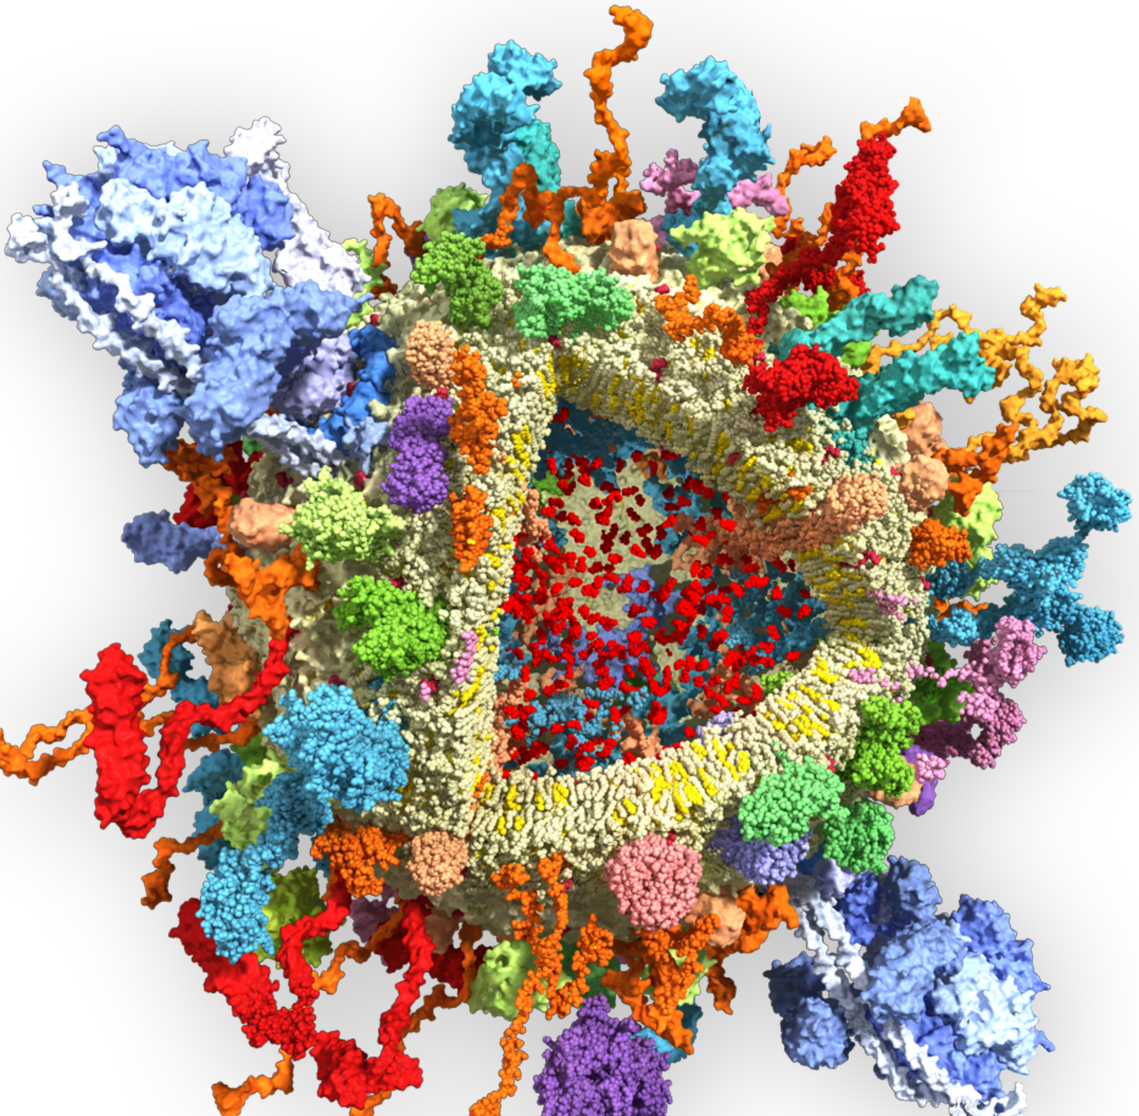

Large Scale Biomolecular Simulations in the Cloud

Carsten Kutzner
Theoretical and Computational
Biophysics Department

MAX PLANCK INSTITUTE
FOR MULTIDISCIPLINARY SCIENCES

Theoretical and Computational Biophysics Department

Dept. head: Helmut Grubmüller

- Our mission:

> Biomolecules such as proteins are the **nanomachines** of life, doing the work in our cells. We want to understand how they work on the basis of **physics**.

- our tool: molecular dynamics (MD) simulation
- proteins, membranes, membrane channels, virus shells, ribosomes, small molecules for computational drug design…
- … using the GROMACS MD package

Molecular Dynamics (MD) Simulation

- Solve Newton's eq. of motion for N atoms
  (N = 10,000, ... 10,000,000)
- Positions $\mathbf{r}_i$, charges q_i, masses m_i, …
- Calculate forces from a potential
  $U(\mathbf{r}_i, q_i, m_i, ...)$

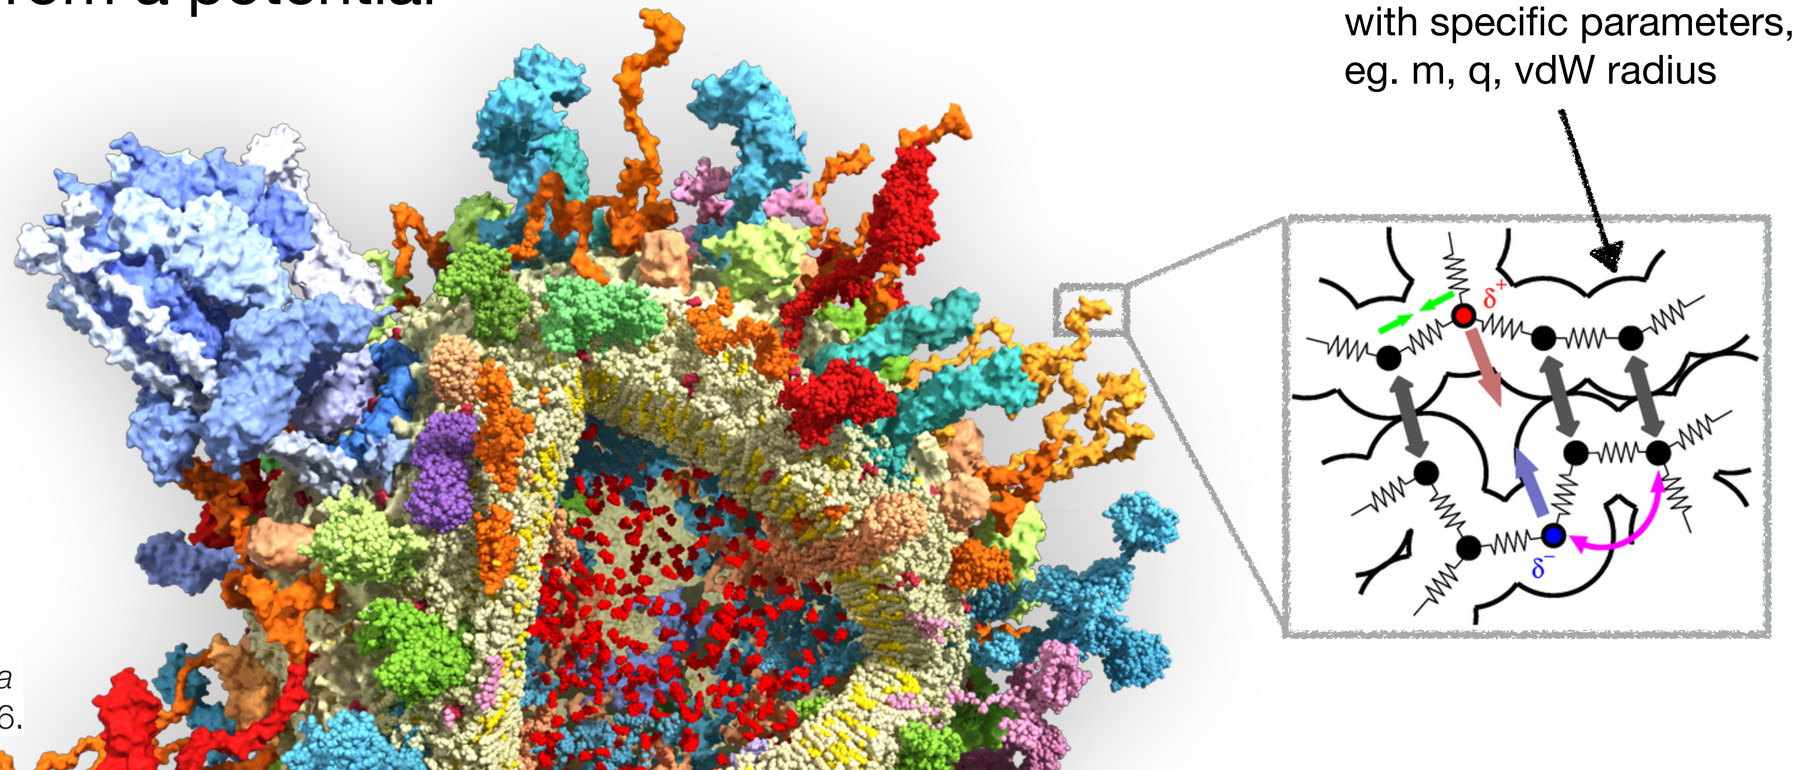

Takamori, Shigeo, et al. *Molecular anatomy of a trafficking organelle.* Cell 127.4 (2006): 831-846.

Challenges of Biomolecular Simulation

- For biologically relevant time scales ➔ μs long trajectories
  ➔ millions of time steps for a single trajectory!
- We are not interested in a single trajectory, but in the average behaviour
- Run ensemble of many slightly different copies for robust statistics

