FAST. FLEXIBLE. FREE.

GROMACS

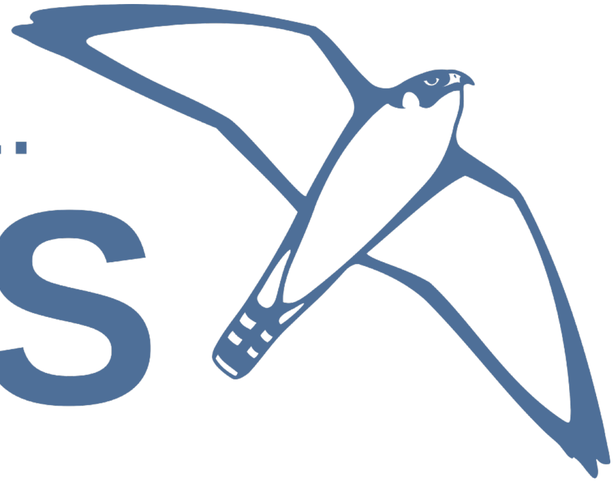

Berk Hess

bioexcel

GROMACS

- GROMACS is an open source (LGPL2) molecular dynamics simulation package
- One of the most popular HPC codes world-wide
  - 10000 citations in 2023, ≈ 4% of HPC usage
- Very efficient code (C++)
- Supports parallelization using OpenMP and MPI
- GPU acceleration using CUDA, OpenCL & SYCL
- Development supported by several EU projects

