bioexcel
Centre of Excellence for Computational Biomolecular Research

Horizon 2020
European Union Funding
for Research & Innovation

GROMACS 2021 overview

Paul Bauer and Berk Hess

pbauer@kth.se, hess@kth.se

What to do with GROMACS?
Biomolecular simulations!

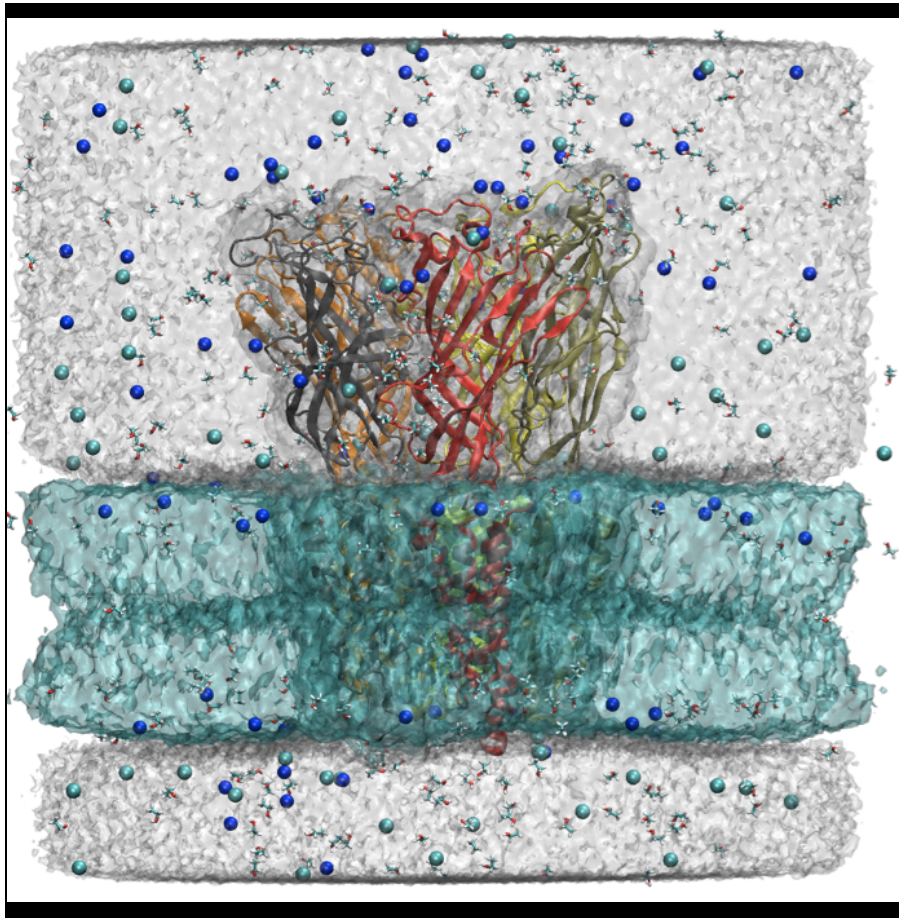

Read the docs! ☺

- All the docs
  - http://manual.gromacs.org/documentation/2021/index.html
- Release notes
  - http://manual.gromacs.org/documentation/2021/release-notes/index.html#major-release

- Previous webinar about GROMACS 2020 capabilities
  - https://bioexcel.eu/webinar-whats-new-in-gromacs-2020-2020-02-20/
- Previous webinar about GROMACS 2019 capabilities
  - https://bioexcel.eu/gromacs-2019-overview-of-the-new-features-and-capabilities-2019-05-17/