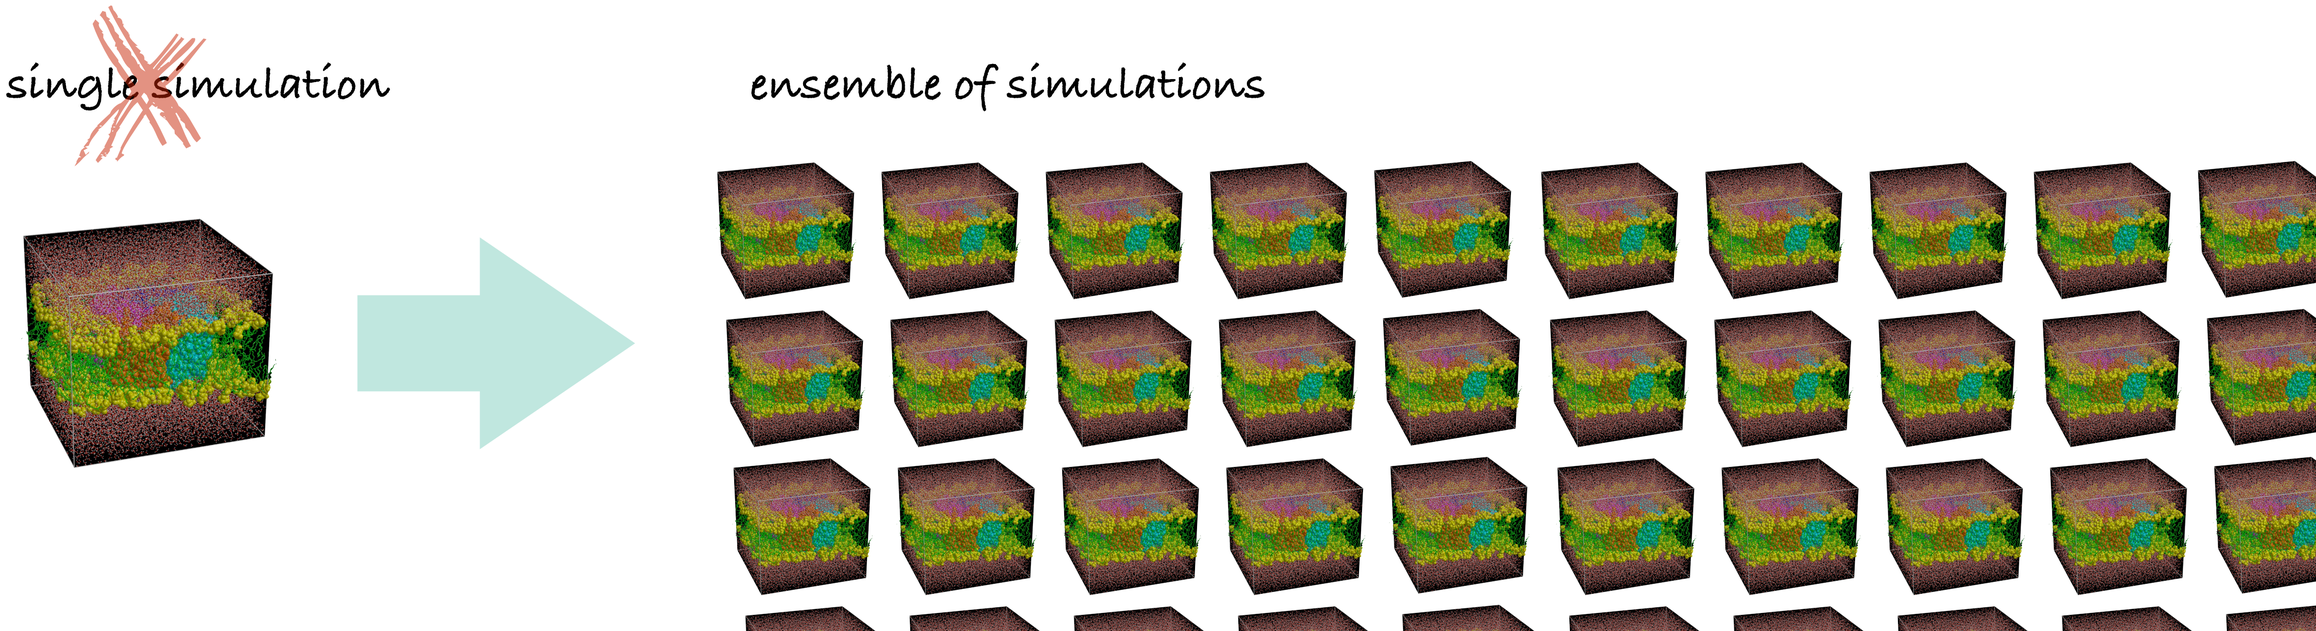

Use Cases Range From HPC to HTC

e.g. the ribosome

HPC
high performance computing

MD simulations

HTC
high throughput computing

e.g. computational drug design

large system: minimize time-to-solution

many small systems: maximize throughput

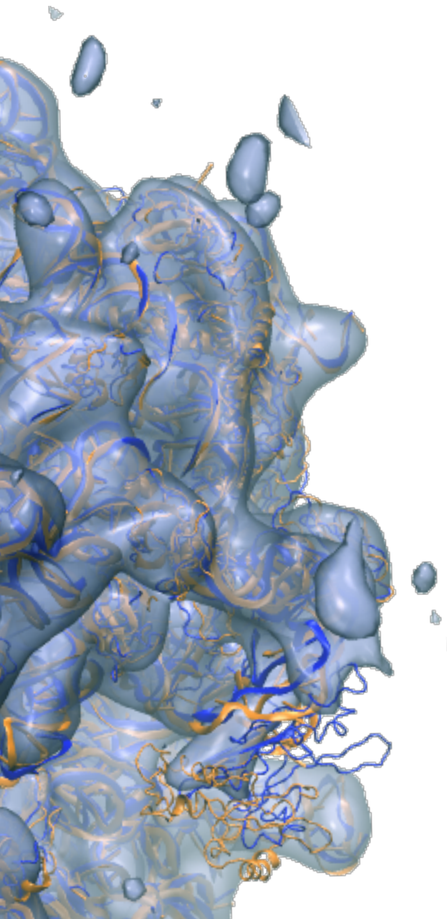

- Compute scenario depends on the questions addressed
- In any case, simulations take weeks to months!
- Where can we run our extensive simulations?

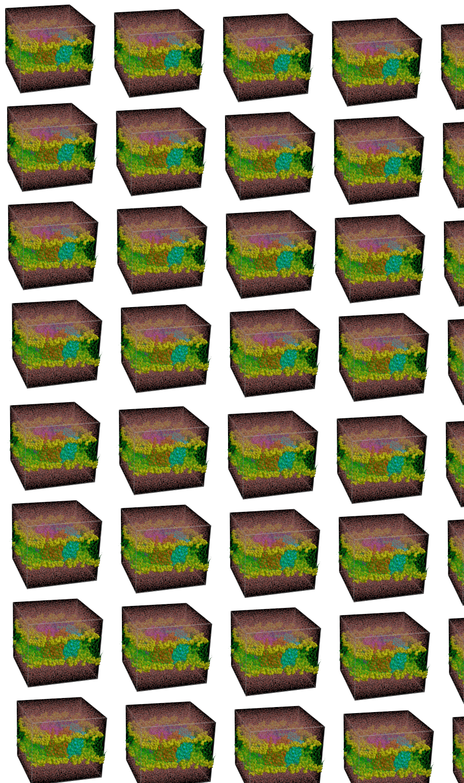

Providers of Computing Resources

National supercomputing centers
LRZ, JSC, HLRS, HLRN, …

Department-level clusters

Institute-wide clusters

University clusters

MPCDF HPC cluster

Folding@Home

What about scientific computing in the Cloud?

Google Cloud Azure Amazon Web Services (AWS) …

Is Cloud Computing Competitive for our Simulations?

Collaborate with Amazon Web Services (AWS) since 09-2020

1. **Can we do HPC in the Cloud?**
   Strong scaling of GROMACS on a cloudy cluster

2. **How expensive is it?**
   Cost-comparison to our in-house cluster tailored towards GROMACS

3. **HTC in the Cloud: How far can we go?**
   Can the cloud outperform an on-premises cluster in terms of throughput?

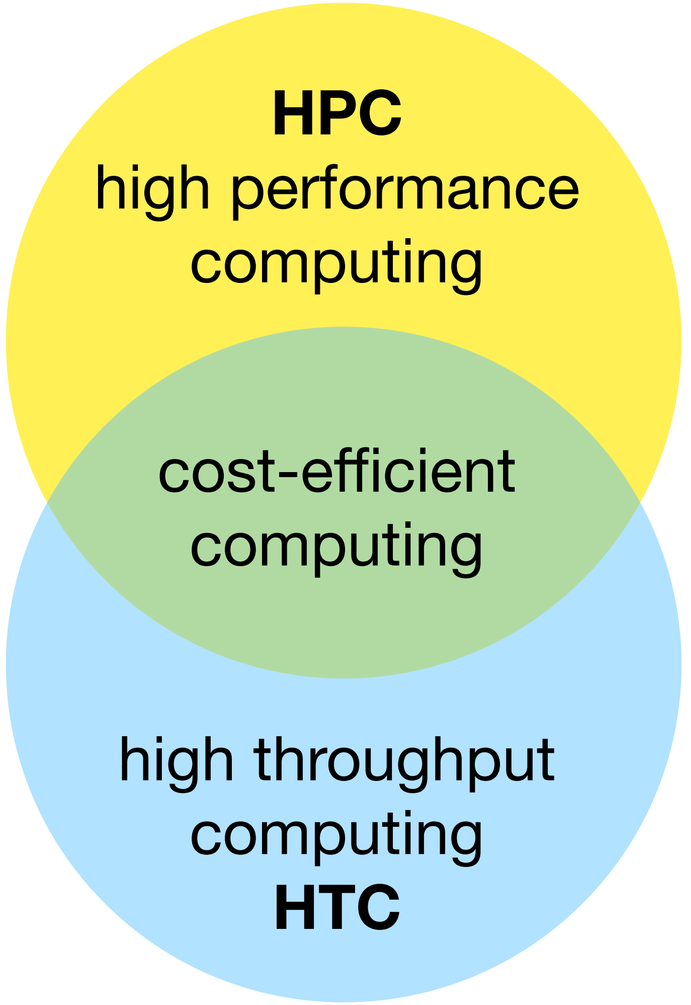

Cloud Computing

- The smallest unit of computing power is an ***instance***, i.e. a virtual machine that runs on part of a node
- AWS offers a zoo of instance types,
  1-192 vCPUs, Intel, AMD, Arm, GPUs, interconnect
- Various **pricing** options, prices and availability vary across regions
- instances come with base OS only, **no preinstalled software**

g5.2xlarge instance

- 4 cores (8 vCPUs)
- 1 NVIDIA A10G GPU
- 32 GB RAM
- 450 GB SSD
- 10 Gbps network

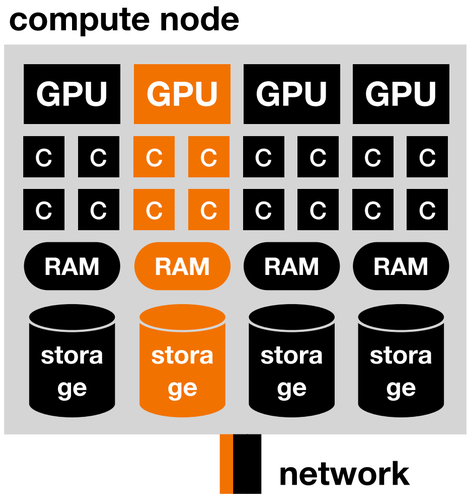

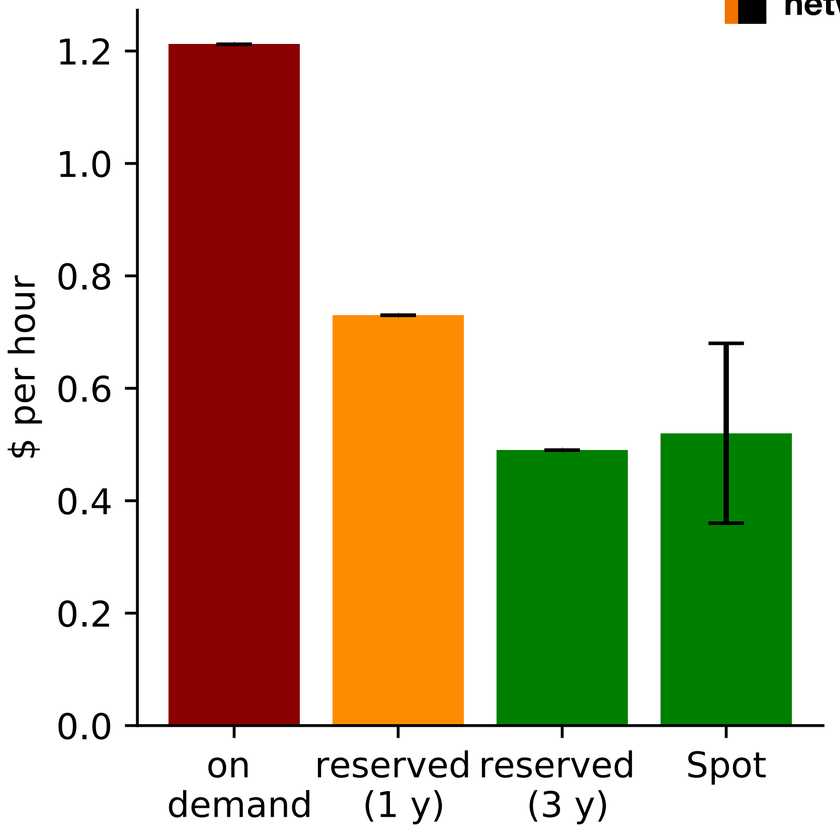

1 Cloud-based HPC
with GROMACS on AWS

Cloud-based HPC

AWS ParallelCluster: A Cloud-Based HPC Cluster

- Set up cluster with **ParallelCluster**
- open source free cluster management tool
- github.com/aws/aws-parallelcluster

