




ENERGY-EFFICIENT HPC APPLICATION EXECUTION USING MERIC

Ondřej Vysocký
IT4Innovations

5. 4. 2022

ENVIRONMENT



| **Use your logins for this training ONLY**

| **All compilations and executions must be done at a compute node allocated followingly**


```
$ qsub -q R603003 -l select=1:ncpus=24 -l walltime=02:00:00 -l msr=1.4.0  
-l hdeem=2.2.7-1 -I -A DD-22-26
```

| `$ /scratch/project/dd-22-26/install_MERIC-RADAR.sh`

| Prepares the necessary data into your \$HOME directory

- | Compiles MERIC
- | Copies RADAR visualizer
- | Copies test application - Kripke

SYSTEM INFO



```
$ bash
```

```
$ cd meric; source set_env
```

```
$ tools/systemInfo
```

SYSTEM INFORMATION

```
CPU name:          Intel(R) Xeon(R) Gold 6240 CPU @ 2.60GHz
Sockets per node:  2
Cores per socket:  18
Threads per core:  1
```

CPU FREQUENCIES

```
Current scaling driver:      acpi-cpufreq
```

```
Current scaling governor:    performance
```

```
Available governors:        conservative ondemand userspace powersave performance
```

```
Hardware controlled P-State: disabled
```

```
Turbo CPU core frequencies:  3900000(2) 3700000(4) 3600000(8) 3600000(12) 3400000(16) 3300000(18)
```

```
kHz(#cores)
```

```
Nominal CPU core frequency:  2600000 kHz
```

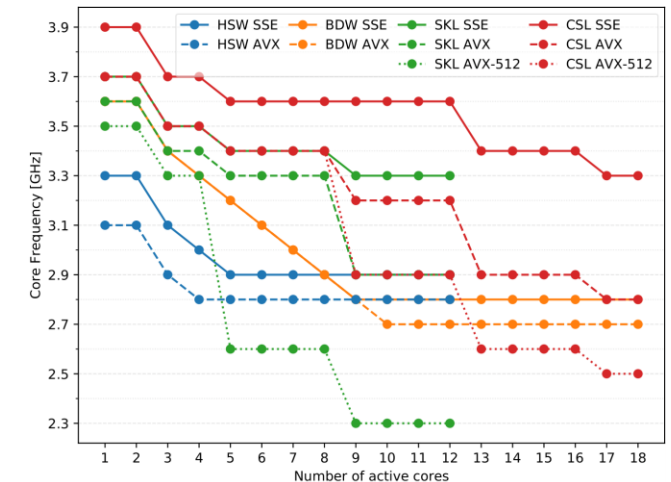
```
Min CPU core frequency:      1000000 kHz
```

```
Available CPU core frequencies:  2601000 2600000 2500000 2400000 2300000 2100000 2000000 1900000
1800000 1700000 1600000 1500000 1300000 1200000 1100000 1000000 kHz
```


```
Max CPU uncore frequency:    2400000 kHz
```

```
Min CPU uncore frequency:    1200000 kHz
```

```
...
```



SYSTEM INFO



...

CPU POWER LIMITS

RAPL time window unit: 976.562 us
PKG max power limit: 376 W
PKG min power limit: 69 W
DRAM max power limit: 26.75 W
DRAM min power limit: 4.5 W


DEFAULT CPU POWER LIMITS

PKG power limit #1: enabled + clamping disabled
PKG power limit #1: 150 W
PKG time window #1: 1 s
PKG power limit #2: enabled + clamping enabled
PKG power limit #2: 180 W
PKG time window #2: 0.000976562 s
DRAM power limit: disabled + clamping disabled
DRAM power limit: 0 W
DRAM time window: 0.000976562 s

AVAILABLE POWER MONITORING SYSTEMS

RAPL
HDEEM


PLAIN RUN



- | LLNL Kripke overview:
- | „Kripke is a simple, scalable, 3D Sn deterministic particle transport code.“
- | Compile Kripke benchmark and run it

```
$ cd $HOME/kripke-benchmark  
$ ./compile_for_0_default.sh  
$ ./run.sh
```


PLAIN RUN ENERGY CONSUMPTION



```
$ bash
$ source set_env
$ KRIPKE_COMMAND="mpirun -n 36 ./kripke --procs
3,3,4 --niter 10 --nest GZD --zones 40,40,40 --
legendre 8 --dset 32"

$ $MERIC_ROOT/tools/energyMeasureStart -e RAPL;
$KRIPKE_COMMAND;
$MERIC_ROOT/tools/energyMeasureStop -e RAPL
$ clearHdeem; sleep 5; startHdeem; $KRIPKE_COMMAND;
stopHdeem
```

INSTRUMENTED RUN




| Run instrumented version of the Kripke

```
$ ./compile_for_1_meric.sh
```

```
$ ./run.sh
```

```
$ ls mericMeasurement/
```


