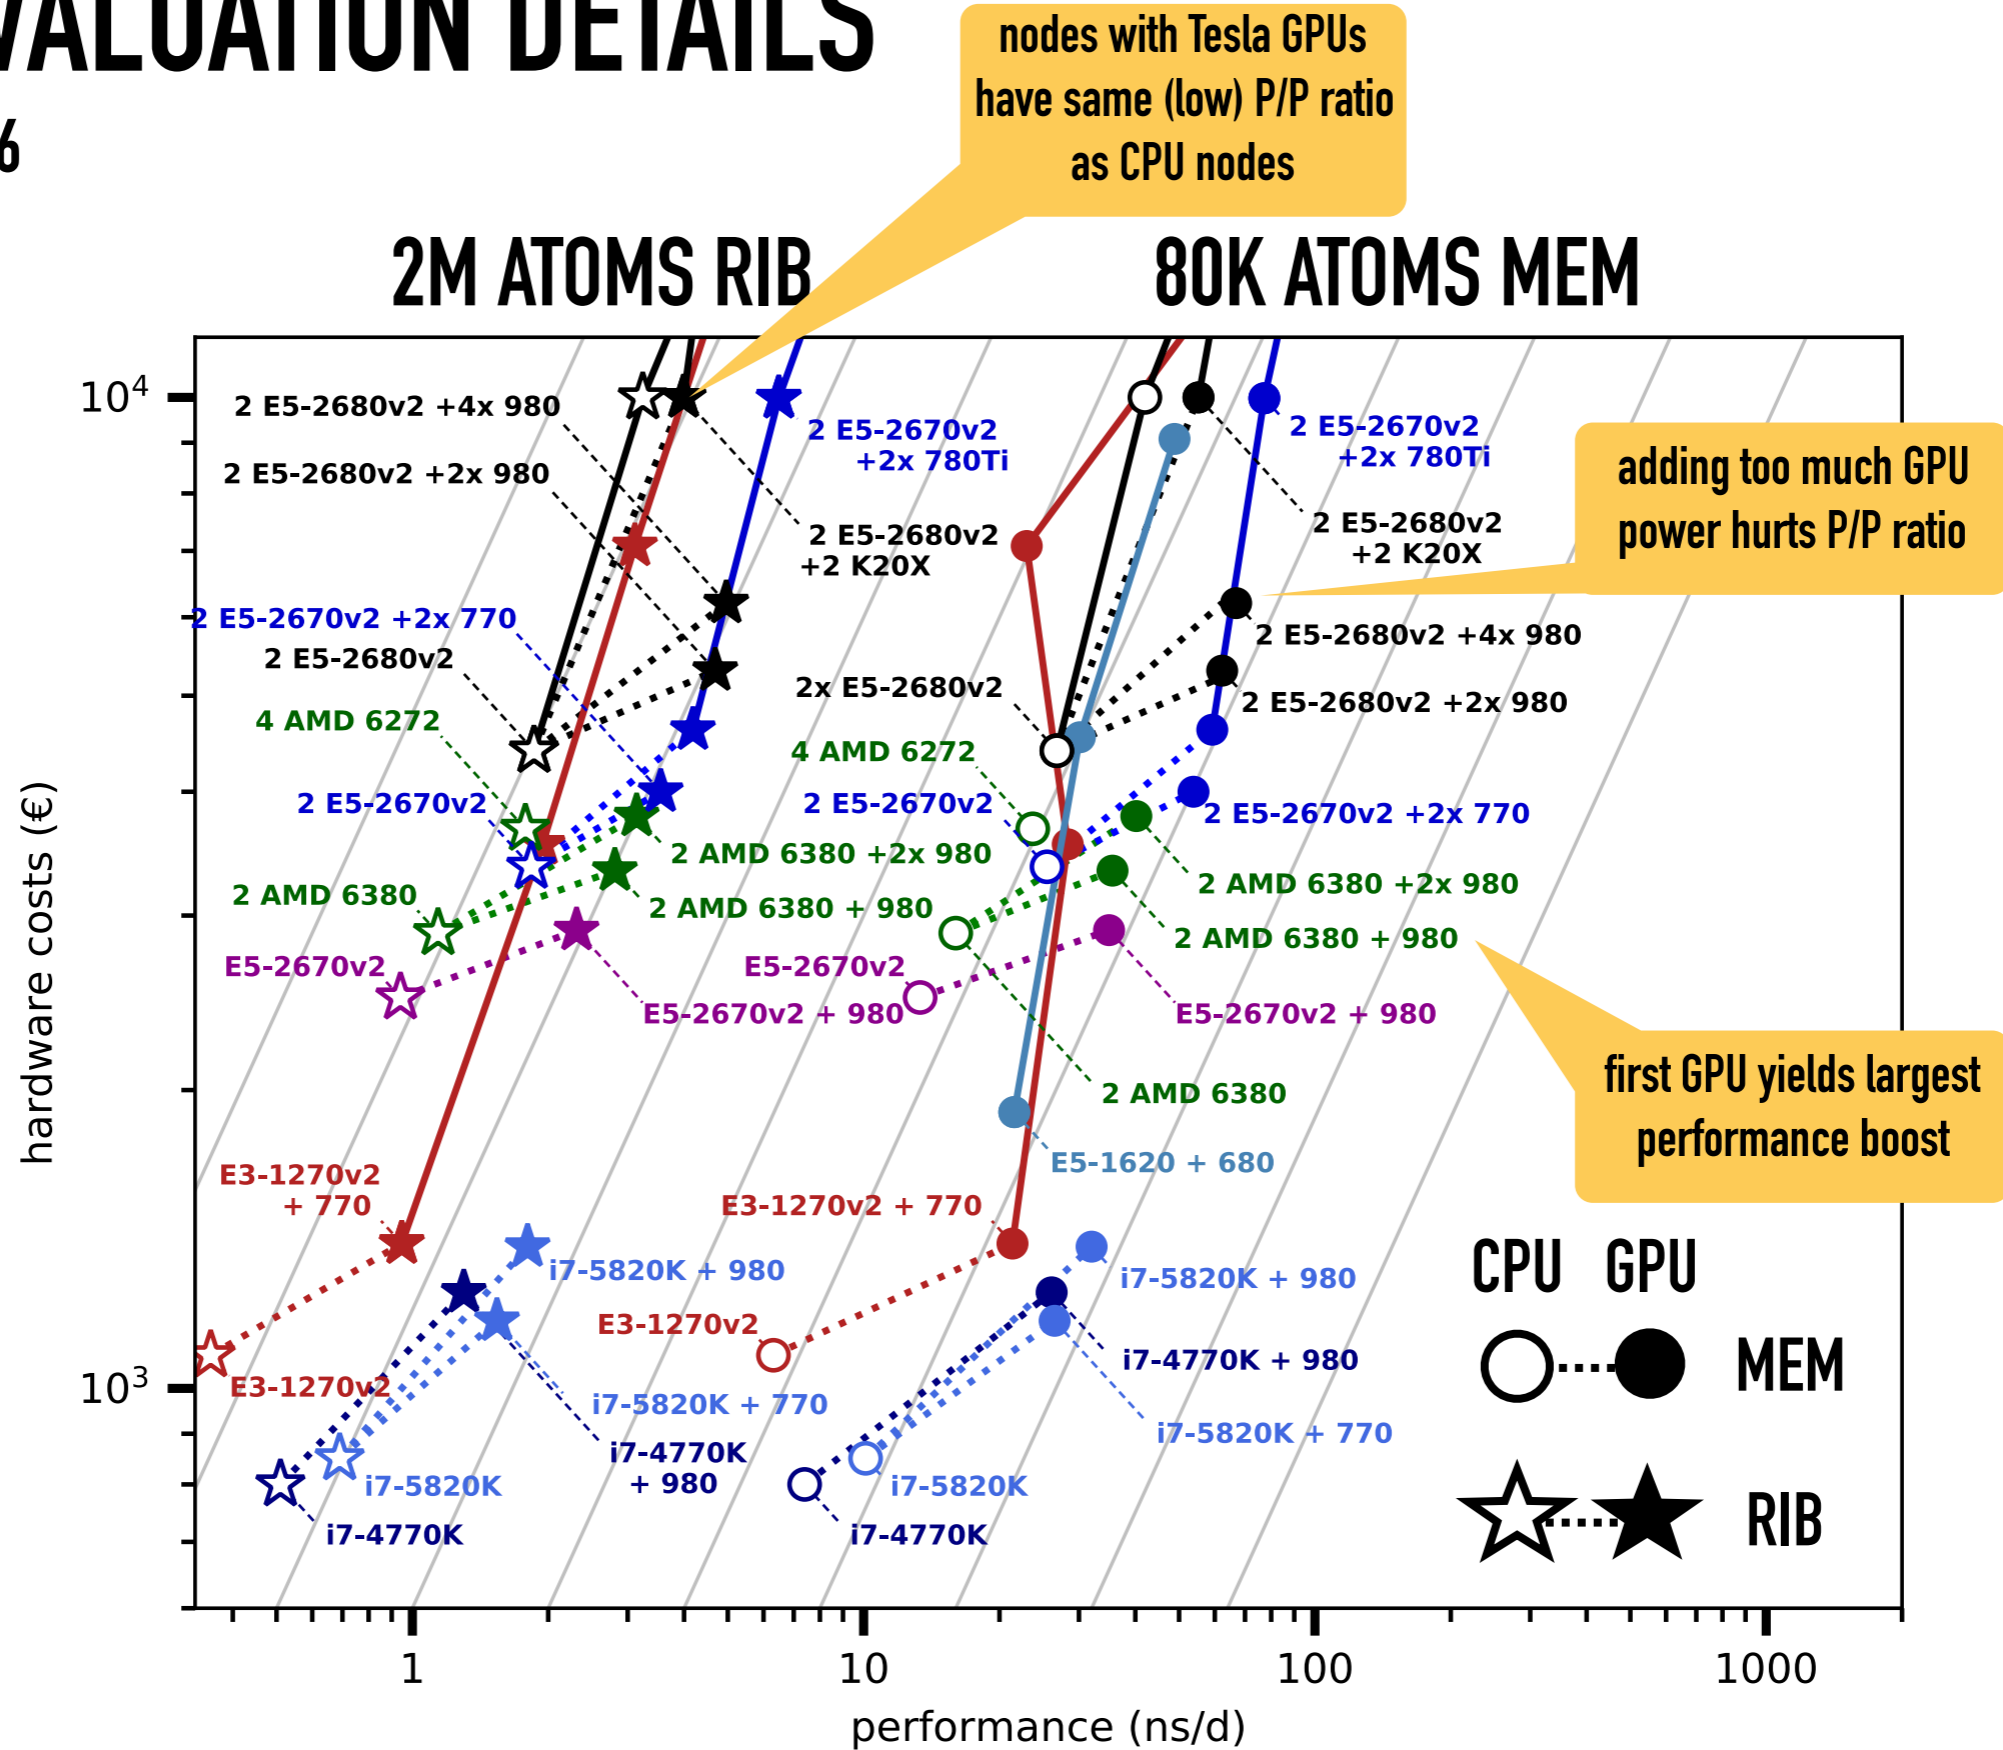
2014 EVALUATION SUMMARY

- ◆ get prices + benchmark GROMACS performance for >50 hardware configurations
- ◆ 12 CPU types + 13 GPU types
- ◆ 2 benchmark systems:
 - ◆ MEM 80k atoms
 - ◆ RIB 2 M atoms
- ◆ on each hardware determined fastest settings for running one simulation
 - ◆ # MPI ranks
 - ◆ # OpenMP threads
 - ◆ # separate PME ranks