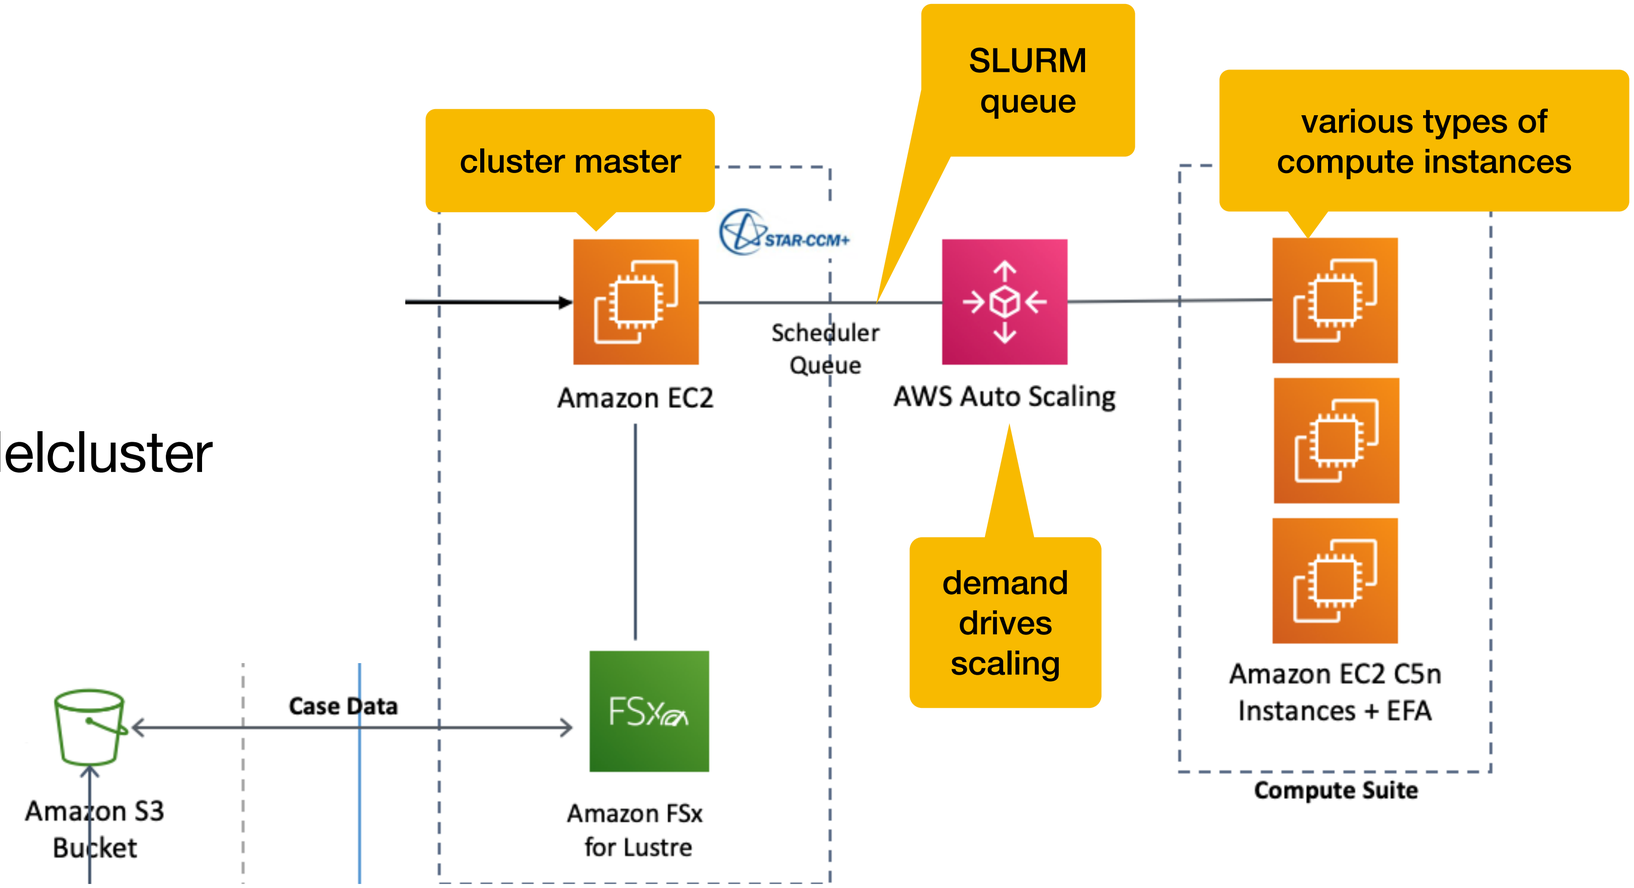

Cloud-based HPC

Benchmarking GROMACS on the Cloud-based Cluster

MEM
81 k atoms
Membrane channel

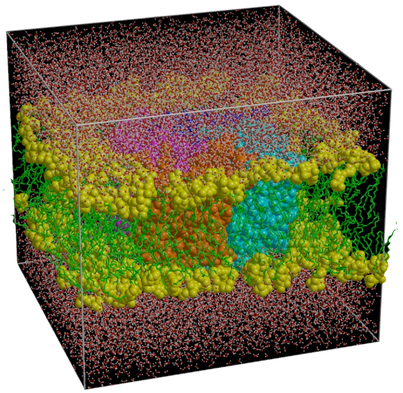

RIB
2.1 M atoms
Ribosome

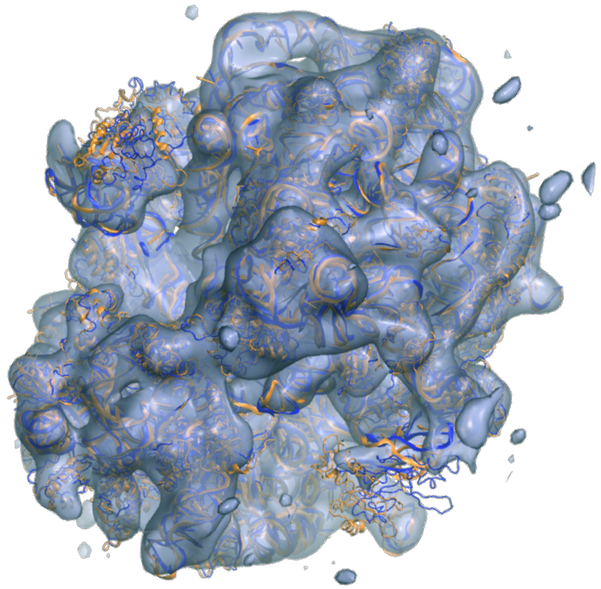

PEP
12.5 M atoms
Peptides

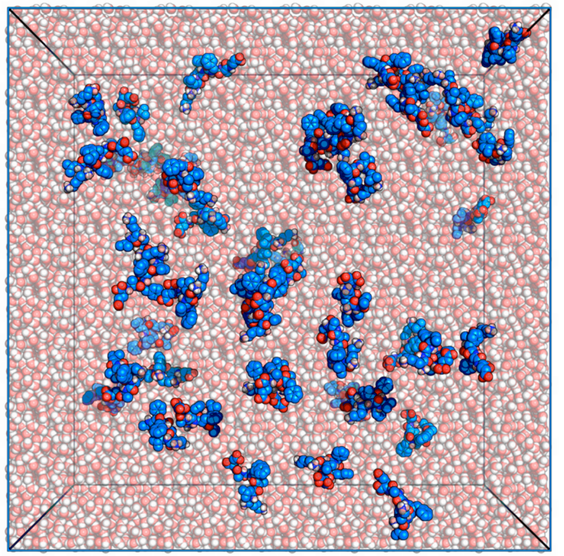

www.mpinat.mpg.de/grubmueller/bench

Cloud-based HPC
with GROMACS on AWS

- **g5:** 4–32 vCPUs AMD EPYC + NVIDIA A10G GPUs
- **hpc6a:** 96 vCPUs AMD EPYC + fast EFA network
- **g4dn:** 4–64 vCPUs Intel + NVIDIA T4 GPU
- **g3:** 4–16 vCPUs Intel + NVIDIA M60 GPU
- **c6i:** 128 vCPUs Intel + fast EFA network
- **c6g:** 4–64 vCPUs ARM Graviton2
- **c5:** 2–96 vCPUs Intel or AMD + optional EFA network
- **p3:** 8–96 vCPUs + 1–8 NVIDIA V100 GPUs
- **p4d:** 96 vCPUs + 8 NVIDIA A100 GPUs

