

Performance analysis with hardware metrics

Likwid-perfctr

Best practices

Energy consumption

Clock speed



- ... are ubiquitous as a starting point for performance analysis (including automatic analysis)
- ... are supported by many tools
- ... are often reduced to cache misses (what could be worse than cache misses?)

Reality:

- Modern parallel computing is **plagued by bottlenecks**
 - There are typical **performance patterns** that cover a large part of possible performance behaviors
 - HPM signatures
 - Scaling behavior
 - Other sources of information
- } **“Performance pattern”**



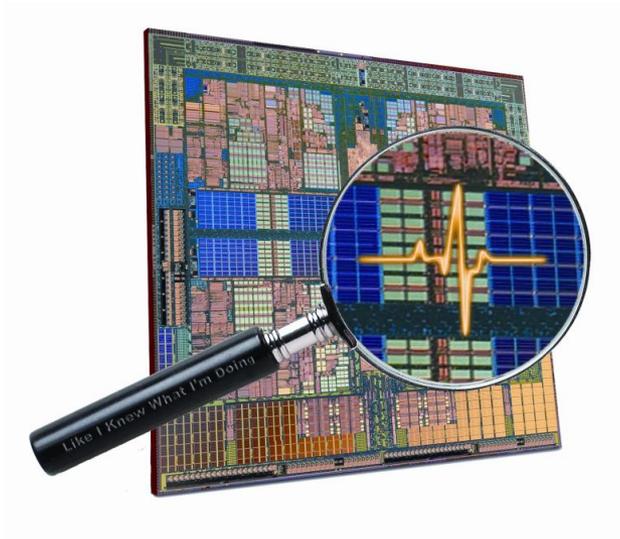
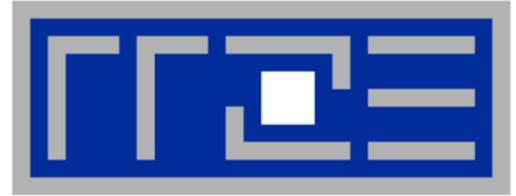
1. **Runtime profile** / Call graph (gprof): Where are the hot spots?
2. **Instrument** hot spots (prepare for detailed measurement)
3. Find **performance signatures**

Possible signatures (selection):

- **Bandwidth** saturation
- **Instruction throughput** limitation (real or language-induced)
- **Latency** impact (irregular data access, high branch ratio)
- **Load imbalance**
- **ccNUMA** issues (data access across ccNUMA domains)
- **SIMD** vs. **scalar** execution
- Pathologic cases (false sharing, expensive operations)

likwid-perfctr
can help here

Goal: Come up with educated guess about a performance-limiting motif
(**Performance Pattern**)



Using hardware performance metrics

`likwid-perfctr`



- How do we find out about the performance properties and requirements of a parallel code?
 - Profiling via advanced tools is often overkill
- A coarse overview is often sufficient
 - **likwid-perfctr** (similar to “perfex” on IRIX, “hpmcount” on AIX, “lipfpm” on Linux/Altix)
 - Simple end-to-end measurement of hardware performance metrics

- Operating modes:

- Wrapper
- Stethoscope
- Timeline
- Marker API

- Preconfigured and extensible metric groups, list with **likwid-perfctr -a**



BRANCH: Branch prediction miss rate/ratio
CACHE: Data cache miss rate/ratio
CLOCK: Clock of cores
DATA: Load to store ratio
FLOPS_DP: Double Precision MFlops/s
FLOPS_SP: Single Precision MFlops/s
FLOPS_X87: X87 MFlops/s
L2: L2 cache bandwidth in MBytes/s
L2CACHE: L2 cache miss rate/ratio
L3: L3 cache bandwidth in MBytes/s
L3CACHE: L3 cache miss rate/ratio
MEM: Main memory bandwidth in MBytes/s
TLB: TLB miss rate/ratio



```
$ likwid-perfctr -C N:0-3 -g FLOPS_DP ./stream.exe
```

```
-----
CPU type:      Intel Core Lynnfield processor
CPU clock:     2.93 GHz
-----
```

```
Measuring group FLOPS_DP
```

```
YOUR PROGRAM OUTPUT
```