GROMACS annual release cycles

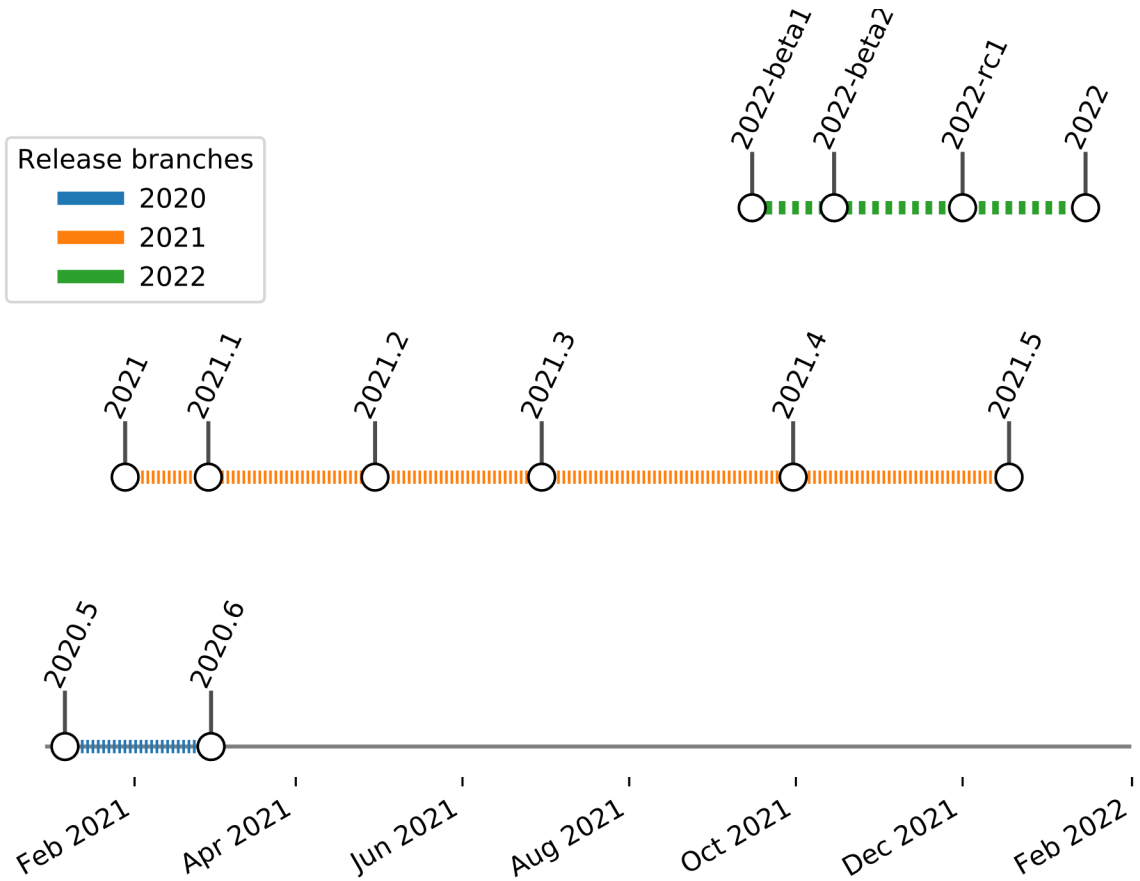

GROMACS impact over time

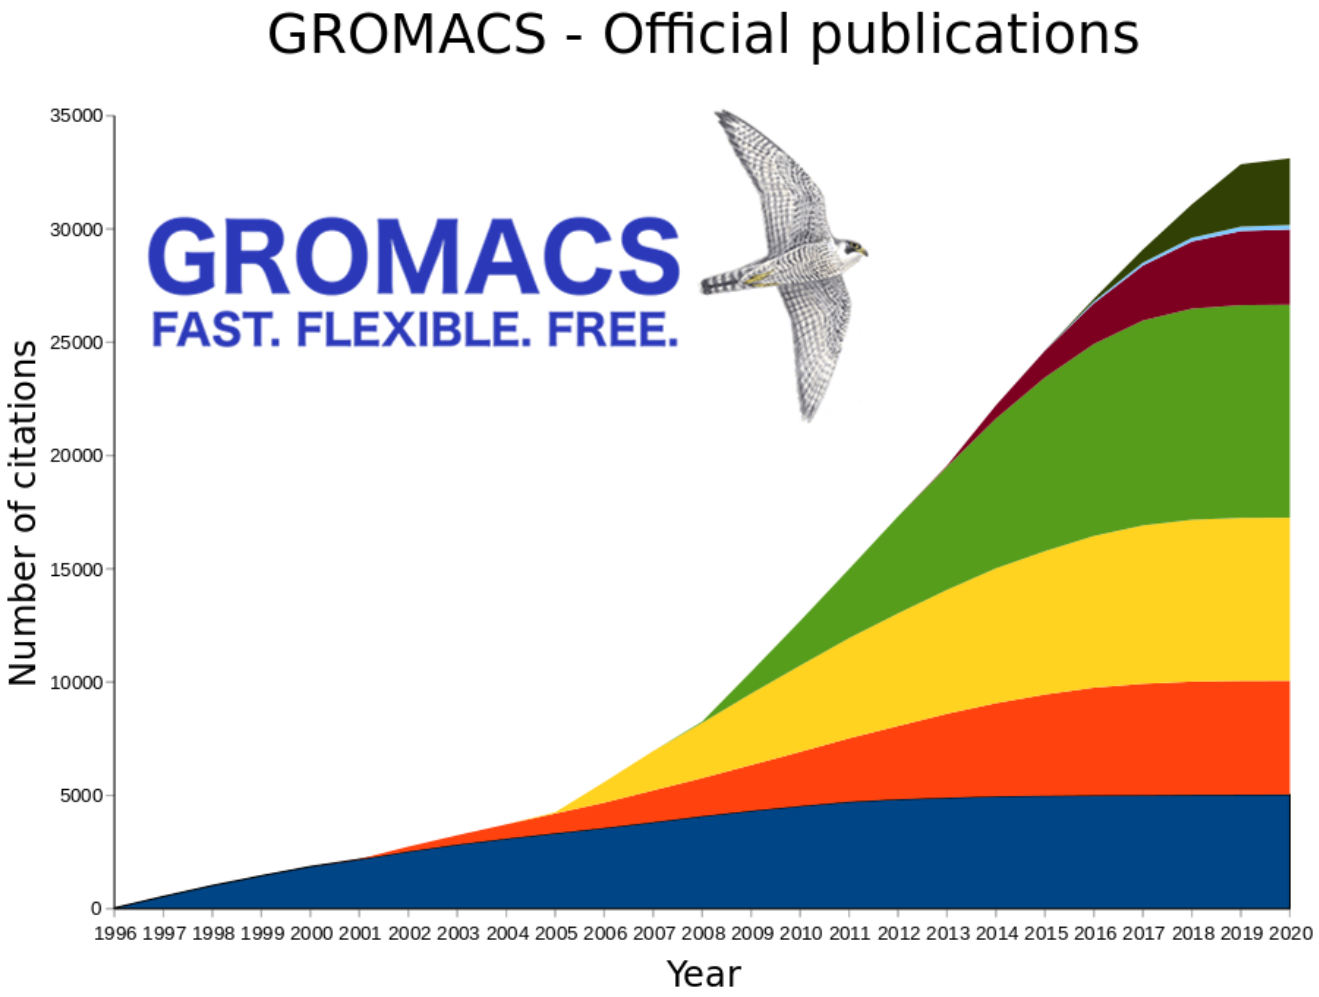

- GROMACS: High performance molecular simulations through multi-level parallelism from laptops to supercomputers
- Tackling Exascale Software Challenges in Molecular Dynamics Simulations with GROMACS.
- GROMACS 4.5: a high-throughput and highly parallel open source molecular simulation toolkit.
- GROMACS 4: Algorithms for highly efficient, load-balanced, and scalable molecular simulation.
- GROMACS: Fast, flexible, and free.
- GROMACS 3.0: a package for molecular simulation and trajectory analysis.
- GROMACS - A Message-Passing parallel Molecular-Dynamics Implementation.

