# Performance Analysis

Application performance: NAMD, WRF

The report contains performance analysis of NAMD and WRF application using input data provide by Bergen research groups.



2010

Alexander Oltu UniBCCS

# Content

# 1 Introduction

These performance studies show how well WRF and NAMD performs and scales on the NOTUR systems. The real input data for these runs have been provided by Bergen research groups and they reflect usual runs on hexagon.

These two applications have used a high amount of the NOTUR CPU-hours in Bergen.

# 2 Performance studies

## 2.1 NAMD

The input data was provided by the Nathalie Reuter group from UIB/UniBCCS. The following two tests have been executed.

prot_water, is a system with 48816 atoms

memb, is a system with 132998 atoms

For both test cases the “numsteps” option have been set to 10000.

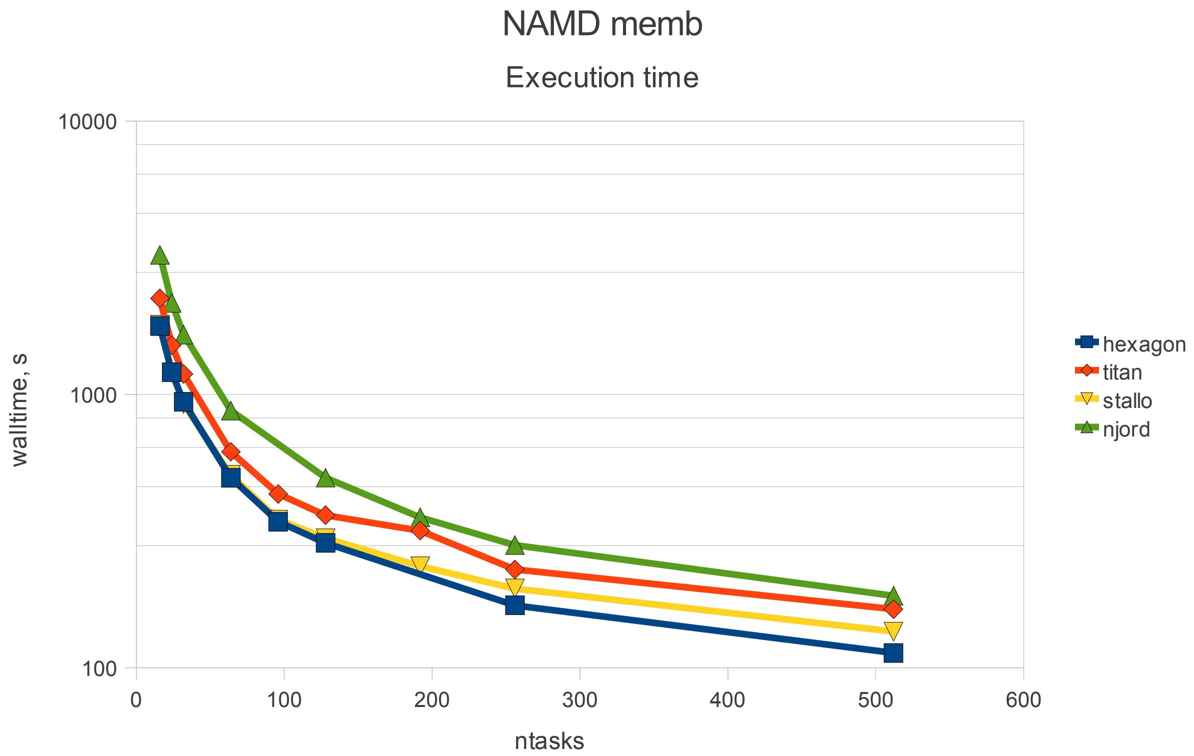


In the “memb” test on the low number of tasks hexagon and stallo has approximately the same performance. Whereas with increasing amount of tasks, i.e. the MPI communications take more and more of execution time, hexagon shows the highest performance. This can also be seen on the next graph.