Event	core 0	core 1	core 2	core 3
INSTR_RETIRED_ANY	1.97463e+08	2.31001e+08	2.30963e+08	2.31885e+08
CPU_CLK_UNHALTED_CORE	9.56999e+08	9.58401e+08	9.58637e+08	9.57338e+08
FP_COMP_OPS_EXE_SSE_FP_PACKED	4.00294e+07	3.08927e+07	3.08866e+07	3.08904e+07
FP_COMP_OPS_EXE_SSE_FP_SCALAR	882	0	0	0
FP_COMP_OPS_EXE_SSE_SINGLE_PRECISION	0	0	0	0
FP_COMP_OPS_EXE_SSE_DOUBLE_PRECISION	4.00303e+07	3.08927e+07	3.08866e+07	3.08904e+07

Always
measured

Configured metrics
(this group)

Metric	core 0	core 1	core 2	core 3
Runtime [s]	0.326242	0.32672	0.326801	0.326358
CPI	4.84647	4.14891	4.15061	4.12849
DP MFlops/s (DP assumed)	245.399	189.108	189.024	189.304
Packed MUOPS/s	122.698	94.554	94.5121	94.6519
Scalar MUOPS/s	0.00270351	0	0	0
SP MUOPS/s	0	0	0	0
DP MUOPS/s	122.701	94.554	94.5121	94.6519

Derived
metrics



Things to look at (in roughly this order)

- **Load balance** (flops, instructions, BW)
- **In-socket memory BW saturation**
- **Shared cache BW saturation**
- **Flop/s, loads and stores per flop metrics**
- **SIMD** vectorization
- **CPI** metric
- **# of instructions**, branches, mispredicted branches

Caveats

- **Load imbalance** may not show in CPI or # of instructions
 - **Spin loops** in OpenMP barriers/MPI blocking calls
 - Looking at “top” or the Windows Task Manager does not tell you anything useful
- **In-socket performance saturation** may have various reasons
- **Cache miss metrics are overrated**
 - If I really know my code, I can often *calculate* the misses
 - Runtime and resource utilization is much more important



- likwid-perfctr counts events on cores; it has no notion of what kind of code is running (if any)