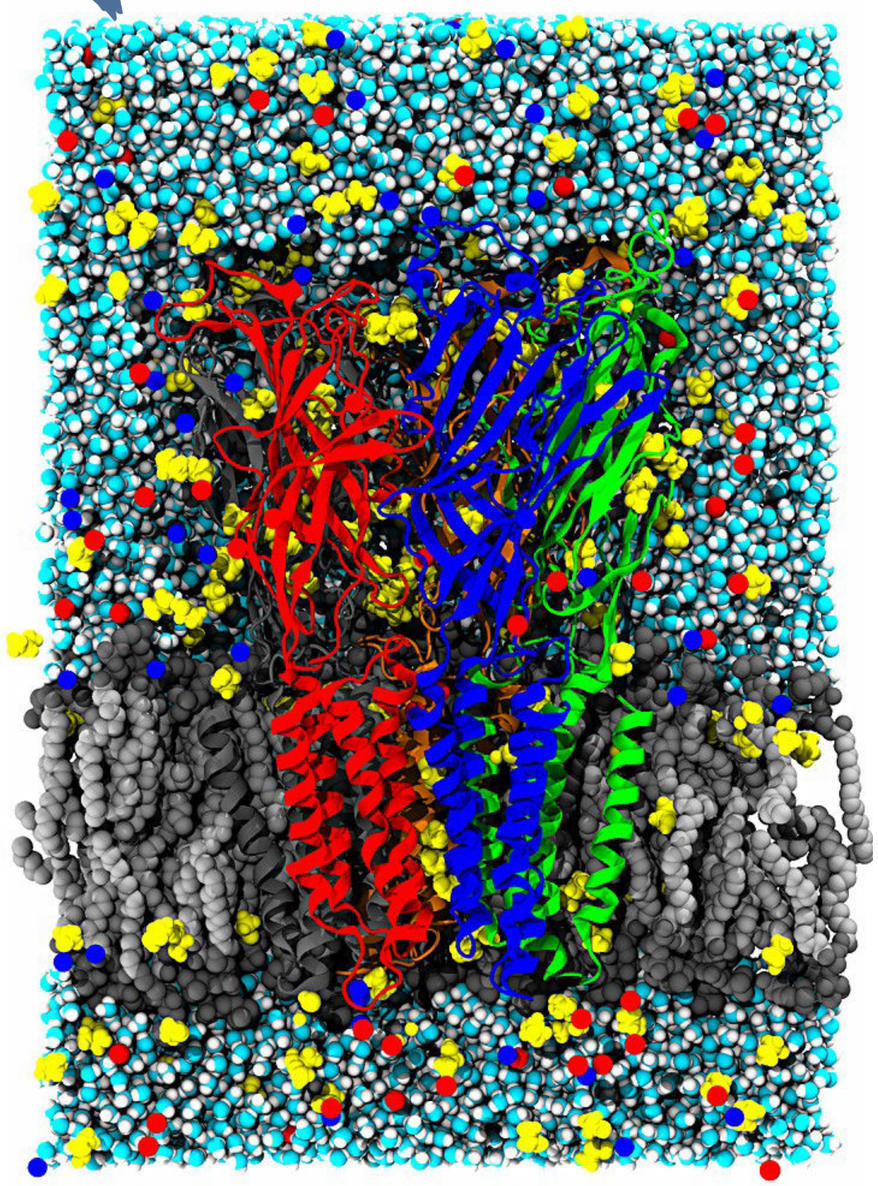

bioexcel

Applications

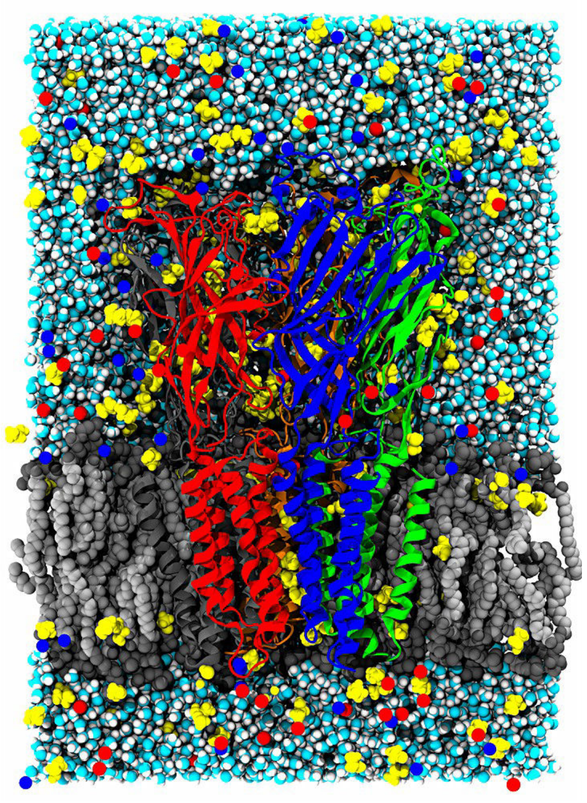

Protein

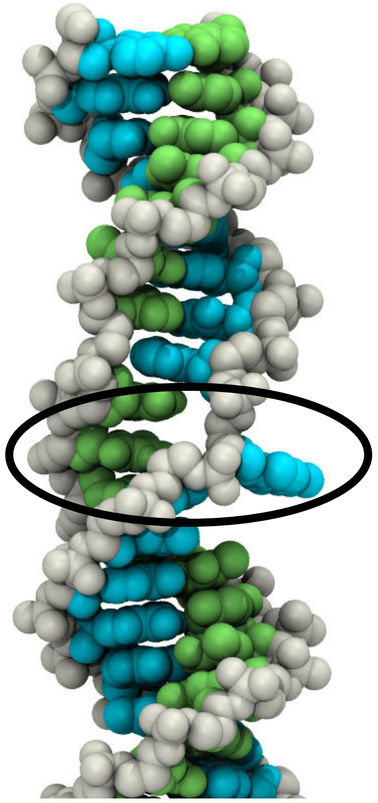

DNA

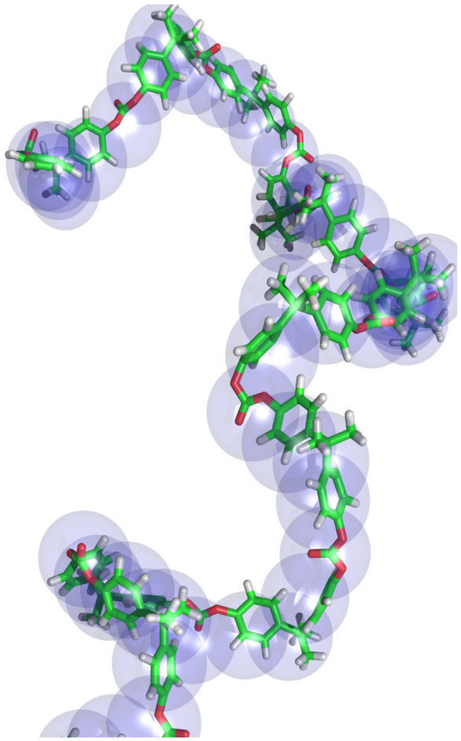

Polymers

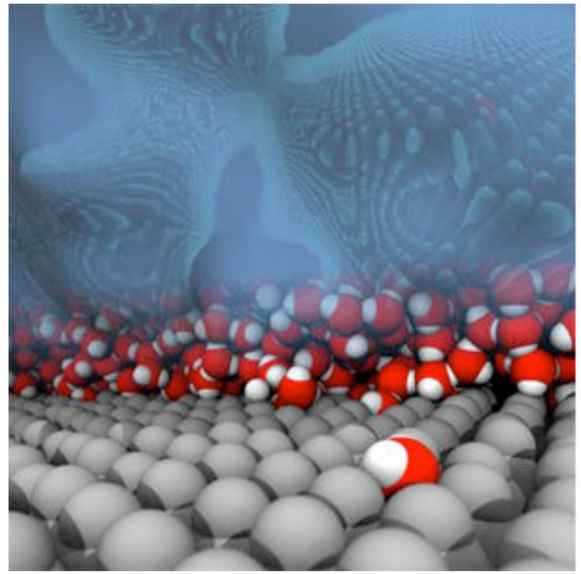

Flow

GROMACS tools

- System preparation, with coordinates taken from the protein data bank, alpha fold or self generated
    - Generates a topology with parameters from the force field chosen by the user
- Preprocessing tool generating a portable run input file
- **The main mdrun simulation engine**
- 100 analysis tools

MD simulation quality

Quality is determined by:

1. Interaction parameters given by the force field
    - Many force field are available in/for GROMACS