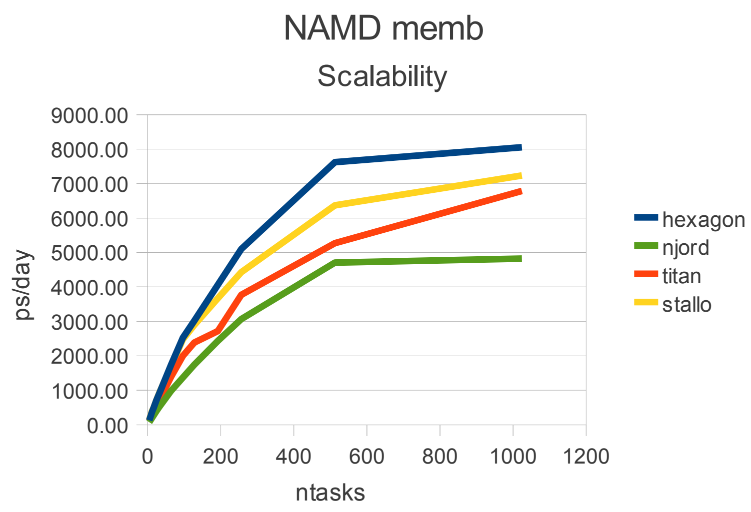


Approximately  the same can be said to the “prot water” test.

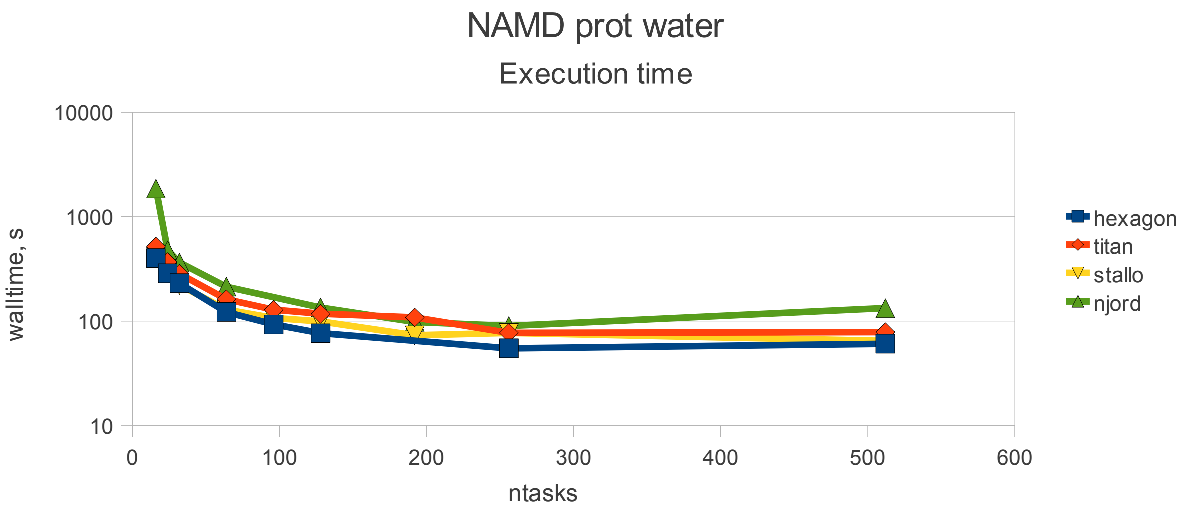


The “prot water” case has a smaller problem size (48816 atoms instead of 132998 of “memb”), the scaling stops  after 256 number of tasks.

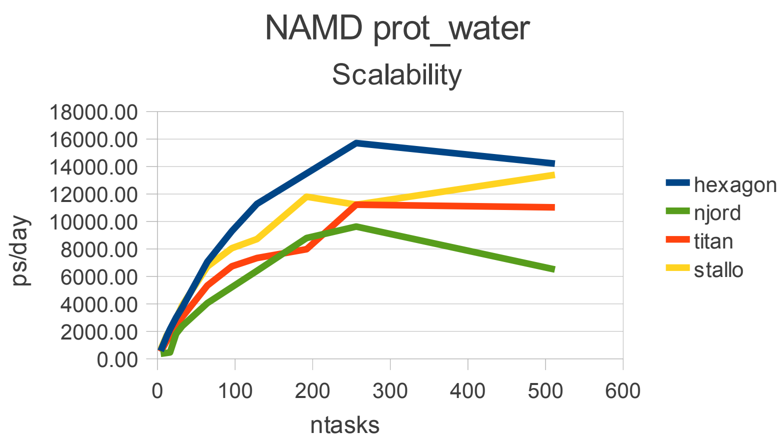

To check for any dependency on file update interval, both test cases have been restarted with values of 200, 500 and 1000 for the XSTFreq, restartfreq, dcdfreq options.

| XSTFreq, restartfreq, dcdfreq | walltime, s | comm, % | walltime, s | comm, % |
|---|---|---|---|---|
| 200 | 956.801959 | 7.8 | 245.756704 | 16.13 |
| 500 | 939.218634 | 6.69 | 230.632549 | 12.95 |
| 1000 | 956.962311 | 7.85 | 245.219785 | 16.01 |

## 2.2 WRF

The input data was provided by the Bjerknes center, Bergen. The data is "Tropical Channel, 50 km".

