

Quantum Espresso 6.7 with GPU support

Webpage

<https://www.quantum-espresso.org/>
<https://gitlab.com/QEF/q-e-gpu/-/releases>

Version

6.7-gpu

Build Environment

- PGI 20.4
- MKL 2019.0.5 (intel 2019 update 5)
- CUDA 10.1

Files Required

- q-e-gpu-qe-gpu-6.7.tar.gz
- openmpi-4.0.2.tar.bz2
- (PBS Pro files under /local/apl/lx/pbs14)

Build Procedure

```
#!/bin/sh

VERSION=6.7
FULLVER=${VERSION}
BASEDIR=/home/users/${USER}/Software/QE/${VERSION}-gpu
TARBALL=${BASEDIR}/q-e-gpu-qe-gpu-${FULLVER}.tar.gz
INSTDIR=/local/apl/lx/espresso67-gpu

# openmpi
WORKDIR=/work/users/${USER}
OMPIVER=4.0.2
OMPITARBALL=/home/users/${USER}/Software/OpenMPI/${OMPIVER}/openmpi-${OMPIVER}.tar.bz2
OMPIROOT=${INSTDIR}/openmpi-${OMPIVER}
PBSROOT=/local/apl/lx/pbs14

PARALLEL=12

# -----
umask 0022

module purge
module load pgi/20.4
module load mkl/2019.0.5

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}

mkdir rccs && cd rccs
```

```

CC=pgcc CXX=pgc++ FC=pgfortran \
./configure --prefix=${OMPIROOT} \
    --with-tm=${PBSROOT} \
    --enable-mpi-cxx \
    --enable-mpi1-compatibility \
    --with-psm2
make -j ${PARALLEL} && make install && make check

# openmpi setting
export OMPI_MCA_btl_openib_allow_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 ${INSTDIR}
if [ -d q-e-gpu-qe-gpu-${FULLVER} ]; then
    mv q-e-gpu-qe-gpu-${FULLVER} q-e-gpu-qe-gpu-erase
    rm -rf q-e-gpu-qe-gpu-erase &
fi

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

export MPIF90=mpif90
export MPIFC=mpif90
export MPIF77=mpif90
export MPICC=mpicc
export MPICXX=mpicxx

# --with-cuda should point cuda bundled with pgi... but i couldn't do it...
FC=pgfortran F90=pgfortran F77=pgfortran CC=pgcc CXX=pgc++ \
./configure --enable-openmp \
    --enable-parallel \
    --with-cuda=/local/apl/lx/cuda-10.1 \
    --with-cuda-cc=60 \
    --with-cuda-runtime=10.1

# force to add cc70 support
sed -i -e "s/cc60/cc60,cc70/" make.inc

make -j${PARALLEL} pw cp
cd test-suite

make run-tests-pw-serial
make run-tests-cp-serial
make clean
make run-tests-pw-parallel
make run-tests-cp-parallel
cd ..

```

Notes

- There were errors on atomic_cmpset* tests of OpenMPI. However, all the QE tests have passed successfully.
- Many functions of pw.x seem to be supported by GPU.
- (GPU version of cp does not seem to be very useful for now.)
 - Limited parts such as fft can be performed by GPUs.
 - Conjugate gradient (cg) is not yet supported.
- OpenMP is enabled in this build.
- It works both on P100 and V100.
- On MPI parallel, single GPU would be assigned to single process.
 - Processes can share single GPU, but it may not be very effective. (Two processes on single GPU might be

advantageous in some situation.)

- Please use mpirun in /local/apl/lx/espresso67-gpu/openmpi-4.0.2 when you perform MPI parallel runs. (Please check sample job script under espresso67-gpu/samples/.)
- Some of calculations might not be supported by GPU. Please check official document or the output file (timing information at the end).
- Very small calculation may not be significantly accelerated by GPU.
- Large calculation may fail due to memory allocation error (probably of GPU memory). This error may be avoided by using multiple GPUs with MPI parallel run.