2. Accuracy of force calculation and integration (nowadays unproblematic)
3. Sampling, **proportional to the speed of the code**

Molecular dynamics

- Basically Newton's equation of motion:

$$m_i \frac{d^2\,\mathbf{x}_i}{d\,t^2} = \mathbf{f}_i(\mathbf{x}) = -\nabla_i V(\mathbf{x})$$

- Particle pair force calculation, within a cut-off, dominates the cost:
  - LJ:

$$V_{ij}(r_{ij}) = 4\epsilon_{ij}\left[\left(\frac{\sigma_{ij}}{r_{ij}}\right)^{12} - \left(\frac{\sigma_{ij}}{r_{ij}}\right)^{6}\right]$$

  - Coulomb:

$$V_{ij}(r_{ij}) = \frac{q_i q_j}{4\pi\epsilon_0 r_{ij}}$$

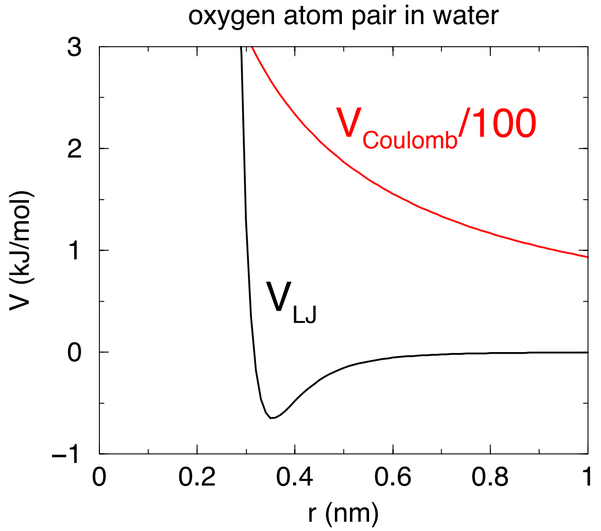

Long-range electrostatics

- 1/r is long range: can not use a plain cut-off
- Particle mesh Ewald:
  - Decompose 1/r into short+long-range:

$$V_{\text{Coulomb}} = V_{direct} + V_{reciprocal} + V_0$$

pair term **long-range all vs all** **constant**

  - The reciprocal part is computed on a grid using a 3D-FFT
    - **This limits parallelisation**
    - **Recent development HeFFTe for GPUs**

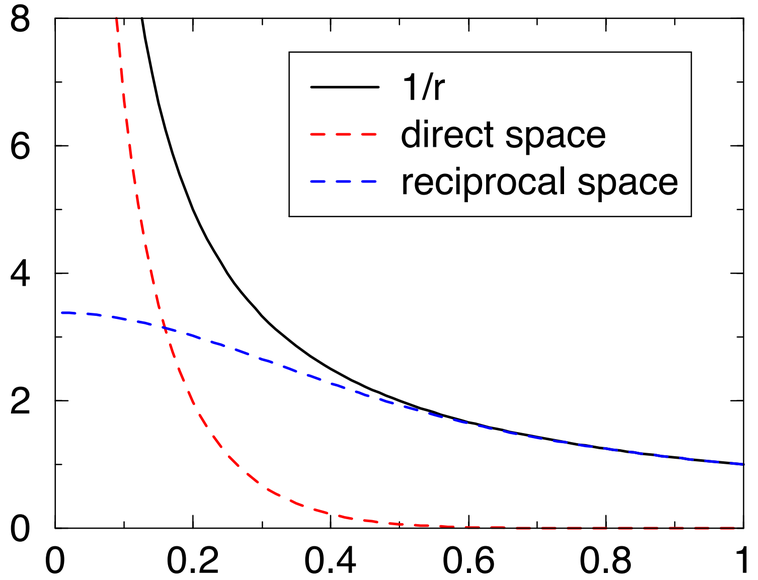

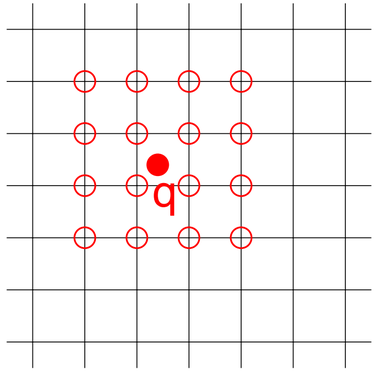

GROMACS parallelization

- Single node:
  - OpenMP shared-memory parallelisation
  - “thread-MPI” + OpenMP parallelisation
    - enables GROMACS to optimise the rank/thread counts
    - enables efficient GPU-GPU comm.
- Multi node:
  - MPI + OpenMP parallelization
  - Dynamic load balancing
  - Multiple-program multiple-data parallel.