Unfortunately the test case can run only on njord and hexagon, on stallo and titan the test fails with a "Segmentation fault" error.

The "simulation speed" in the following graph is the model time step, 150 seconds, divided by average time per time step.

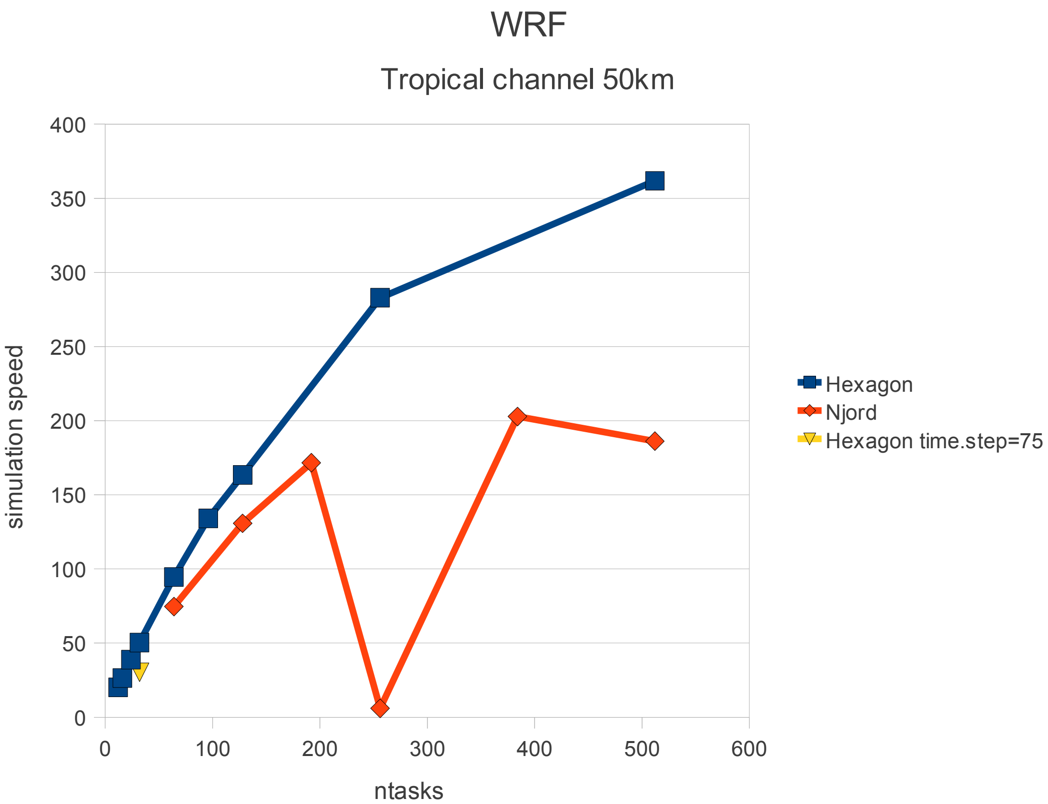


The default "time_step" was set 150. Additional two runs have been performed on hexagon using 32 tasks with time_step equal to 75 and 300. Time step of 75 resulted in slower run and time step 300 can't run at all, with crashing WRF after "Writing auxhist1_d01_2005-02-02_06:35".

