

Performance HPC Linux
Bull Echirolles

Workbook Profiling

C.Berthelot
Christophe.Berthelot@bull.net

Copyright (©) Bull S.A.S. 2008

- Equipment list
- OpenMP profiling and PAPI
- Profiling with HPCToolkit 1/2
- Profiling with HPCToolkit 2/2
- Profiling with vtune
- MPI profiling with MPIBULL2

- Equipment list
 - Hardware
 - Software
- OpenMP profiling and PAPI
- Profiling with HPCToolkit 1/2
- Profiling with HPCToolkit 2/2
- Profiling with vtune
- MPI profiling with MPIBULL2

Hardware

NovaScale 3045

- ▶ Number of processor: 4 CPUs Itanium 2 (8 cores)
- ▶ Frequency: 1.6Ghz
- ▶ Memory size: 32G

Networking equipment

- ▶ Infiniband network

Software

- ▶ Bull Linux AS4 V5.1 FIX10
- ▶ Compiler: Fortran 10.1
- ▶ Compiler: C 10.1
- ▶ MPI: mpibull2
- ▶ NPB benchmark

modules:

```
>module list  
Currently Loaded Modulefiles:  
1) oscar-modules/1.0.3
```

- Equipment list
- OpenMP profiling and PAPI
 - OpenMP profiling
 - OpenMP and PAPI
- Profiling with HPCToolkit 1/2
- Profiling with HPCToolkit 2/2
- Profiling with vtune
- MPI profiling with MPIBULL2

OpenMP Profiling

Compile

- ▶ `cd $HOME/TP/MandelOpenMP_PAPI`
- ▶ Edit Makefile to add openmp option to enables analysis
- ▶ `export OMP_NUM_THREADS=8`
- ▶ Run `bench mandel`
- ▶ Look profiling

OpenMP and PAPI

Add PAPI api

- ▶ Add PAPI api to calculate MFLOPS/threads
- ▶ Compile and link with papi and perfmon library
- ▶ Run bench mandel

Profiling with HPCToolkit 1/2

Compile

- ▶ `cd $HOME/TP/NPB3.2.1/NPB3.2-SER/`
- ▶ Compile with `-g` benchmark `cg`

Run

- ▶ run with `hpcrun ./cg.W` with default events (`PAPI_TOT_CYC`)

Analyze

- ▶ `hpcprof -e ./cg.W cg.W.PAPI_TOT_CYC.<hostname>.pid.`
`-n | more`

Profiling with HPCToolkit 2/2

Compile

- ▶ `cd $HOME/TP/NPB3.2.1/NPB3.2-SER/`
- ▶ Compile with `-g benchmark cg`

Run

- ▶ `bloop ./cg.W > cg.W.psxml`
- ▶ `hpcrun -e IA64_INST_RETIRED -e L3_MISSES -e PAPI_TOT_CYC ./cg.W > profile`
- ▶ `hpcquick -I<SRC_PATH> -S cg.W.psxml -P profile -n`
- ▶ `hpcview -o experiment-db hpcquick.xml`

Analyze

- ▶ `hpcviewer`

First use of Vtune: sampling

Run vtserver on a node

- ▶ Create a lsf job with a sleep and submit with ccc_msub
- ▶ Find allocate node
- ▶ Log on the allocate node
- ▶ Run vtserver -no-auth

First Sampling

- ▶ Create a sampling activity to look the system
- ▶ Time: 30s

Script

```
#!/bin/bash
#BSUB -J Vtune
#BSUB -i /dev/null
#BSUB -n 1
vtserver --no-auth &
sleep 40
killall vtlistenerd
```

Sampling

```
vtl activity -d 30 -c sampling,,NODE run
```

Benchmark BT from NAS

Compile

- ▶ `cd $HOME/TP/NPB3.2.1/NPB3.2-MPI`
- ▶ `make bt NPROCS=4 CLASS=B`

Run

- ▶ Create a lsf job to run bt.B.4 on one node and submit it with `ccc_bsub`

Record

- ▶ Time

Use Vtune to profile BT

Create sampling

- ▶ Time: more than BT time
- ▶ Events: FP_OPS_RETIRED, CPU_CYCLES-ALL and IA64_INST_RETIRED-THIS

Profile

- ▶ Run activity and vtune

Save results

- ▶ Save informations line for the hot spot

Script

```
#!/bin/bash
#BSUB -J BT-4
#BSUB -i /dev/null
#BSUB -n 4
vtserver -t --no-auth &
pdsh -w gallium3 vtl activity -c sampling,, $SLURM_NODELIST
-o "-ec en=FP_OPS_RETIRED en=CPU_OP_CYCLES-ALL
    en=IA64_INST_RETIRED-THIS" run &
srun -n 4 ./bt.B.4
killall vtlistenerd
```

Command

- ▶ vtl show
- ▶ vtl view aX::rY -processes
- ▶ vtl view aX::rY -pid 0XXXXX
- ▶ vtl view a1::r1 -pid 0XXXXX -mn bt.B.4 -hf

MPI profiling with MPIBULL2

Compile

- ▶ `cd $HOME/TP/NPB3.2.1/NPB3.2-MPI`
- ▶ Modify `conf/make.def` to add `mpianalyzer` lib
- ▶ `make clean; make bt NPROCS=9 CLASS=B`

Run

- ▶ Export variables
- ▶ Run it

Run readpfc

- ▶ Create outputs



Architect of an Open World™



- ▶ (c) Copyright Bull. All rights reserved
 - ✓ Users Restricted Rights - Use, duplication or disclosure restricted.
 - ✓ Any copy of these documents should keep all copyright, logos and other proprietary notices contained herein.
 - ✓ This publication may include technical inaccuracies or typographical errors.
 - ✓ This publication is provided "AS IS" without any warranty either expressed or implied including but not limited to the implied warranties of merchantabilities or fitness of the described product.
 - ✓ Course Material Licensing Terms : No sublicensing rights.
 - ✓ For other licensing needs, please contact Bull