equally sized domains

with dynamic load balancing

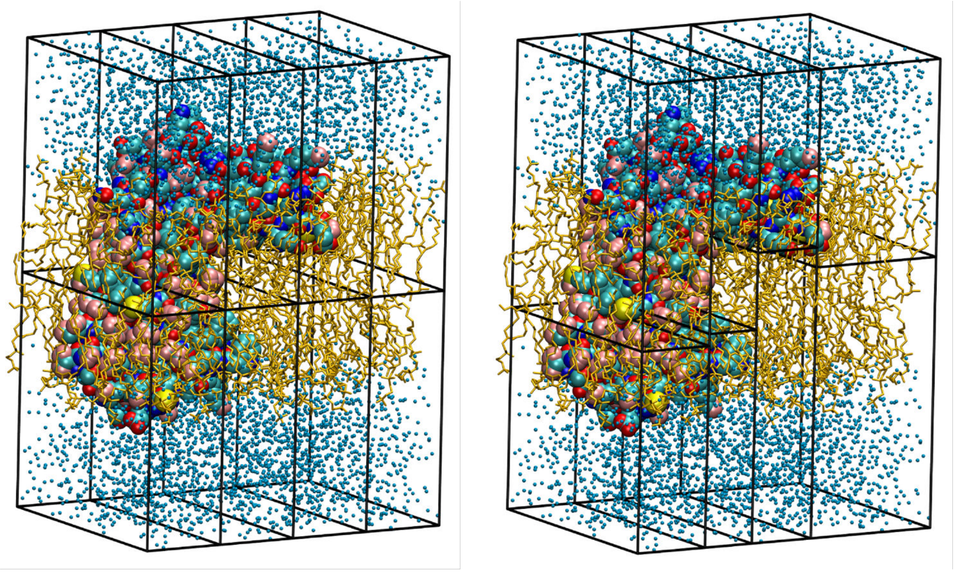

MD performance

- Bio-molecular systems have a fixed size, usually 50000 - 300000 atoms
- This means we have a strong scaling problem
  - more cores → less work per core, shorter iteration times

- A time step in atomistic simulations is usually 2 femtoseconds
- GROMACS can execute a time step in less than 1 milliseconds
  - > 100 million steps per day, or > 200 **ns/day**

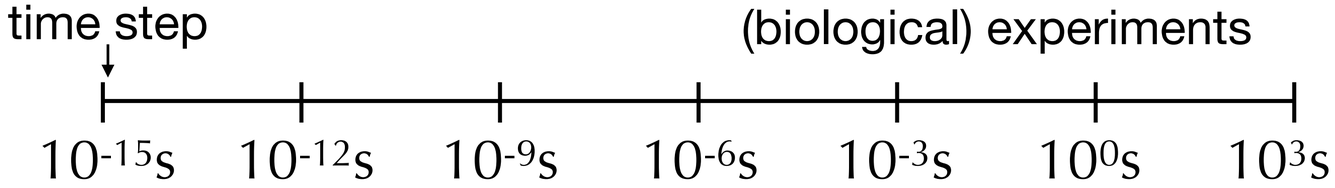

Performance on LUMI

- LUMI is challenging because of AMD GPUs:
  - No CUDA support
  - Software stack issues
- GROMACS solution:
  - SYCL code, convert to HIP using OpenSYCL
  - Use HeFFTe for parallel FFT over GPUs