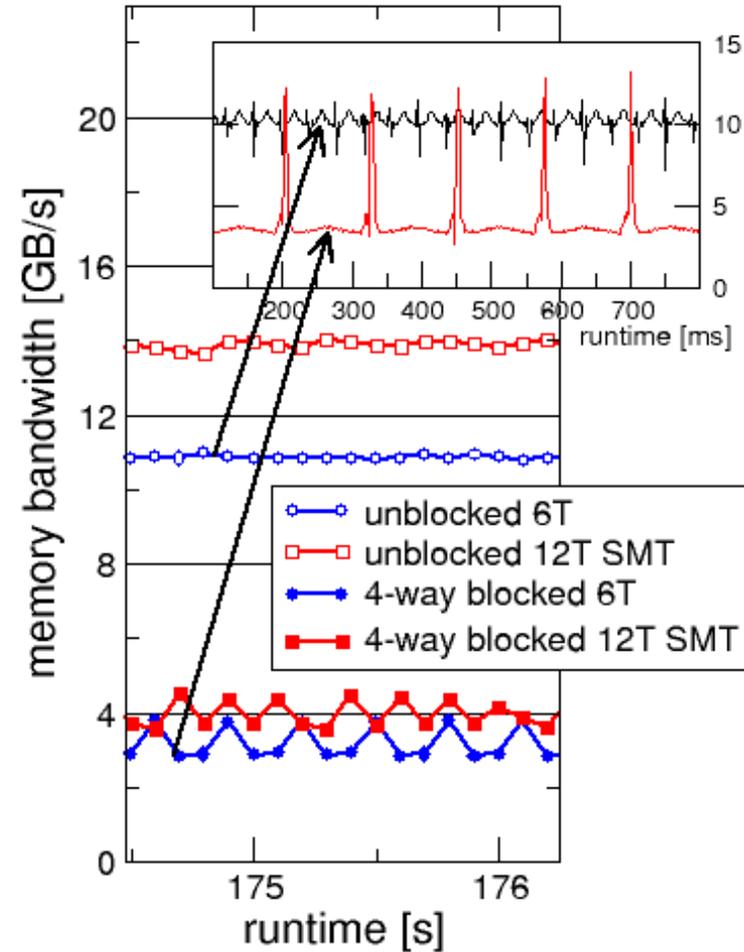
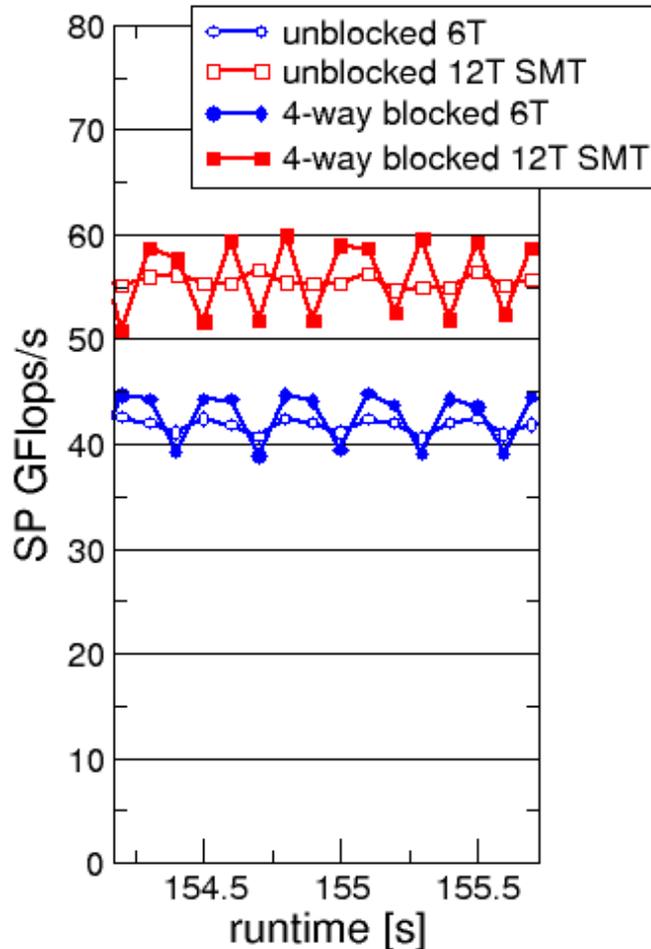
This allows you to “listen” to what is currently happening, without any overhead:

```
likwid-perfctr -c N:0-11 -g FLOPS_DP -S 10
```

- It can be used as cluster/server monitoring tool
- A frequent use is to measure a certain part of a long running parallel application from outside

- likwid-perfctr supports time-resolved measurements of full node:

```
$ likwid-perfctr -c N:0-11 -g MEM -t 50ms > out.txt
```





- The marker API can restrict measurements to code regions
- The API only turns counters on/off. The configuration of the counters is still done by likwid-perfctr
- Multiple named regions support, accumulation over multiple calls
- Inclusive and overlapping regions allowed

```
#include <likwid.h>
. . .
LIKWID_MARKER_INIT;           // must be called from serial region
#pragma omp parallel
{
    LIKWID_MARKER_THREADINIT; // only reqd. if measuring multiple threads
}
. . .
LIKWID_MARKER_START("Compute");
. . .
LIKWID_MARKER_STOP("Compute");
. . .
LIKWID_MARKER_START("Postprocess");
. . .
LIKWID_MARKER_STOP("Postprocess");
. . .
LIKWID_MARKER_CLOSE;         // must be called from serial region
```

- Activate macros with `-DLIKWID_PERFMON`
- Run `likwid-perfctr` with `-m` option to activate markers



SHORT PSTI

EVENTSET

```
FIXC0 INSTR_RETIRED_ANY
FIXC1 CPU_CLK_UNHALTED_CORE
FIXC2 CPU_CLK_UNHALTED_REF
PMC0 FP_COMP_OPS_EXE_SSE_FP_PACKED
PMC1 FP_COMP_OPS_EXE_SSE_FP_SCALAR
PMC2 FP_COMP_OPS_EXE_SSE_SINGLE_PRECISION
PMC3 FP_COMP_OPS_EXE_SSE_DOUBLE_PRECISION
UPMC0 UNC_QMC_NORMAL_READS_ANY
UPMC1 UNC_QMC_WRITES_FULL_ANY
UPMC2 UNC_QHL_REQUESTS_REMOTE_READS
UPMC3 UNC_QHL_REQUESTS_LOCAL_READS
```