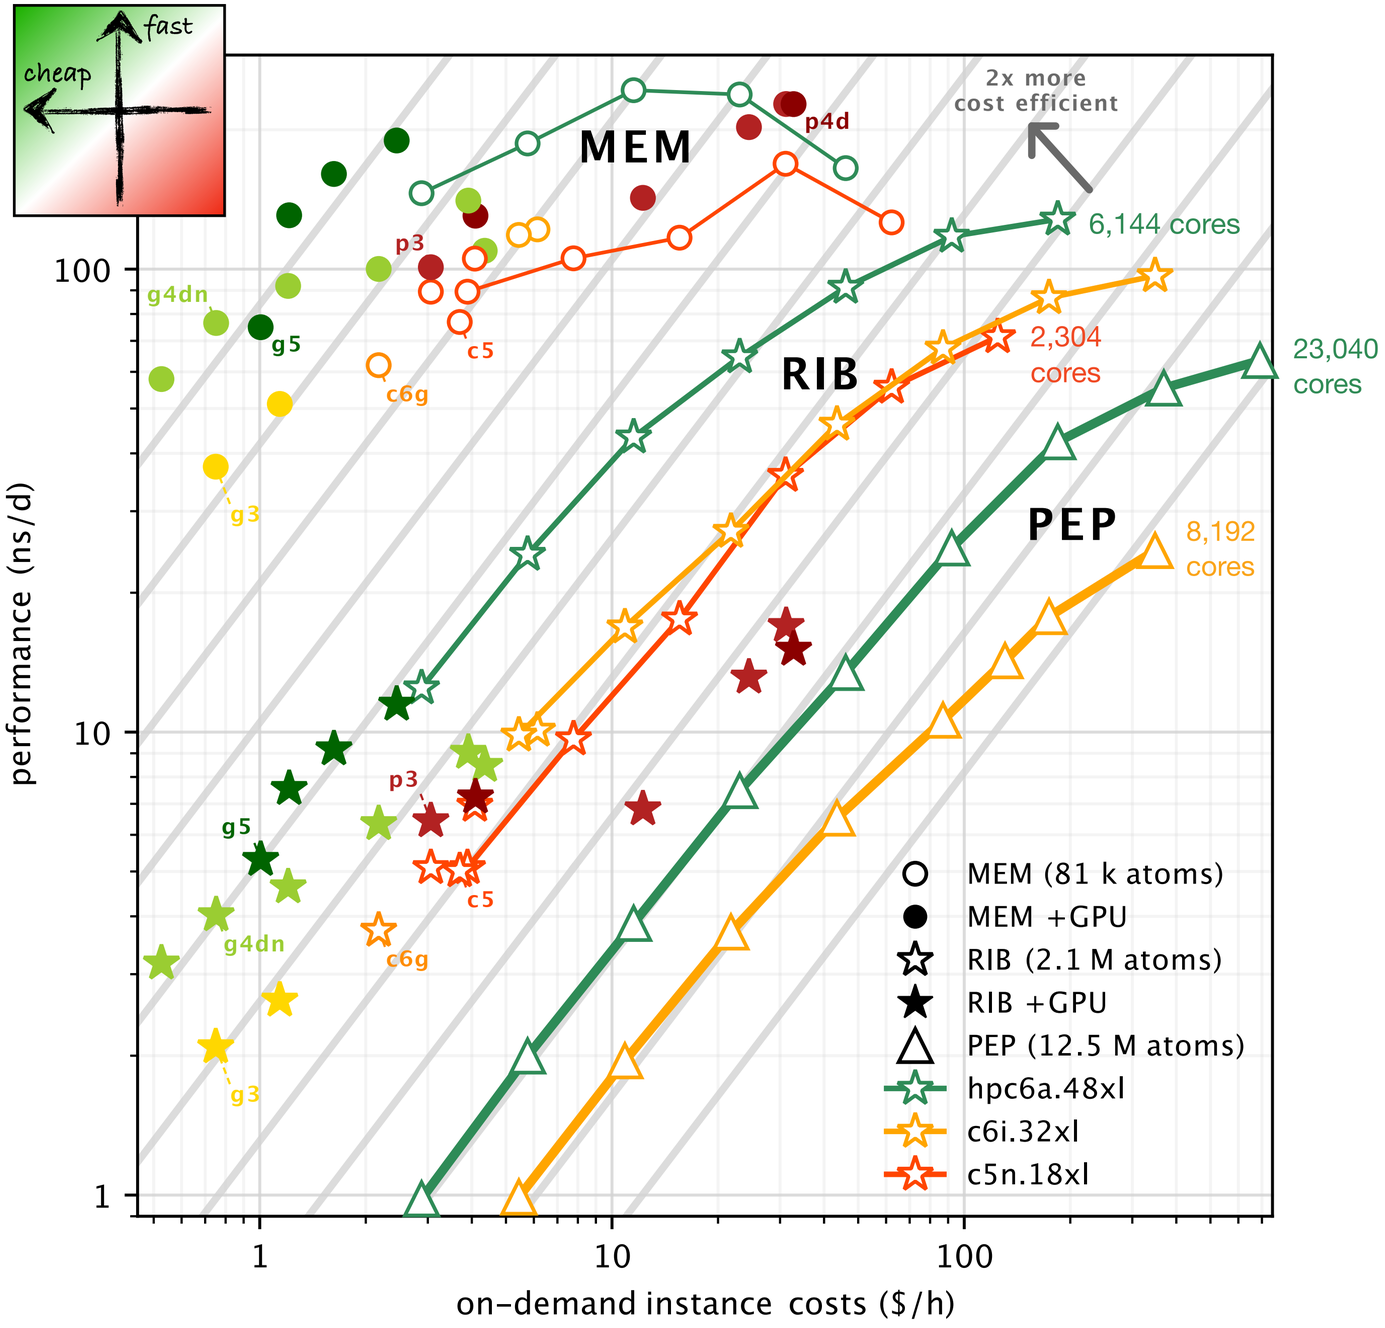

Cloud-based HPC
with GROMACS on AWS

Results I:

- **strong scaling** on up to 23,040 cores
- **cost-efficiency** depends critically on selected instance type
- Optimal for GROMACS:
  - single instances: **g5** and **g4dn**
  - scaling: **hpc6a**

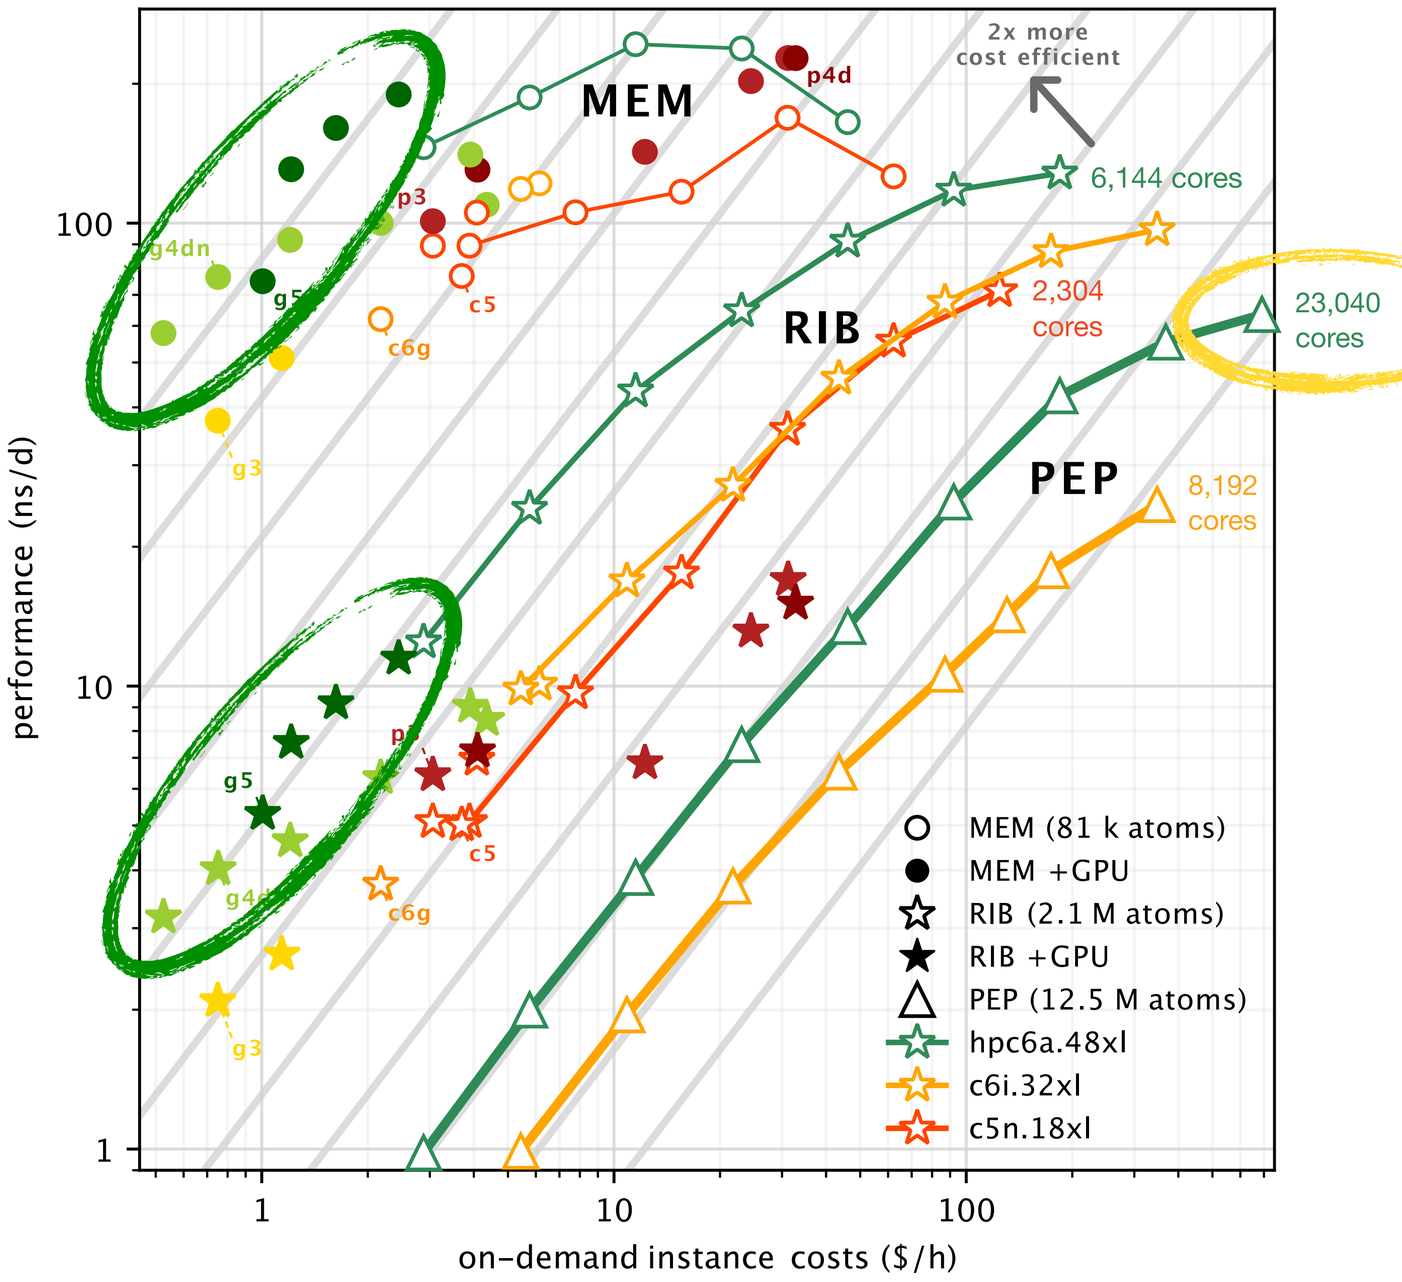

2 Isn’t a local cluster cheaper?

Cost-comparison for biomolecular simulation

Costs to Produce a µs of RIB Trajectory

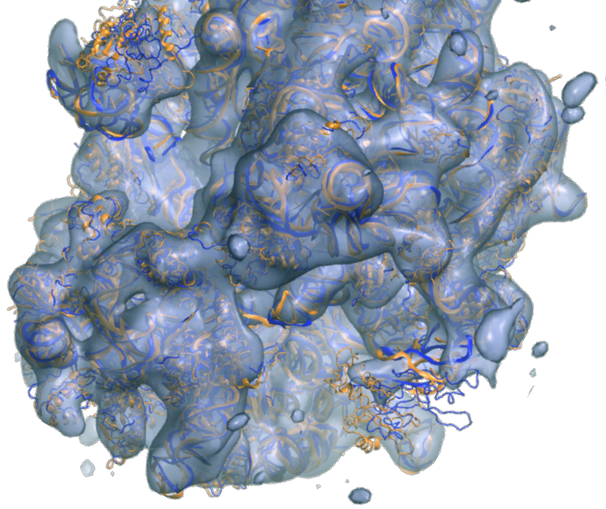

- Compute costs for operating a 500 node on-premises cluster over 3 years

- Nodes tuned for throughput and cost-efficiency with GROMACS!
  - consumer GPUs (e.g. GeForce)
  - reasonably priced CPUs
  - no HPC interconnect
  - minimum RAM
  - dense packing (~1 GPU per U)