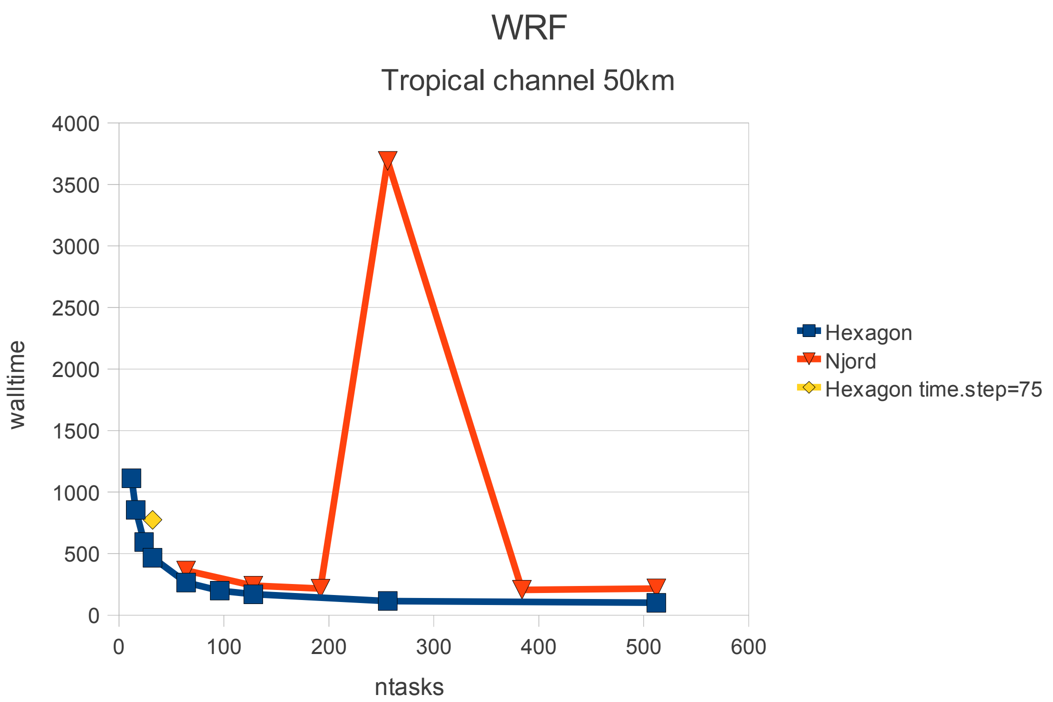


The big performance decrease with 256 tasks on the njord system can be related to that the job has been suspended and resumed.

To measure stallo and titan performance the official WRF benchmark "CONUS, 12km" was executed. The test on the titan system failed with a "Segmentation fault" error.

The importance of a good interconnect for the WRF application can be clearly seen in the following graph.

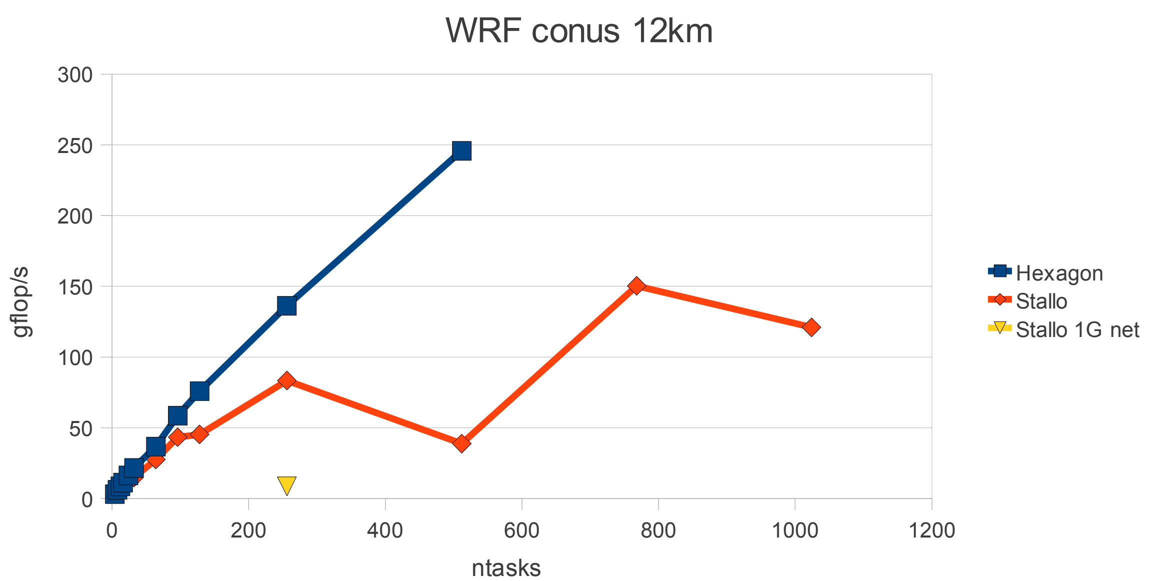


As an additional comparison the graph has one run using 32 tasks on the stallos 1Gbit interconnect.

