CPU Energy Meter: A Tool for Energy-Aware Algorithms Engineering

Proc. TACAS 2020, https://doi.org/10/f4bk

Dirk Beyer

Joint work with Philipp Wendler

LMU Munich, Germany

LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

SoSy-Lab Software Systems

Why measure energy?

- Verification certainly consumes a lot of energy
- "Green" verification needed

Why measure energy?

- Verification certainly consumes a lot of energy
- “Green” verification needed
- We do not even know how much energy is consumed!
- You cannot improve what you cannot measure!

How to measure energy?

Suggestion: Buy smart energy meter, hook up to computer, etc.

- Needs additional hardware and installation effort (costs)
- Typically low measurement resolution
- Complex setup
- Often not usable (shared machines, data center, etc.)

$\Rightarrow$ Unlikely to be adopted in research practice

How to measure energy: RAPL

- Intel **R**unning **A**verage **P**ower **L**imit [2]
- API for accessing energy-consumption counters
- Available in common Intel CPUs
- Measurements per CPU and per CPU component (cores, GPU, memory controller)
- Resolution $\sim 10^{-5}$ J (e.g., 10 mW for 1 ms)
- No official statements on precision and accuracy, but experiments found good accuracy

CPU Energy Meter

`github.com/ sosy-lab / cpu-energy-meter`

- Easy to use command-line tool for reading RAPL values
- BSD License
- Available as Debian/Ubuntu package,
  no manual configuration necessary
- Manual execution possible
- Intended to be used by scripts / benchmarking frameworks
- Integrated in BenchExec
  (will be used automatically if installed
  and whole CPUs are used)

Demo 1: Manual

```
cpu-energy-meter
<Ctrl+C>
```

Output:

```
+-------------------------------------+
| CPU Energy Meter             Socket 0 |
+-------------------------------------+
Duration                     1.085165 sec
Package                      1.318359 Joule
Core                         0.161438 Joule
Uncore                       0.048279 Joule
DRAM                         1.252380 Joule
PSYS                         9.194458 Joule
```

Demo 2: Calculating Pi

Start measurement, calculate pi, end measurement:

```
cpu-energy-meter &
echo "scale=2000; a(1)*4" | bc -l > /dev/null
kill -INT %1
```

Output:

```
+-------------------------------------+
| CPU Energy Meter             Socket 0 |
+-------------------------------------+
Duration                    1.601558 sec
Package                    18.213501 Joule
Core                       15.999451 Joule
Uncore                      0.042236 Joule
DRAM                        1.764526 Joule
PSYS                       40.531677 Joule
```

Reminder: Measures whole CPU energy, not only single process!

Demo 2: BenchExec integration

```
runexec -- /bin/sh -c \
  'echo "scale=2000; a(1)*4" | bc -l > /dev/null'
```

Output:

```
2021-03-17 15:56:46 - INFO - Starting command /bin/sh -c 'echo "scale=2000; a(1
2021-03-17 15:56:46 - INFO - Writing output to output.log
starttime=2021-03-17T15:56:46.158262+01:00
returnvalue=0
walltime=1.6216954150004312s
cputime=1.620933245s
cputime-cpu0=1.619139808s
cputime-cpu1=0.000238686s
cputime-cpu2=0.001383994s
cputime-cpu3=0.000170757s
memory=1073152B
blkio-read=0B
blkio-write=0B
cpuenergy=18.772644J
cpuenergy-pkg0-core=16.334778J
cpuenergy-pkg0-dram=1.902161J
cpuenergy-pkg0-package=18.772644J
cpuenergy-pkg0-psys=41.596741J
cpuenergy-pkg0-uncore=0.189575J
```

Integration in BenchExec [1]

Click here to select …	status	cputime (s)	walltime (s)	memory (MB)	cpuenergy (J)
total	2964	73800	56600	750000	800000
correct results	1	avg: 270, max: 956, median: 81.7, min: 22.7, stdev: 296			
correct true	902	7880	3390	172000	74400

Energy results available just like time, memory, etc.