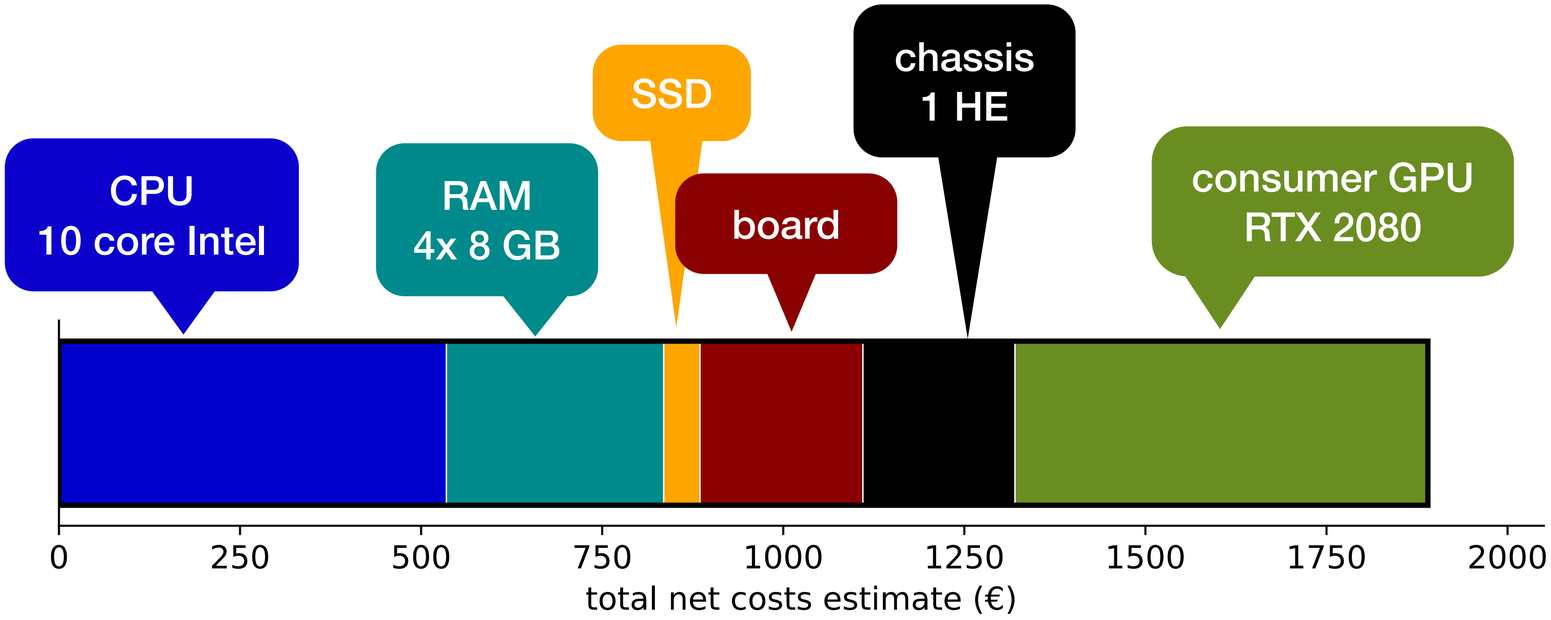

3-year warranty included

Cluster Node Tuned for GROMACS

Fixed Costs + 1-Year Variable Costs Estimate

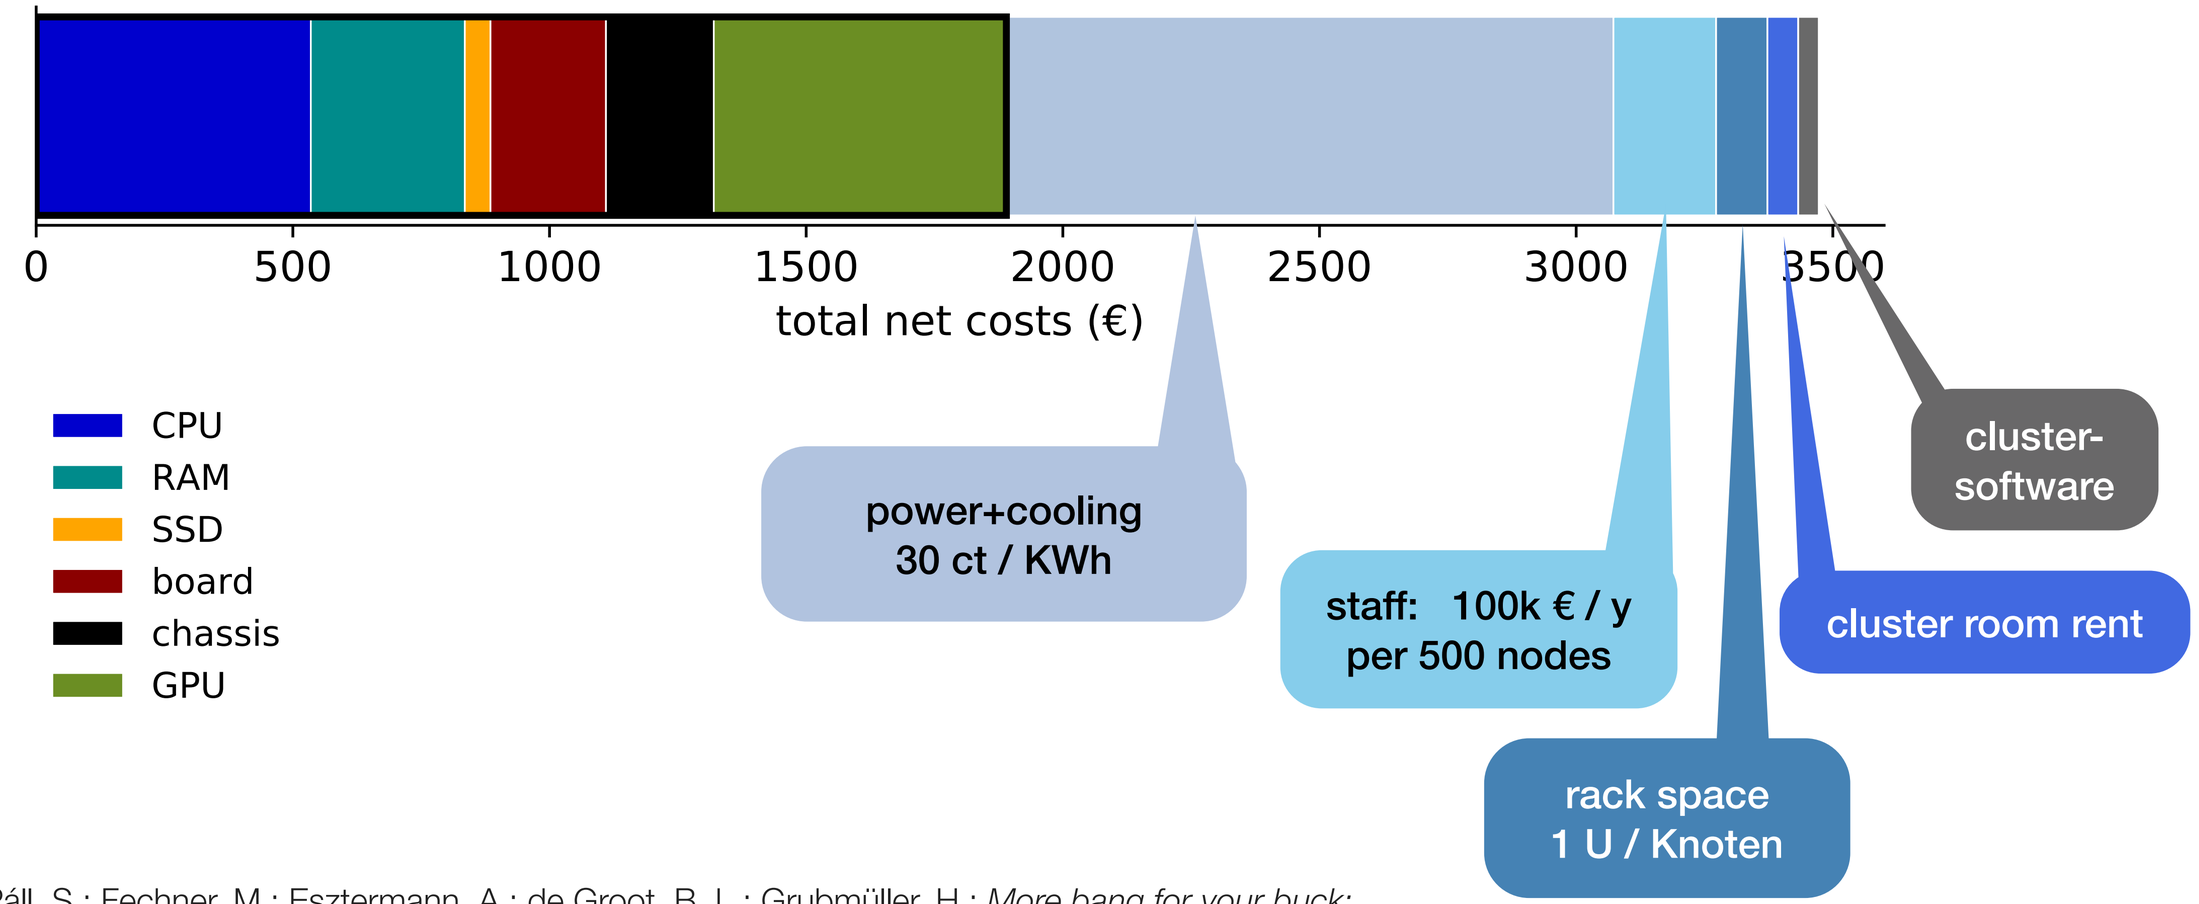

Kutzner, C.; Páll, S.; Fechner, M.; Esztermann, A.; de Groot, B. L.; Grubmüller, H.: *More bang for your buck: Improved use of GPU nodes for GROMACS 2018*. J. Comp. Chem. 40 (27), pp. 2418 - 2431 (2019)