Active projects around GROMACS

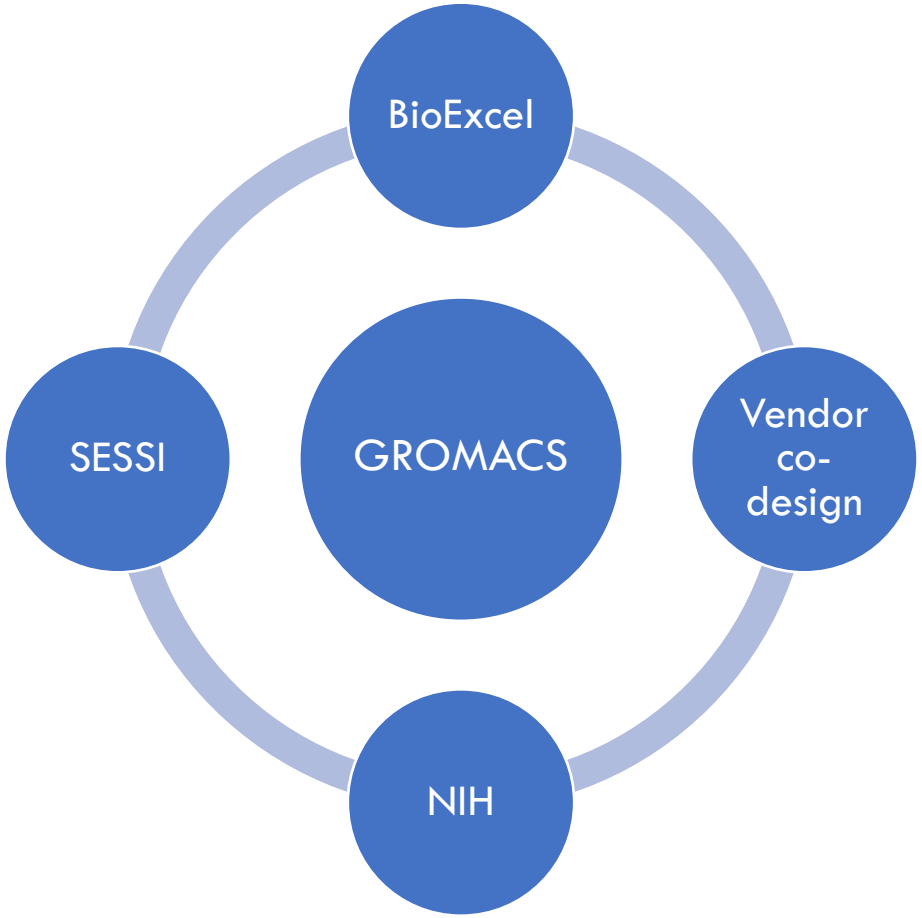

New features

- Support for multiple time stepping
- Stochastic cell rescaling
- SYCL as accelerator framework
- Free energy perturbation with AWH
- PME offloading to GPU for free energy simulations
- ARM SVE and Fujitsu A64FX (contribution by RIST)
- New nonbonded interaction API with NB-LIB (in collaboration with PRACE)
- New logo!

Multiple time stepping

- Improve performance by calculating forces less frequently
- r-RESPA reversible and symplectic integrator
- Split the potential in fast and slowly varying parts
- Simple to implement with leap-frog integrator:
  - Every n steps:

$$v_{i+1/2} = v_{i-1/2} + \frac{\Delta t}{m}(F_{\text{fast}} + n\,F_{\text{slow}})$$

  - Every other step:

$$v_{i+1/2} = v_{i-1/2} + \frac{\Delta t}{m}F_{\text{fast}}$$

- Constraint virial is more complicated

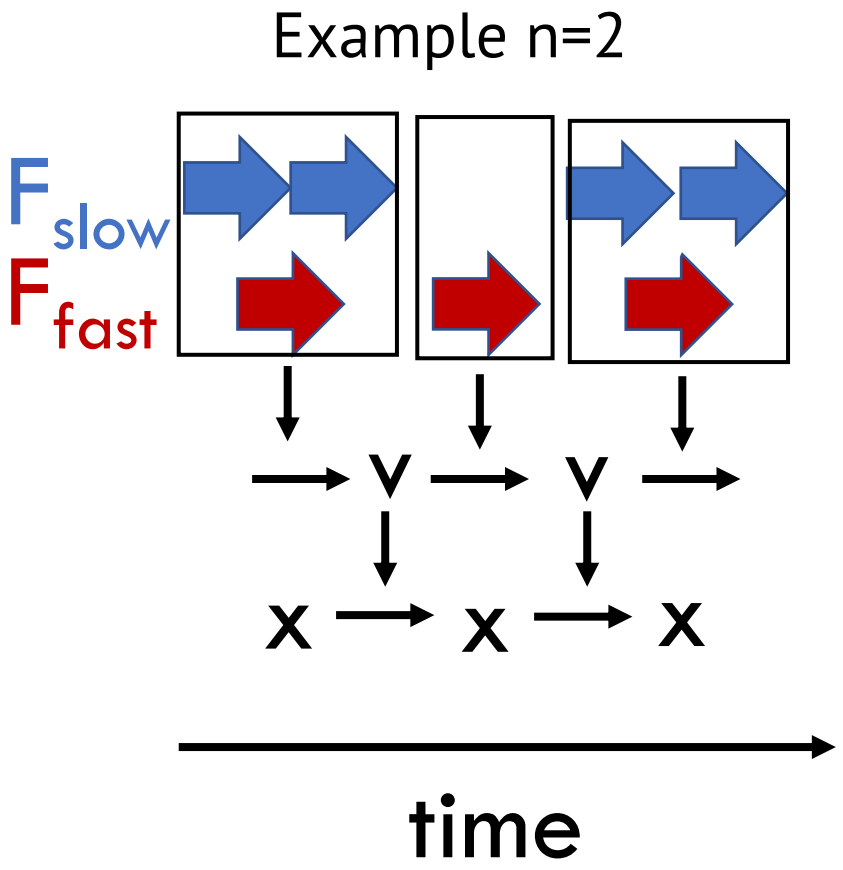

Multiple time stepping options

- mdp file options:
  - `mts = yes` (default = no)
  - `mts-levels = 2` (only 2 supported)
  - `mts-level2forces = longrange-nonbonded` (= PME grid forces, default)
  - `mts-level2factor = 2` (default)
- Small to moderate performance gain in most cases
- Higher performance gain at high parallelization
- Can be used for COM pull forces: high gain for parallel simulations