2014 EVALUATION DETAILS

GROMACS 4.6


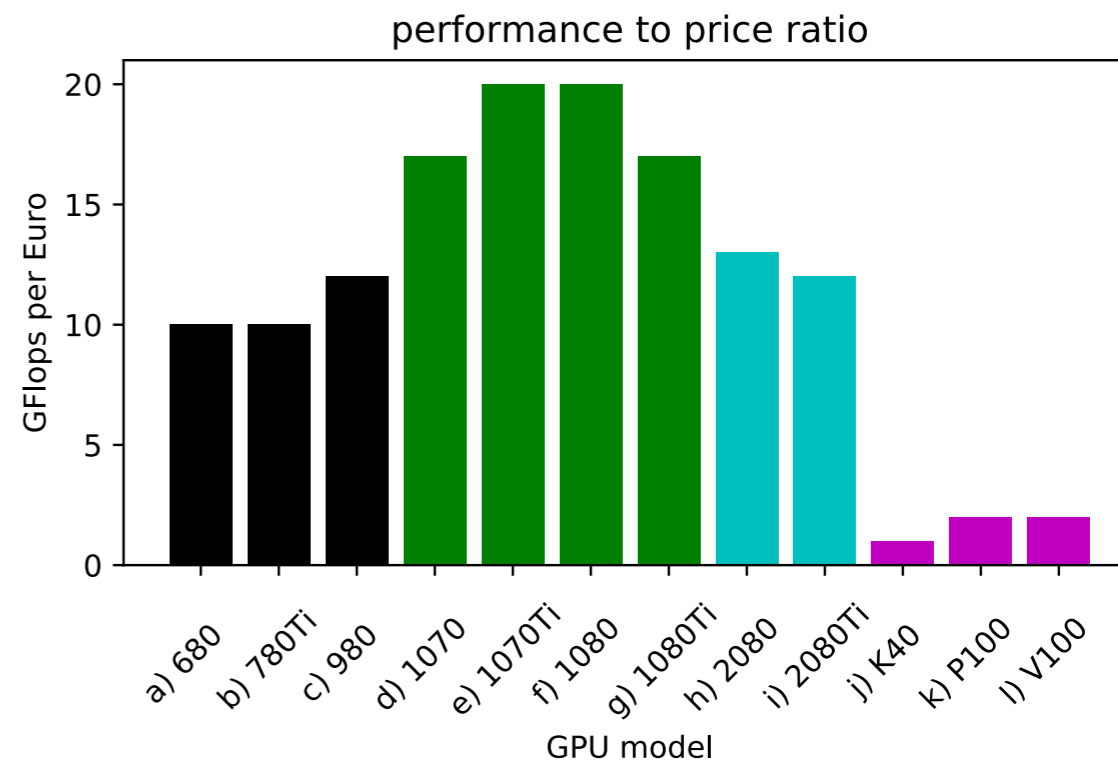
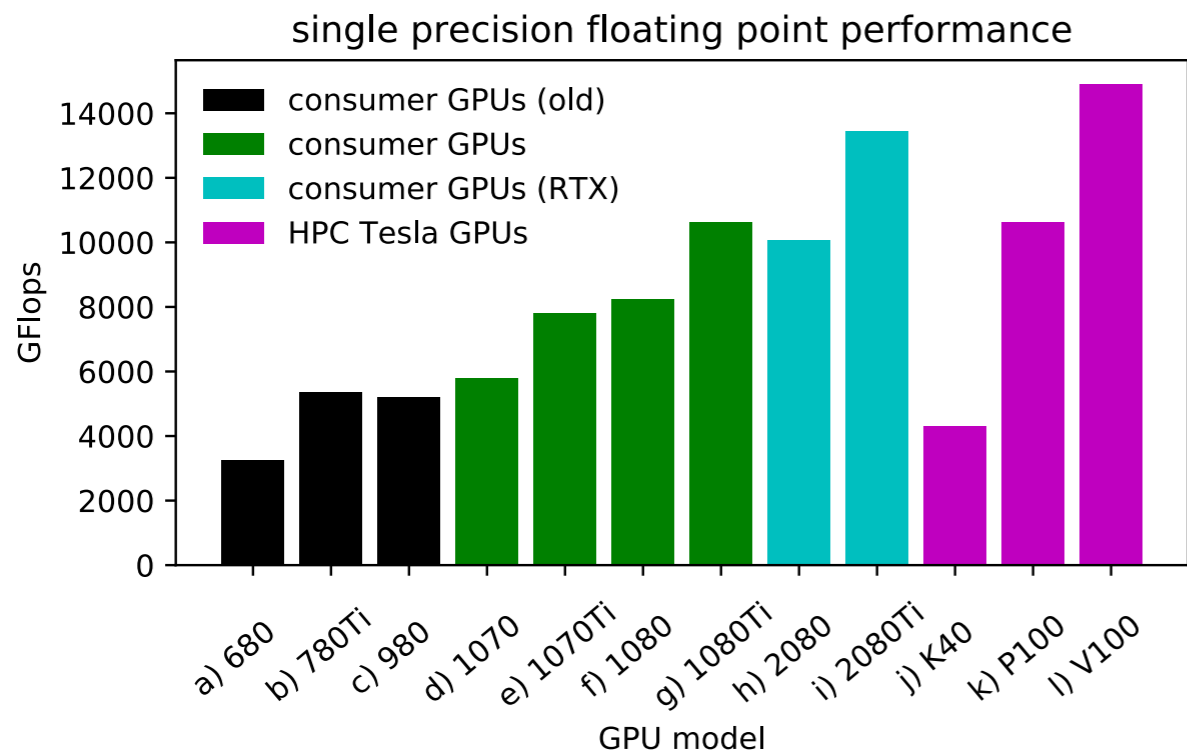


2018

WHAT'S NEW?


- 1. Hardware:**
GPUs with higher performance
- 2. Software:**
PME can be offloaded to the GPU
(among many other features!)
- 3. Benchmarks:**
(Slight) change of protocol