METRICS

```
Runtime [s] FIXC1*inverseClock
CPI FIXC1/FIXC0
Clock [MHz] 1.E-06*(FIXC1/FIXC2)/inverseClock
DP MFlops/s (DP assumed) 1.0E-06*(PMC0*2.0+PMC1)/time
Packed MUOPS/s 1.0E-06*PMC0/time
Scalar MUOPS/s 1.0E-06*PMC1/time
SP MUOPS/s 1.0E-06*PMC2/time
DP MUOPS/s 1.0E-06*PMC3/time
Memory bandwidth [MBytes/s] 1.0E-06*(UPMC0+UPMC1)*64/time;
Remote Read BW [MBytes/s] 1.0E-06*(UPMC2)*64/time;
```

LONG

Formula:

```
DP MFlops/s = (FP_COMP_OPS_EXE_SSE_FP_PACKED*2 + FP_COMP_OPS_EXE_SSE_FP_SCALAR)/ runtime.
```

- Groups are architecture-specific
- They are defined in simple text files
- Code is generated on recompilation of likwid
- `likwid-perfctr -a` prints a list of groups
- An extended documentation (help text) is available for every group

Example: triangular dense MVM



Multiplication of upper triangular matrix with vector

```
#define N 10000 // large enough to use memory
#define ROUNDS 10

[...]

// Calculation loop
#pragma omp parallel
{
    for (int k = 0; k < ROUNDS; k++) {
        #pragma omp for private(current,j)
        for (int i = 0; i < N; i++) {
            current = 0;
            for (int j = i; j < N; j++)
                current += mat[(i*N)+j] * bvec[j];
            cvec[i] = current;
        }
        while (cvec[N>>1] < 0) {dummy();break;}
    }
}
```



LIKWID marker API

```
#include <likwid.h>
[...]
LIKWID_MARKER_INIT;
#pragma omp parallel
{
    LIKWID_MARKER_THREADINIT;
}
#pragma omp parallel
{
    for (int k = 0; k < ROUNDS; k++) {
        LIKWID_MARKER_START("Compute");
        #pragma omp for private(current,j)
        for (int i = 0; i < N; i++) {
            current = 0;
            for (int j = i; j < N; j++)
                current += mat[(i*N)+j] * bvec[j];
            cvec[i] = current;
        }
        LIKWID_MARKER_STOP("Compute");
        while (cvec[N]>>1 < 0) {dummy();break;}
    }
}
LIKWID_MARKER_CLOSE;
```

Example: triangular dense MVM



```
$ likwid-perfctr -C 0,1,2 -g L2 -m ./a.out
```

```
-----  
CPU type:      Intel Core SandyBridge EN/EP processor  
CPU clock:    3.09 GHz  
-----
```

```
=====  
Group 1: Region Compute  
=====
```

