

Quantum Espresso 6.8 with GPU support

Webpage

<https://www.quantum-espresso.org/>

(source code was merged into main branch)

Version

6.8

Build Environment

- PGI 20.4
- CUDA 10.1 (bundled with PGI 20.4)
- MKL 2020.0.2 (intel 2020 update 2)
- OpenMPI 3.1.6

Files Required

- q-e-qe-6.8.tar.gz
 - patch_extlibs_makefile

```
--- install/extlibs_makefile.org    2021-12-15 12:32:45.000000000 +0900
+++ install/extlibs_makefile      2021-12-15 12:33:50.000000000 +0900
@@ -106,6 +106,7 @@
     --with-cuda-runtime=$(CUDA_RUNTIME) \
     --disable-parallel \
     --enable-cuda-env-check=no; \
+   sed -i -e "s/cc60/cc60,cc70/" make.inc include/configure.h install/make_lapack.inc; \
make all
touch ../install/libcuda_devxlib # do not download and configure again
```

- openmpi-3.1.6.tar.bz2
- (PBS Pro files under /local/apl/lx/pbs14)

Build Procedure

OpenMPI 3.1.6

```
#!/bin/sh

VERSION=6.8
FULLVER=${VERSION}
BASEDIR=/home/users/${USER}/Software/QE/${VERSION}
INSTDIR=/local/apl/lx/espresso68-gpu

# nvhpc openmpi awares cuda, but not tm (PBS)...
WORKDIR=/work/users/${USER}
OMPIVER=3.1.6
OMPITARBALL=/home/users/${USER}/Software/OpenMPI/${OMPIVER}/openmpi-${OMPIVER}.tar.bz2
OMPIROOT=${INSTDIR}/openmpi-${OMPIVER}
PBSROOT=/local/apl/lx/pbs14

PARALLEL=12

#export CUDA_HOME=/local/apl/lx/nvhpc-21.9/Linux_x86_64/21.9/cuda
export CUDA_HOME=/local/apl/lx/cuda-10.1

# -----
umask 0022

module purge
```

```

#module load nvhpc/21.9-nompi
module load pgi/20.4
module load mkl/2020.0.2

export LANG=C
export LC_ALL=C

ulimit -s unlimited

# build openmpi first
cd ${WORKDIR}
if [ -d openmpi-${OMPIVER} ]; then
  mv openmpi-${OMPIVER} openmpi-erase
  rm -rf openmpi-erase &
fi

tar jxf ${OMPITARBALL}
cd openmpi-${OMPIVER}

export CFLAGS="-fPIC"
export FCFLAGS="-fPIC"
export CXXFLAGS="-fPIC"
export LDFLAGS="-fPIC"

mkdir rccs && cd rccs
CC=pgcc CXX=pgc++ FC=pgf90 \
  ./configure --prefix=${OMPIROOT} \
    --with-tm=${PBSROOT} \
    --enable-mpi-cxx \
    --with-cuda=${CUDA_HOME} \
    --with-psm2
make -j ${PARALLEL} && make install && make check

```

■ tests

- atomic tests failed

QE

```

#!/bin/sh

VERSION=6.8
FULLVER=${VERSION}
BASEDIR=/home/users/${USER}/Software/QE/${VERSION}
TARBALL=${BASEDIR}/q-e-qe-${FULLVER}.tar.gz
INSTDIR=/local/apl/lx/espresso68-gpu

PATCH0=${BASEDIR}/patch_extlibs_makefile

# nvhpc openmpi awares cuda, but not tm (PBS)...
OMPIVER=3.1.6
OMPIROOT=${INSTDIR}/openmpi-${OMPIVER}

PARALLEL=12

#export CUDA_HOME=/local/apl/lx/nvhpc-21.9/Linux_x86_64/21.9/cuda
export CUDA_HOME=/local/apl/lx/cuda-10.1

# -----
umask 0022

module purge
#module load nvhpc/21.9-nompi
module load pgi/20.4
module load mkl/2020.0.2

```

```

export LANG=C
export LC_ALL=C

ulimit -s unlimited

# openmpi setting
export OMPI_MCA_btlib_ompi_ib=1
export CPATH="${OMPIROOT}/include:${CPATH}"
export LIBRARY_PATH="${OMPIROOT}/lib:${LIBRARY_PATH}"
export LD_LIBRARY_PATH="${OMPIROOT}/lib:${LD_LIBRARY_PATH}"
export PATH="${OMPIROOT}/bin:${PATH}"

# qe build
cd ${INSTDIR}
if [ -d q-e-qe-${FULLVER} ]; then
  mv q-e-qe-${FULLVER} q-e-qe-erase
  rm -rf q-e-qe-erase &
fi

tar xzf ${TARBALL}
cd q-e-qe-${FULLVER}
mv *.[a-zA-Z]* ./
cd ../ && rmdir q-e-qe-${FULLVER}

export MPIF90=mpif90

patch -p0 < ${PATCH0}
# complicated...
sed -i -e 's/cc$(GPU_ARCH)/cc60,cc70/' install/Makefile.lib_eigsolve

./configure --enable-openmp \
  --enable-openacc \
  --with-scalapack=no \
  --with-cuda=${CUDA_HOME} \
  --with-cuda-cc=60 \
  --with-cuda-runtime=10.1

# force to add curand to library list
sed -i -e "s/cusolver/cusolver,curand/" make.inc
# add cc70 (is it really ok?)
sed -i -e "s/cc60/cc60,cc70/" \
  make.inc \
  install/make_lapack.inc \
  install/make_wannier90.inc \
  include/configure.h

make -j${PARALLEL} all # neb may fail...

cd test-suite
make run-tests-serial
make clean
make run-tests-parallel
cd ..

```

Notes

- beef.in, beef-spin.in tests were failed (openacc version does not support beef). All the other tests were passed successfully.
- openacc was tentatively enabled
- **Though V100 flag (cc70) was specified, binaries do not work on V100. Please use P100 (jobtype=gpu).**
 - The reason is not yet clear. Even when `--with-cuda-cc=70` was specified (manipulations about cc were removed, of course), binaries didn't work at all on V100.
 - At least this is not due to cuFFT issue <https://gitlab.com/QEF/q-e/-/issues/315> (cufft version is ok and setting environment variable didn't change anything).
 - CUDA or GPU driver version might be related to this issue. Upgrading driver might fix the bug.

- By changing nvhpc 21.9 to pgi 20.4, MPI error (occasional?) and V100 error disappeared.
- OpenMPI 4.x failed on some of its own tests. We thus employ 3.x. (NVIDIA SDK's openmpi is also 3.x one.)
- benchmark result of ausurf system (final WALL value; -nk 1 -nb 1 -nt 1)
 - cpu 16 cores (same cpu as jobtype=small): 209.05s
 - P100*2 (cpu 2 cores and 2 gpus on single node): 39.70s
 - it does not work on P100*1 (16 GB) due to insufficient memory error.
 - V100*1 (cpu 1 core and 1 gpu on single node): 49.88s