

I/O libraries Exercises

All the source code for this practice session are located under `/lrz/sys/courses/PPHPS19-advanced/I0-LIB`

Output files for all exercises can use the `$SCRATCH` or `$WORK_MPIIO` file space in CoolMUC3.

The following module must be used to the different I/O libraries.

- For HDF5 :

```
module load hdf5/mpi/1.8.20_SLES12_3
or
module load hdf5/serial/1.8.20_SLES12_3
```

Compilation and Linkage can be done with

```
h5pfc -o hyperslab_by_col-f90 hyperslab_by_col-test.f90
h5fc -o hyperslab h5_hyperslab_1.f90
```

For Parallel version you can replace the Makefile with Makefile-f90 or Makefile-C, then you use make command.

- For SIONLib:

```
module load sionlib/1.7.2
```

Compilation and Linkage can be done with a Makefile.

Example:

```
include Makefile.defs
LDLFLAGS += `sionconfig --libs --64 --mpi -be`
CFLAGS += `sionconfig --cflags --64 --mpi -be`
# -- Rules --
all: sionlib-parallel-example
clean:
    rm -rf sionlib-parallel-example a.* *.o
sionlib-parallel-example.o: sionlib-parallel-example.c
    $(MPICC) sionlib-parallel-example.c -c $(CFLAGS)
sionlib-parallel-example: sionlib-parallel-example.o
    $(MPICC) sionlib-parallel-example.o -o sionlib-parallel-example $(LDLFLAGS)
```

Perform the following exercises

A. Serial HDF5

The program `h5_hyperslab_1.f90` write and read a hyperslab. A subset is write with the value "5", that is selected using:

```
start = (/ 1, 2 /)
count = (/ 3, 4 /)
stride = (/ 1, 1 /)
```

A.1 Compile and run the program. Use `h5dump` to verify the data.

A.2 Change the source code to write a subset of 4x2 that starts in 2,3 with the value "10". Verify the file data using h5dump.

For running:

```
module load intel
module load mpi.intel/2017
module use -a /lrz/sys/share/modules/extfiles
module load hdf5/serial/1.8.20_SLES12
export LD_LIBRARY_PATH=/lrz/sys/io_tools/zip/2.1.1-SLES12_3/lib:
$LD_LIBRARY_PATH
$PATH_EXECUTABLE/hyperslab
```

B. Parallel HDF5

The program `hyperslab_by_col.f90` does parallel I/O in a contiguous hyperslab (from tutorial PHDF5 <http://www.hdfgroup.org/HDF5/Tutor/phycecont.html>). The example provided only work up to 8 processes. Change the dataset dimensions to work up to 64 processes.

The example is explained bellow to help you understand the problem.

The example has a dataset of 5 (rows) x 8 (columns). Since a contiguous hyperslab in Fortran 90 consists of consecutive columns, each process will be writing an equal number of columns to the dataset. The size of the hyperslab to write to the dataset can be defined as follows:

$$\text{count}(1) = \text{dimsf}(1)$$

$$\text{count}(2) = \text{dimsf}(2) / \text{number_of_processes}$$

where, $\text{dimsf}(1)$ is the number of rows in the dataset and $\text{dimsf}(2)$ is the number of columns

The offset for the hyperslab dimension would be different for each process:

$$\text{offset}(1) = 0$$

$$\text{offset}(2) = k * \text{count}(2)$$

where, $\text{offset}(1) = 0$ indicates to start at the beginning of the column. "k" is the process id number.

$\text{count}(2)$ is the number of columns to be written by each hyperslab.

The number of processes that could be used in this example are 1, 2, 4, or 8.

By using 4 processes, the `offset` and `count` parameters for Process 1 would look like Figure 1.




Figure 1: offset and count on file logical view for Process 1.

For running:

```
module load intel
module load mpi.intel/2017
module load hdf5/mpi/1.8.20_SLES12
export LD_LIBRARY_PATH=/lrz/sys/io_tools/zip/2.1.1-SLES12_3/lib:
$LD_LIBRARY_PATH
```

```
cd $SCRATCH/HDF5-Test
mpiexec -np 4 $PATH_EXECUTABLE/Hyperslab_by_row
```

Exercises: SIONlib

A. The program `sionpar.c` uses SIONLib library in its parallel version. In this example, each MPI process writes and reads to/from a shared file.

A.1 The following variable controls the I/O pattern in the sion file.

- `BUFSIZE`: define the size of buffer local used for a MPI task.
- `chunksize`: define the size of chunk on the sion file.
- `fsblksize`: define the block size of the file system.

Compile and run the program. Analyze the metadata information of the `parfile.sion` (output file) with `siondump`. What are the values reported for the following metrics?

1. file usage (user data/chunk size)
2. file usage (used fsblocks/chunk size)
3. file usage (user data/used fsblocks)

Change the source code to increase the percentage of these metrics to 100%. These metrics are related with the variables `BUFSIZE`, `chunksize` and `fsblksize`.

A.2 Change the source code to write into 2 files instead of a shared file.

For running:

```
module load intel
module load mpi.intel/2017
module use -a /lrz/sys/share/modules/extfiles
module load sionlib/1.7.2
cd $SCRATCH/SIONLib-Test
mpiexec -np 16 $PATH_EXECUTABLE/sionlib-parallel
```