

Optimization workshop

Dmitry Sergeev
Dmitry Mishura
Maxim Perminov
Intel

Software and Services Group

Copyright © 2011, Intel Corporation. All rights reserved.

Intel® and other Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States or other countries.

*Other brands and names are the property of their respective owners.



General optimization stages

- Compile & run the code
- Collect profile using IPTU
 - Collect basic profile
 - Collect bandwidth related events
 - Collect cache related events
- Collect profile using Intel Trace Analyzer

Compile and run the code

```
$ ssh compiler
```

```
$ mpienv -m impi-4.0.1
```

Modify your makefile for using Intel's compilers & mpi (mpiicc, mpiifort, icc & ifort for serial or OpenMP apps)

```
$ make
```

Submit the batch job to the queue:

```
$ sbatch -n <CPUs> -p test impi ./<application>
```

```
$ squeue
```

Modify your workload, to make the job running no longer than 10 minutes!

Collect profile using IPTU : basic

Create bash script in the app's folder ([get_sa.sh](#)):

```
#!/bin/bash
sa=10000000
out_dir=$1
shift
if [ "$PMI_ID" = "0" ]; then
/opt/intel/ptu40_005_lin_intel64/bin/vtsarun $out_dir -s \
-ec CPU_CLK_UNHALTED.THREAD:sa=$sa \
-ec INST_RETIRED.ANY:sa=$sa \
-ec DTLB_MISSES.ANY:sa=$sa \
-ec MEM_LOAD_RETIRED.LLC_MISS:sa=$sa \
-- "$@"
else
"$@"
fi
```

Collect profile using IPTU : basic

Create bash script in the app's folder (**get_ss.sh**):

```
#!/bin/bash
out_dir=$1
shift
if [ "$PMI_ID" = "0" ]; then
/opt/intel/ptu40_005_lin_intel64/bin/vtssrun $out_dir -- "$@"
else
"$@"
fi
```

Add execute attributes:

```
$ chmod +x get_sa.sh
```

```
$ chmod +x get_ss.sh
```

Collect profile using IPTU: basic

Submit the batch job to collect profile:

```
$ sbatch -n <CPUs> -p test impi ./get_sa.sh <prof_dir1> ./<application>
```

```
$ sbatch -n <CPUs> -p test impi ./get_ss.sh <prof_dir2> ./<application>
```

Look at <prof_dir> for collected profile:

```
$ls -la ./<prof_dir>
```

```
$/opt/intel/ptu40_005_lin_intel64/bin/vtsaview <prof_dir>
```

Options: -g m; -g p; -a e; -sea p

DO NOT cancel the job!

Software and Services Group

Copyright © 2011, Intel Corporation. All rights reserved.

Intel® and other Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States or other countries.

*Other brands and names are the property of their respective owners.

Collect profile using IPTU: bandwidth events

Total memory bandwidth

Modify the ./get_sa.sh script → ./get_sa_bnd1.sh:

```
#!/bin/bash
sa=10000000
out_dir=$1
shift
if [ "$PMI_ID" = "0" ]; then
/opt/intel/ptu40_005_lin_intel64/bin/vtsarun $out_dir -s \
-ec CPU_CLK_UNHALTED.THREAD:sa=$sa \
-ec INST_RETIRED.ANY:sa=$sa \
-ec OFFCORE_RESPONSE_0.DATA_IN.ANY_DRAM :sa=100000 \
-- "$@"
else
"$@"
fi
```


Collect profile using IPTU: bandwidth events

Remote DRAM bandwidth

Modify the `./get_sa.sh` script → `./get_sa_bnd2.sh`:

```
#!/bin/bash
sa=10000000
out_dir=$1
shift
if [ "$PMI_ID" = "0" ]; then
/opt/intel/ptu40_005_lin_intel64/bin/vtsarun $out_dir -s \
-ec CPU_CLK_UNHALTED.THREAD:sa=$sa \
-ec INST_RETIRED.ANY:sa=$sa \
-ec OFFCORE_RESPONSE_0.REMOTE_DRAM :sa=100000 \
-- "$@"
else
"$@"
fi
```

Collect profile using IPTU: bandwidth events

Submit the batch job to collect profile:

```
$ sbatch -n <CPUs> -p test impi ./get_sa_bnd1.sh <prof_dir>  
./<application>
```

```
$ sbatch -n <CPUs> -p test impi ./get_sa_bnd2.sh <prof_dir>  
./<application>
```

Look at <prof_dir> for collected profile:

```
$ls -la ./<prof_dir>
```

```
$/opt/intel/ptu40_005_lin_intel64/bin/vtsaview <prof_dir> -g m -sea e  
( ( 64 * <events> ) / CLK ) * 2.93 = <total read bandwidth> GB/s
```

PEAK: 36 GB/s

DO NOT cancel the job!

Collect profile using Intel Trace Analyzer

Recompile the application using `-trace` switch:

```
$mpicc -trace <app_files>
```

Submit the job to the queue as usual:

```
$sbatch -n <CPUs> -p test impi ./<application>
```

The trace profile will be located into app's folder (*.srt).