Total Costs For 3 Years of Operation

Cloud vs. Own Cluster

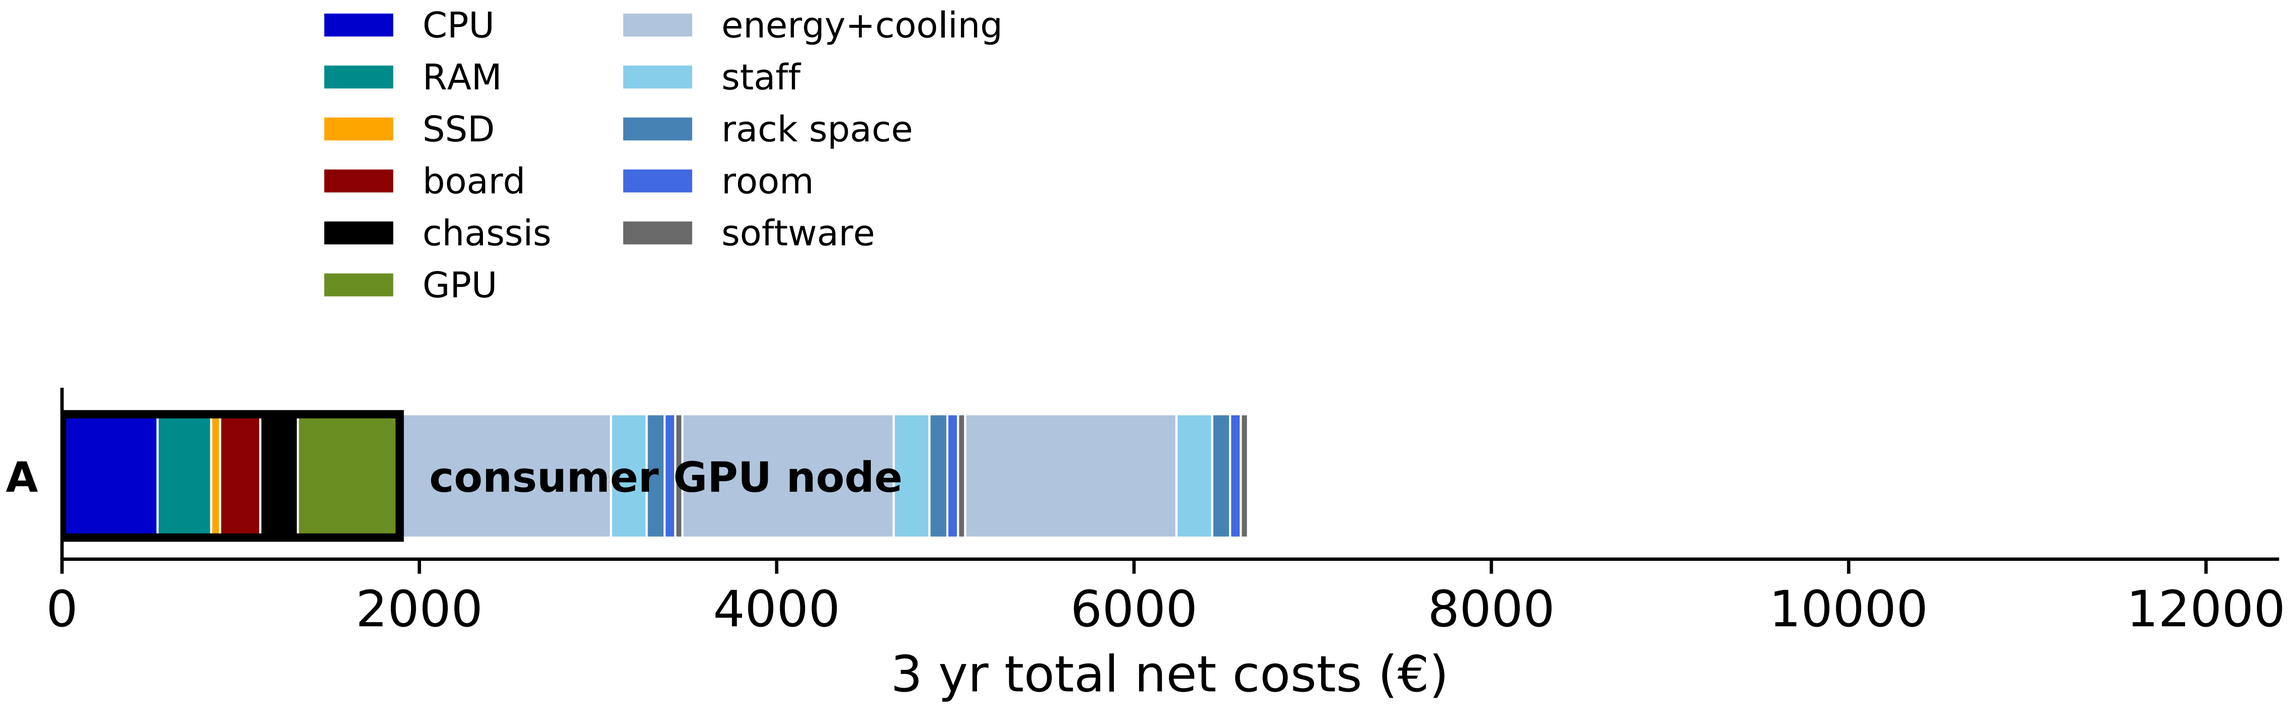

Total Costs For 3 Years of Operation

Cloud vs. Own Cluster

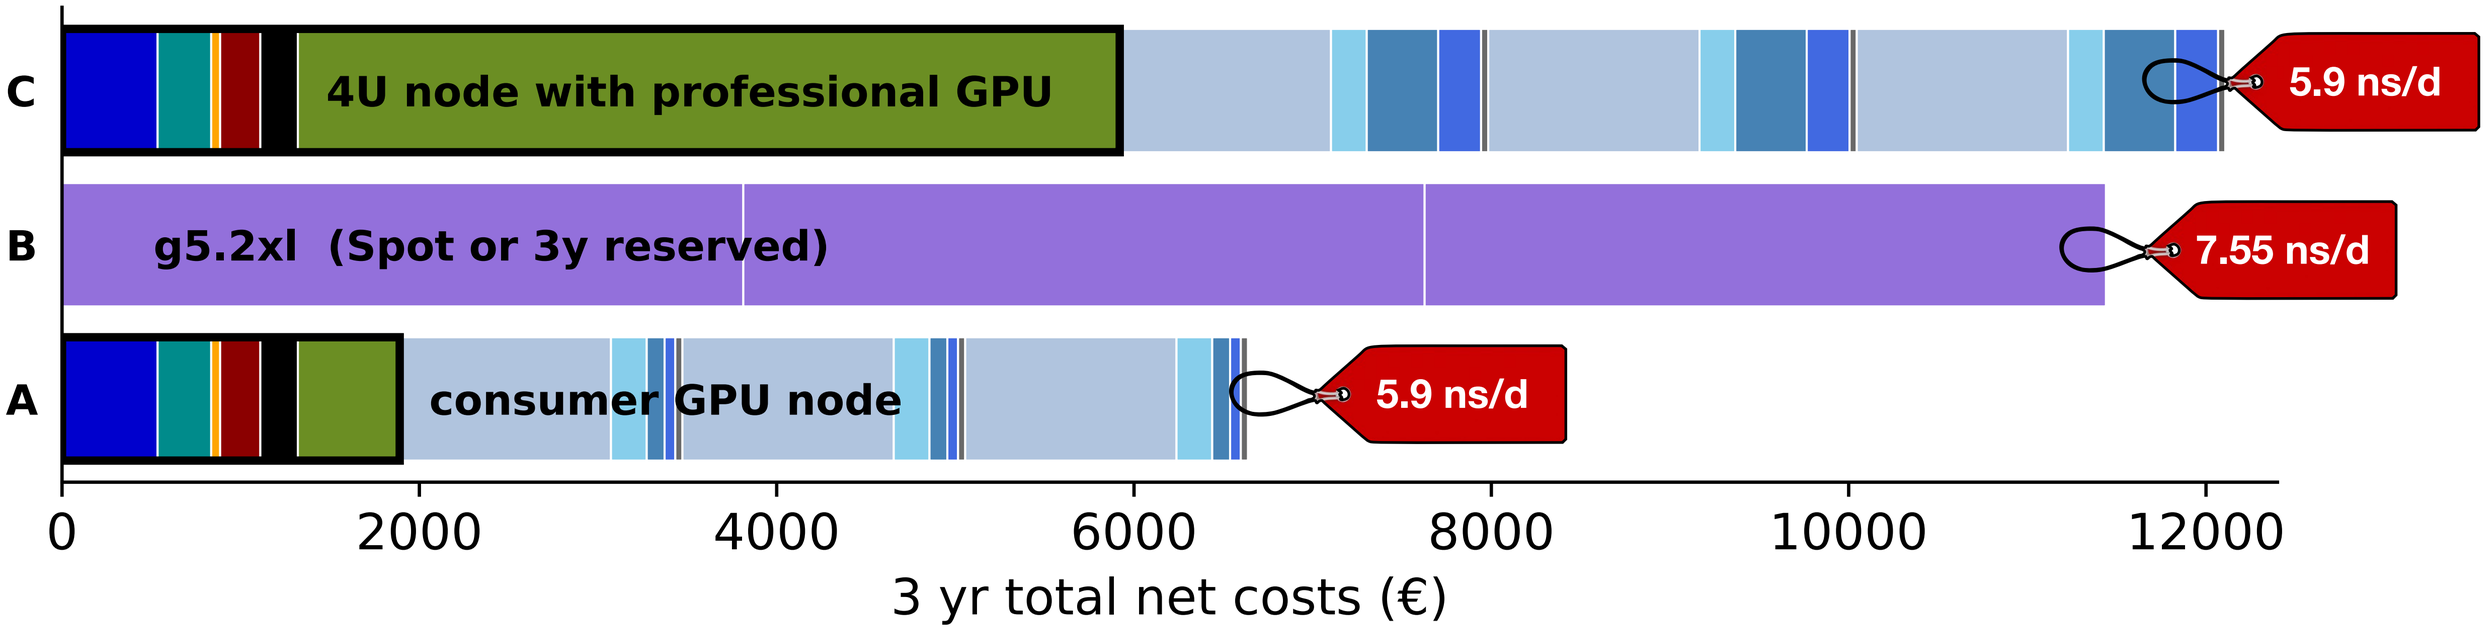

Isn’t a local cluster cheaper?

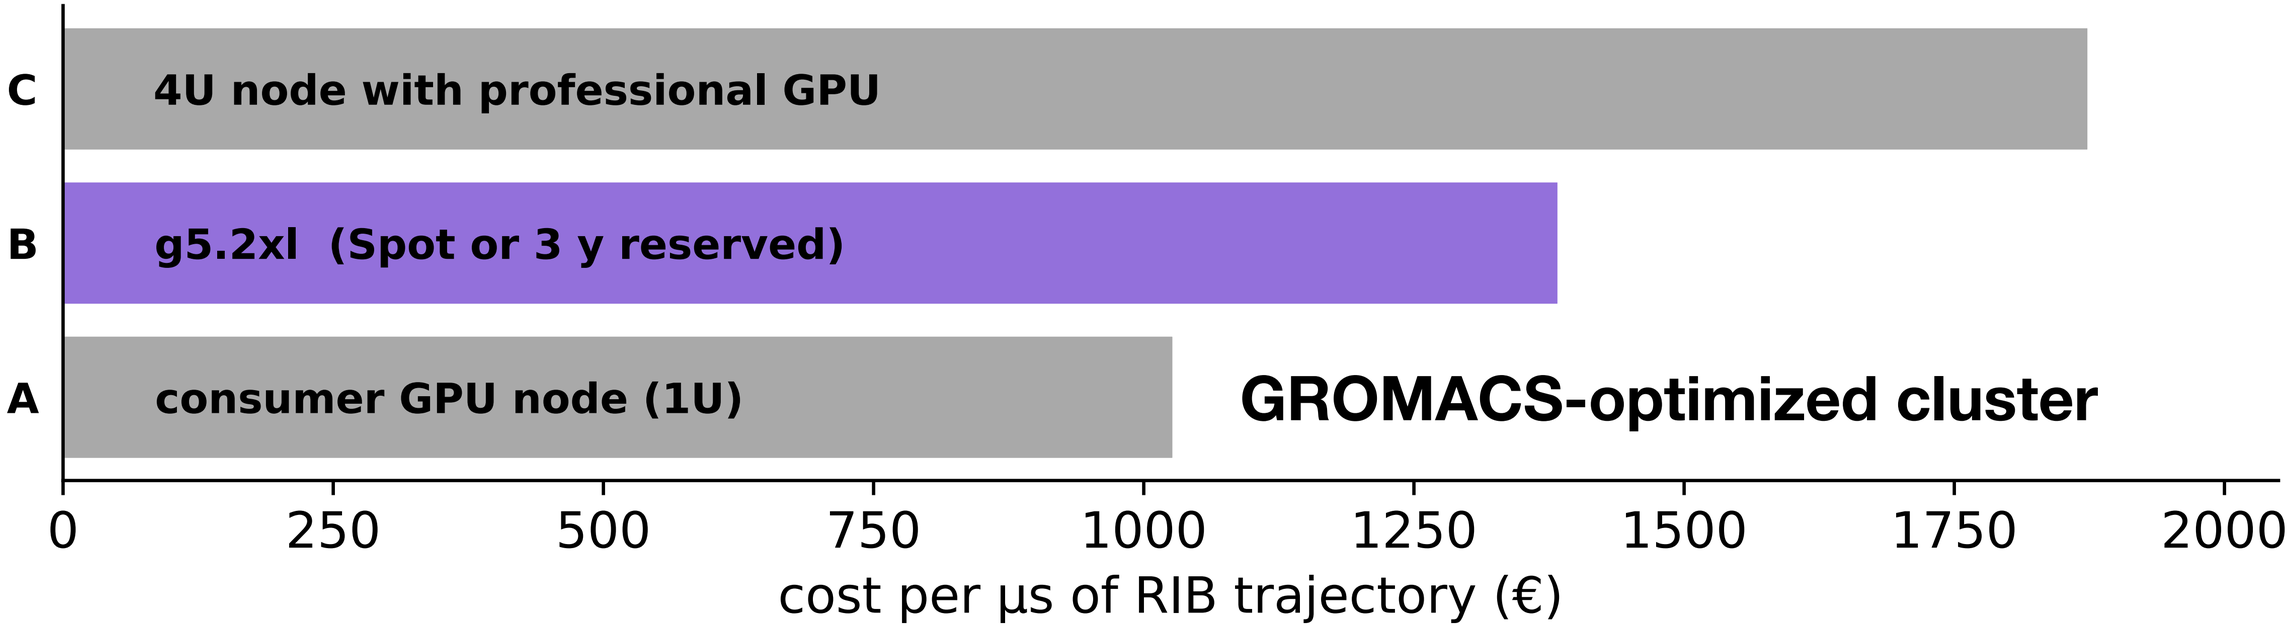

Results II:

- Cloud is also cost effective if appropriate instances are chosen
  (For GROMACS: g5.2xl and similar Spot GPU instances)

3 Cloud-based HTC

How much throughput can we get?

