LAMMPS 29Oct20 with GPU

Webpage

http://lammps.sandia.gov/

Version

29Oct20

Build Environment

- Intel Parallel Studio XE 2018 update 4
- CUDA 11.1 Update 1
- cmake 3.16.3

Files Required

- lammps-29Oct20.tar.gz (renamed lammps-stable.tar.gz)
- (some of files will be downloaded during installation)

Build Procedure

```
#!/bin/sh

VERSION=29Oct20
NAME=lammps-${VERSION}
INSTALL_PREFIX=/local/apl/lx/lammps${VERSION}-CUDA

BASEDIR=/home/users/${USER}/Software/LAMMPS/${VERSION}
LAMMPS_TARBALL=${BASEDIR}/${NAME}.tar.gz

WORKDIR=/work/users/${USER}
LAMMPS_WORKDIR=${WORKDIR}/${NAME}

GPU_ARCH=sm_60
CUDA_ROOTDIR=/local/apl/lx/cuda-11.1
VMD_MOLFILE_INC=/local/apl/lx/vmd193/lib/plugins/include

PARALLEL=12

#----------------------------------------------------------------
umask 0022
export LANG=C

module purge
module load intel_parallelstudio/2018update4
module load cuda/11.1
module load cmake/3.16.3

export CC=mpiicc
export CXX=mpiicpc
export FC=mpiifort
export MPICC=mpiicc
export MPICXX=mpiicpc
export MPIFC=mpiifort

cd ${WORKDIR}
if [ -d ${NAME} ]; then
  mv ${NAME} lammps_erase
  rm -rf lammps_erase &
fi

tar zxf ${LAMMPS_TARBALL}

cd ${NAME}
mkdir build && cd build

# Disabled PKGs:
# ADIOS, QUIP, QMMM, VTK: noavail
# MSCG: gsl too old
# KOKKOS: not sure what this is
# MESSAGE: ZeroMQ support not enabled

cmake ../cmake \
  -DLAMMPS_MACHINE=rccs-cuda \
  -DENABLE_TESTING=on \
  -DCMAKE_INSTALL_PREFIX=${INSTALL_PREFIX} \
  -DCMAKE_C_COMPILER=mpiicc \
  -DCMAKE_CXX_COMPILER=mpiicpc \
  -DCMAKE_Fortran_COMPILER=mpiifort \
  -DCMAKE_MPI_C_COMPILER=mpiicc \
  -DCMAKE_MPI_CXX_COMPILER=mpiicpc \
  -DCMAKE_MPI_Fortran_COMPILER=mpiifort \
  -DBUILD_SHARED_LIBS=on \
```

```
-DBUILD_TOOLS=on \
-DBUILD_MPI=on \
-DBUILD_OMP=on \
-DFFT=MKL \
-DFFT_SINGLE=on \
-DDOWNLOAD_LATTE=on \
-DDOWNLOAD_SCAFACOS=on \
-DDOWNLOAD_VORO=on \
-DDOWNLOAD_PLUMED=on \
-DDOWNLOAD_EIGEN3=on \
-DMOLFILE_INCLUDE_DIRS=${VMD_MOLFILE_INC} \
-DWITH_JPEG=yes \
-DWITH_PNG=yes \
-DWITH_GZIP=yes \
-DPKG_OPT=on \
-DPKG_USER-OMP=on \
-DPKG_USER-INTEL=on \
-DPKG_GPU=on \
-DGPU_API=cuda \
-DGPU_ARCH=${GPU_ARCH} \
-DCUDA_TOOLKIT_ROOT_DIR=${CUDA_ROOTDIR} \
-DPKG_KOKKOS=off \
-DPKG_ASPHERE=on \
-DPKG_BODY=on \
-DPKG_CLASS2=on \
-DPKG_COLLOID=on \
-DPKG_COMPRESS=on \
-DPKG_CORESHELL=on \
-DPKG_DIPOLE=on \
-DPKG_GRANULAR=on \
-DPKG_KSPACE=on \
-DPKG_LATTE=on \
-DPKG_MANYBODY=on \
-DPKG_MC=on \
-DPKG_MESSAGE=on \
-DPKG_MISC=on \
-DPKG_MLIAP=on \
-DPKG_MOLECULE=on \
-DPKG_MPIIO=on \
-DPKG_PERI=on \
-DPKG_POEMS=on \
-DPKG_PYTHON=on \
-DPKG_QEQ=on \
-DPKG_REPLICA=on \
-DPKG_RIGID=on \
-DPKG_SHOCK=on \
-DPKG_SNAP=on \
-DPKG_SPIN=on \
-DPKG_SRD=on \
-DPKG_USER-ATC=on \
-DPKG_USER-AWPMD=on \
-DPKG_USER-BOCS=on \
-DPKG_USER-CGDNA=on \
-DPKG_USER-CGSDK=on \
-DPKG_USER-COLVARS=on \
-DPKG_USER-DIFFRACTION=on \
-DPKG_USER-DPD=on \
-DPKG_USER-DRUDE=on \
-DPKG_USER-EFF=on \
-DPKG_USER-FEP=on \
-DPKG_USER-H5MD=on \
-DPKG_USER-LB=on \
-DPKG_USER-MANIFOLD=on \
-DPKG_USER-MEAMC=on \
-DPKG_USER-MESODPD=on \
-DPKG_USER-MESONT=on \
-DPKG_USER-MGPT=on \
-DPKG_USER-MISC=on \
-DPKG_USER-MOFFF=on \
-DPKG_USER-MOLFILE=on \
-DPKG_USER-NETCDF=on \
-DPKG_USER-PHONON=on \
-DPKG_USER-PLUMED=on \
-DPKG_USER-PTM=on \
-DPKG_USER-QTB=on \
-DPKG_USER-REACTION=on \
-DPKG_USER-REAXC=on \
```

```
  -DPKG_USER-SCAFACOS=on \
  -DPKG_USER-SDPD=on \
  -DPKG_USER-SMD=on \
  -DPKG_USER-SMTBQ=on \
  -DPKG_USER-SPH=on \
  -DPKG_USER-TALLY=on \
  -DPKG_USER-UEF=on \
  -DPKG_USER-YAFF=on \
  -DPKG_VORONOI=on \
  -DBLAS_LIBRARIES="-mkl" \
  -DCMAKE_BUILD_TYPE=Release

make -j ${PARALLEL}

# to avoid error on COMPILER_SUPPORTSfast=2:INTERNAL= line
# (cuda version has additional problem...)
sed -i -e "s/line\.split.*/line\.rsplit('=',1)/" \
       -e "s/\].split.*/].split(':',1)/" \
    ../unittest/python/python-capabilities.py \
    ../unittest/python/python-pylammps.py

make test # will put error...
make install

cp -a ../examples ${INSTALL_PREFIX}

# no extra tests... current test suite is designed for developers, not for us
cd ${INSTALL_PREFIX}
for f in etc/profile.d/*; do
  ln -s $f .
done
```

Notes

- Built on ccgpup. (Access to external web site is now allowed!)
- Both of P100 and V100 are available for this binary (lmp_rccs-cuda).
- Unittest results are copied to /local/apl/lx/lammps29Oct20/Testing. There are very limited number of unittests for GPU build.
- Please also check information about CPU version.