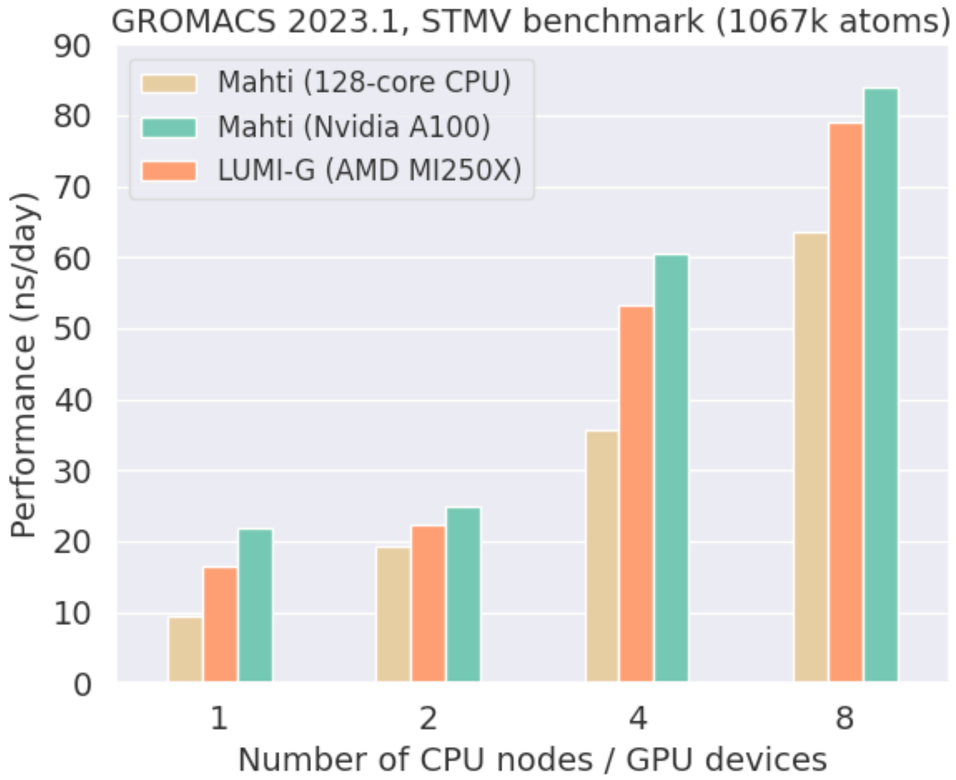

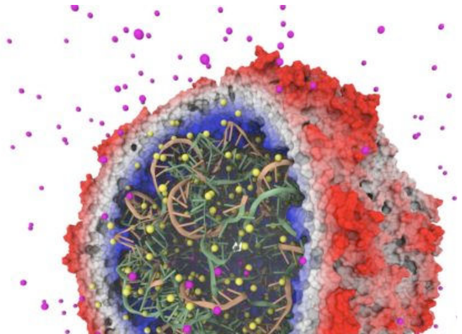

GROMACS highlights

Heterogeneous parallelization and GPU acceleration

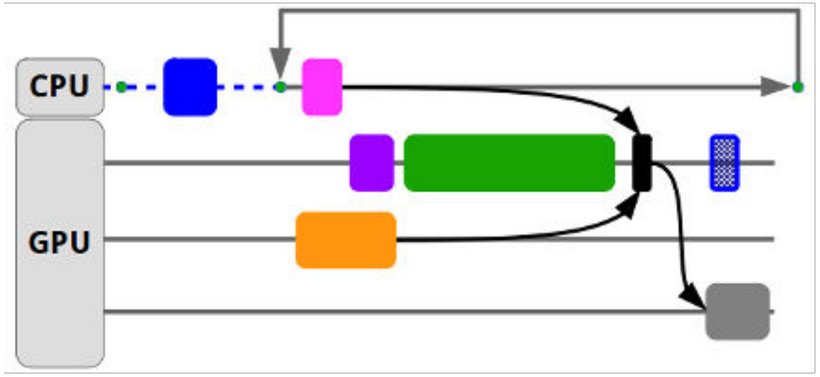

Accelerating sampling with the AWH method

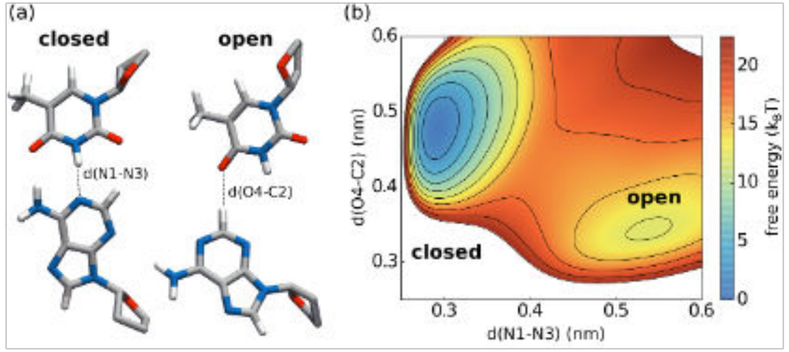

gmxapi Python interface for GROMACS

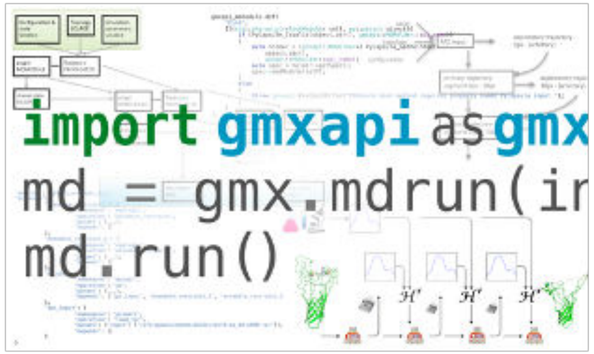

Hybrid Quantum-Classical simulations (QM/MM) with CP2K interface

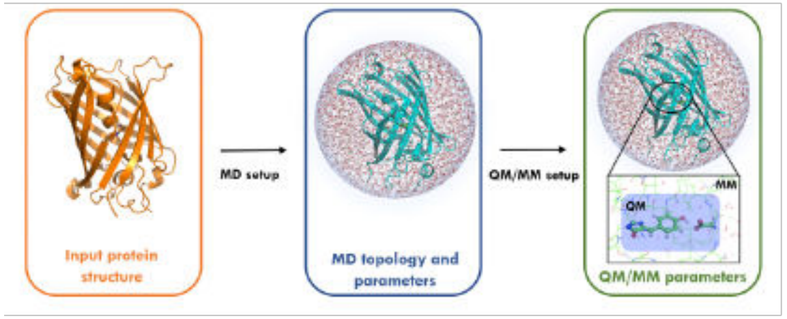

