

POP

POP 2.0.1 is the Los Alamos National Laboratory's Parallel Ocean Program. POP is a component of the NCAR Community Climate System Model. The benchmark includes two configurations, of which we will only be running the

7-day 1-degree production benchmark.

Building POP requires GNU make, MPI, a C compiler, a Fortran90 compiler, and the NetCDF libraries.

POP can be downloaded from <http://web.ncar.teragrid.org/~bmayer/benchmarks/pop.tar.gz>

1 Procedure

Follow these steps to build the POP benchmark.

1. Set the environment variable ARCHDIR to the correct value for your environment and run type. This will cause the file **input_templates/\$ARCHDIR.gnu** to be copied by the setup script and the compile settings in that file to be used at compile time. Predefined options can be found in the **input_templates/.gnu*** files. If a **\$ARCHDIR.gnu** file does not exist for the desired architecture it will have to be created prior to compilation.
2. Use the **setup** script to create the run directory

For the 1 degree case type:

```
./setup_run_dir run_dir.x1
```

This will create a directory for the 1-degree run called **run_dir.x1** and copy the files necessary to compile and run the model into that directory.

1. Change into the **data** subdirectory. The **Makefile** in this directory builds a data conversion utility to convert the benchmark data into native binary format. Edit the **Makefile** to set the correct compile options for your environment.

For the 1-degree case type:

```
make x1 RUNDIR=./run_dir.x1
```

This will convert the necessary data files to native binary format and copy all needed data and source files into the run directory for the desired case.

1. Change into the run directory for the desired case (**./run_dir.x1 for 1-degree**) and edit the file **\$ARCHDIR.gnu** and change any necessary compile time options (compilers, optimization flags, link flags etc).
2. Type **gmake** to build the **pop** benchmark. The executable will be placed in the current directory and will be named **pop**.
3. The **run** directory should now contain the executable, the input namelist file (**pop_in**) and all the data necessary to run **pop**. Edit the **pop_in** file to set the number of processors you wish to use by setting the lines:

```
nprocs_clinic=NP
```

```
nprocs_tropic=NP
```

where NP is the desired number of processors

1. Use **mpirun** to run the POP benchmark, e.g.:

```
mpirun --np 32 ./pop
```

Note: Each time **pop** is to be run with a different number of processors the namelist input file **pop_in** must be edited to reflect the desired number of processors. For certain processor numbers the values of **max_blocks_clinic** and **max_blocks_tropic** may need to be adjusted in the **domain_size.F90** file (located in the **run** directory for each test case). If the values are too small, a message will be printed to standard out at run time stating what the values should be set to. The values should be adjusted accordingly and the code will need to be recompiled prior to running. The number of blocks assigned to each processor can be adjusted by changing the block sizes in the file **domain_size.F90**. the number of blocks will be :

```
(nx_global/block_size_x)*(ny_global/block_size_y).
```

On some systems the best performance is found using one block per processor. The value can be tuned to find the best performance on the target system. The values should be adjusted accordingly and the code recompiled prior to each run.

2 Validation

Numerical validation can be checked by examining the value of "mean K.E." printed out at the end of the run. This number should agree to at least 5 decimal digits with the baseline values obtained on NCAR's IBM POWER4 (**bluesky**) system, which are
0.965524462029654 for the 1-degree run
0.261029444867348 for the 0.1-degree run

3 Data

Timing information is written to standard output at the end of the run. The most important timing number is timer number 11 which represents the wallclock execution time less initialization overhead.

The POP production benchmark (1-degree, 7-day simulation) should be run using from 1 to the maximum number of processors in the system. All of the output should be saved and the timer number 11 value at the end of each run should be recorded in the benchmark results spreadsheet.