Evaluation

- Can energy measurements give new insights when comparing algorithms/tools?
- Experiment:
  - Official SV-COMP'19 results
  - Energy usage of whole CPU
    (results for CPU components similar)

Power usage varies across tools and runs

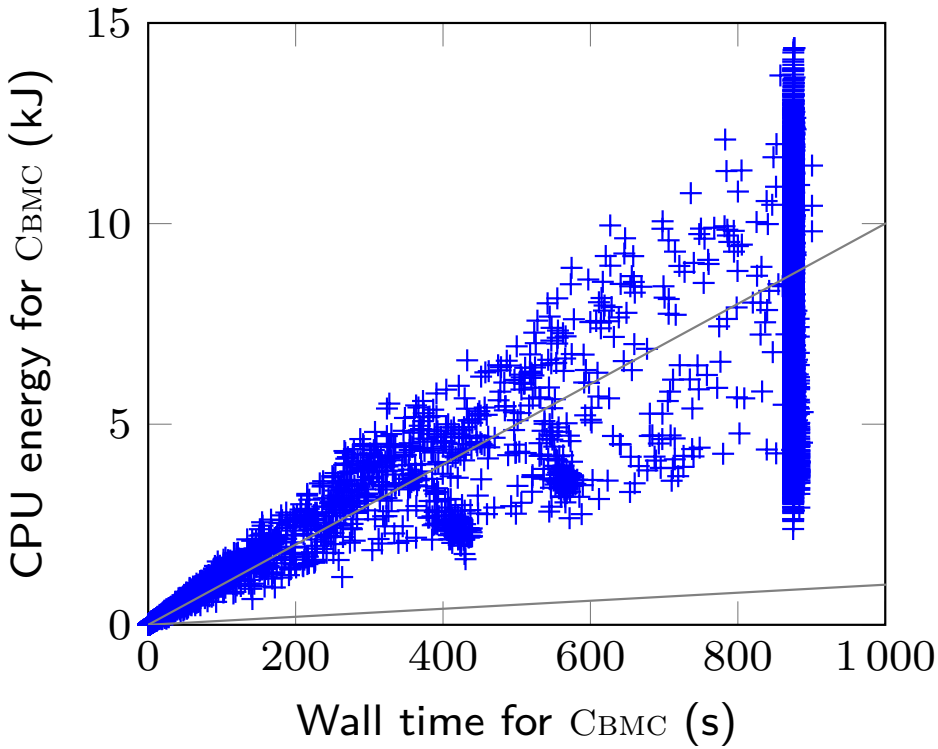

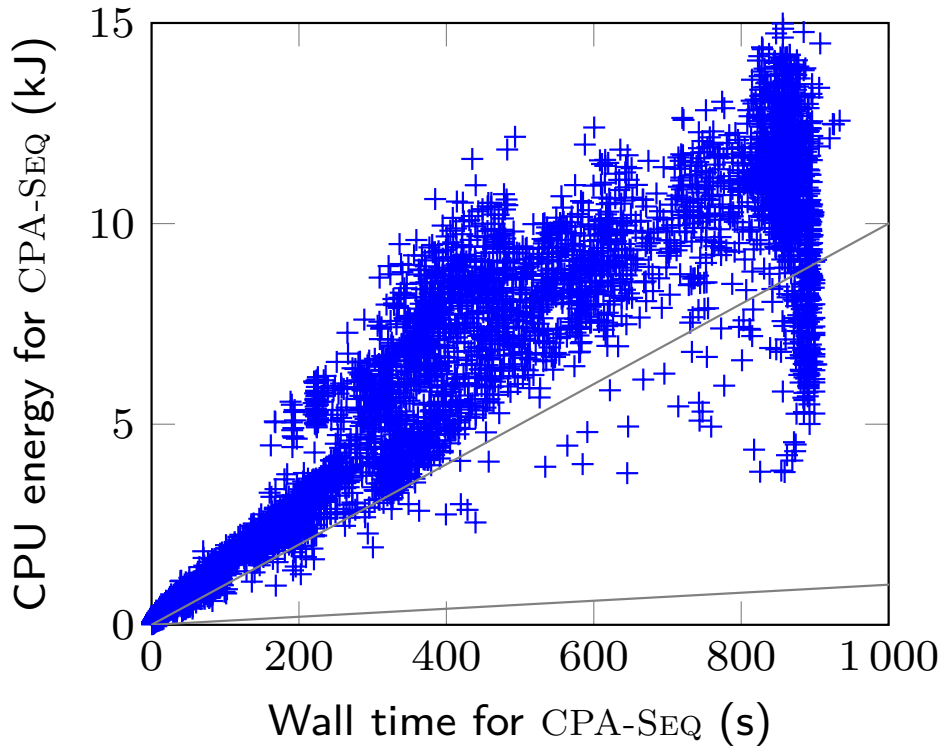

Lines through origin represent average power
(gray lines: 1 W and 10 W)