Applying the Accelerated Weight Histogram method to alchemical transformations

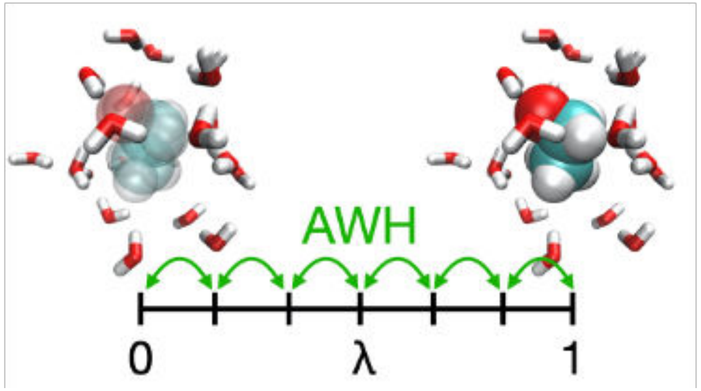

Ions in Action! Studying ion channels by Computational Electrophysiology

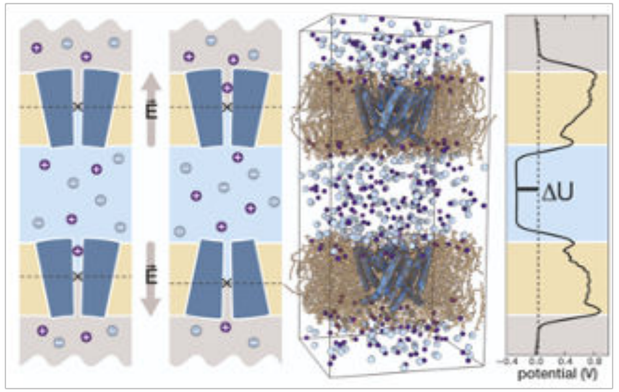

Code development

- Open development
- Issue tracking, code review & CI/CD through GitLab
- Yearly major release cadence
- Quarterly bug-fix releases
- Co-design efforts with all major vendors:

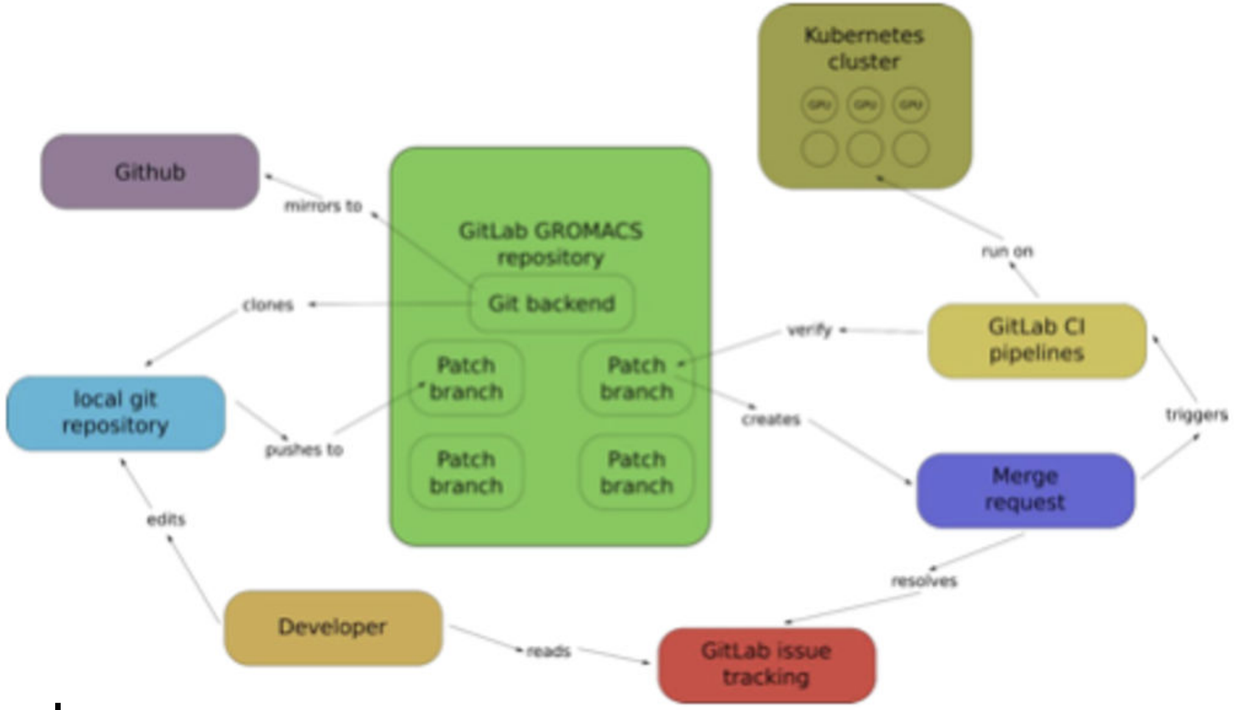

NVIDIA

intel

AMD

More information / developers

- Website: www.gromacs.org
- Development: gitlab.com/gromacs
- Tutorials: tutorials.gromacs.org
- Forums (user/developer discussion, announce): gromacs.bioexcel.eu

Project leaders: Erik Lindahl, Berk Hess

Current contributors:
Mark Abraham, Andrey Alekseenko, Vladimir Basov, Cathrine Bergh, Eliane Briand, Ania Brown, Mahesh Doijade, Giacomo Fiorin, Stefan Fleischmann, Sergey Gorelov, Gilles Gouaillardet, Alan Gray, M. Eric Irrgang, Farzaneh Jalalypour, Joe Jordan, Carsten Kutzner, Justin A. Lemkul, Magnus Lundborg, Pascal Merz, Vedran Miletic, Dmitry Morozov, Julien Nabet, Szilárd Páll, Andrea Pasquadibisceglie, Michele Pellegrino, Hubert Santuz, Roland Schulz, Tatiana Shugaeva, Alexey Shvetsov, Philip Turner, Alessandra Villa, Sebastian Wingbermuehle.

FAST. FLEXIBLE. FREE.
GROMACS

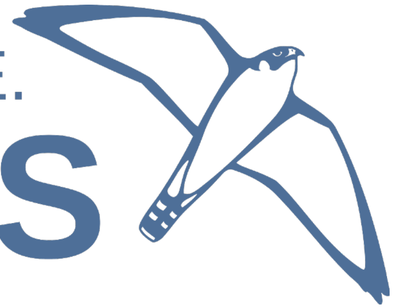

bioexcel