- Original idea was to replace dt=4 with hydrogens replaced by virtual sites with:
  - `mts-level2forces = longrange-nonbonded nonbonded pairs dihedrals`
- But a hydrogen get unstable every ~100 ns with dt=2 fs (dt=1.8 fs might work)

Multiple time stepping improvements

- Improvements considered for GROMACS 2022 or later:
  - Support for update on GPU combinded with MTS
  - Support for SD integrator
  - Support for more special forces in MTS
  - Support for more than 2 MTS levels

Free-energy with AWH

- Conventional procedure for alchemical free-energy calculations:
  - Add coupling parameter λ to interpolate between state A and B
  - Run many independent simulations for different λ values
  - Use `gmx bar` to calculate the free-energy difference with Bennett's acceptance ratio

- Now you can use the accelerated weight histogram method
  - `awh = yes`
  - `awh1-dim1-coord-provider = fep-lambda`
  - `awh1-dim1-diffusion = 0.001` (seems to work well)
- You get the free-energy efficiently from a single simulation
- You can parallelize with multiple walkers using mdrun -multidir

BioExcel webinar on AWH + FEP coming soon

New features

Stochastic cell rescaling

Algorithm description here: https://doi.org/10.1063/5.0020514

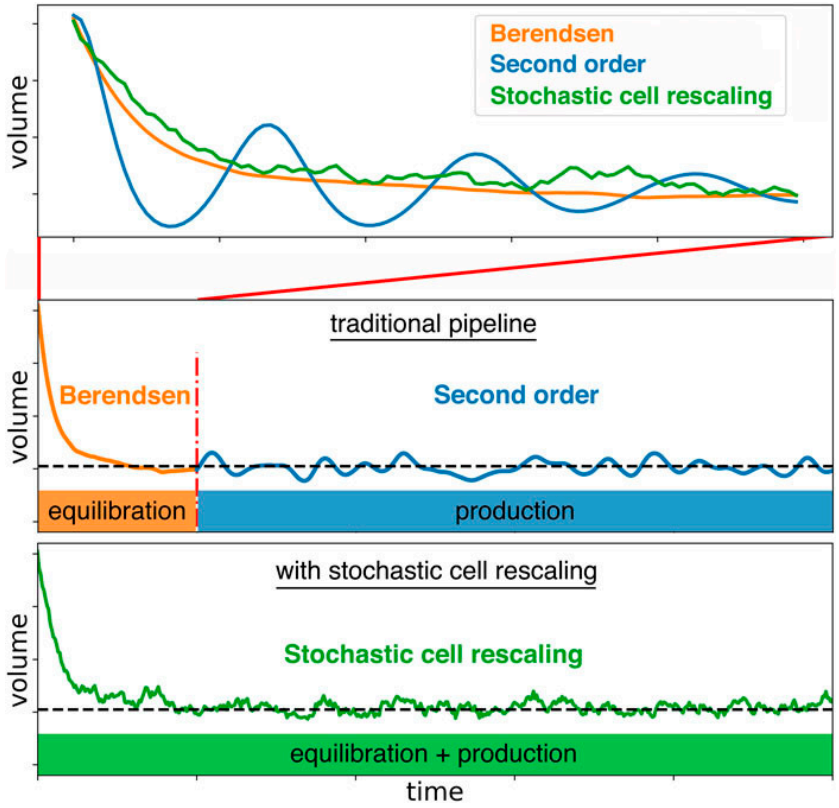

FIG. 1. Graphical representation of protocols for a constant-pressure simulation. In the traditional pipeline (middle panel), an equilibration run using a first-order Berendsen barostat[28] is followed by a production run using a second-order barostat.[2–24] Indeed, using a second-order barostat on a non-equilibrated system might lead to oscillations and instabilities (upper panel). The stochastic cell rescaling algorithm introduced here relaxes straight to the correct volume and then produces correct fluctuations (lower panel). It can thus be used both for equilibration and production runs.