Cloud-based HTC

Global compute resources accelerate computational drug design

- **Aim:** Minimize time-to-solution for an ensemble of 20,000 jobs (5,000–100,000 atoms)
- **Approach:** Run all jobs at the same time, each on a separate Spot instance, wherever there is capacity
- **Solution: AWS Cyclone**-based setup
  - distributes jobs to AWS Batch queues around the globe
  - open source, community-supported https://github.com/awslabs/aws-cyclone-solution

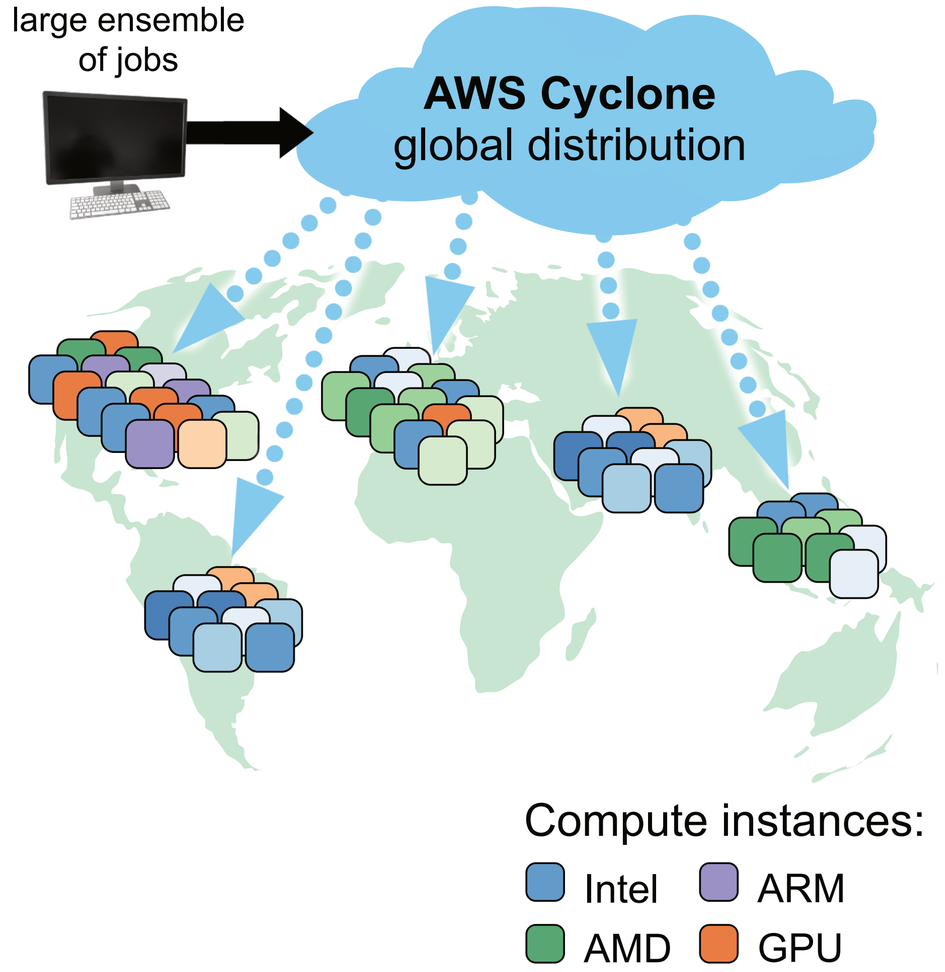

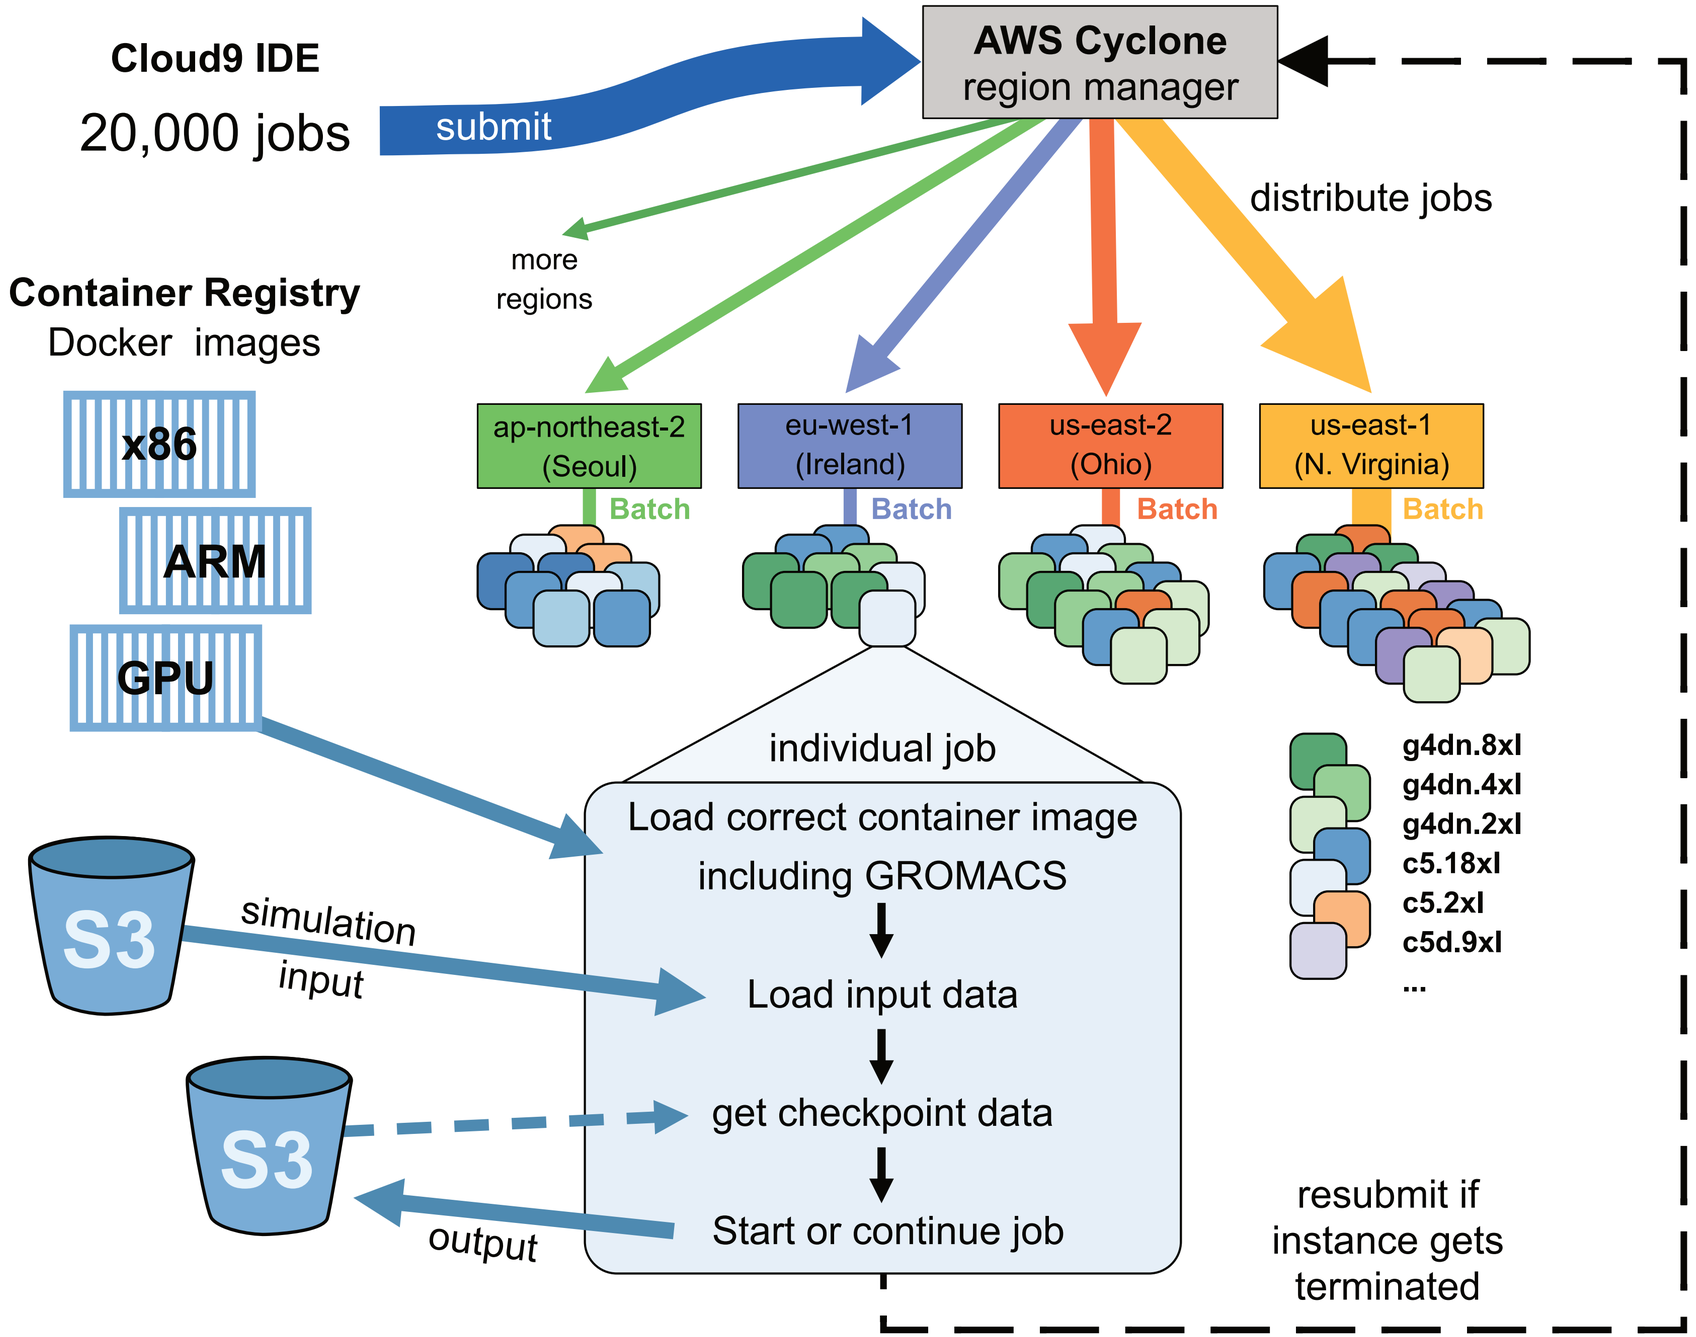
Cloud9 IDE
20,000 jobs
submit
AWS Cyclone
region manager
distribute jobs
more
regions
Container Registry
Docker images
x86
ARM
GPU
ap-northeast-2
(Seoul)
eu-west-1
(Ireland)
us-east-2
(Ohio)
us-east-1
(N. Virginia)
Batch
Batch
Batch
Batch
g4dn.8xl
g4dn.4xl
g4dn.2xl
c5.18xl
c5.2xl
c5d.9xl
...
individual job
Load correct container image
including GROMACS
Load input data
get checkpoint data
Start or continue job
S3
simulation
input
S3
output
resubmit if
instance gets
terminated