MEASUREMENT OUTPUT



| See MERIC output in `mericMeasurement/`

```
# CALLTREE;init_4
# Job info
JOB ID,603210.1
Runtime of function [s],44.0138
Function start timestamp,1649155824.06406
# CALLTREE;init_4
# SAMPLES - HDEEM (Blade):
Energy consumption [J],18073.6
MAX power [W],491.25
MIN power [W],304.875
AVG power [W],410.635
ID of first sample,452
ID of last sample,44464
# SAMPLES - HDEEM (CPU0):
Energy consumption [J],6856.33
MAX power [W],188.25
MIN power [W],118.875
AVG power [W],155.777
ID of first sample,47
ID of last sample,4447
...
# CALLTREE;init_4
# COUNTERS - RAPL:
RAPL_RAM_0 [J],656.51973
RAPL_RAM_1 [J],652.61183
RAPL_PCKG_0 [J],6605.80060
RAPL_PCKG_1 [J],6603.47150
RAPL Energy consumption [J],14518.40366
# CALLTREE;init_4
# COUNTERS - MSR:
# RATIOS:
CPU_FREQ_0 [Hz],3148992452.62181
CPU_FREQ_1 [Hz],3149181571.83058
CPU_FREQ_10 [Hz],3149219341.03368
CPU_FREQ_11 [Hz],3149220776.26192
CPU_FREQ_12 [Hz],3149161361.77316
CPU_FREQ_13 [Hz],3149157158.45573
...
```


RADAR VISUALIZER



| For GUI you need to be logged in with -X or use VNC

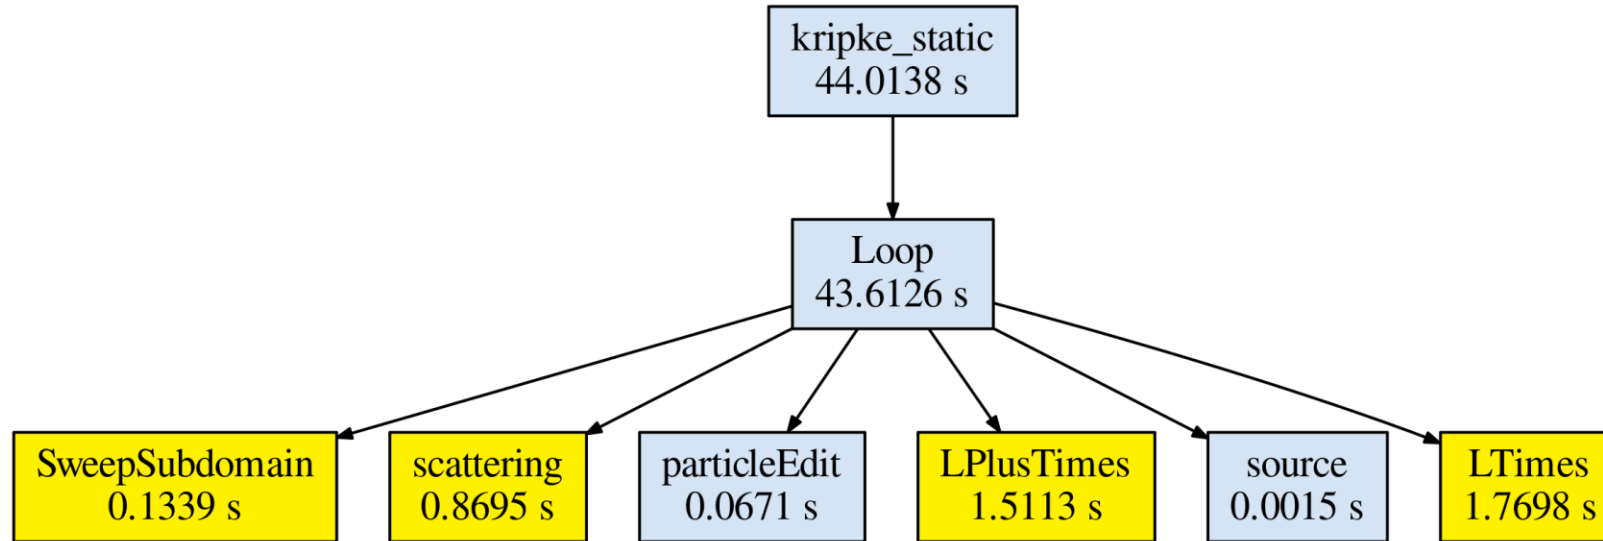

```
$ ssh -X -i /path/to/id_rsa u<LOGIN>@barbora.it4i.cz
```

```
$ cd $HOME/radar/
```


```
$ source set_env_barbora.sh
```

```
$ ./runRadarGUI_analyze.py
```

MERIC ANALYSIS



RADAR VISUALIZER



| Data location: \$HOME/kripke-benchmark/mericMeasurement

RADAR configuration (on login2.karolina.it4i.cz)

Regions | Data parameters

Main region
kripke_static

Phase region
None

Filter nested regions by tlr
0,00000000 s
Apply filter
Show region tree

Nested regions

- kripke_static
- Loop
- LPlusTimes
- LTimes
- particleEdit
- scattering
- source
- SweepSubdomain

Select All
Next

RADAR configuration (on login2.karolina.it4i.cz)

Regions | Data parameters

y label

Time-energy variables

Time
Job info Runtime of function [s]

Energy
COUNTERS - RAPL: RAPL Energy

Baseline 80 W

Parameters

Parameter 1	Role	Default value	Multiplier	Unit
CF	funcLabel	2600000000	1.0	
UnCF	xLabel	2400000000	1.0	

Previous Run RADAR analysis