Published in: Mattia Bernetti; Giovanni Bussi; *J. Chem. Phys.* **153,** 114107 (2020)
DOI: 10.1063/5.0020514

New features

SYCL support

- Added experimental support for SYCL
- Will be developed over the year to allow native offloading to Intel GPUs
- Investigation if it is possible to use hipSYCL for AMD GPU support
- Development branch for those happy to experiment here:
  - https://manual.gromacs.org/2021-sycl/download.html

New features

PME on GPU for free energy calculations

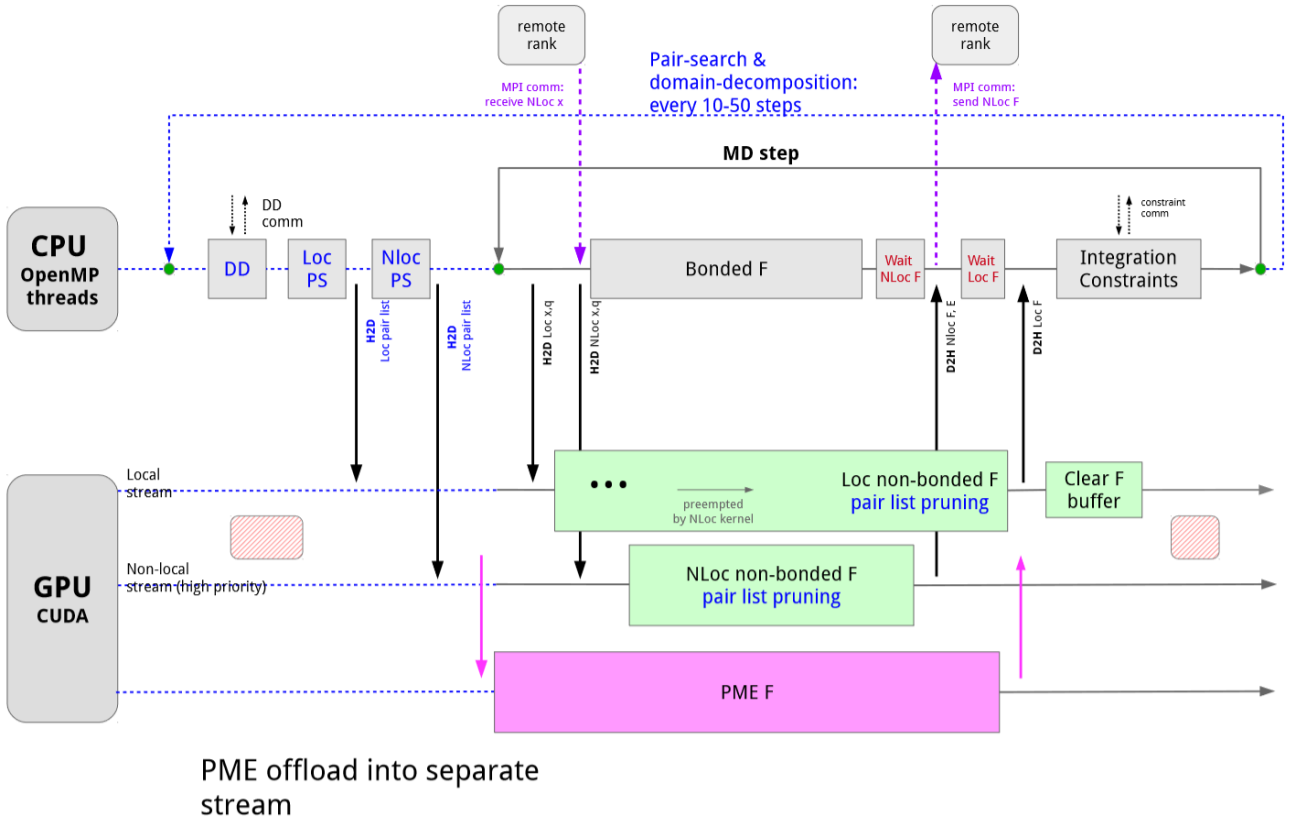

PME offload into separate stream

New features

ARM Scalable Vector Extensions (SVE) support

- SIMD kernels added for SVE on ARM
- Needs compatible compiler and toolchain
- For now, only possible to compile in one vector size determined at configure time
- Provides possibility to support A64FX chips
- Contributed by Research Organization for Science Information and Technology (RIST).