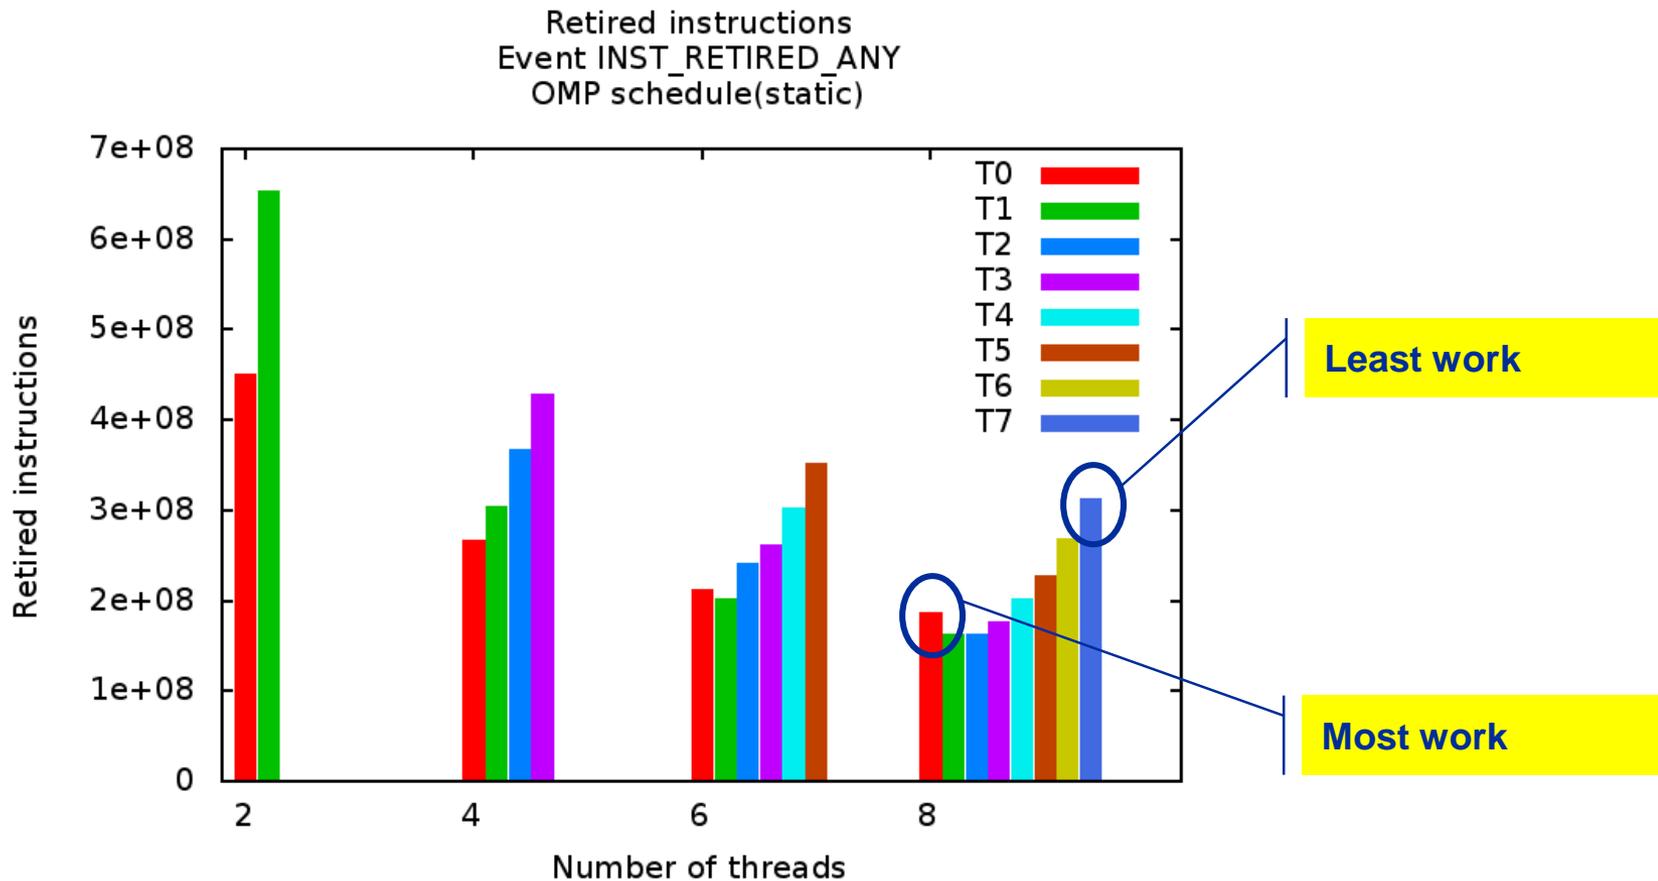
Region Info	Core 0	Core 1	Core 2
RDTSC Runtime [s]	0.161382	0.161365	0.161365
call count	10	10	10