# 3 Conclusion

For both applications the best scaling is offered by the hexagon system.

The NAMD “memb” with 132998 atoms can scale well on hexagon up to 1024 cores, whereas the optimal run is around 512 cores.

The NAMD “prot_water” with 48816 atoms do not scale beyond 256 cores, for this problem size the optimal number of cores will be around 128.

The stallo and titan machines are suitable for the smaller NAMD cases with lower atoms count or when a user wants to use up to a maximum of 100 cores.

The WRF “Tropical channel, 50km” will not give a big decrease in walltime when running it on more than 256 cores on hexagon.

In general, the WRF application requires a high speed, low latency interconnect.

Unfortunately, it was not possible to run any of the WRF benchmarks on the titan system.

# 4 Appendix A – build environment and results

## 4.1 NAMD

### 4.1.1 Hexagon

```
module switch PrgEnv-pgi PrgEnv-gnu
module load fftw/2.1.5.2 # compiles with fftw v2
export CC=cc && export CXX=CC && export FC=ftn
export MPICXX=CC

tar xf charm*.tar
cd charm*
./build charm++ mpi-crayxt --no-build-shared --with-production
cd mpi-crayxt/tests/charm++/megatest && make pgm
aprun -n4 ./pgm #(in another window with qsub -I -l mppwidth=4 :))

cd ../../../../../
vim arch/CRAY-XT.fftw (FFTDIR=$(FFTW_DIR)/../)
vim arch/CRAY-XT.tcl (TCLDIR=/work/apps/tcl/8.5.8-gnu and TCLLIB=-L$(TCLDIR)/lib -ltcl8.5 -ldl)
./config CRAY-XT-g++ --charm-arch mpi-crayxt
cd CRAY-XT-g++/
make -j4
aprun -n4 ./namd2
aprun -n4 ./namd2 src/alanin
```

### 4.1.2 Njord

```
Currently Loaded Modulefiles:
  1) modules            3) hpnssh/5.5p1-13v9   5) ploticus/2.41       7) tcltk/8.4.16
  2) local/1.0.3        4) hdf/5-1.6.5         6) ipm/0.982
```

FFTW:

```
./configure --prefix /home/ntnu/aoltu/benchmarks/notur2010/namd/NAMD_2.7b3_Source/fftw --enable-static --disable-shared --enable-mpi --disable-threads
 --enable-type-prefix --enable-float
```

make && make install

TCL 8.4:

```
./configure --prefix /home/ntnu/aoltu/benchmarks/notur2010/namd/NAMD_2.7b3_Source/tcl
--enable-static --disable-shared --disable-threads
make && make install
```

Charm++

```
MAKE=gmake  ./build charm++ mpi-sp -j4 -L$IPM_LIBPATH -lipm
```

NAMD

```
./config AIX-POWER-xlC --charm-arch mpi-sp --cxx-opts "-L$IPM_LIBPATH
-lipm" --cc-opts "-L$IPM_LIBPATH -lipm"

cd AIX-POWER-xlC && make -j4
```

### 4.1.3 Titan

```
Currently Loaded Modulefiles:
  1) modules            3) intellib/10.1        5) intel/11.1
  2) local/1.0.0        4) ipm/0.982            6) openmpi/1.4.2.intel

cd charm-6.2.1
env MPICXX=mpicxx ./build charm++ mpi-linux-x86_64 --no-build-shared --with-production
cd ..
wget http://www.ks.uiuc.edu/Research/namd/libraries/fftw-linux-x86_64.tar.gz
tar xzf fftw-linux-x86_64.tar.gz
```

```
mv linux-x86_64 fftw
wget http://www.ks.uiuc.edu/Research/namd/libraries/tcl-linux-x86_64.tar.gz
tar xzf tcl-linux-x86_64.tar.gz
mv linux-x86_64 tcl
./config Linux-x86_64-g++ --charm-arch mpi-linux-x86_64
cd Linux-x86_64-g++/
make -j4
```

### 4.1.4 Stallo

```
 Currently Loaded Modulefiles:
  1) intel-compiler/11.1   2) openmpi/1.3.3         3) system-manual/1.0     4) gold/2.1.5.0
5) netcdf/3.6.2          6) ipm/0.982

env MPICXX=mpicxx ./build charm++ mpi-linux-x86_64 --no-build-shared --with-production
```

The rest is same as for titan.