New features

NB-LIB nonbonded interaction API

- First release containing working example for test cases
- Able to generate small programs to test algorithms
- Supports new topology and interaction type formats

Removed features

- Nothing new! But still some missing support from group scheme deprecation
  - Membrane embedding still not working
  - Still no support for user supplied tables

New requirements

- C++17 compatible compilers (gcc-7, clang-7)
- C++14 compatible CUDA nvcc (starting at 9.0)
- Cmake > 3.13.0

Deprecations and announcements

- deffnm will be removed
- OpenCL will be removed once replacements are in place
- Legacy SIMD architecture support
  - Intel KNC (MIC
  - Sparc64 HPC ACE
  - Armv7
  - Power 7
- mdrun-only build
- HWLOC API version 1
- Legacy API
- Constant-acceleration MD
- Reading .pdo files in gmx wham

Active co-design projects

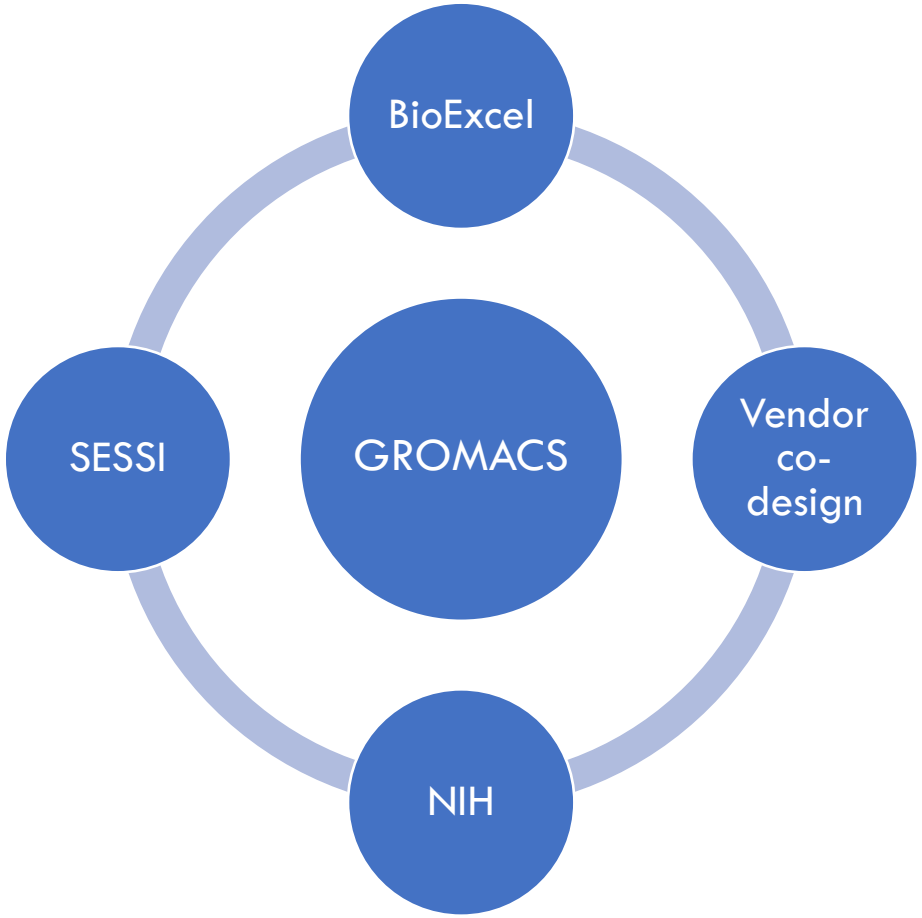

NVIDIA

arm

STREAM
High Performance Computing

intel

AMD

Long term plans for GROMACS

- multiple time stepping to replace virtual-sites
- Modularization of core
- Encapsulated low-level layers
- Python and C++ API
- Extensible force calculation
- Task parallelism
- Auto-tuning
- FMM

- Ensemble strategies for free-energy surfaces
- Modernized preparation and analysis suites
- Improved UI and UX
- New input and output formats
- Container-based CI testing
- Container distribution

Acknowledgments

- Paul Bauer
- Szilard Pall
- Aleksei Yupinov
- Sander Pronk
- Christian Blau
- Magnus Lundborg
- Roland Schulz
- Teemu Murtola
- Mark Abraham
- Berk Hess
- Erik Lindahl
- Jonathan Vincent
- Alan Gray
- Anca Hamararu
- Joe Jordan
- Pascal Merz
- Michael Shirts
- Eric Irrgang
- Peter Kasson
- David van der Spoel
- Justin Lemkul
- Kevin Boyd

Audience Q&A session

- Please use the Questions function in GoToWebinar application
  - If you *don't have audio*, please mention that in the question.

- Any other questions or points to discuss after the live webinar? Join the discussions at http://ask.bioexcel.eu.