Event	Counter	Core 0	Core 1	Core 2
INSTR_RETIRED_ANY	FIXC0	2.626800e+08	3.187585e+08	3.780255e+08
CPU_CLK_UNHALTED_CORE	FIXC1	4.972802e+08	4.961411e+08	4.933711e+08
CPU_CLK_UNHALTED_REF	FIXC2	4.972801e+08	4.961404e+08	4.933714e+08
L1D_REPLACEMENT	PMC0	5.490278e+07	3.927353e+07	2.364295e+07
L1D_M_EVICT	PMC1	2.920200e+04	2.876600e+04	2.861000e+04
ICACHE_MISSES	PMC2	4.649000e+03	4.984000e+03	5.321000e+03

Example: triangular dense MVM



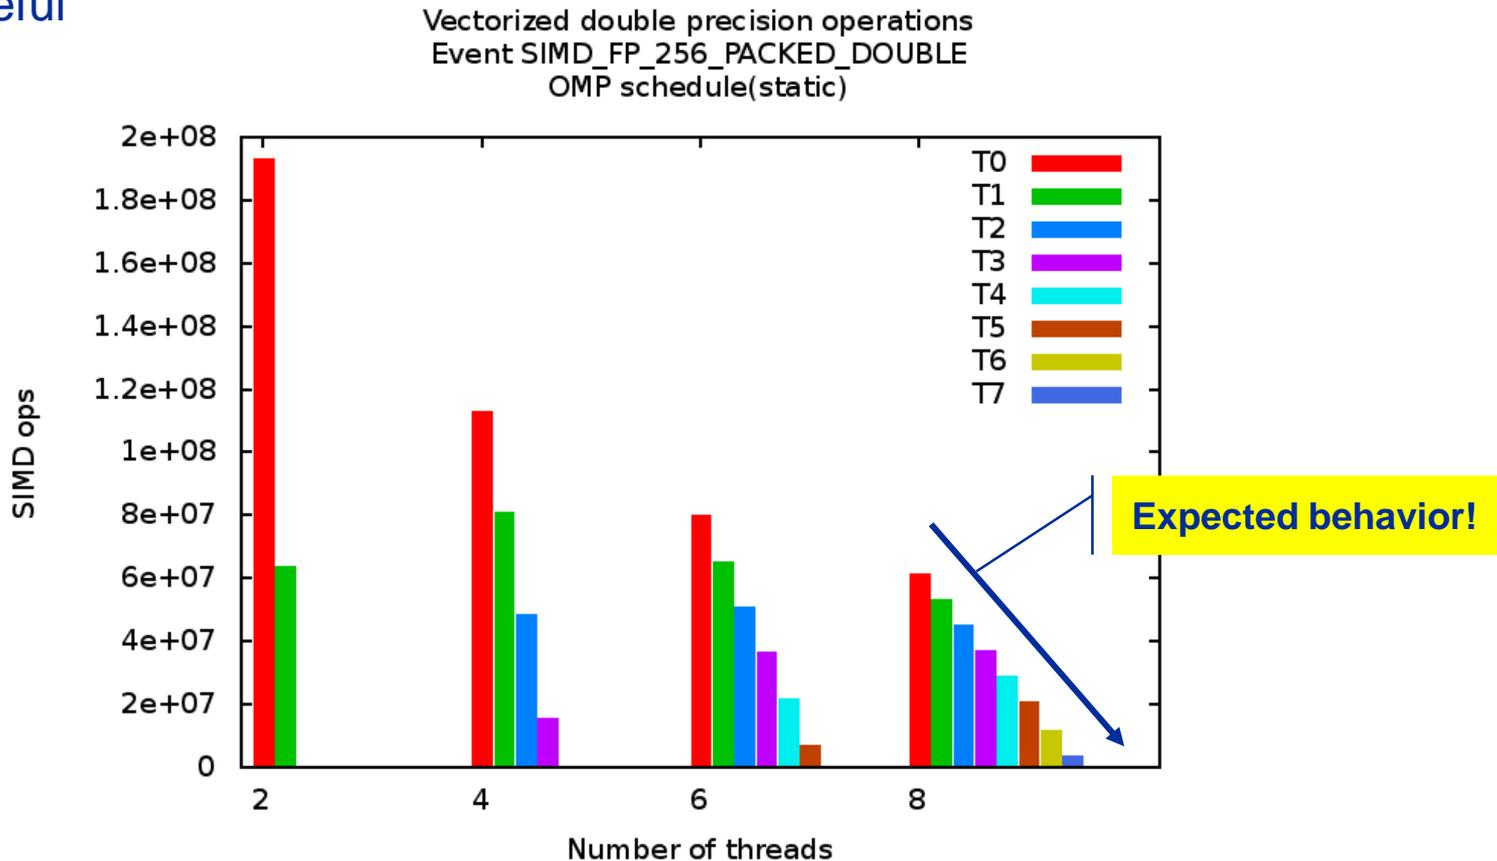
- Retired instructions are misleading
- Waiting in implicit OpenMP barrier issues many instructions