### 4.1.5 Results

| | hexagon | | | | njord | | | titan | | | stallo | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| nb (ps) | walltime | Comm | ps/day | speedup | walltime | ps/day | speedup | walltime | ps/day | speedup | walltime | Comm, | ps/day | speedup |
| **4** | | | | | | | | | | | | | | |
| 10 **memb** | 6717.137104 | 1.61 | 128.63 | 1.00 | 11008.88574 | 78.48 | 1.00 | 7997.98975 | 108.03 | 1.00 | 5895.346415 | 8.47 | 146.56 | 1.00 |
| 10 **prot_water** | 1486.80777 | 3.07 | 581.11 | 1.00 | 2414.296875 | 357.87 | 1.00 | 1745.19824 | 495.07 | 1.00 | 1279.192822 | 10.33 | 675.43 | 1.00 |
| **12** | | | | | | | | | | | | | | |
| 10 **memb** | 2338.73088 | 3.53 | 369.43 | 2.87 | | | | 2910.39111 | 296.87 | 2.75 | 2194.596982 | 13.02 | 393.69 | 2.69 |
| 10 **prot_water** | 528.154285 | 5.96 | 1635.89 | 2.82 | | | | 657.865601 | 1313.34 | 2.65 | 500.794209 | 16.92 | 1725.26 | 2.55 |
| **16** | | | | | | | | | | | | | | |
| 10 **memb** | 1777.765542 | 4.11 | 486.00 | 3.78 | 3243.848389 | 266.35 | 3.39 | 2246.57227 | 384.59 | 3.56 | 1799.869698 | 16.73 | 480.03 | 3.28 |
| 10 **prot_water** | 400.867301 | 6.33 | 2155.33 | 3.71 | 1856.479492 | 465.40 | 1.30 | 518.057007 | 1667.77 | 3.37 | 413.99586 | 19.52 | 2086.98 | 3.09 |
| **24** | | | | | | | | | | | | | | |
| 10 **memb** | 1205.503433 | 4.86 | 716.71 | 5.57 | 2161.424805 | 399.74 | 5.09 | 1514.63696 | 570.43 | 5.28 | 1204.58925 | 15.73 | 717.26 | 4.89 |
| 10 **prot_water** | 287.245401 | 9.2 | 3007.88 | 5.18 | 477.051727 | 1811.12 | 5.06 | 368.728912 | 2343.18 | 4.73 | 294.688227 | 21.48 | 2931.91 | 4.34 |
| **32** | | | | | | | | | | | | | | |
| 10 **memb** | 939.218634 | 6.69 | 919.91 | 7.15 | 1656.099854 | 521.71 | 6.65 | 1188.15576 | 727.18 | 6.73 | 923.490034 | 16.54 | 935.58 | 6.38 |
| 10 **prot_water** | 230.632549 | 12.95 | 3746.22 | 6.45 | 363.507233 | 2376.84 | 6.64 | 283.669006 | 3045.80 | 6.15 | 222.234185 | 21.35 | 3887.79 | 5.76 |
| **64** | | | | | | | | | | | | | | |
| 10 **memb** | 495.818371 | 9.18 | 1742.57 | 13.55 | 873.35 | 989.29 | 12.61 | 616.926208 | 1400.49 | 12.96 | 507.671644 | 19.9 | 1701.89 | 11.61 |
| 10 **prot_water** | 122.193899 | 14.31 | 7070.73 | 12.17 | 213.684937 | 4043.34 | 11.30 | 161.886932 | 5337.06 | 10.78 | 128.66586 | 26.52 | 6715.07 | 9.94 |
| **96** | | | | | | | | | | | | | | |
| 10 **memb** | 342.05405 | 10.51 | 2525.92 | 19.64 | | | | 432.495514 | 1997.71 | 18.49 | 348.749884 | 20.54 | 2477.42 | 16.90 |
| 10 **prot_water** | 92.92371 | 21.71 | 9297.95 | 16.00 | | | | 128.307083 | 6733.84 | 13.60 | 107.183792 | 36.6 | 8060.92 | 11.93 |
| **128** | | | | | | | | | | | | | | |
| 10 **memb** | 286.345096 | 12.7 | 3017.34 | 23.46 | 496.308807 | 1740.85 | 22.18 | 361.772186 | 2388.24 | 22.11 | 297.936572 | 22.96 | 2899.95 | 19.79 |
| 10 **prot_water** | 76.652273 | 25.62 | 11271.68 | 19.40 | 135.116089 | 6394.50 | 17.87 | 117.754417 | 7337.30 | 14.82 | 99.123396 | 43.67 | 8716.41 | 12.91 |
| **192** | | | | | | | | | | | | | | |
| 10 **memb** | | | | | 355.256561 | 2432.05 | 30.99 | 317.476074 | 2721.46 | 25.19 | 235.215082 | 30.79 | 3673.23 | 25.06 |
| 10 **prot_water** | | | | | 98.164841 | 8801.52 | 24.59 | 108.448746 | 7966.90 | 16.09 | 73.193444 | 44.76 | 11804.34 | 17.48 |
| **256** | | | | | | | | | | | | | | |
| 10 **memb** | 169.307029 | 23.37 | 5103.15 | 39.67 | 281.410187 | 3070.25 | 39.12 | 228.913086 | 3774.36 | 34.94 | 195.079257 | 36.87 | 4428.97 | 30.22 |
| 10 **prot_water** | 55.000121 | 38.18 | 15709.06 | 27.03 | 89.70874 | 9631.17 | 26.91 | 76.984947 | 11222.97 | 22.67 | 76.993198 | 53.34 | 11221.77 | 16.61 |
| **512** | | | | | | | | | | | | | | |
| 10 **memb** | 113.346371 | 35.19 | 7622.65 | 59.26 | 183.567978 | 4706.70 | 59.97 | 163.786377 | 5275.16 | 48.83 | 135.665331 | 49.96 | 6368.61 | 43.46 |
| 10 **prot_water** | 60.766768 | 62.4 | 14218.30 | 24.47 | 132.947891 | 6498.79 | 18.16 | 78.400505 | 11020.34 | 22.26 | 64.541783 | 66.33 | 13386.68 | 19.82 |
| **1024** | | | | | | | | | | | | | | |
| 10 **memb** | 107.259763 | 58.79 | 8055.21 | 62.62 | 179.254761 | 4819.96 | 61.41 | 127.275063 | 6788.45 | 62.84 | 119.440568 | 65.59 | 7233.72 | 49.36 |
| 10 **prot_water** | 557.667257 | 91.59 | 1549.31 | 2.67 | 1806.513306 | 478.27 | 1.34 | 386.985962 | 2232.64 | 4.51 | 490.076372 | 90.24 | 1762.99 | 2.61 |