1. HARDWARE DEVELOPMENTS— GPUS 2014...2018

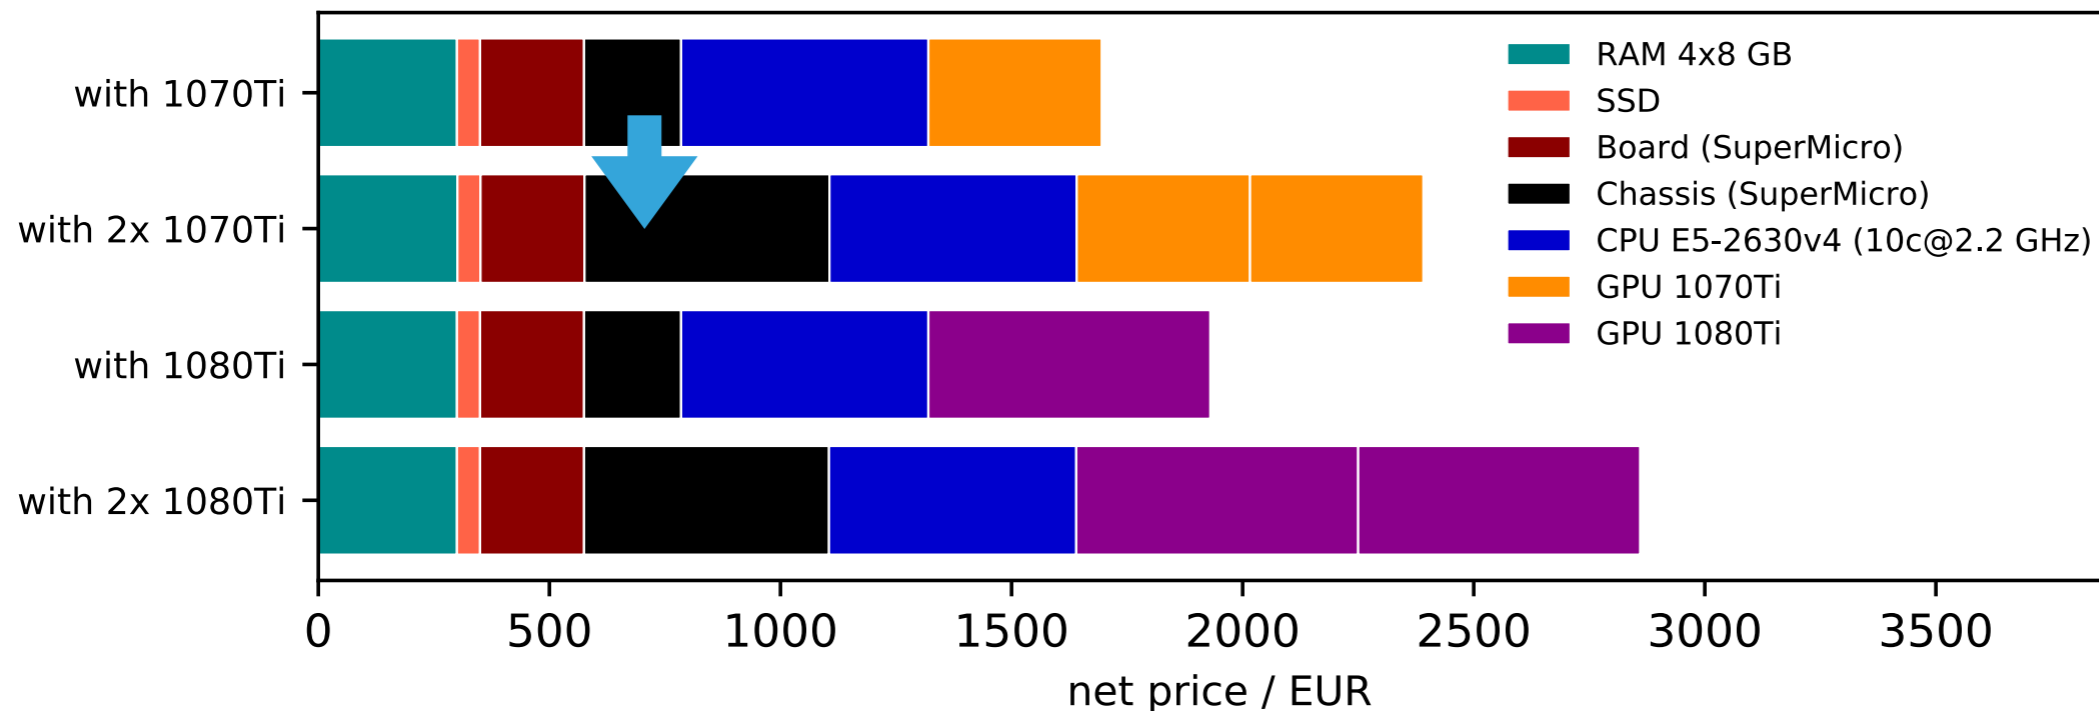


1. HARDWARE: BUILDING A GPU NODE

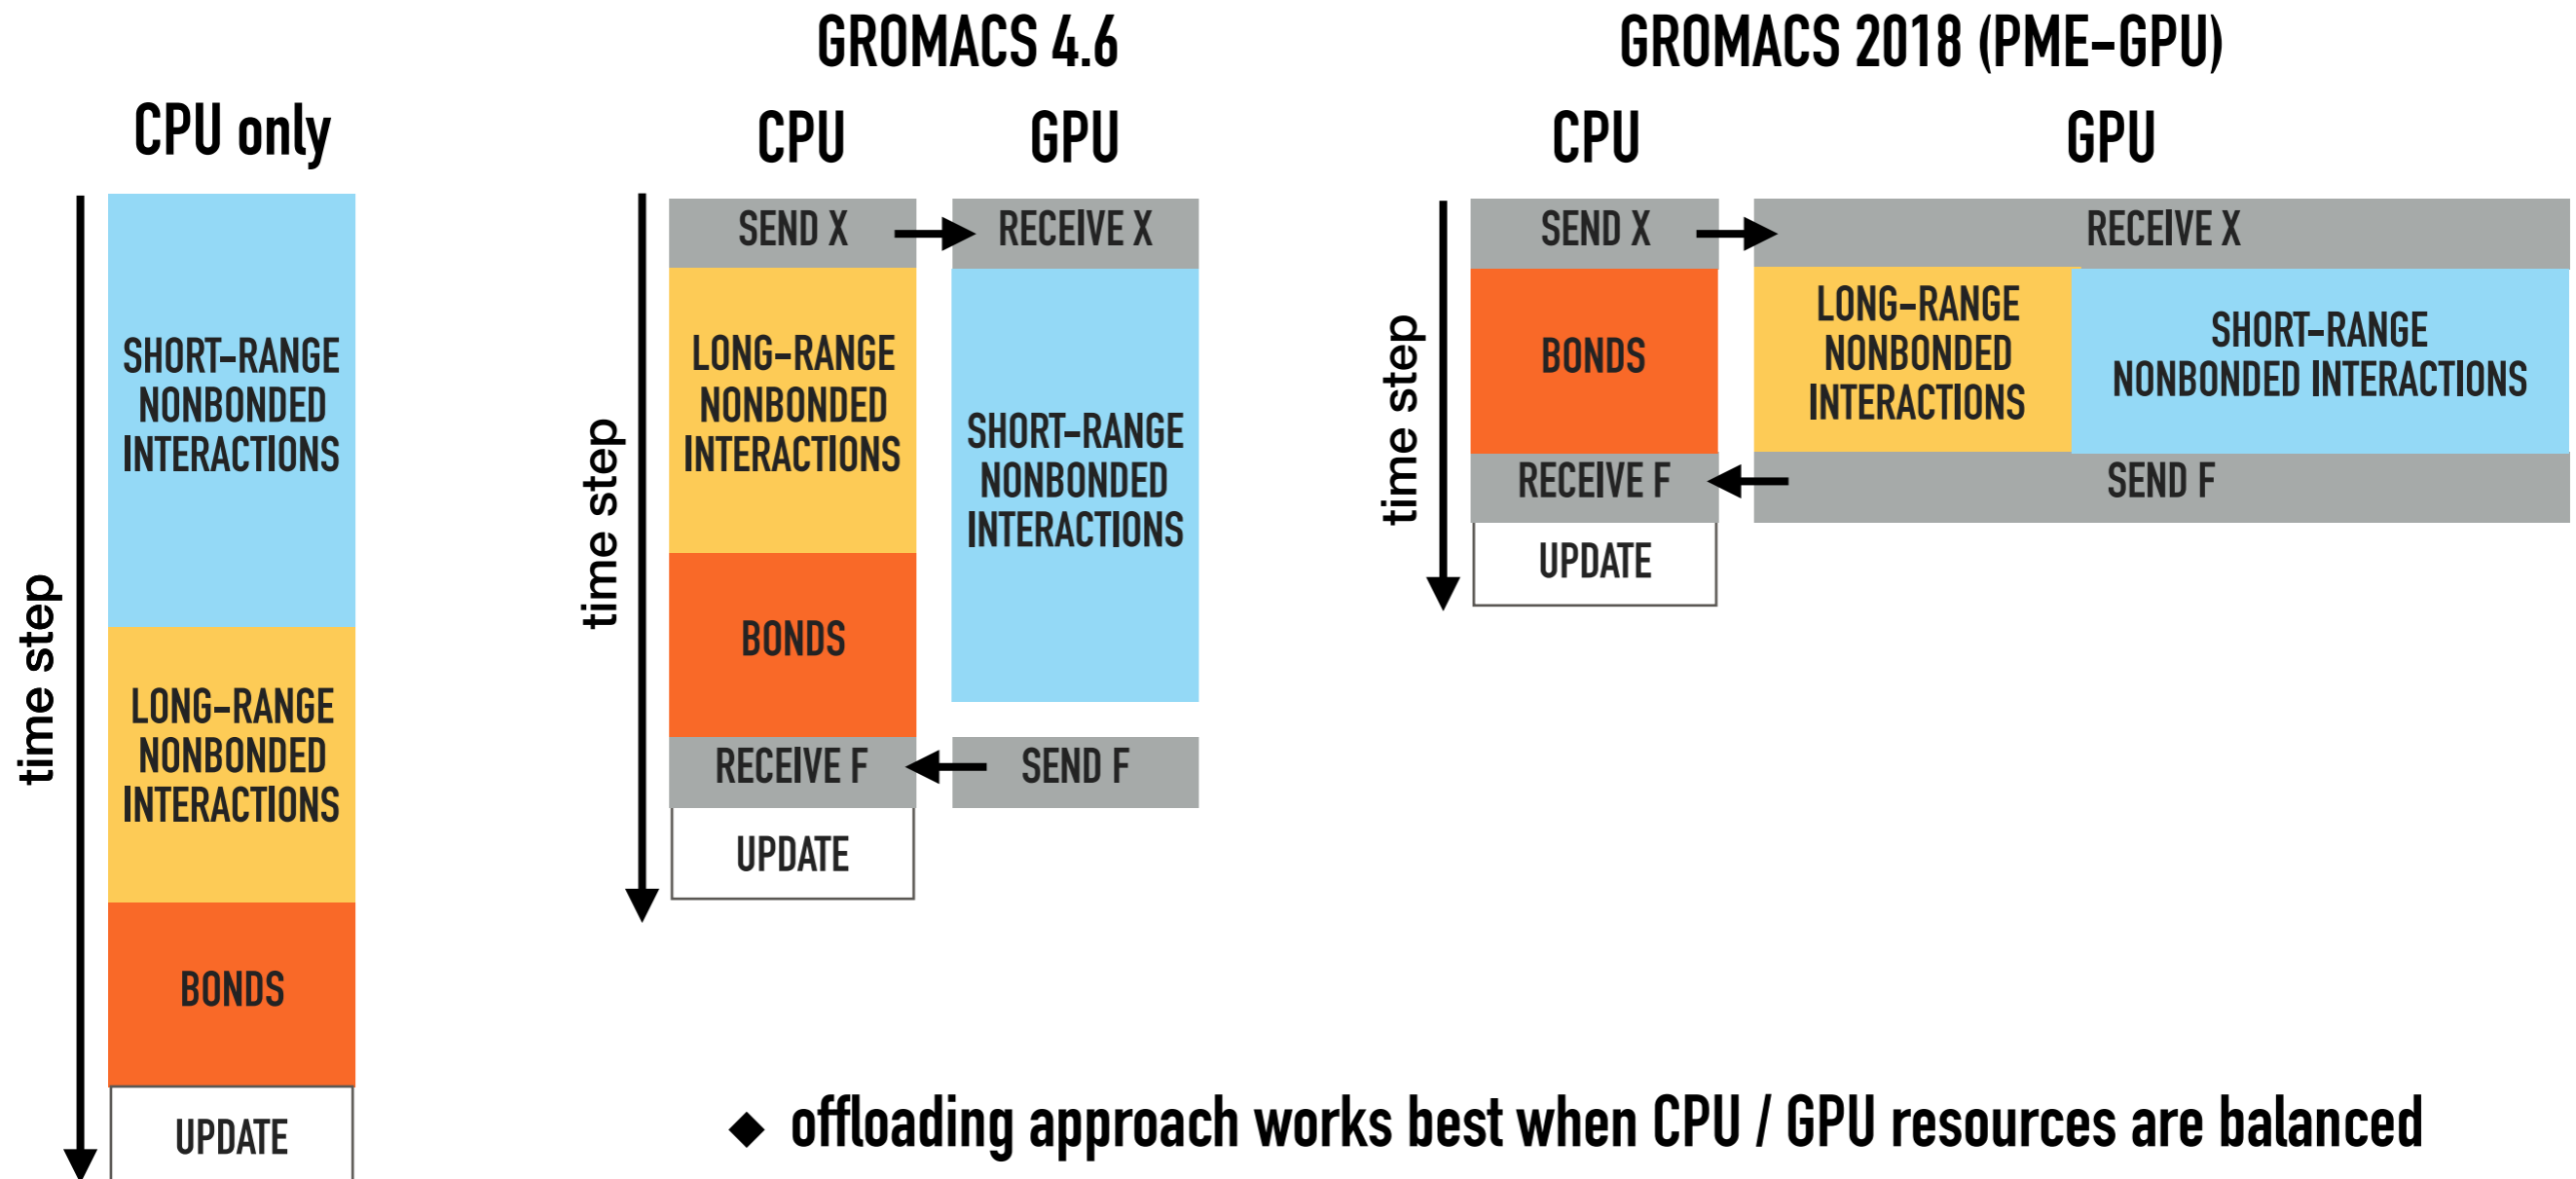
why not exclusively use GPUs with optimal performance to price (P/P) ratio?