Cloud-based HTC

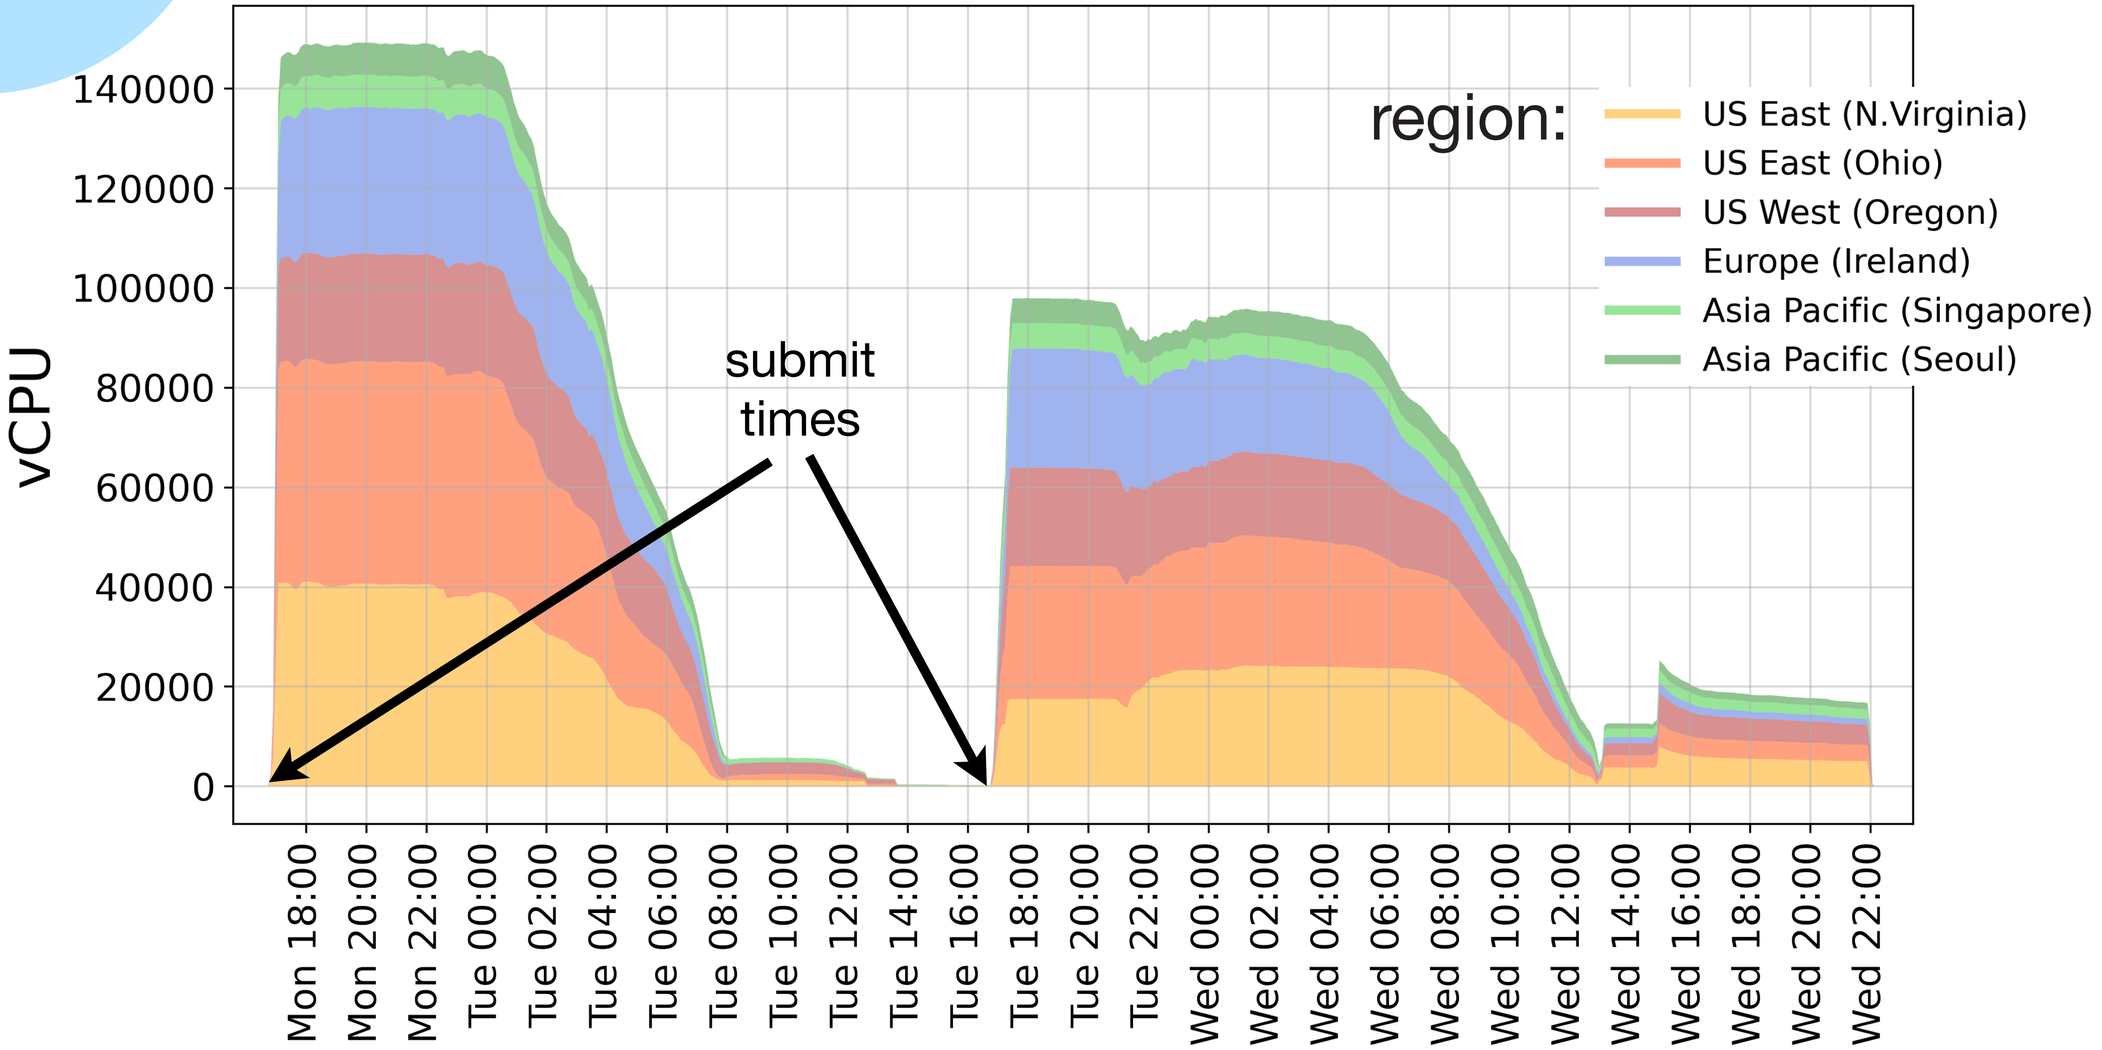

Cloud-based HTC

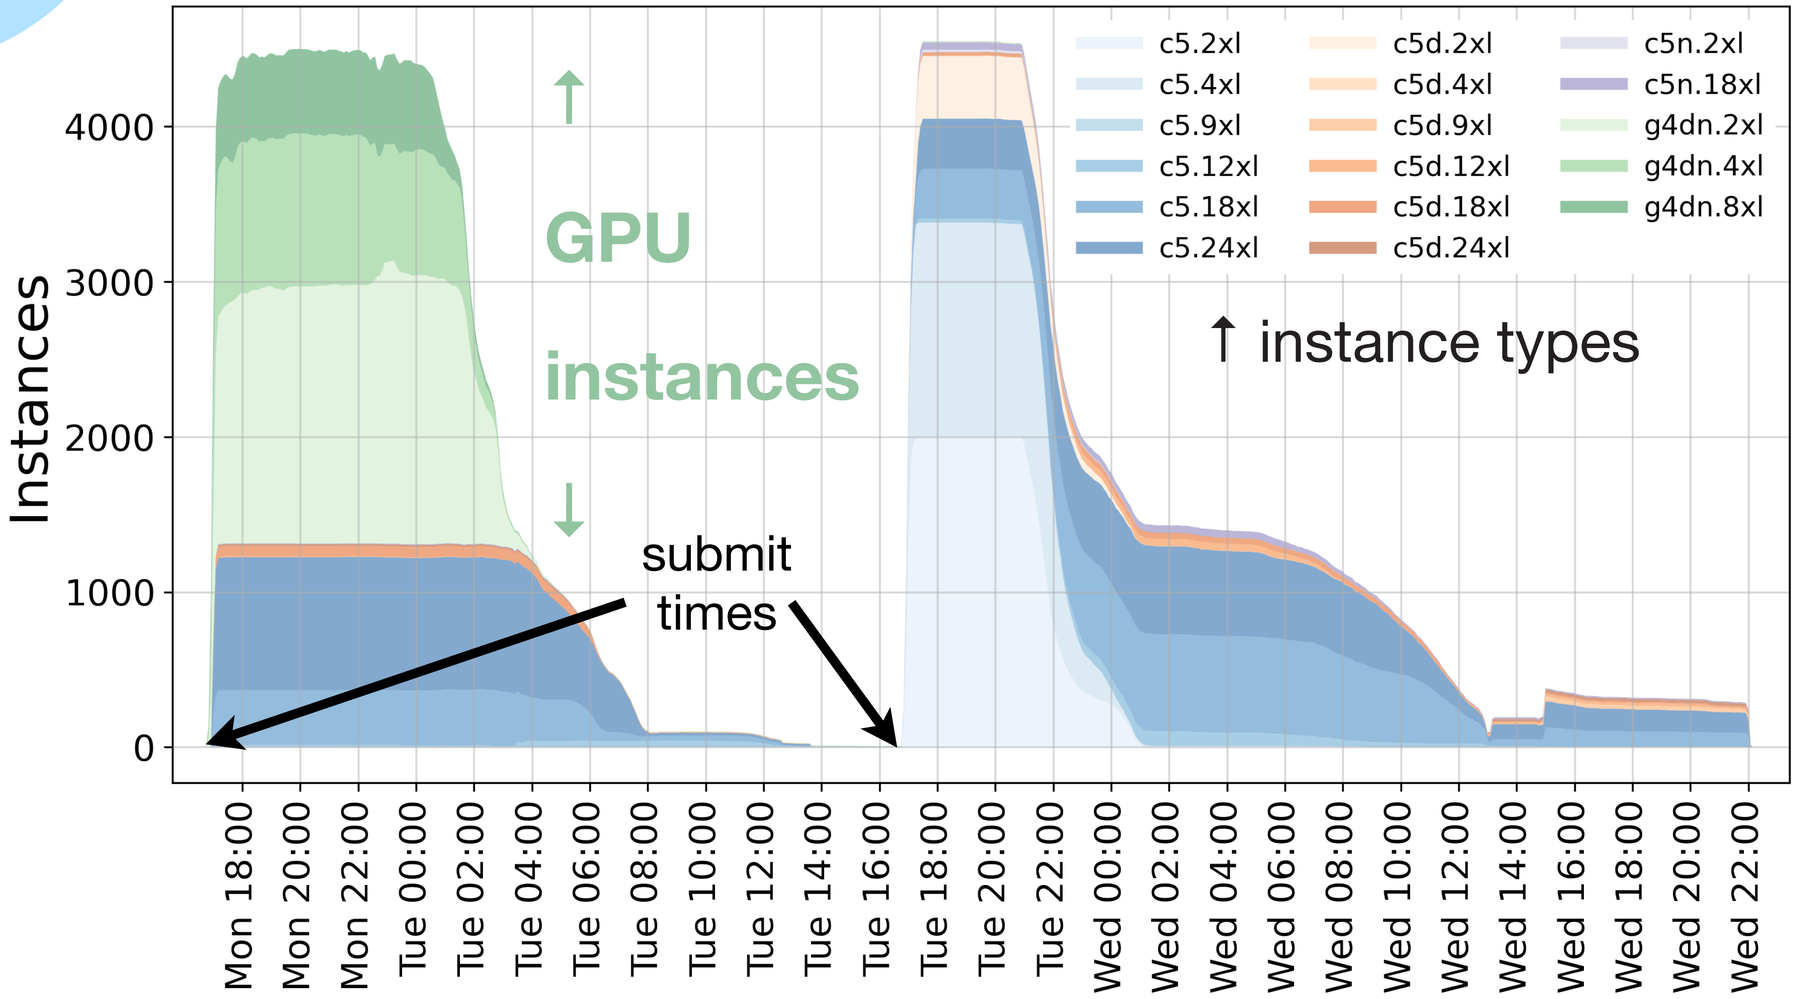

Cloud-based HTC

Global compute resources accelerate computational drug design

Results III:

- With AWS Cyclone, our **20,000 jobs** acquired
  - **140,000 cores**
  - **3,000 GPUs**
  - **4,000 instances** shortly after submission

- in total ~1.8 M core hours
  i.e. about what a power user spends on MPCDF clusters per month

- Whole ensemble **finished in two days** (each job 3 h – 14 h),
  though < 24 h seems possible!

3 questions

regarding biomolecular simulations in the Cloud

1. **Can we do HPC in the Cloud?**
   Demonstrated GROMACS strong scaling to >23,000 cores
   ➔ *ParallelCluster*

2. **Is Cloud computing expensive?**
   With the right instances, Cloud computing can be as cost effective as an own cluster

3. **Throughput: How far can we go?**
   Quickly deploy a few thousand Spot instances globally
   Even large MD ensembles can potentially be run overnight
   ➔ *Cyclone*

C Kutzner, C Kniep, A Cherian, L Nordstrom, H Grubmüller, BL de Groot, V Gapsys: *GROMACS in the Cloud: A Global Supercomputer to Speed Up Alchemical Drug Design*, J. Chem. Inf. Model. 62 (7), 1691–1711 (2022)

Acknowledgments

People involved in this collaboration

MAX PLANCK INSTITUTE
FOR MULTIDISCIPLINARY SCIENCES

MPI

- Dr Vytas Gapsys – Computational drug design
- Prof Bert de Groot – Computational Biomolecular Dynamics Group head
- Prof Helmut Grubmüller – Theoretical and Computational Biophysics Dept. head

AWS

- Christian Kniep – Container-based workflow setup
- Ludvig Nordstrom – AWS Cyclone
- Austin Cherian – Strong scaling benchmarks
- Torsten Bloth, Olly Perks, Stephen Sachs, Cristian Măgherușan-Stanciu, Bruno Silva, Johannes Schulz – Advice and support
- Agata Jablonka (AWS public sector) – compute time for cloud-based simulations