Example: triangular dense MVM



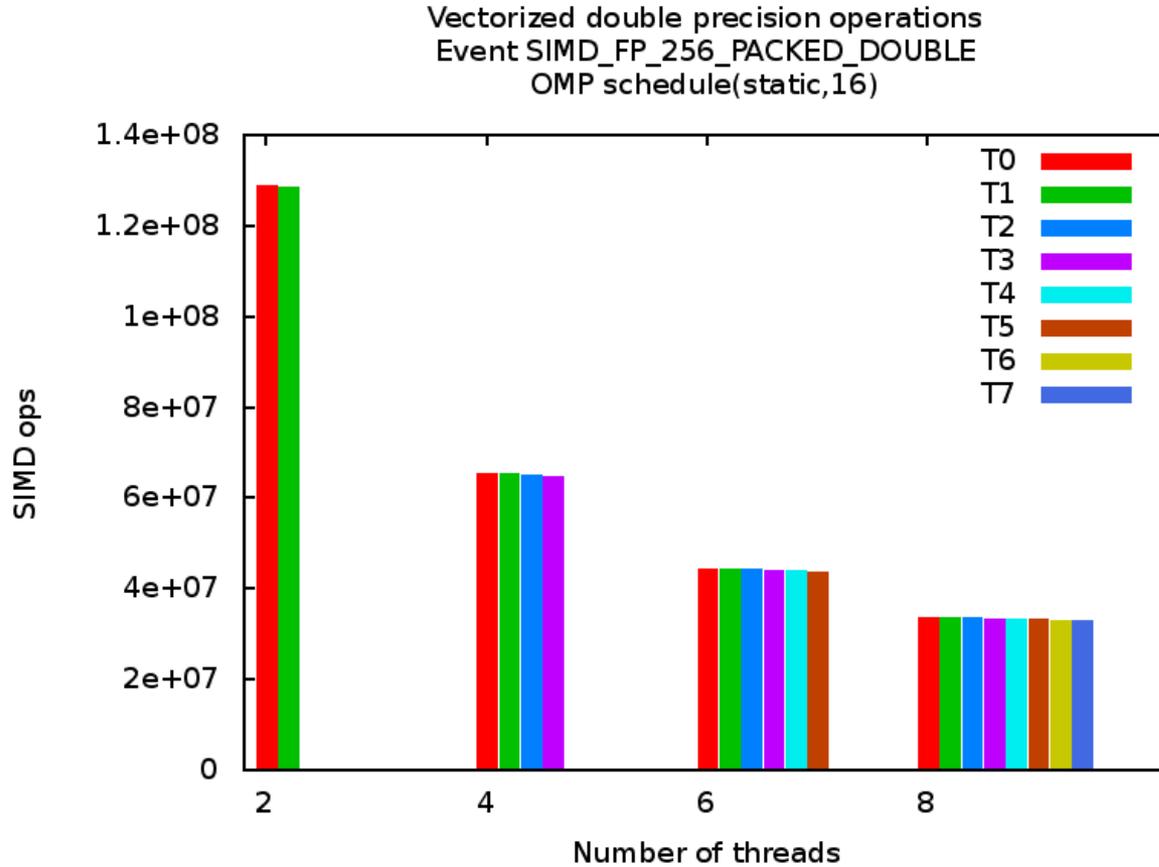
- Floating point instructions are reliable \leftrightarrow useful work metric!
- But floating point instr. counters suffer from miscounting since SandyBridge
 - Still useful