Breakdown of node price example

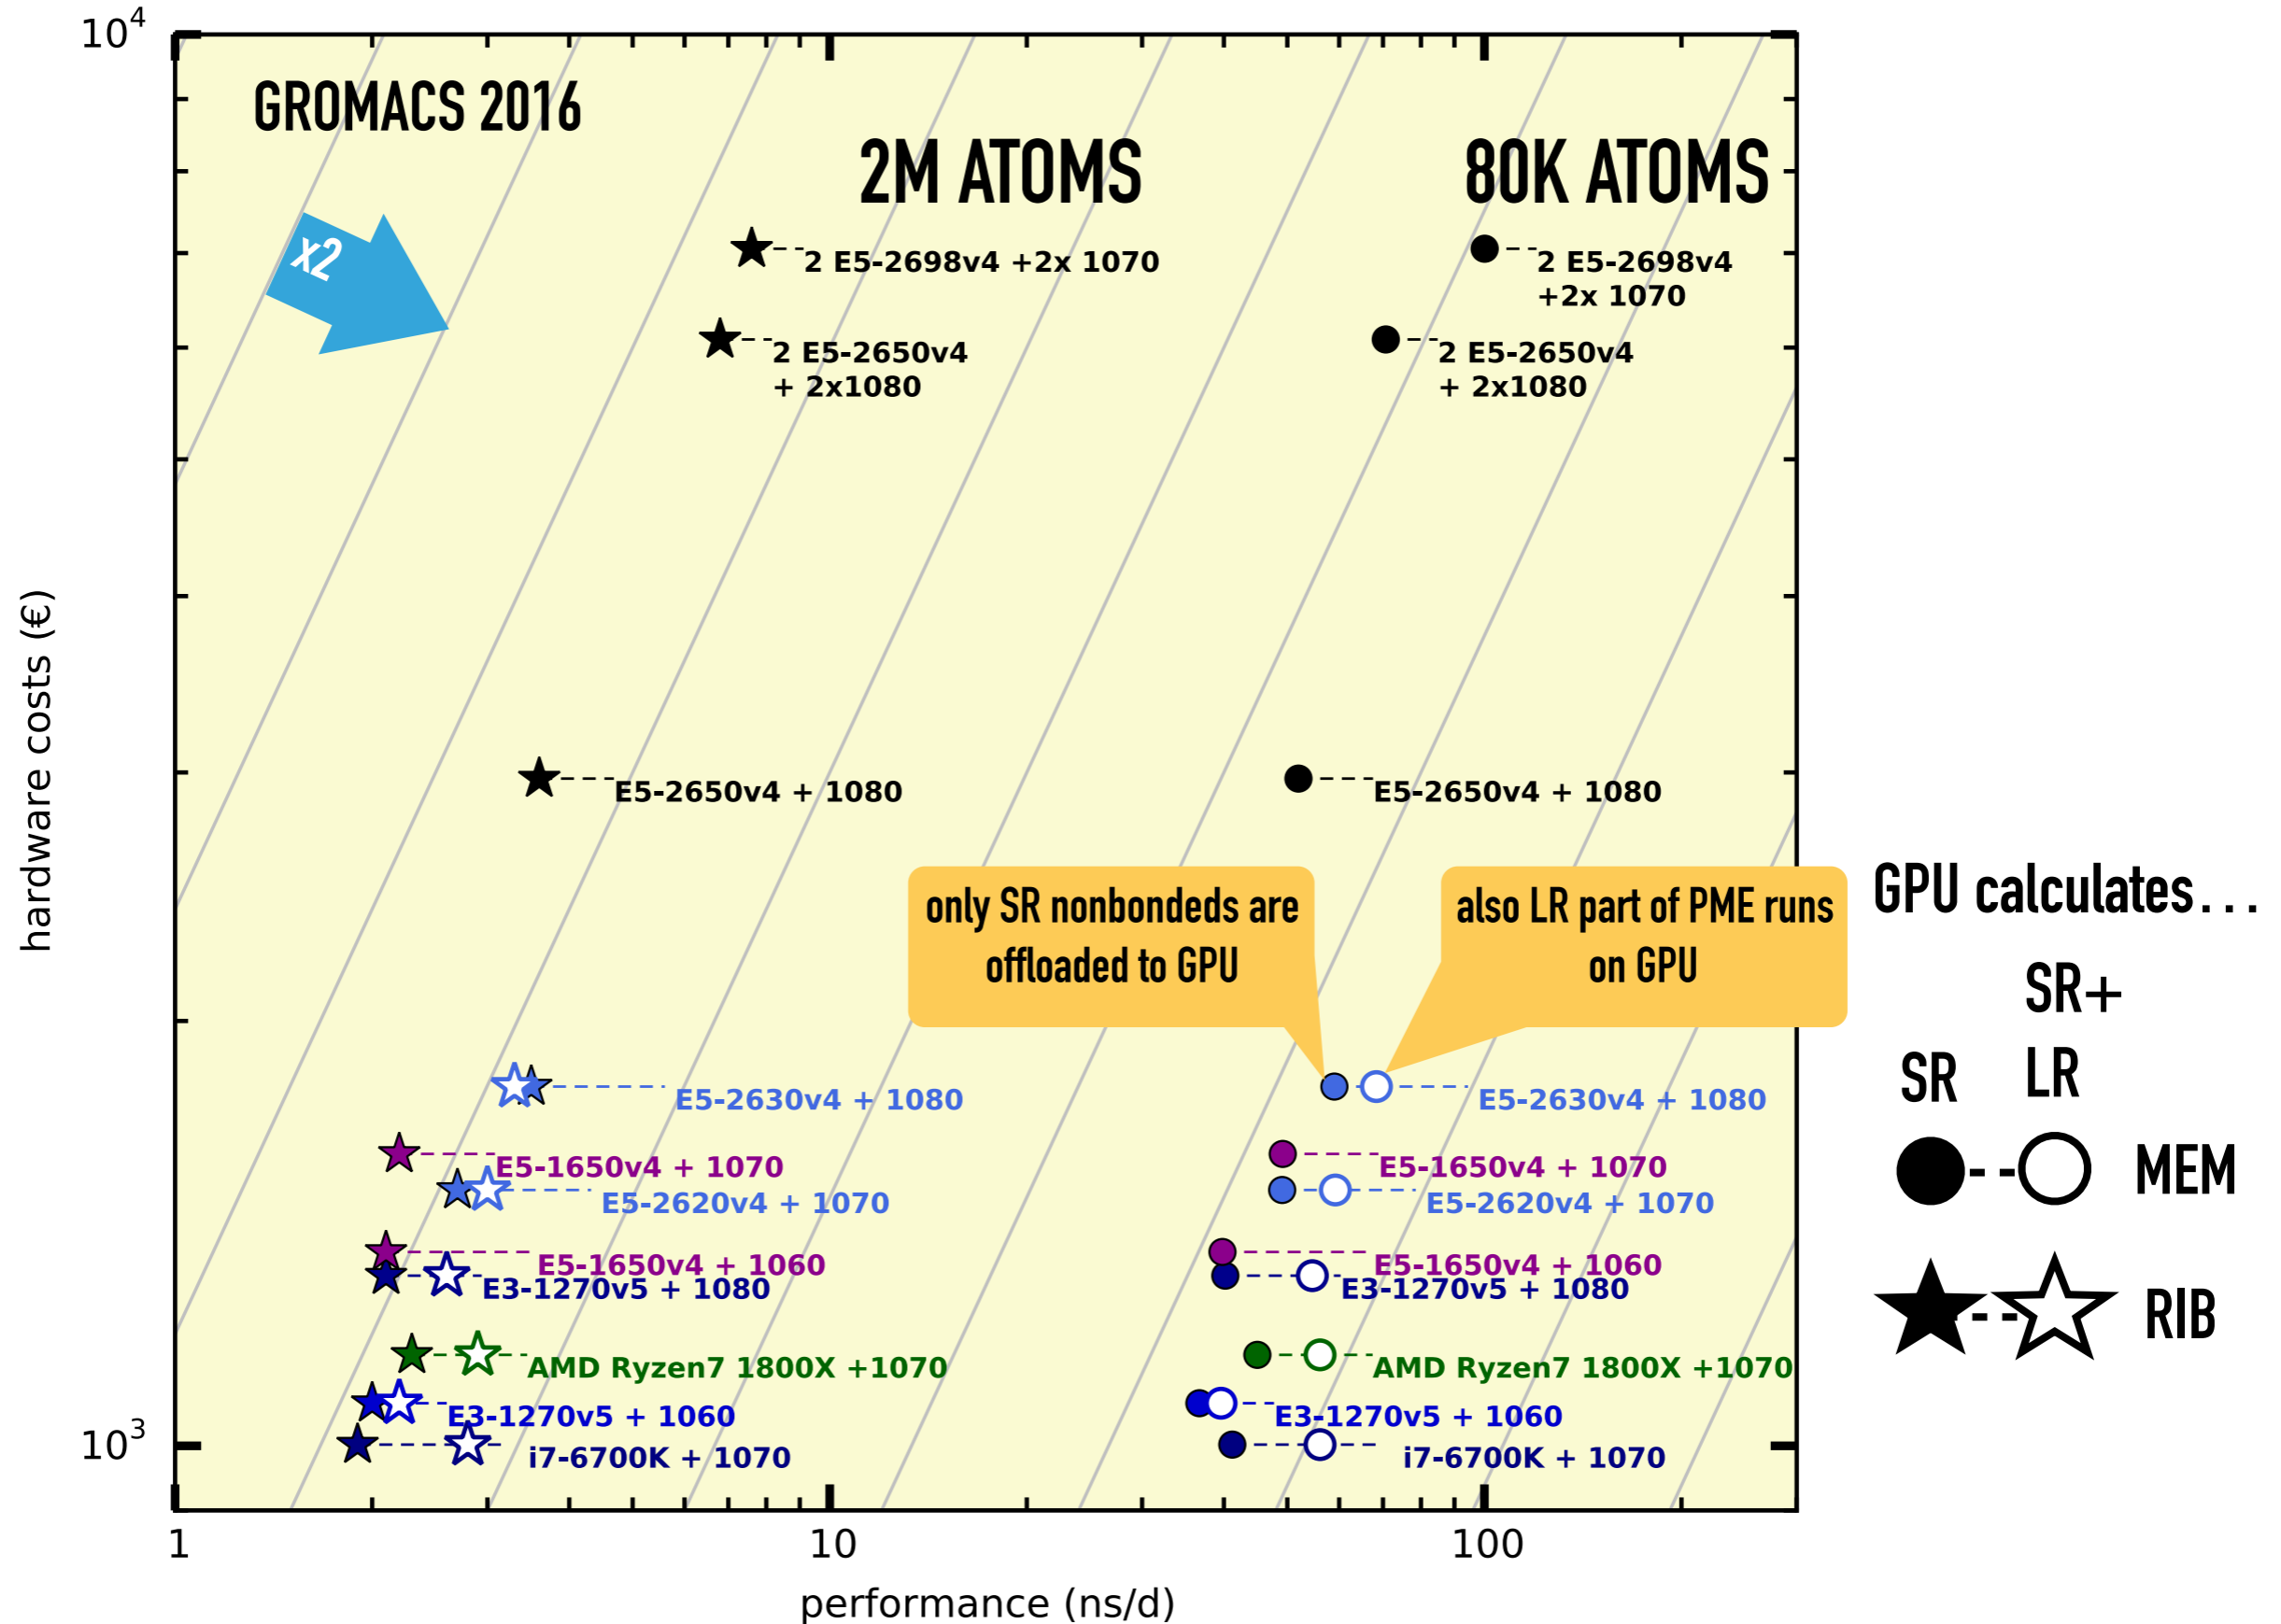


2. SOFTWARE: GPU OFFLOADING SCHEMES



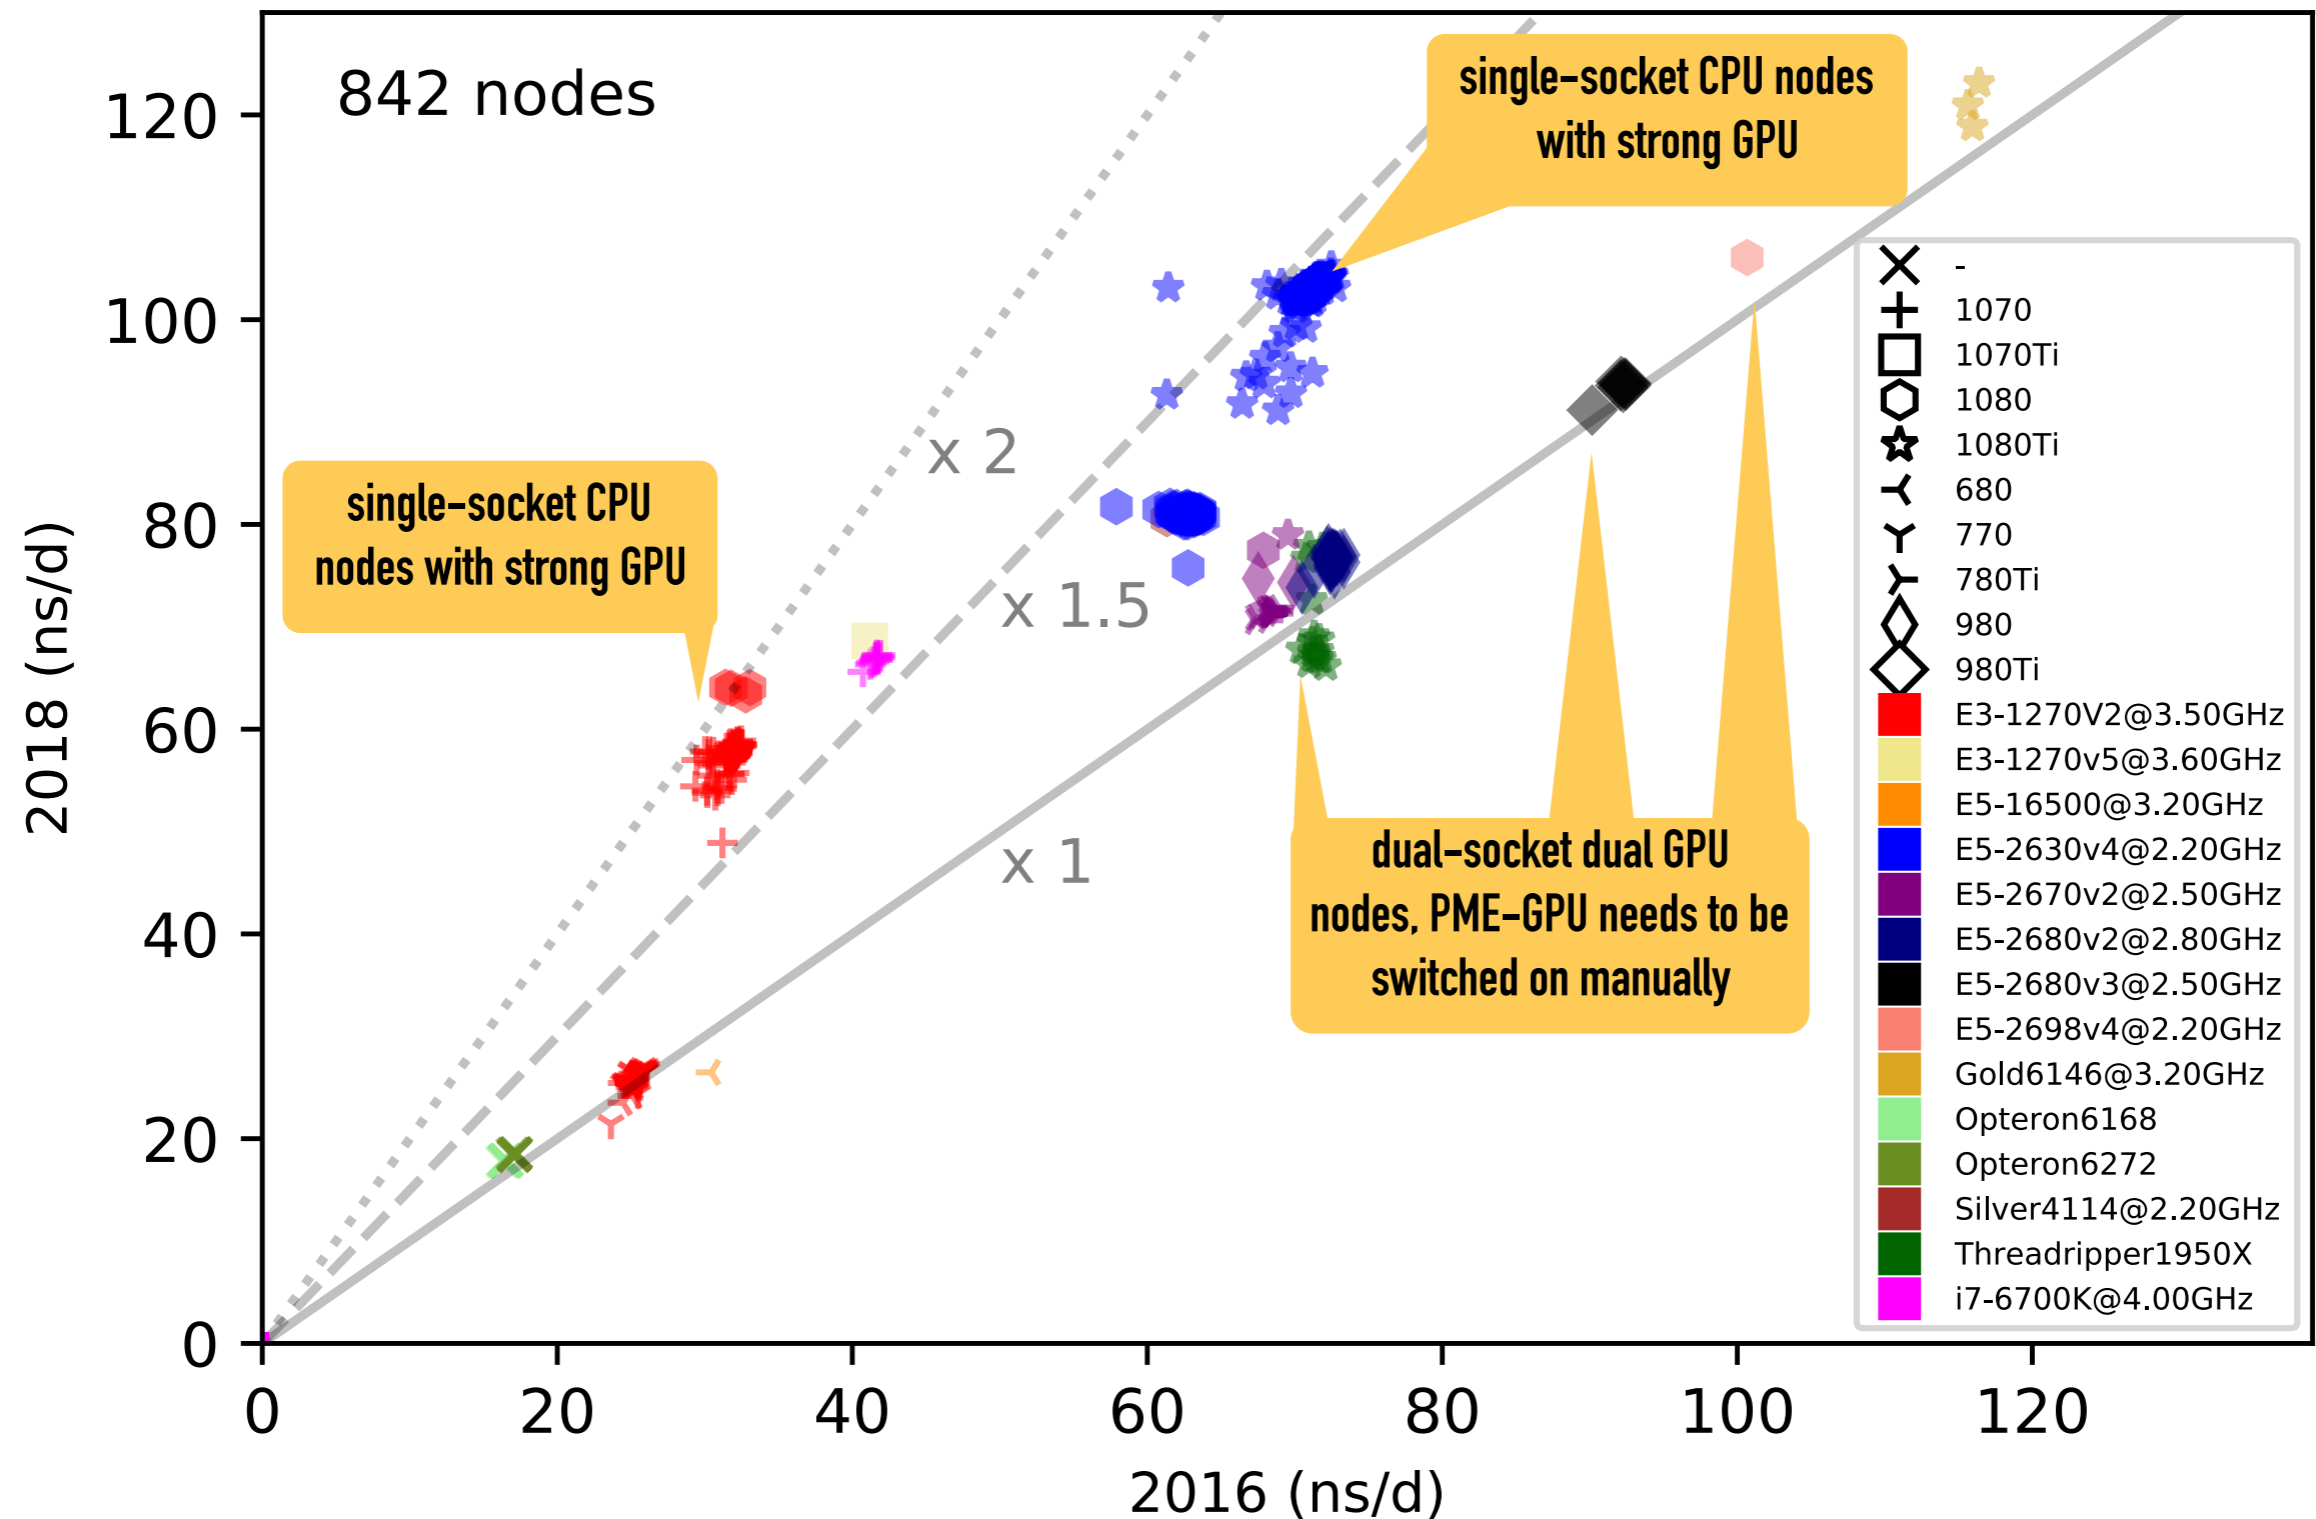
- ◆ offloading approach works best when CPU / GPU resources are balanced
- ◆ if a run is CPU-bound, more GPU power won't shorten the time step