High values for CPU time do not imply high values for energy

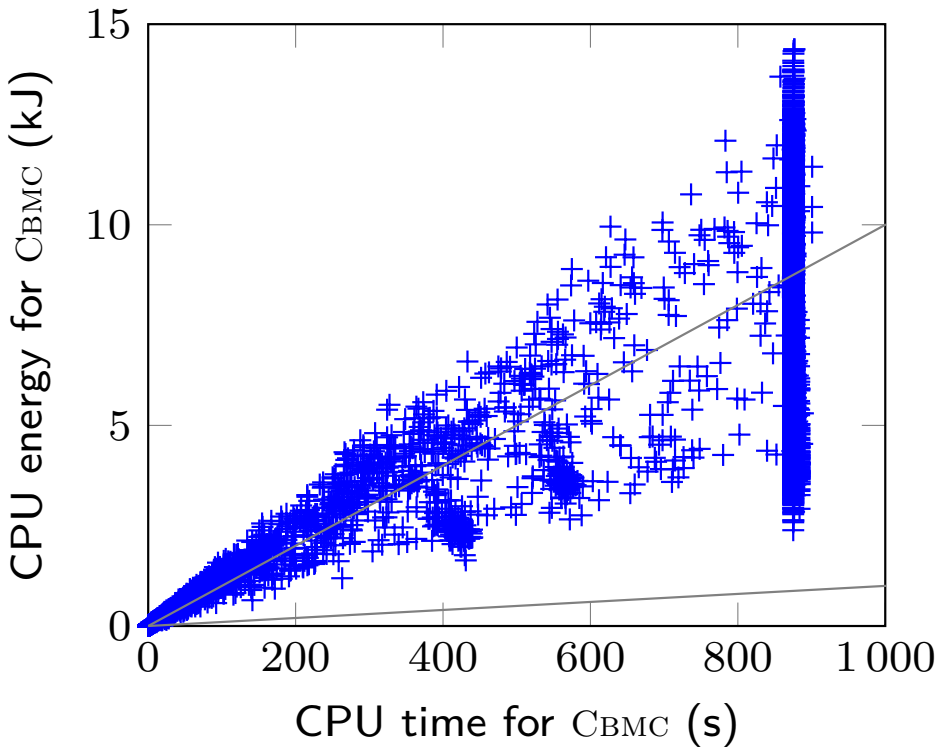

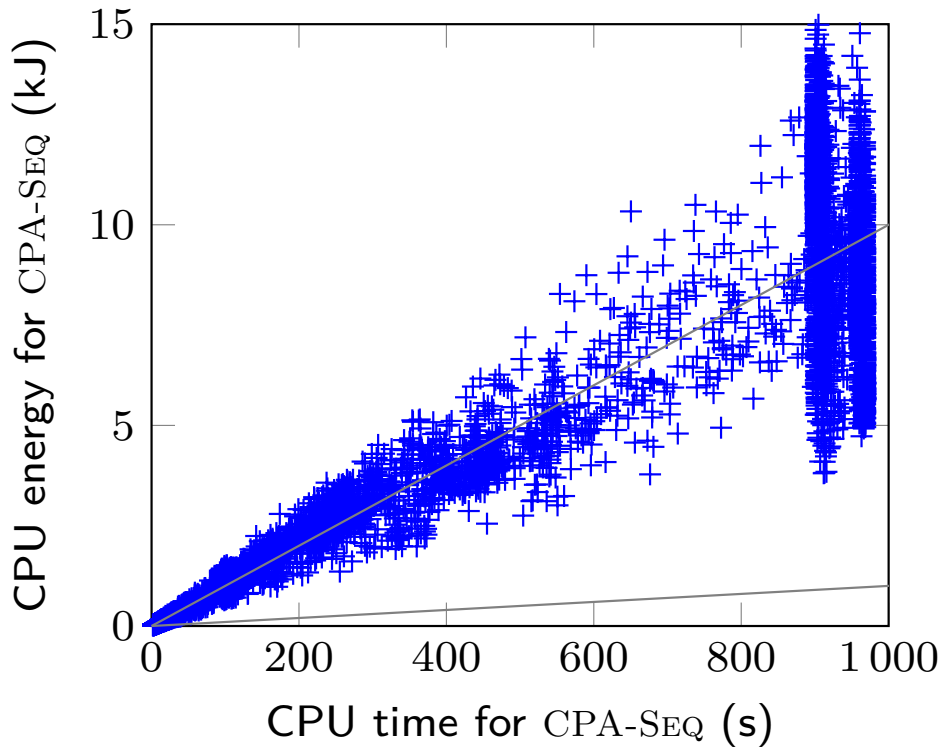

"CPU-Power" usage between 2 W and 16 W

Comparing different tools regarding CPU time can lead to different conclusions than energy-based comparisons

