

Gromacs (GPU) Best Practices

1. Introduction:

The following best practices document is provided as courtesy of the HPC Advisory Council.

2. Application Description:

GROMACS is a versatile package to perform molecular dynamics, i.e. simulate the Newtonian equations of motion for systems with hundreds to millions of particles.

It is primarily designed for biochemical molecules like proteins, lipids and nucleic acids that have a lot of complicated bonded interactions, but since GROMACS is extremely fast at calculating the nonbonded interactions (that usually dominate simulations) many groups are also using it for research on non-biological systems, e.g. polymers. More information on Gromacs can be found at the following web site: <http://www.gromacs.org>.

3. Version Information:

This guideline is based on gromacs-5.0.4. The source code can be downloaded from <http://www.gromacs.org/Downloads>.

4. Prerequisites:

4.1 Hardware:

The instructions from this best practice have been tested on the HPC Advisory Council, Dell™ PowerEdge™ R730 32-node cluster.

- Dual Socket Intel® Xeon® 14-core CPUs E5-2697 V3 @ 2.60 GHz
- Mellanox ConnectX-4 EDR 100Gb/s InfiniBand adapters
- Mellanox Switch-IB SB7700 36-Port 100Gb/s EDR InfiniBand switches
- NVIDIA K40, K80 GPUs

4.2 Software:

- a. OS: Red Hat Enterprise Linux 6.5
- b. GNU Compiler for Linux; 4.7.x or higher
- c. Other:
 - CUDA-6.5 or 7.0
 - Cuda-aware MPI - MVAPICH2, OpenMPI, or Platform MPI
 - Cmake 2.8.x or higher

5. Building Gromacs

5.1 CPU only version

CPU version can be built with either single or double precision.

```
module load gnu/4.9.3
```

```
module load cmake/3.2.3
```

```
source /opt/intel/composer_xe_2015.3.187/bin/compilervars.sh intel64
```

```
source /opt/intel/impi/5.0.3.048/bin64/mpivars.sh
```

```
BASEDIR=$PWD
```

```
rm -fr $BASEDIR/build
```

```
mkdir $BASEDIR/build
```

```
cd $BASEDIR/build
```

```
FLAGS="-DGMX_GPU=OFF "
```

```
FLAGS+="-DGMX_DOUBLE=ON "
```

```
FLAGS+="-DCMAKE_INSTALL_PREFIX=$BASEDIR/install/CPU-impi.DP "
```

```
cmake .. $FLAGS -DGMX_MPI=ON -DGMX_CPU_ACCELERATION=AVX2_256 -DGMX_FFT_LIBRARY=fftw3 -DCMAKE_C_COMPILER=mpiicc -DCMAKE_CXX_COMPILER=mpicpc -DGMX_PREFER_STATIC_LIBS=ON -DGMX_FFT_LIBRARY=mk1
```

```
make -j 32
```

```
make install
```

5.2 GPU version

GPU version is available in only single precision.

Intel 2015 compilers did not support Cuda-6.x, but Cuda-7.x.

```
module load gnu/4.9.3
```

```
module load cmake/3.2.3
```

```
source /opt/intel/composer_xe_2015.3.187/bin/compilervars.sh intel64
```

```
CVER=7.0
```

```
module use /application/hpcx-v1.3.336-icc-MLNX_OFED_LINUX-3.0-1.0.1-redhat6.5-x86_64/modulefiles
```

```
module load hpcx-ompi-mellanox-v1.8_vanilla_cuda$CVER
```

```
export OMPI_MPICXX=icpc
```

```

export OMPI_MPICC=icc
export OMPI_MPIF77=ifort
export OMPI_MPIF90=ifort
export PATH=/usr/local/cuda-$CVER/bin:$PATH
export CFLAGS="-I/usr/local/cuda-$CVER/include"
export LDFLAGS="-L/usr/local/cuda-$CVER/lib64"
export LD_LIBRARY_PATH=/usr/local/cuda-$CVER/
lib64:$LD_LIBRARY_PATH
BASEDIR=$PWD
rm -fr $BASEDIR/build
mkdir $BASEDIR/build
cd $BASEDIR/build
FLAGS="-DGMX_GPU=ON "
FLAGS+="-DGMX_DOUBLE=OFF "
FLAGS+="-DCUDA_TOOLKIT_ROOT_DIR=/usr/local/
cuda-$CVER "
FLAGS+="-DCMAKE_INSTALL_PREFIX=$BASEDIR/
install/GPU-mpi.SP "
cmake .. $FLAGS -DGMX_MPI=ON -DGMX_
CPU_ACCELERATION=AVX2_256 -DGMX_FFT_
LIBRARY=fftw3 -DCMAKE_C_COMPILER=icc
-DCMAKE_CXX_COMPILER=icpc -DGMX_PREFER_
STATIC_LIBS=ON -DGMX_FFT_LIBRARY=mkl
make -j 32
make install

```

5.3 Build OpenMPI with CUDA support

```

export HPCX_HOME=/application/hpcx-v1.3.336-icc-
MLNX_OFED_LINUX-3.0-1.0.1-redhat6.5-x86_64
module use $HPCX_HOME/modulefiles
module load hpcx
source /opt/intel/composer_xe_2015.3.187/bin/com-
pilervars.sh intel64
CVER=7.0
export PATH=/usr/local/cuda-$CVER/bin:$PATH
export LD_LIBRARY_PATH=/usr/local/cuda-$CVER/
lib64:$LD_LIBRARY_PATH
./configure --prefix=${HPCX_HOME}/ompi-mellanox-
v1.8_vanilla_cuda$CVER \
    --with-slurm --with-pmi \
    --with-cuda=/usr/local/cuda-$CVER
make all
make install

```

6. Running Gromacs with CPU or with GPU

CPU benchmark

```

-----
module load gnu/4.9.3
source /opt/intel/composer_xe_2015.3.187/bin/com-
pilervars.sh intel64
source /opt/intel/impi/5.0.3.048/bin64/mpivars.sh
mpiexec -IB -genv MV2_USE_APM 0 -genv I_MPI_OFA_
ADAPTER_NAME mlx5_0 -genv I_MPI_OFA_NUM_
PORTS 1 -perhost 28 -ppn 28 -envall -np 112 <path to
Gromacs executables>/mdrun_mpi_d -s topol

```

GPU benchmark

```

-----
module load gnu/4.9.3
source /opt/intel/composer_xe_2015.3.187/bin/com-
pilervars.sh intel64
export LD_LIBRARY_PATH=/usr/local/cuda-7.0/bin:$LD_
LIBRARY_PATH
module use /application/hpcx-v1.3.336-icc-MLNX_
OFED_LINUX-3.0-1.0.1-redhat6.5-x86_64/modulefiles
module load hpcx-ompi-mellanox-v1.8_vanilla_cuda7.0
mpirun -hostfile hostfile.txt -np 8 --report-bindings
-mca btl_sm_use_knem 1 -mca coll_fca_enable 0 -mca
coll_hcoll_enable 0 -mca mtl ^mxm -mca pml ob1 -mca
btl openib,sm,self -mca btl_openib_if_include mlx5_0:1
--bind-to none <path to Gromacs executables>/mdrun_
mpi -s topol-pme

```



350 Oakmead Pkwy, Sunnyvale, CA 94085
Tel: 408-970-3400 • Fax: 408-970-3403
www.hpcadvisorycouncil.com