2. SOFTWARE: PME ON GPU



2. SOFTWARE: PME ON GPU

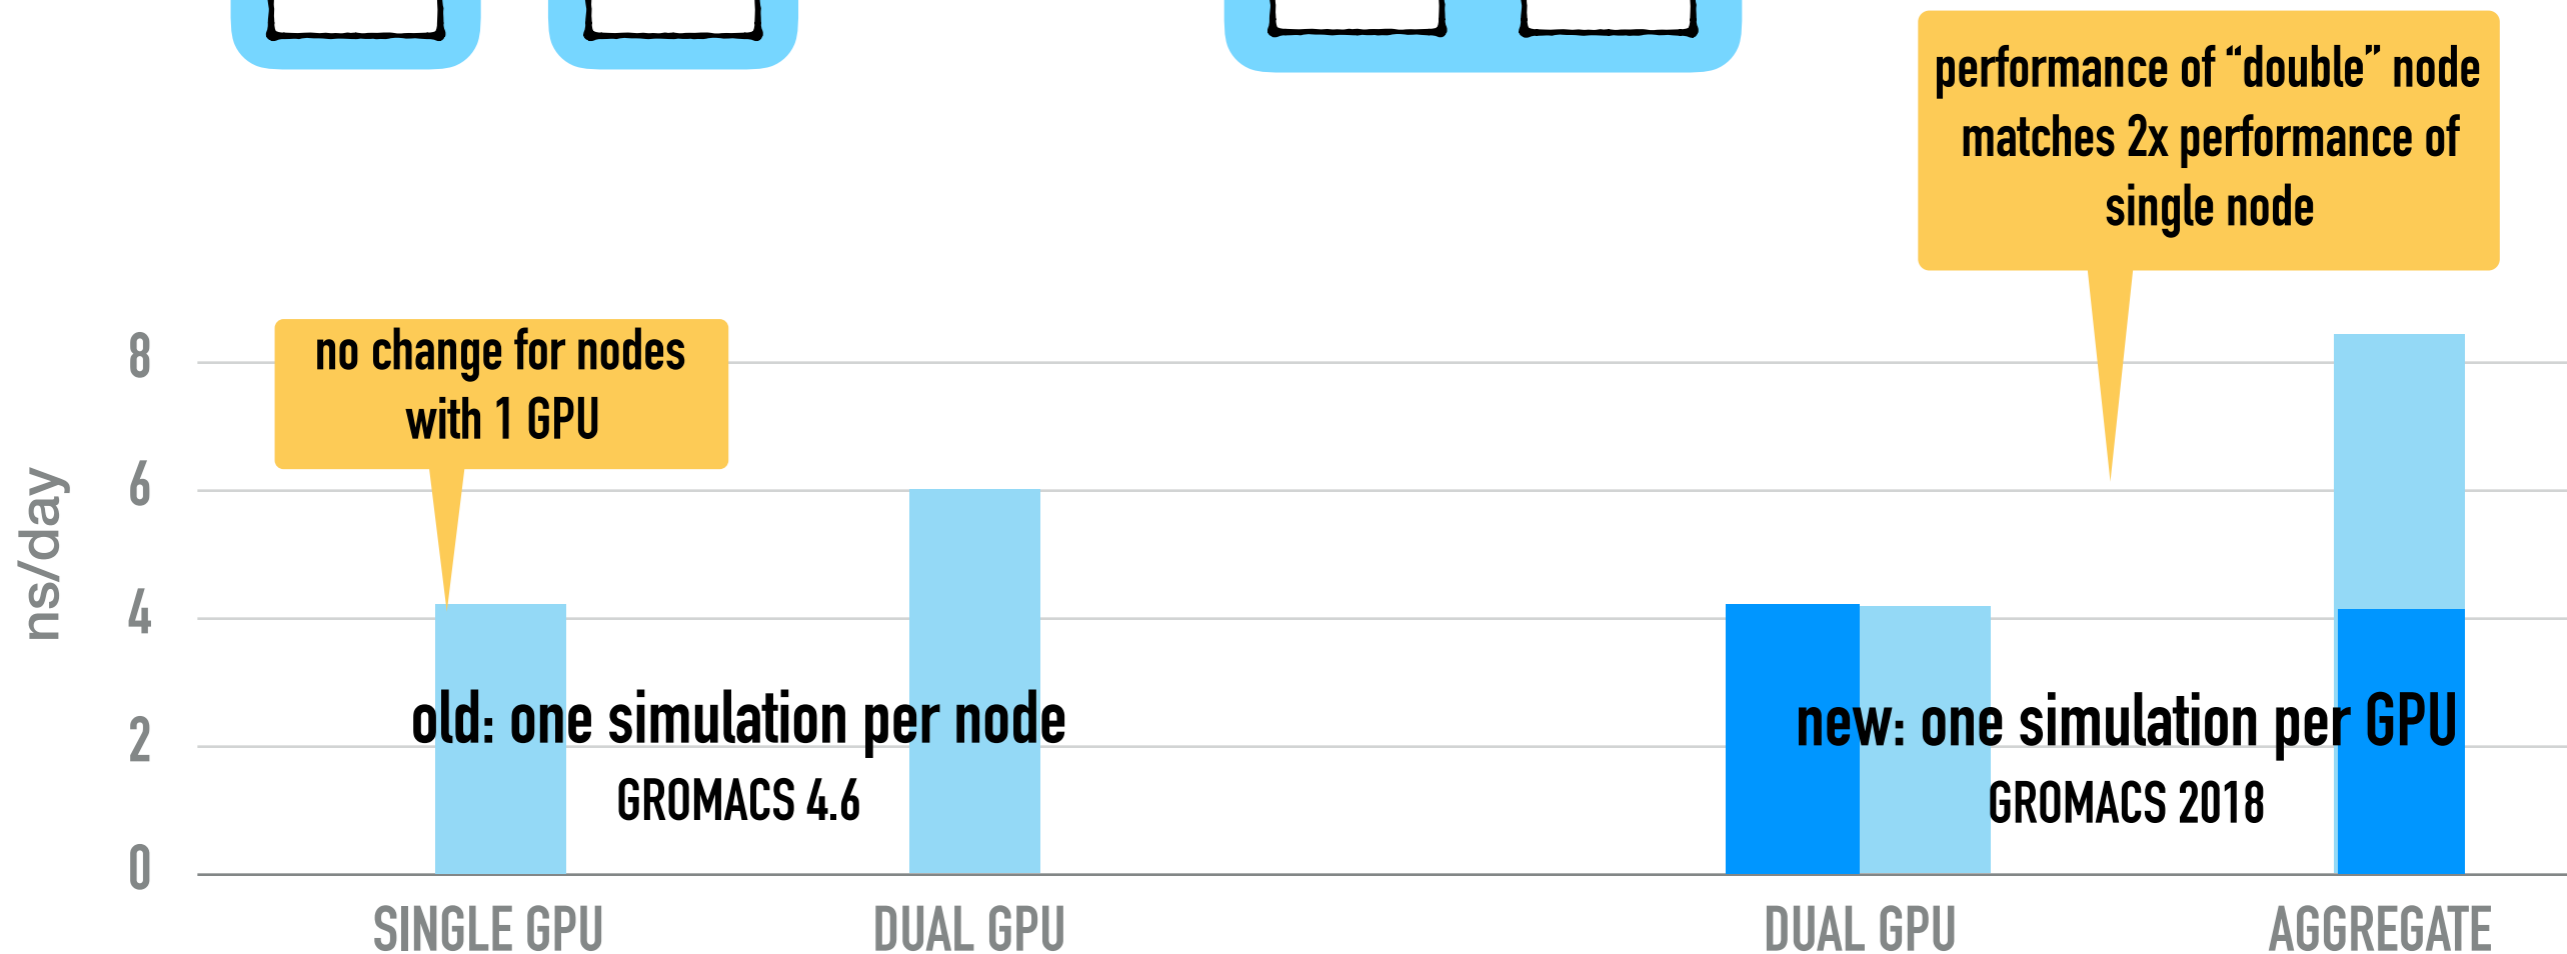

cluster health check using 80k atoms MEM benchmark



3. BENCHMARK PROTOCOL CHANGE - WHY?

We don't want to penalize the aggregation of compute power
(which may offer price and rack space savings!)

How do we measure the performance of a node? Ideally we get:



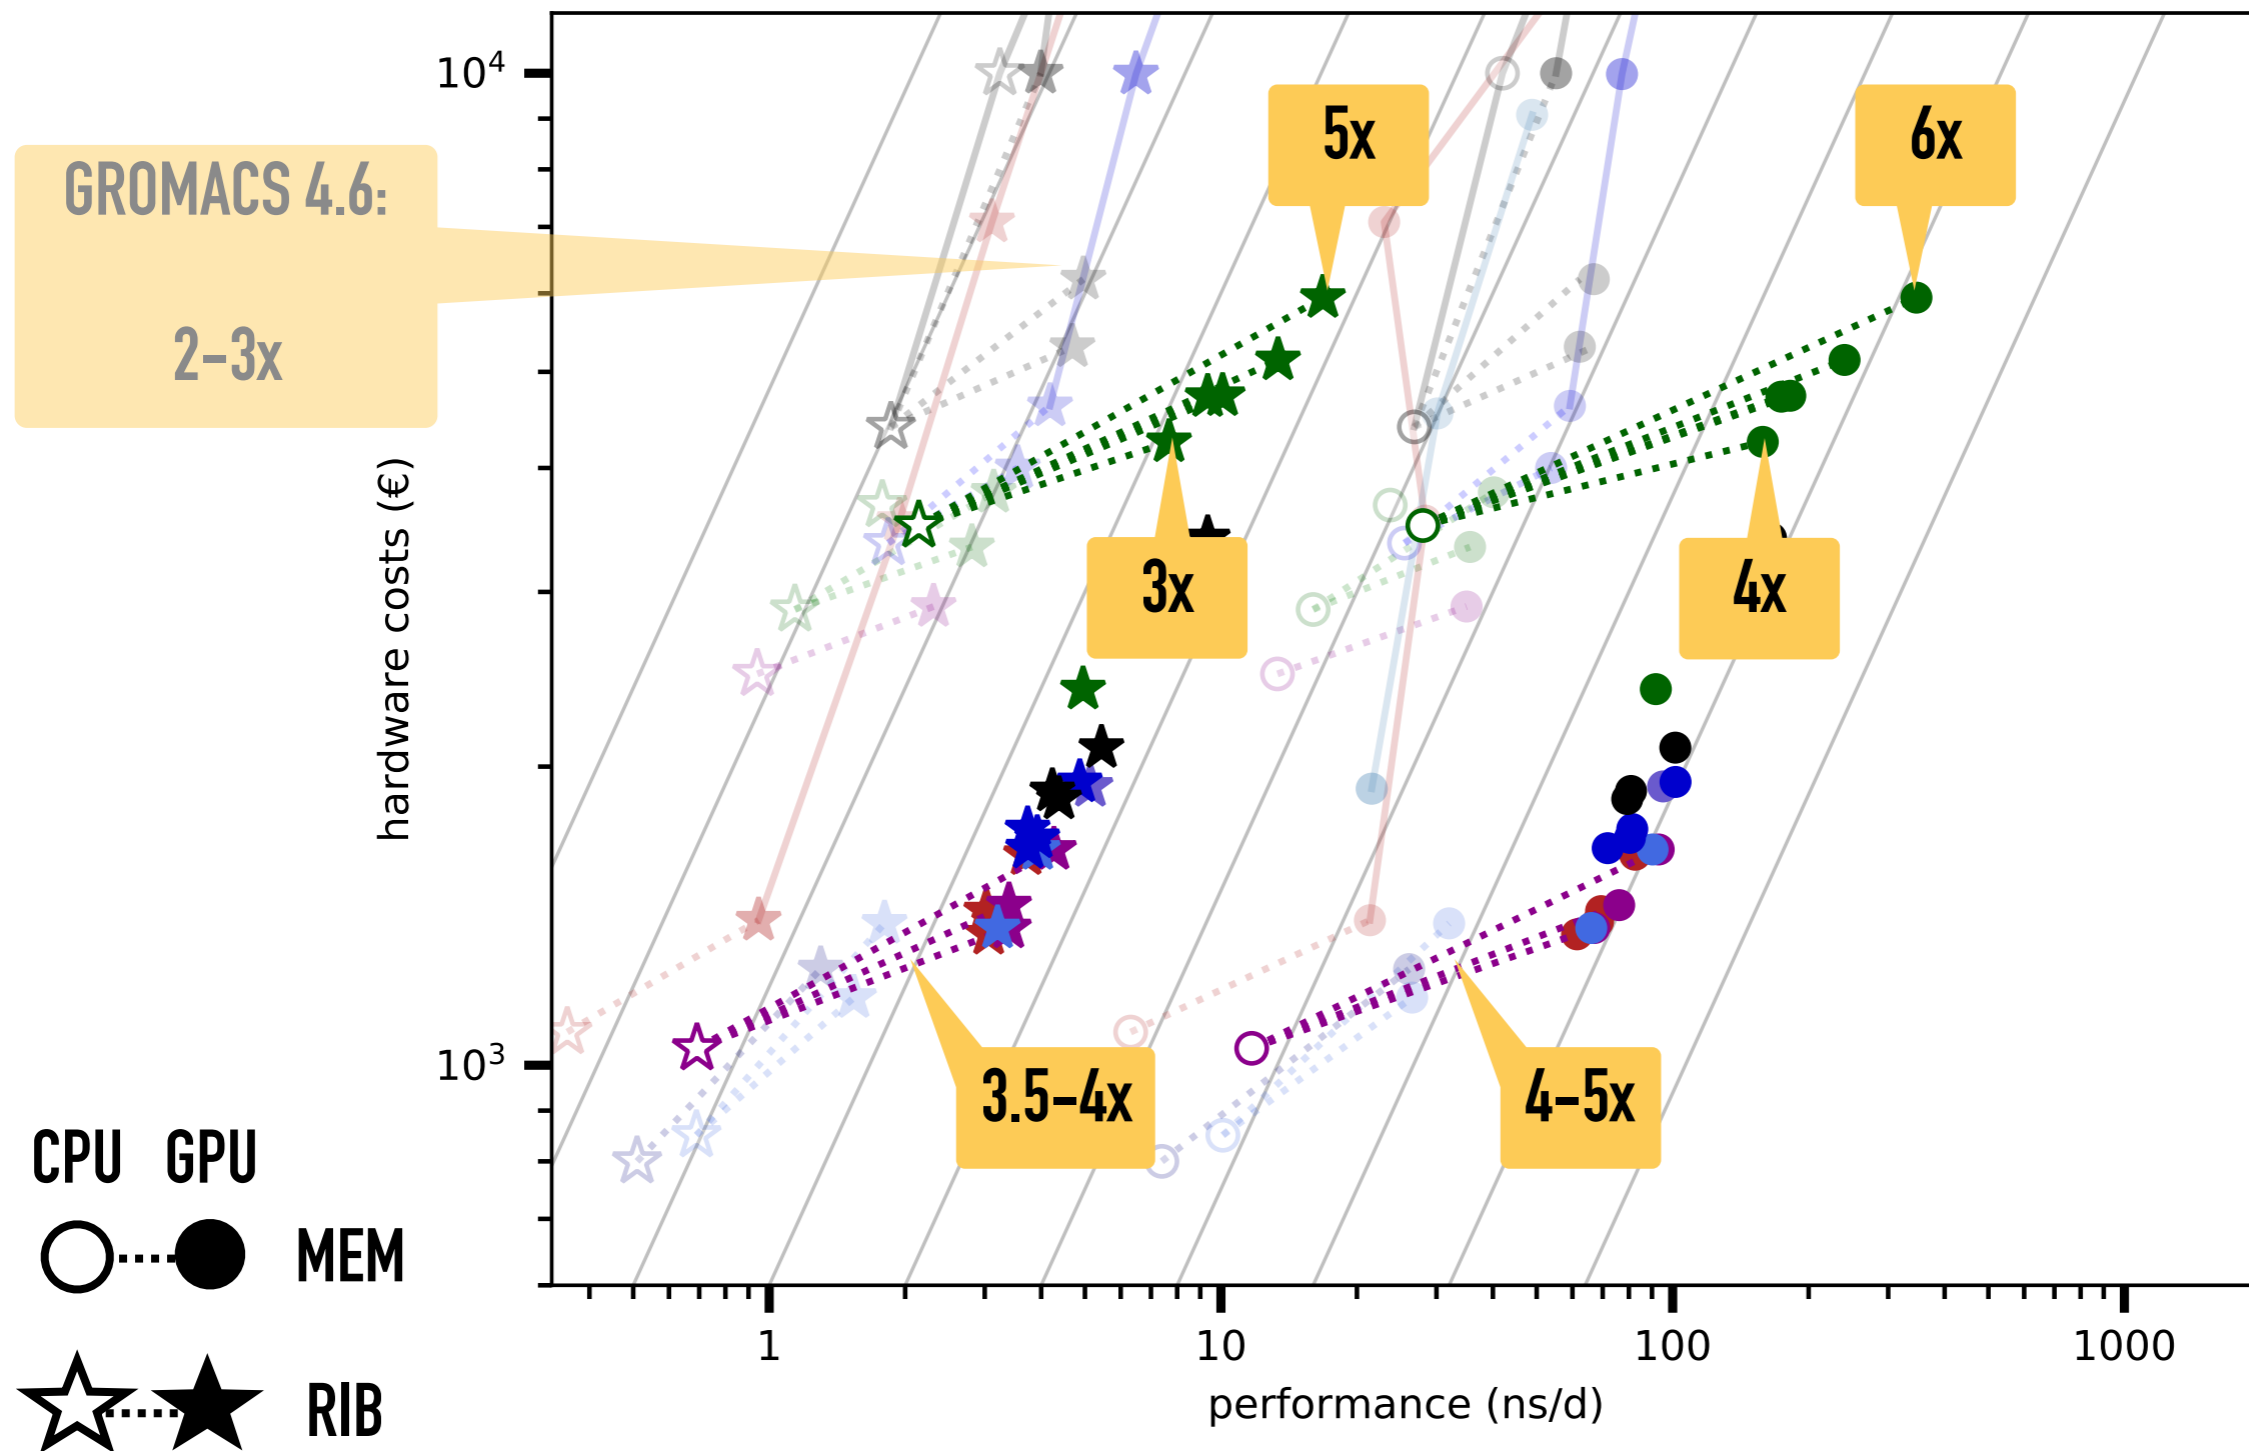
RESULTS 2018

4-YEAR PERFORMANCE TO PRICE DEVELOPMENT

2014...2018


2M ATOMS

80K ATOMS




ENERGY EFFICIENCY

ADD ENERGY COSTS TO THE BILL



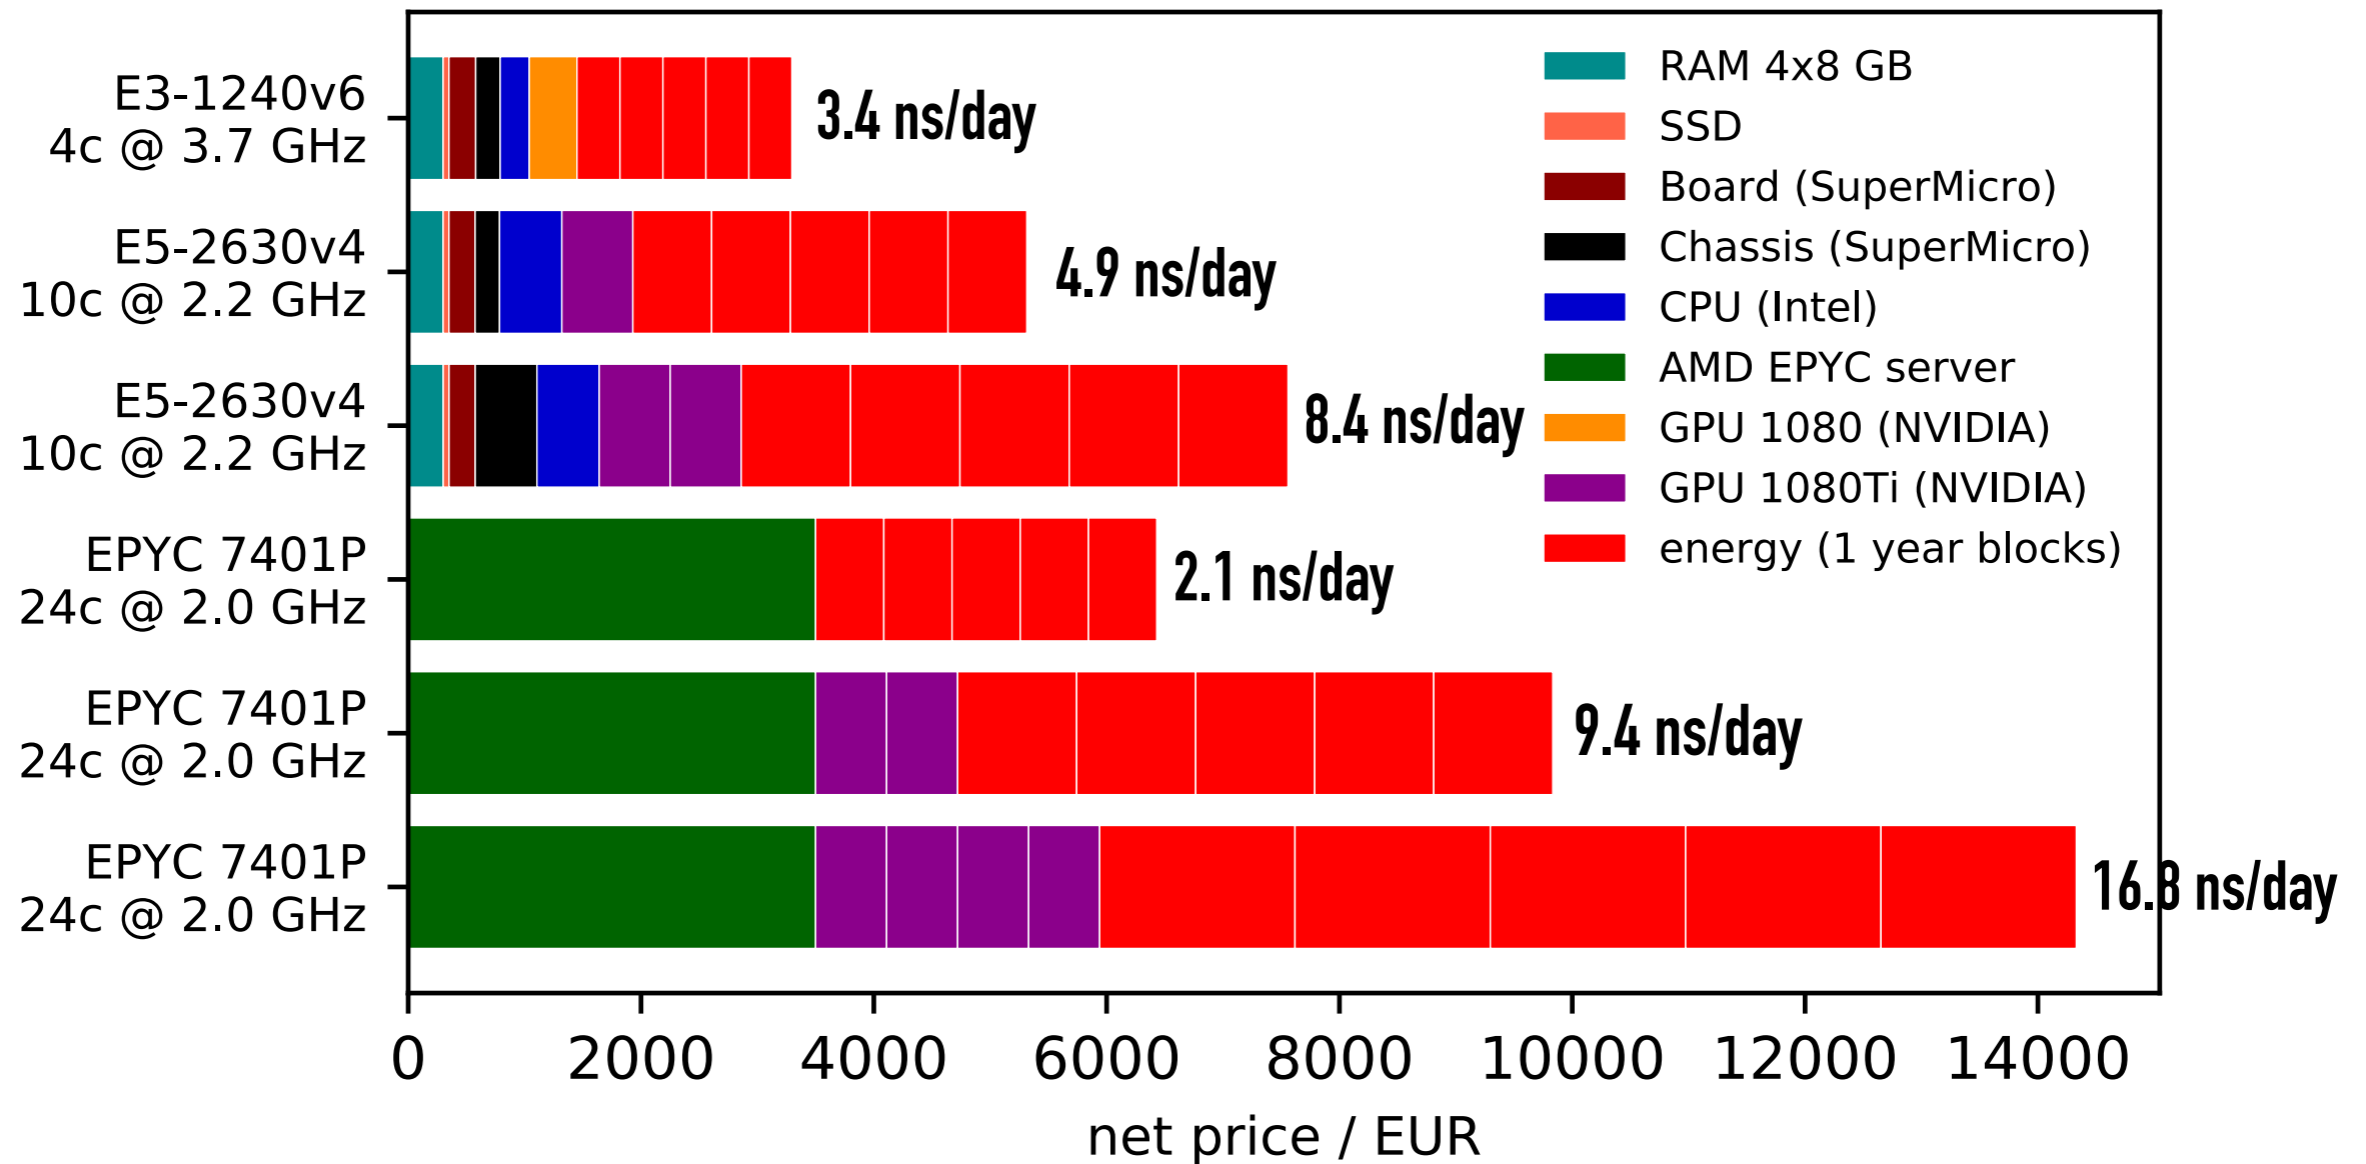
ADD ENERGY COSTS TO THE BILL

node costs taking into account energy + cooling (0.2 EUR / kWh) RIB



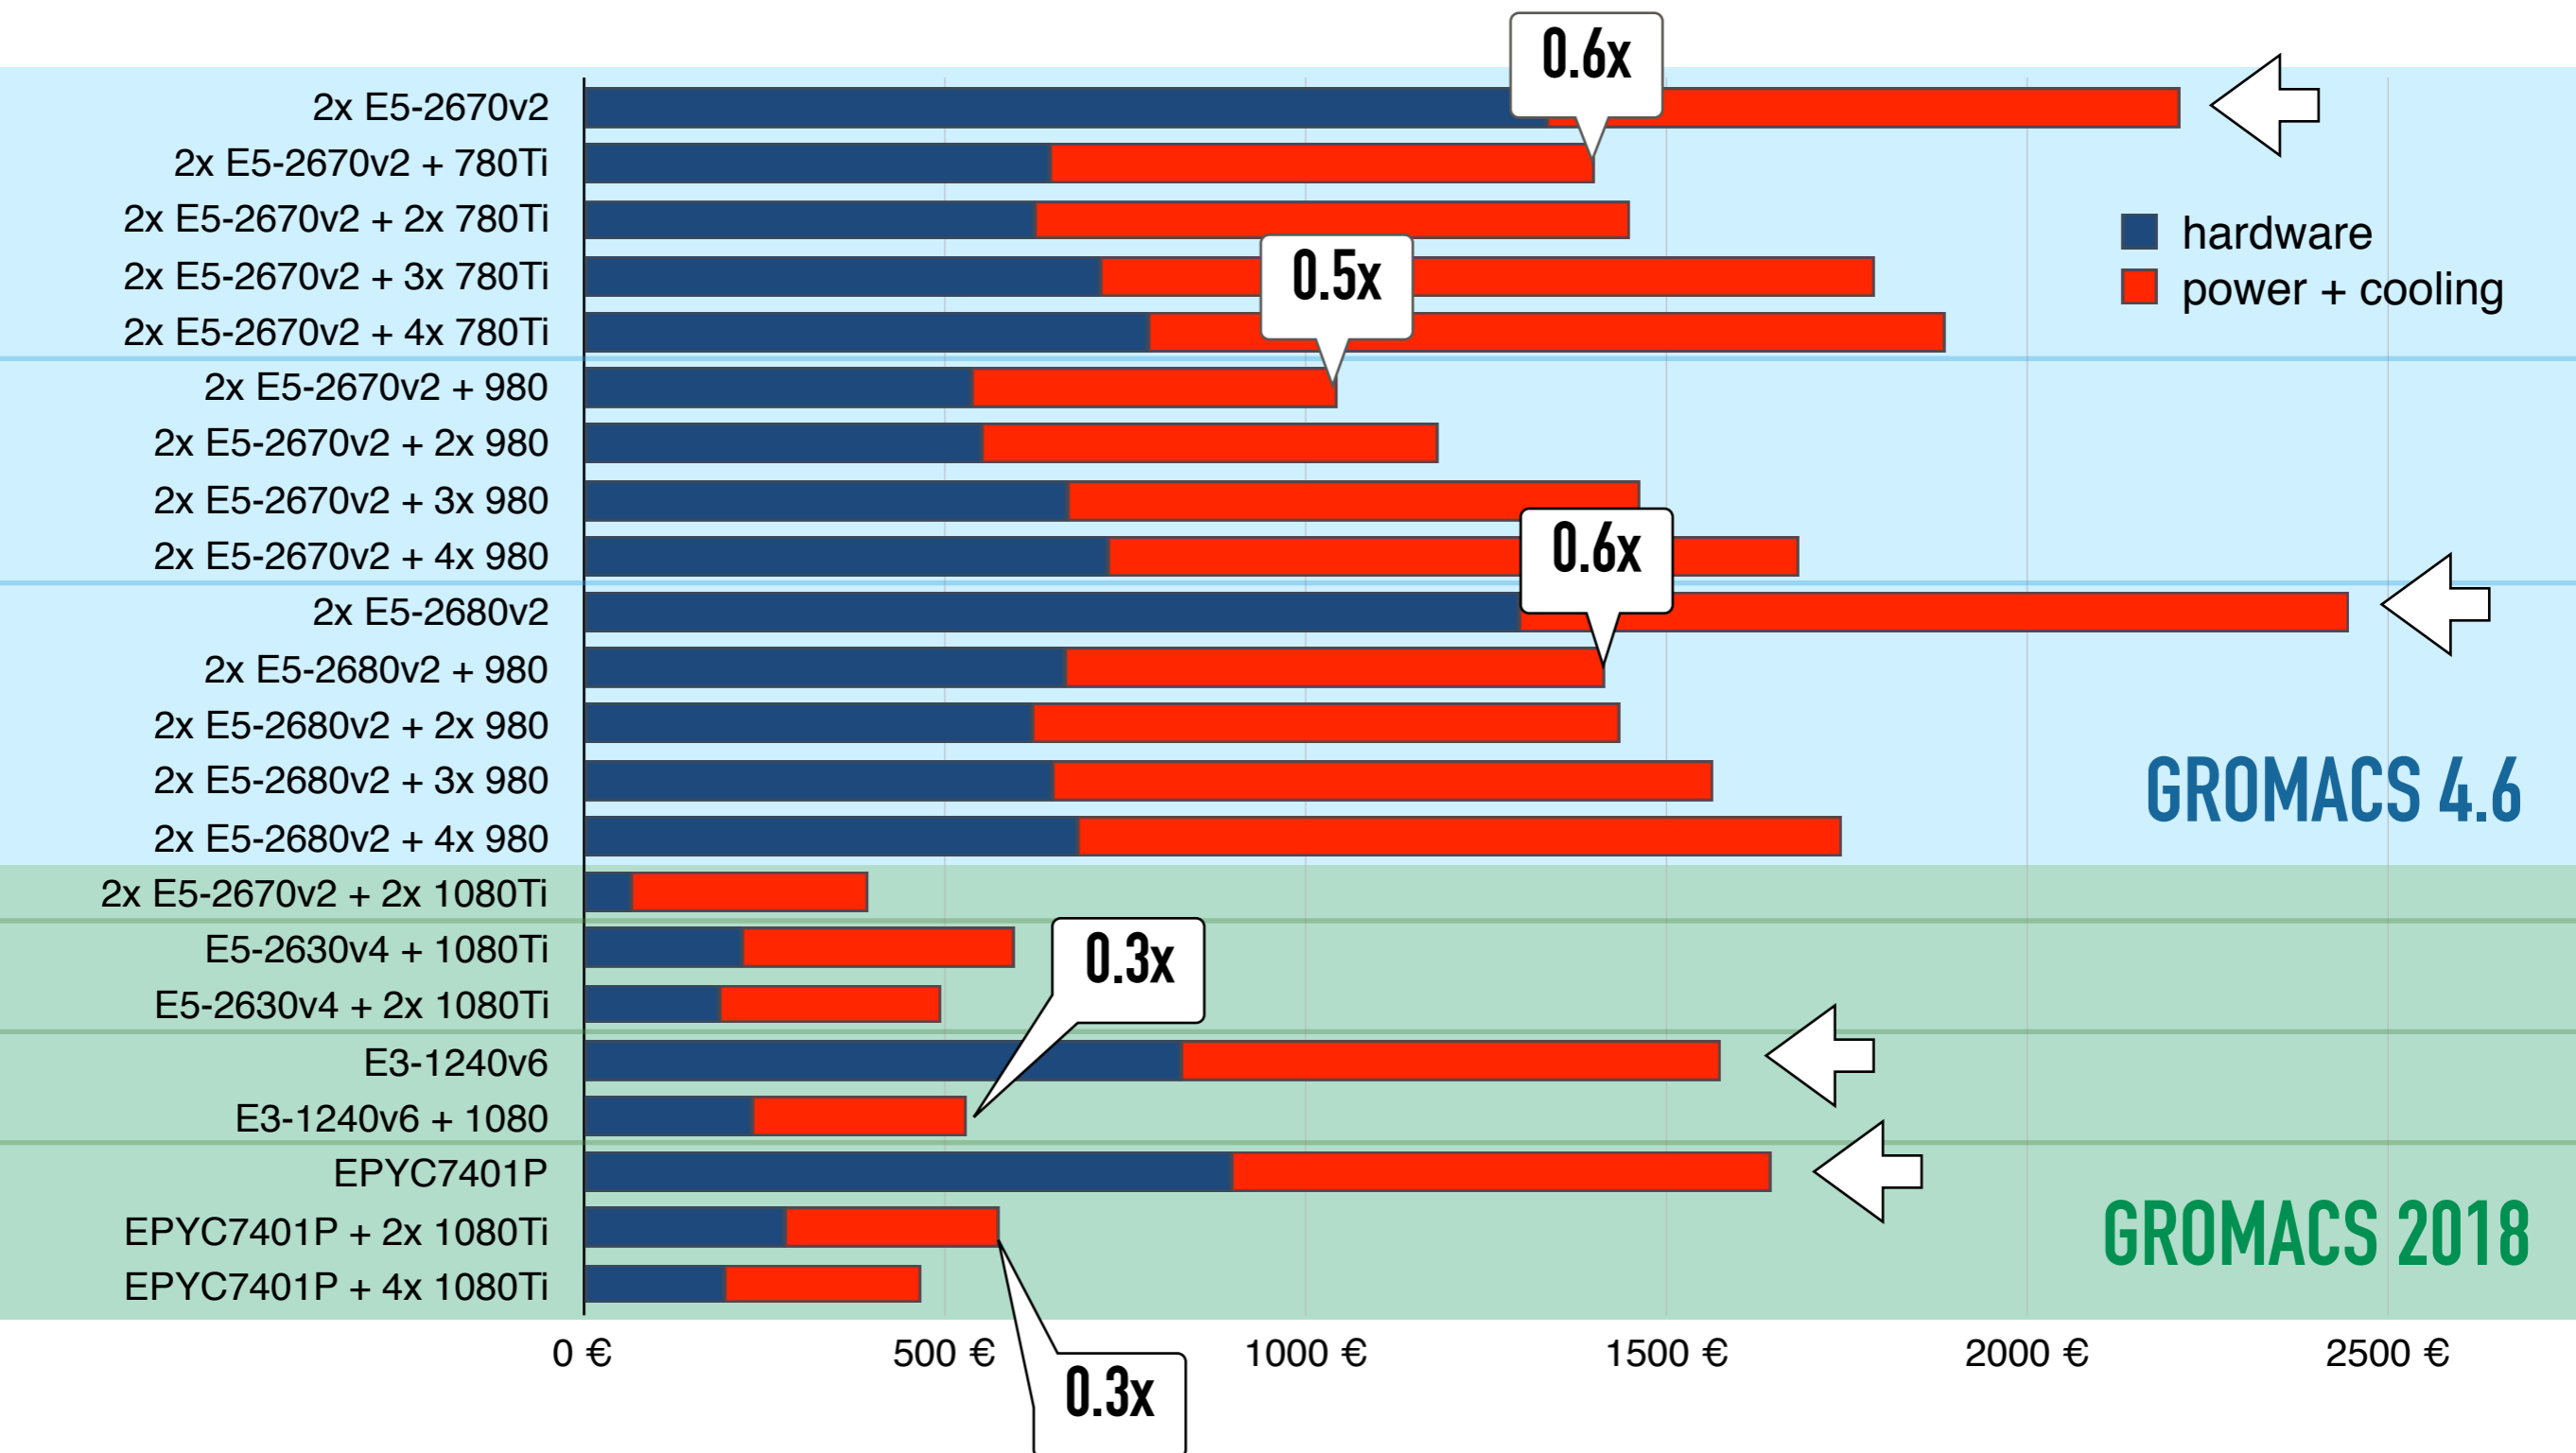
ADD ENERGY COSTS TO THE BILL

node costs taking into account energy + cooling (0.2 EUR / kWh) RIB



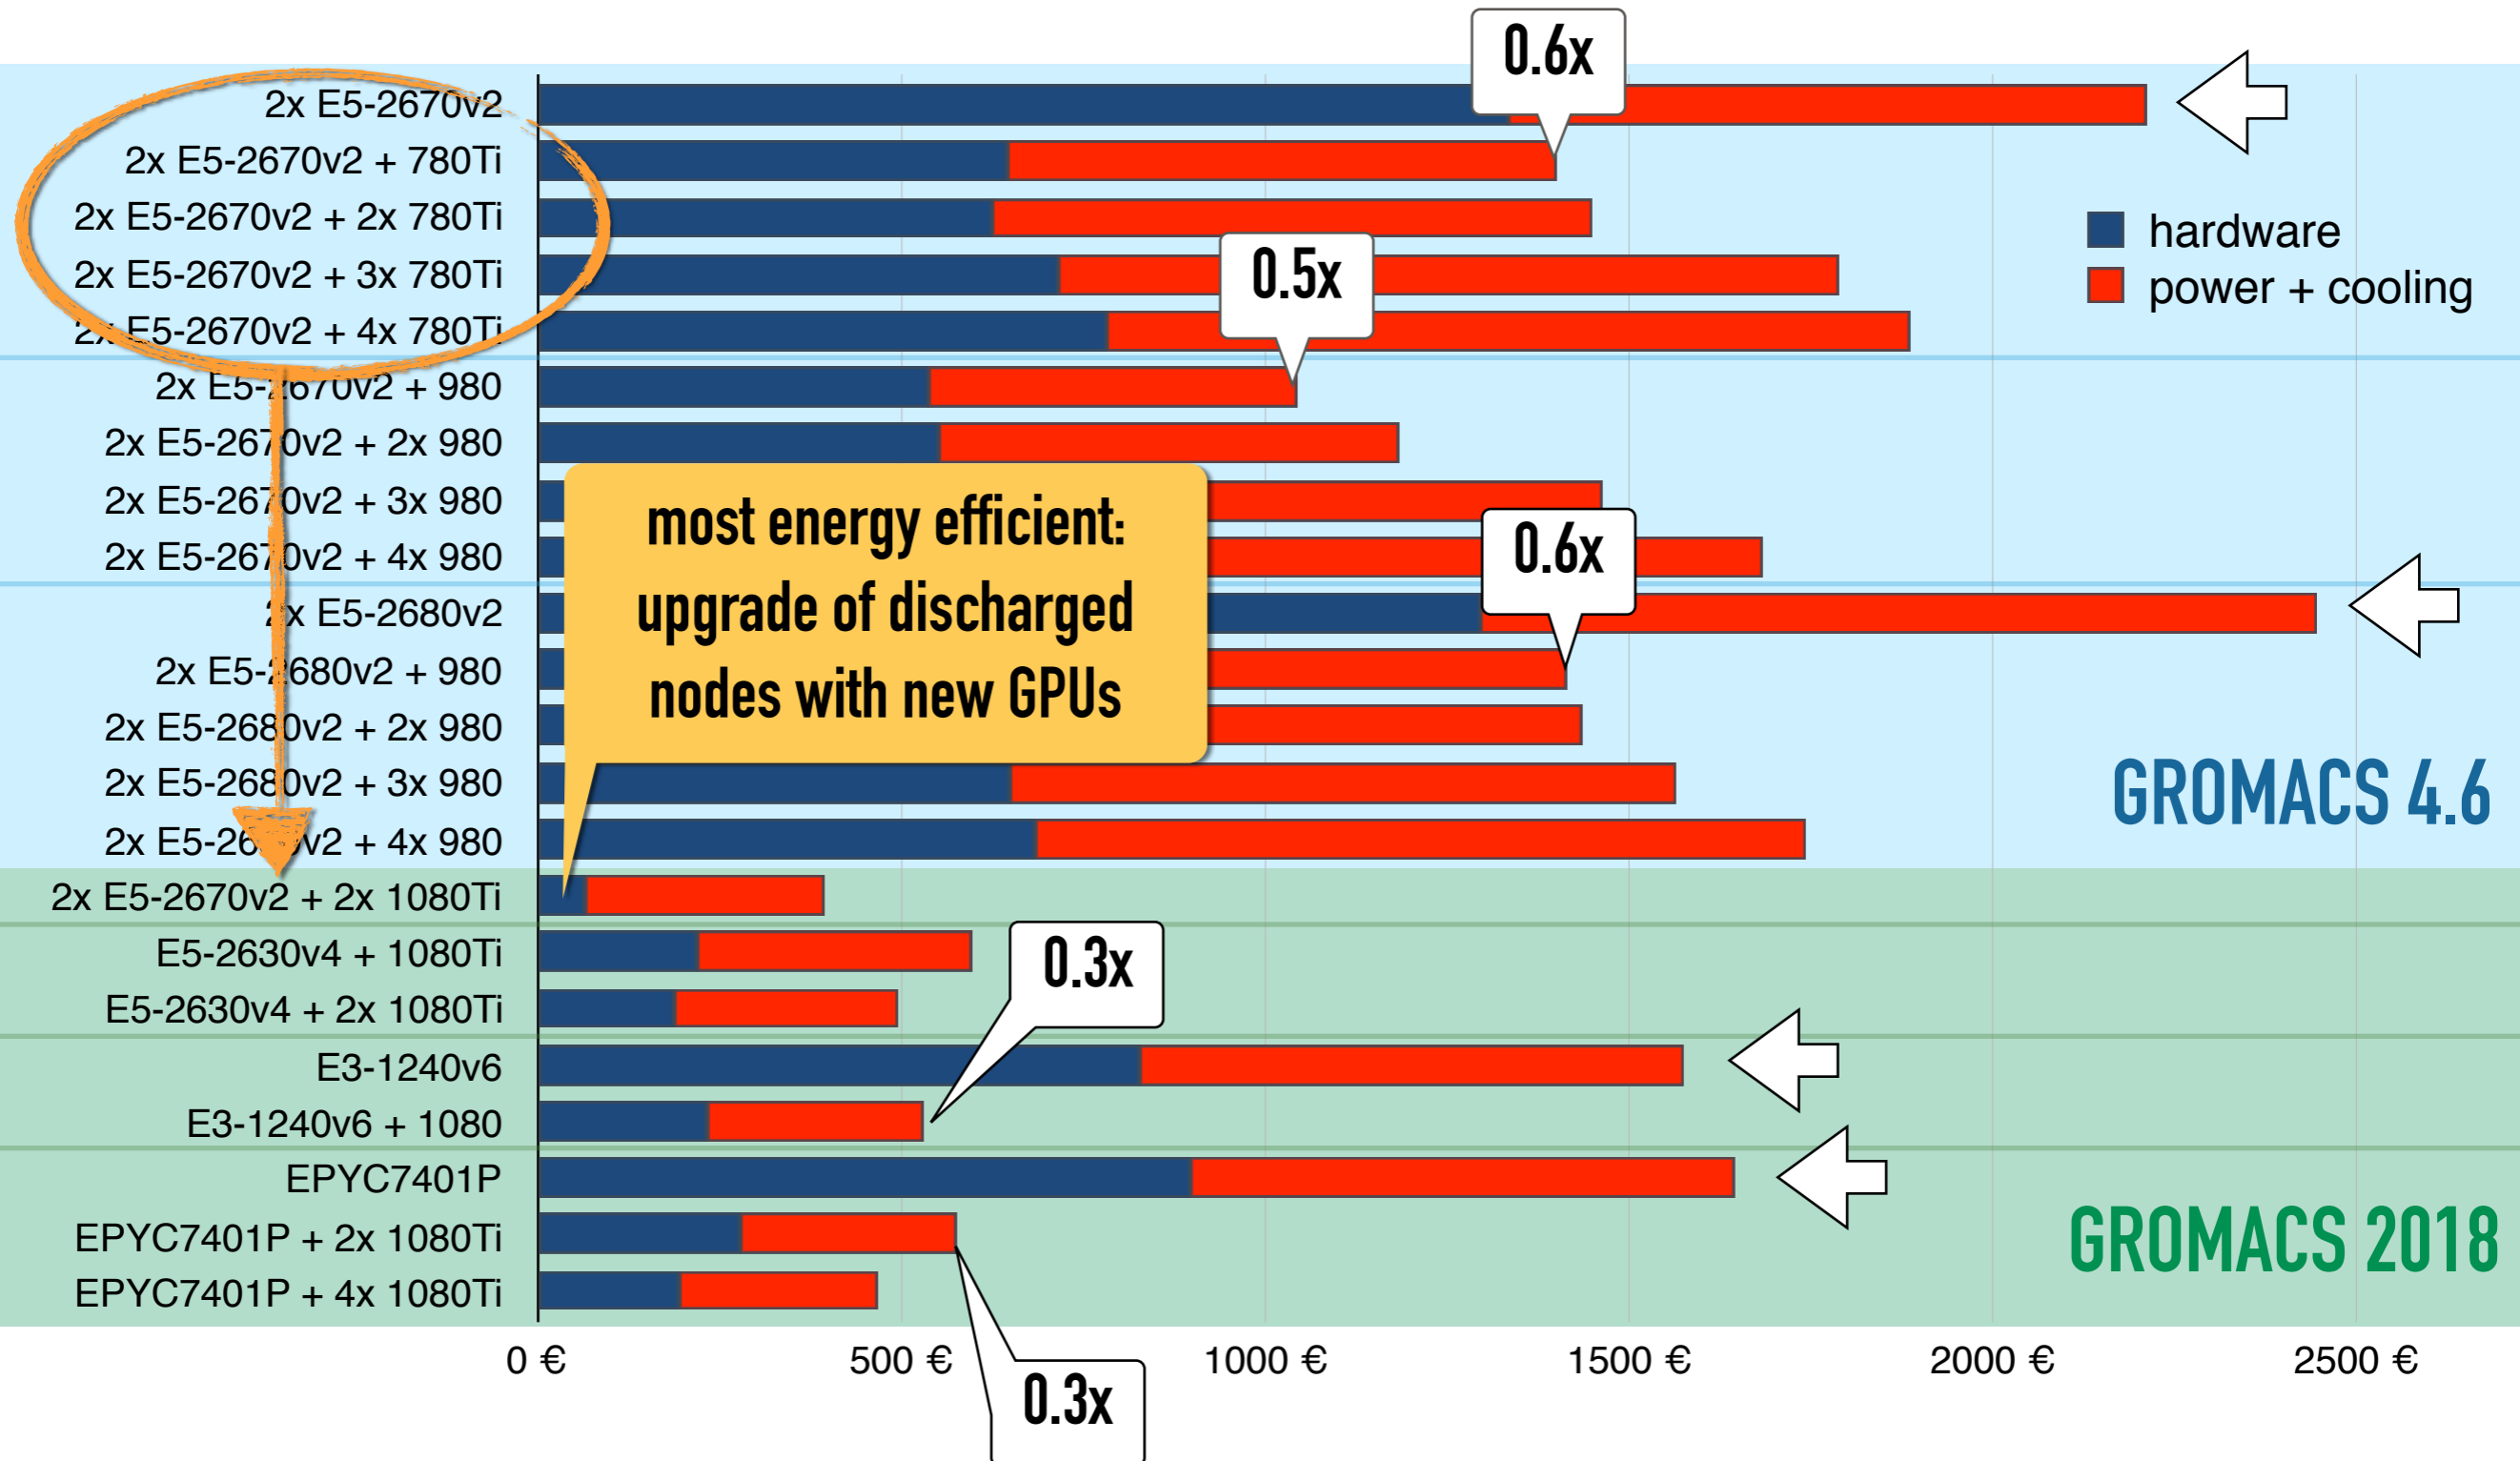
ENERGY EFFICIENCY

trajectory costs per microsecond RIB for 5 years of operation



