

Gromacs 2023.2 with GPU support

Webpage

<http://www.gromacs.org/>

Version

2023.2

Build Environment

- GCC 9.2.1 (gcc-toolset-9)
- Open MPI 4.1.5 (CUDA-aware)
- CUDA 12.1 Update 1

Files Required

- gromacs-2023.2.tar.gz
- regressiontests-2023.2.tar.gz
- fftw-3.3.8.tar.gz

Build Procedure

```
#!/bin/sh
#PBS -l select=1:ncpus=16:mpiprocs=8:ompthreads=2:ngpus=2
#PBS -l walltime=03:00:00

VERSION=2023.2
INSTALL_PREFIX=/apl/gromacs/${VERSION}-CUDA

BASEDIR=/home/users/${USER}/Software/Gromacs/${VERSION}/
GROMACS_TARBALL=${BASEDIR}/gromacs-${VERSION}.tar.gz
REGRESSION_TARBALL=${BASEDIR}/regressiontests-${VERSION}.tar.gz
WORKDIR=/gwork/users/${USER}
REGRESSION_PATH=${WORKDIR}/regressiontests-${VERSION}

PARALLEL=12
export LANG=C

FFTW_VER=3.3.8
FFTW_PATH=${BASEDIR}/fftw-${FFTW_VER}.tar.gz

#-----
umask 0022

module -s purge
module -s load gcc-toolset/9
module -s load openmpi/4.1.5/gcc9-cuda12.1u1
module -s load cuda/12.1u1

#export CUDA_VISIBLE_DEVICES=0,1
unset OMP_NUM_THREADS

cd ${WORKDIR}
if [ -d gromacs-${VERSION} ]; then
  mv gromacs-${VERSION} gromacs_erase
  rm -rf gromacs_erase &
fi

if [ -d regressiontests-${VERSION} ]; then
  mv regressiontests-${VERSION} regressiontests_erase
  rm -rf regressiontests_erase &
fi
```

```

tar xzf ${GROMACS_TARBALL}
tar xzf ${REGRESSION_TARBALL}
cd gromacs-${VERSION}

# single precision, no MPI
mkdir rccs-s
cd rccs-s
cmake .. \
  -DCMAKE_INSTALL_PREFIX=${INSTALL_PREFIX} \
  -DCMAKE_VERBOSE_MAKEFILE=ON \
  -DCMAKE_C_COMPILER=gcc \
  -DCMAKE_CXX_COMPILER=g++ \
  -DGMX_MPI=OFF \
  -DGMX_GPU=CUDA \
  -DGMX_DOUBLE=OFF \
  -DGMX_THREAD_MPI=ON \
  -DGMX_USE_CUFFTMP=OFF \
  -DGMX_BUILD_OWN_FFTW=ON \
  -DGMX_BUILD_OWN_FFTW_URL=${FFTW_PATH} \
  -DREGRESSIONTEST_DOWNLOAD=OFF \
  -DREGRESSIONTEST_PATH=${REGRESSION_PATH}
make -j${PARALLEL} && make check && make install
cd ..

# single precision, with MPI
mkdir rccs-mpi-s
cd rccs-mpi-s
cmake .. \
  -DCMAKE_INSTALL_PREFIX=${INSTALL_PREFIX} \
  -DCMAKE_VERBOSE_MAKEFILE=ON \
  -DCMAKE_C_COMPILER=mpicc \
  -DCMAKE_CXX_COMPILER=mpicxx \
  -DGMX_MPI=ON \
  -DGMX_GPU=CUDA \
  -DGMX_DOUBLE=OFF \
  -DGMX_THREAD_MPI=OFF \
  -DGMX_USE_CUFFTMP=OFF \
  -DGMX_BUILD_OWN_FFTW=ON \
  -DGMX_BUILD_OWN_FFTW_URL=${FFTW_PATH} \
  -DREGRESSIONTEST_DOWNLOAD=OFF \
  -DREGRESSIONTEST_PATH=${REGRESSION_PATH}
make -j${PARALLEL} && make check && make install
cd ..

```

Notes

- (Please also check notes about CPU version)
- CUDA-aware Open MPI was employed for this version.
- cuFFTMp was not enabled. Multiple GPUs can't be assigned for PME reciprocal part.
 - NVSHMEM is required to enable NVSHMEM. NVSHMEM in NVIDIA HPC SDK can't be used due to the incompatibility. It is difficult to build and use our own version of NVSHMEM due to its license.
- When CUDA-aware HPC-X 2.11 was employed, some regression test occasionally failed due to timeout (probably hung). Although this error does not happen for CUDA-aware Open MPI 4.1.5 version so far, there can be the same problem...