| | **memb** | | **prot_water** | |
|---|---|---|---|---|
| **XSTFreq, restartfreq, dcdfreq** | **walltime, s** | **comm, %** | **walltime, s** | **comm, %** |
| 200 | 956.801959 | 7.8 | 245.756704 | 16.13 |
| 500 | 939.218634 | 6.69 | 230.632549 | 12.95 |
| 1000 | 956.962311 | 7.85 | 245.219785 | 16.01 |

## 4.2 WRF

### 4.2.1 Hexagon

```
tar xf /work/aoltu/sources/WRFV3.2.TAR.gz
cd WRFV3/
cd tools/
mv gen_allocs.c{,.orig} && mv reg_parse.c{,.orig}
wget http://www.mmm.ucar.edu/wrf/src/fix/reg_parse.c && wget
http://www.mmm.ucar.edu/wrf/src/fix/gen_allocs.c
export WRFIO_NCD_LARGE_FILE_SUPPORT=1

module load ipm
module load netCDF libfast
export CC=cc && export CXX=CC && export FC=ftn
./configure arw

configure.wrf:
LDFLAGS_LOCAL   =       -L/work/apps/ipm/0.982-pgi/lib -lipm
OPTERON_TYPE    =      -tp barcelona-64
SFC             =       ftn
SCC             =       cc
CCOMP           =       cc
DM_FC           =       ftn
DM_CC           =       cc -I$(MPICH_DIR)/include -DMPI2_SUPPORT

./compile em_real >& compile-em_real.$(date +%F).log
```