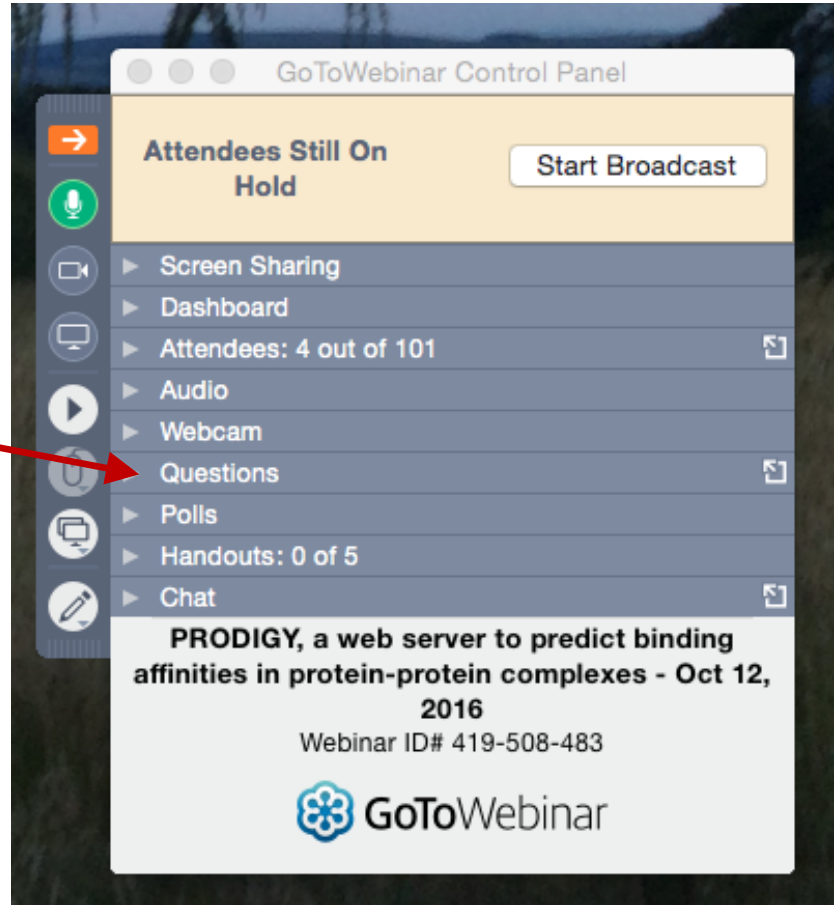

Next webinar occasion

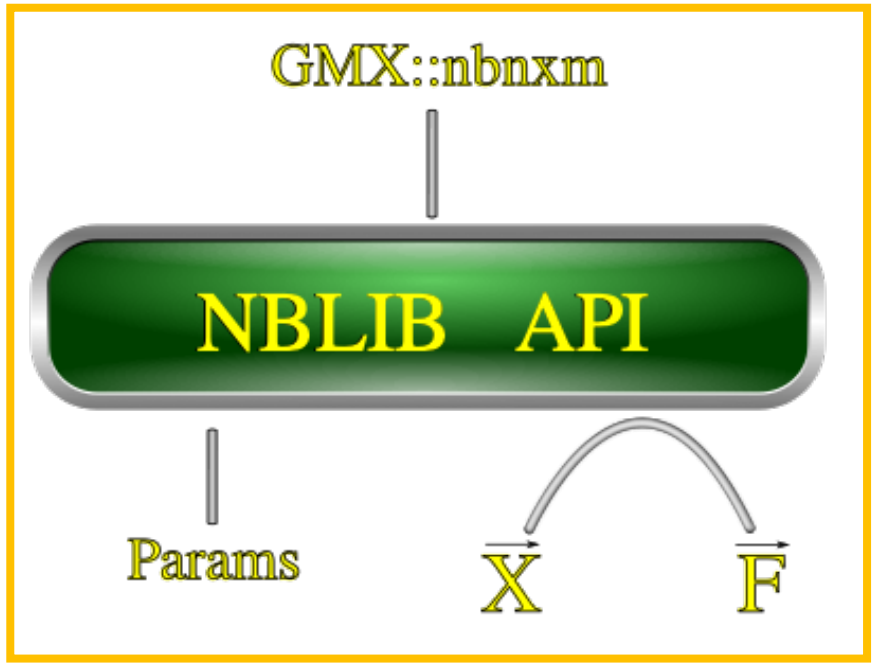

NB-LIB: A performance portable library for computing forces and energies of multi-particle systems

by
Joe Jordan and Sebastian Keller

See https://bioexcel.eu/