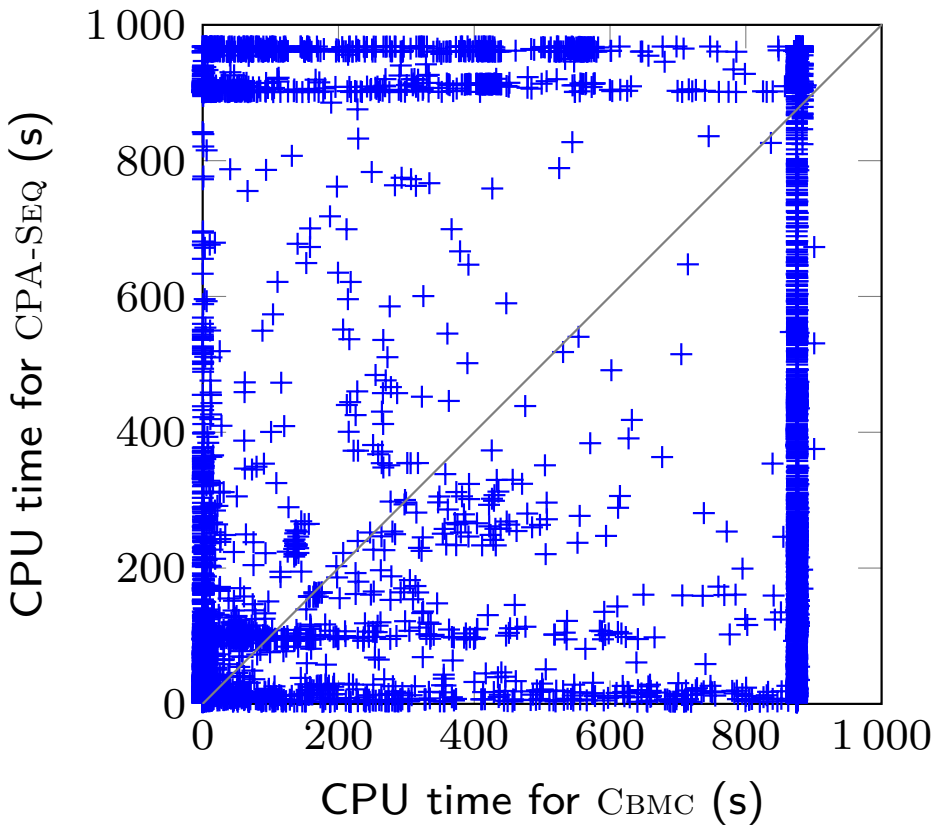

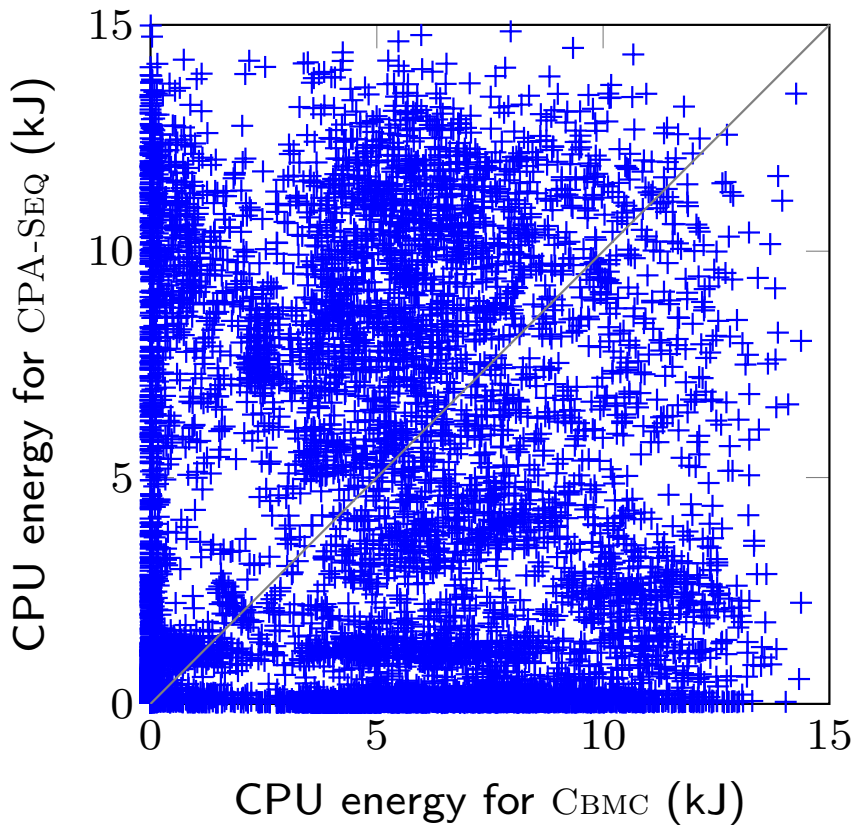

Summary

- Energy-aware research is important
- Energy measurements more difficult
  than time measurements
- Time is not a good proxy for energy
- At least measure energy as far as possible
  (and discuss in evaluation)
- Add efficiency rankings to competitions
- BenchExec + CPU Energy Meter leave no excuse ;-)

References I

Beyer, D., Löwe, S., Wendler, P.: Reliable benchmarking: Requirements and solutions. Int. J. Softw. Tools Technol. Transfer **21**(1), 1–29 (2019).

https://doi.org/10.1007/s10009-017-0469-y

Intel: Intel 64 and IA-32 architectures software developer's manual, vol. 3B, chap. 14.9. Intel (December 2017), available at https://software.intel.com/en-us/articles/intel-sdm