


VI-HPS

SOFTWARE



0.00 <<time step loop>>
0.00 updatedt
6.62 updatex
372.85 updateien
0.00 gene
0.00 <<iteration loop>>
293.65 genbc



PRODUCTIVITY

FAST SOLUTIONS

☒ PAPI_L1_DCM
☒ PAPI_L1_ICM
☐ PAPI_L2_DCM
☒ PAPI_L2_ICM
☒ PAPI_L2_TCM
☐ PAPI_L2_TCM

Demo example code: NPB-MZ-MPI / BT

VI-HPS Team

- The NAS Parallel Benchmark suite (MPI+OpenMP version)
 - Available from

<http://www.nas.nasa.gov/Software/NPB>

- 3 benchmarks in Fortran77
 - Configurable for various sizes & classes
- Move into the NPB3.3-MZ-MPI root directory

```
% cd Tutorial; ls
bin/      common/  jobscript/  Makefile  README.install  SP-MZ/
BT-MZ/    config/    LU-MZ/      README    README.tutorial  sys/
```

- Subdirectories contain source code for each benchmark
 - plus additional configuration and common code
- The provided distribution has already been configured for the tutorial, such that it's ready to “make” one or more of the benchmarks and install them into a (tool-specific) “bin” subdirectory

- What does it do?
 - Solves a discretized version of unsteady, compressible Navier-Stokes equations in three spatial dimensions
 - Performs 200 time-steps on a regular 3-dimensional grid
- Implemented in 20 or so Fortran77 source modules
- Uses MPI & OpenMP in combination
 - 4 processes with 4 threads each should be reasonable
 - don't expect to see speed-up when run on a laptop!
 - bt-mz_W.4 should run in around 13 seconds on a laptop
 - bt-mz_B.4 is more suitable for dedicated HPC compute nodes
 - Each class step takes around 10-15x longer

- Type “make” for instructions

```
% make
=====
=      NAS PARALLEL BENCHMARKS 3.3      =
=      MPI+OpenMP Multi-Zone Versions    =
=      F77                                =
=====
```

To make a NAS multi-zone benchmark type

```
make <benchmark-name> CLASS=<class> NPROCS=<nprocs>
```

where <benchmark-name> is “bt-mz”, “lu-mz”, or “sp-mz”
<class> is “S”, “W”, “A” through “F”
<nprocs> is number of processes

[...]

```
*****
* Custom build configuration is specified in config/make.def *
* Suggested tutorial exercise configuration for LiveISO/DVD:  *
*      make bt-mz CLASS=W NPROCS=4                            *
*****
```

Hint: the recommended build configuration is available via

% make suite

- Specify the benchmark configuration
 - benchmark name: **bt-mz**, lu-mz, sp-mz
 - the number of MPI processes: **NPROCS=4**
 - the benchmark class (S, W, A, B, C, D, E): **CLASS=W**

```
% make bt-mz CLASS=W NPROCS=4
cd BT-MZ; make CLASS=W NPROCS=4 VERSION=
make: Entering directory 'BT-MZ'
cd ../sys; cc -o setparams setparams.c
../sys/setparams bt-mz 4 W
mpif77 -c -O3 -fopenmp bt.f
[...]
cd ../common; mpif77 -c -O3 -fopenmp timers.f
mpif77 -O3 -fopenmp -o ../bin/bt-mz_W.4 \
bt.o initialize.o exact_solution.o exact_rhs.o set_constants.o \
adi.o rhs.o zone_setup.o x_solve.o y_solve.o exch_qbc.o \
solve_subs.o z_solve.o add.o error.o verify.o mpi_setup.o \
../common/print_results.o ../common/timers.o
Built executable ../bin/bt-mz_W.4
make: Leaving directory 'BT-MZ'
```

- Launch as a hybrid MPI+OpenMP application

```
% cd bin
% OMP_NUM_THREADS=4 mpiexec -np 4 ./bt-mz_W.4
NAS Parallel Benchmarks (NPB3.3-MZ-MPI) - BT-MZ MPI+OpenMP Benchmark
Number of zones:    4 x    4
Iterations:    200    dt:    0.000800
Number of active processes:    4
Total number of threads:    16    (  4.0 threads/process)

Time step    1
Time step    20
Time step    40
[...]
Time step   160
Time step   180
Time step   200
Verification Successful

BT-MZ Benchmark Completed.
Time in seconds = 5.57
```

Hint: save the benchmark output (or note the run time) to be able to refer to it later

- The tutorial steps are similar and repeated for each tool
- Edit [config/make.def](#) to adjust build configuration
 - Modify specification of compiler/linker: [MPIF77](#)
- Make clean and build new tool-specific executable

```
% make clean  
% make bt-mz CLASS=W NPROCS=4  
Built executable ../bin.$(TOOL)/bt-mz_W.4
```

- Change to the directory containing the new executable before running it with the desired tool configuration

```
% cd bin.$(TOOL)  
% export ...  
% OMP_NUM_THREADS=4 mpiexec -np 4 ./bt-mz_W.4
```

```
#           SITE- AND/OR PLATFORM-SPECIFIC DEFINITIONS
#-----
# Items in this file may need to be changed for each platform.
#-----
...
#-----
# The Fortran compiler used for MPI programs
#-----
MPIF77 = mpif77
```

Default (no instrumentation)

```
# Alternative variants to perform instrumentation
#MPIF77 = psc_instrument -u user,mpi,omp -s ${PROGRAM}.sir mpif77
#MPIF77 = tau_f90.sh
#MPIF77 = scalasca -instrument mpif77
#MPIF77 = vtf77 -vt:hyb -vt:f77 mpif77
#MPIF77 = scorep --user mpif77
```

Hint: uncomment one of these
alternative compiler wrappers
to perform instrumentation

```
# PREP is a generic preposition macro for instrumentation preparation
#MPIF77 = $(PREP) mpif77

# This links MPI Fortran programs; usually the same as ${MPIF77}
FLINK    = $(MPIF77)
...
```