### 4.2.2 Njord

```
tar xf /work/aoltu/sources/WRFV3.2.TAR.gz
cd WRFV3/
cd tools/
mv gen_allocs.c{,.orig} && mv reg_parse.c{,.orig}
wget http://www.mmm.ucar.edu/wrf/src/fix/reg_parse.c && wget
http://www.mmm.ucar.edu/wrf/src/fix/gen_allocs.c

module load netcdf
module load ploticus ipm
export NETCDF=/usr/local/netcdf/netcdf-3.6.2
./configure arw

configure.wrf:
LDFLAGS_LOCAL   =       -lmass -lmassv -L/usr/local/ipm/ipm-0.982/lib -lipm

./compile em_real >& compile-em_real.$(date +%F).log
```

### 4.2.3 Titan

```
module load netcdf.intel ipm papi openmpi/1.3.3.intel ibmpich intel
export NETCDF=/site/netcdf/AuthenticAMD-x86_64/ifort

configure.wrf:
LDFLAGS_LOCAL   =       -L/site/VERSIONS/ipm-0.982/lib -lipm
CFLAGS_LOCAL    =       -w -O3

./compile em_real >& compile-em_real.$(date +%F).log
```

### 4.2.4 Stallo

```
module load netcdf
module load ipm

export NETCDF=/global/apps/netcdf/3.6.2
```

```
configure.wrf:LDFLAGS_LOCAL    =       -ip -L/global/apps/ipm/0.982/lib -lipm
./compile em_real >& compile-em_real.$(date +%F).log
```

### 4.2.5 Results

Tropical channel

| | Hexagon | | | Njord | | | Hexagon time.step=75 | | |
|---|---|---|---|---|---|---|---|---|---|
| **nodes** | **Sim speed** | **walltime, s** | **comm, %** | **Sim speed** | **walltime, s** | **comm, %** | **Sim speec** | **walltime, s** | **comm, %** |
| **12** | 20.185681354 | 1112.48044 | 11.77 | | | | | | |
| **16** | 26.492012482 | 856.130858 | 13.67 | | | | | | |
| **24** | 38.850984717 | 595.393459 | 17.78 | | | | | | |
| **32** | 50.359465867 | 467.723182 | 20.01 | | | | 30.308725 | 775.3837 | 17.65 |
| **64** | 94.608393783 | 266.478776 | 29.98 | 74.6191439 | 363.441247 | 33.91 | | | |
| **96** | 134.15579781 | 199.108223 | 38.89 | | | | | | |
| **128** | 163.37110836 | 170.157562 | 44.16 | 130.79806 | 240.372272 | 49.1 | | | |
| **192** | | | | 171.605668 | 216.152691 | 60.82 | | | |
| **256** | 282.90783994 | 116.07438 | 60.17 | 5.9907309 | 3691.143565 | 97.94 | | | |
| **384** | | | | 202.778882 | 206.795823 | 77.74 | | | |
| **512** | 361.74825698 | 101.752277 | 75.82 | 186.206794 | 217.899975 | 83.5 | | | |
| **1024** | | | | 107.548314 | 370.17078 | 91.38 | | | |
| **1536** | | | | 63.1301972 | 537.588955 | 95.08 | | | |
| **2048** | | | | 40.8472037 | 743.964248 | 95.82 | | | |

CONUS 12km

| **nodes** | **Hexagon Gflop/s** | **Stallo Gflop/s** | **Stallo 1G net Gflop/s** |
|---|---|---|---|
| **4** | 3.0823934965 | | |
| **8** | 6.1113171049 | | |
| **12** | 8.6553518609 | | |
| **16** | 11.194575303 | | |
| **24** | 16.454470552 | | |
| **32** | 21.552965037 | 14.953753966 | |
| **64** | 36.753355869 | 27.605082945 | |
| **96** | 58.591856402 | 43.398324112 | |
| **128** | 75.911436657 | 45.347223582 | |
| | | 66.666371317 | |
| **256** | 136.29447142 | 83.444237423 | 8.626224456 |
| **512** | 245.77593036 | 38.829189497 | |
| **768** | | 150.31390314 | |
| **1024** | | 120.97824514 | |