ENERGY EFFICIENCY

trajectory costs per microsecond RIB for 5 years of operation




CONCLUSIONS

- ◆ compared to CPU nodes or nodes with Tesla GPUs, nodes with consumer GPUs yield significantly higher trajectory output per invested Euro
 - ◆ taking into account raw node price: **2–3 x for GROMACS 4.6**, and **3–6 x for GROMACS 2018**
 - ◆ including energy costs: about **2 x for GROMACS 4.6**, and **3 x for GROMACS 2018**
- ◆ PME on GPU...
 - ◆ ...moves the optimal hardware balance even more towards the GPU side (4–8 cores / 1080Ti)
 - ◆ ...allows to cheaply upgrade old nodes with state-of-the-art GPUs!



OUTLOOK

- ◆ upgrade old E5-2670v2 nodes (2x 10 cores @ 2.8 GHz) with 2 or 4 GTX 1080Ti.
- ◆ benchmark configurations with AMD GPUs (e.g. VEGA64)
- ◆ want to compare your own hardware and contribute to benchmarking?
<https://www.mpibpc.mpg.de/grubmueller/bench> has various benchmark .tprs for download (CC licensed, also FE benchmarks)
- ◆ keep an updated list of benchmark results for the most recent GROMACS version



ACKNOWLEDGEMENTS



- ◆ the Grubmüller department
 - ◆ Martin Fechner, Ansgar Esztermann
- ◆ Szilard Pall (KTH)
- ◆ Markus Rampp (MPCDF)