



MPI Runtime Error Detection with MUST

NPB Hands-On

At the 30th VI-HPS Tuning Workshop

Joachim Protze
IT Center RWTH Aachen University
January 2019

Content

- Motivation
- MPI usage errors
- Hands-on
- Examples: Common MPI usage errors
 - Including MUST's error descriptions
- Correctness tools
- MUST usage
- **NPB Hands-on**

```
$ cp -r ~nct00009/tutorial/must ~/must-examples
$ source ~/must-examples/load-must.sh
```

Hands On – Build for MUST

- Go into the NPB directory
- Edit config/make.def or copy must.def from the must-example-dir
- Disable any other tool (i.e. use mpif77, unset PREP)
- Use intel tool chain
- Build:

```
FFLAGS = = -g -xCORE-AVX512 -mtune=skylake -xHost -O3 $(OPENMP)
...
MPIF77 = mpif77
```

```
% module load intel
% make clean
% make bt-mz NPROCS=6 CLASS=B
=====
=   NAS PARALLEL BENCHMARKS 3.3   =
=   MPI+OpenMP Multi-Zone Versions =
=   F77                           =
=====
```

```
cd BT-MZ; make CLASS=B NPROCS=6
make[1]: Entering directory
...
mpif77 -O3 -g -openmp -extend-source -o ../bin/bt-mz_B.6
bt_scorep_user.o ...
```

```
$ cp -r ~nct00009/tutorial/must ~/must-examples
$ source ~/must-examples/load-must.sh
```

Hands On - Prepare Job

- Create and edit the jobscript

```
$ cp -r ~nct00009/tutorial/must ~/must-examples
$ source ~/must-examples/load-must.sh
$ cd bin
$ cp ~/must-examples/must*.sbatch .
```

- Jobscript:

```
...
module use ~nct00009/.modules
module load must/intel
...
export OMP_NUM_THREADS=6
CLASS=B
NPROCS=6
...
mustrun --must:mpiexec mpiexec -n $NPROCS -t $OMP_NUM_THREADS $EXE
```

MUST needs one extra process!
We use 6 processes * 6 threads + 1 tool
process

```
$ cp -r ~nct00009/tutorial/must ~/must-examples  
$ source ~/must-examples/load-must.sh
```

Hands On – Executing with MUST

- Submit the jobscript:

```
sbatch must.sbatch
```

- Job output should read like:

```
...  
[MUST] Using prebuilt infrastructure at ../modules/mode1-layer2  
...  
[MUST] Executing application:  
  
NAS Parallel Benchmarks (NPB3.3-MZ-MPI) - BT-MZ MPI+OpenMP Benchmark  
...  
Total number of threads: 36 ( 6.0 threads/process)  
Calculated speedup = 35.84  
  
Time step 1  
...  
Verification Successful  
...  
[MUST] Execution finished, inspect "(...)/MUST_Output.html"!
```



```
$ cp -r ~nct00009/tutorial/must ~/must-examples
$ source ~/must-examples/load-must.sh
```

BT – MUST Results

- Open the MUST output: <Browser> MUST_Output.html

Rank(s)	Type	Message
0-3	Warning	<p>You requested 3 threads by OMP_NUM_THREADS=3 but the requested thread level MPI_THREAD_FUNNELED from the mpi library but thr library provides no thread support. This is ok as long as your application doesn't make use of OpenMP</p> <p>MPI_Init_thread (1st occurrence) called from: #0 MAIN_@bt.f:90 #1 main@bt.f:319</p>
0-3	Error	<p>There are 1 communicators that are not freed when MPI_Finalize was issued, a quality application should free all MPI resources before calling MPI_Finalize. Listing information for these communicators:</p> <p>-Communicator 1: Communicator created at reference 1 size=4</p> <p>Representative location: MPI_Comm_split (1st occurrence) called from: #0 MAIN_@bt.f:90 #1 main@bt.f:319</p> <p>References of a representative process: reference 1 rank 2: MPI_Comm_split (1st occurrence) called from: #0 MAIN_@bt.f:90 #1 main@bt.f:319</p>

BT-MZ should evaluate the "provided" thread level and don't use threads.

Resource leak:
A communicator created with MPI_Comm_split is not freed

Stacktraces in MUST

- We use an external lib for stacktraces
- This lib has officially no support for Intel compiler
 - But: in most cases it's compatible to icc compiled C applications
- ifort compiled FORTRAN applications can lead to segfault:
 - Use MUST w/o stacktraces for fortran applications
 - Use GNU compiler to build your application and use MUST w/ stacktraces
- Supposed your application has no faults you won't need stacktraces 😊

From
Representative location: MPI_Init_thread (1st occurrence) called from: #0 MAIN_@bt.f:90 #1 main@bt.f:319
Representative location: MPI_Comm_split (1st occurrence) called from: #0 MAIN_@bt.f:90 #1 main@bt.f:319

Rank(s)	Type	Message	From	References
	Information	MUST detected no MPI usage errors nor any suspicious behavior during this application run.		



Thank You