Example: triangular dense MVM



- Changing static schedule to static,16 → load balanced



Measuring energy consumption

likwid-powermeter and likwid-perfctr -g ENERGY



Implements Intel RAPL (Running average power limit) interface (Sandy Bridge++)

```
$ likwid-powermeter -i
```

```
-----  
CPU name:      Intel(R) Xeon(R) CPU E5-2695 v3 @ 2.30GHz  
CPU type:      Intel Xeon Haswell EN/EP/EX processor  
CPU clock:     2.30 GHz  
-----
```

```
Base clock:    2300.00 MHz  
Minimal clock: 1200.00 MHz
```

Turbo Boost Steps:

C0 3300.00 MHz

C1 3300.00 MHz

C2 3100.00 MHz

C3 3000.00 MHz

C4 2900.00 MHz

[...]

C13 2800.00 MHz

Info for RAPL domain **PKG**:

Thermal Spec Power: 120 Watt

Minimum Power: 70 Watt

Maximum Power: **120 Watt**

Maximum Time Window: 46848 micro sec

Info for RAPL domain **DRAM**:

Thermal Spec Power: 21.5 Watt

Minimum Power: 5.75 Watt

Maximum Power: **21.5 Watt**

Maximum Time Window: 44896 micro sec

Energy with likwid-perfctr in wrapper mode



```
$ likwid-perfctr -g ENERGY -C S0:0-1 ./a.out
```

[...]

Metric	Core 0	Core 1
Runtime (RDTSC) [s]	2.362547e+01	2.362547e+01
Runtime unhaltd [s]	3.347379e+01	3.348318e+01
Clock [MHz]	3.284280e+03	3.284278e+03
CPI	6.105052e-01	6.116845e-01
Temperature [C]	33	33
Energy [J]	1.423615e+03	0
Power [W]	6.025761e+01	0
Energy PP0 [J]	0	0
Power PP0 [W]	0	0
Energy DRAM [J]	89.424216	0
Power DRAM [W]	3.785076e+00	0

Package (socket)

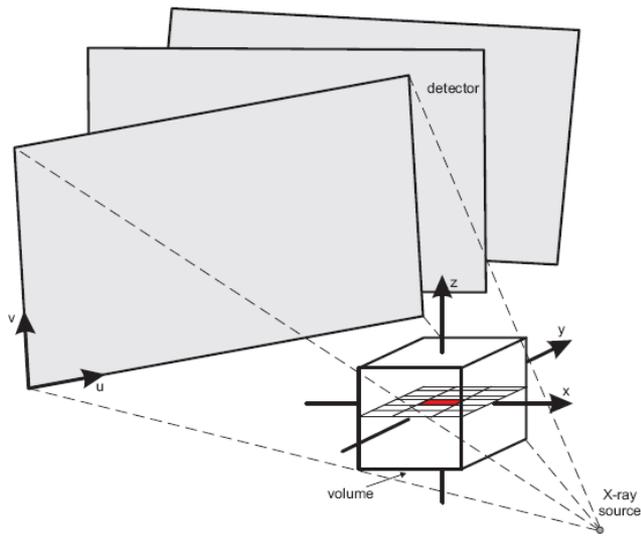
Cores only (if available)

DRAM (DIMMs)

[...]

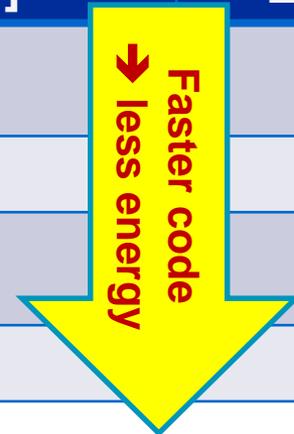
Example:

A medical image reconstruction code on Sandy Bridge



Sandy Bridge EP (8 cores, 2.7 GHz base freq.)

Test case	Runtime [s]	Power [W]	Energy [J]
8 cores, plain C	90.43	90	8110
8 cores, SSE	29.63	93	2750
8 cores (SMT), SSE	22.61	102	2300
8 cores (SMT), AVX	18.42	111	2040





- The “**Turbo Mode**” feature makes reliable benchmarking harder
 - CPU can change clock speed at its own discretion
- Clock speed reduction may **save a lot of energy**

- So how do we set the clock speed? → LIKWID to the rescue!

```
$ likwid-setFrequencies -l
Available frequencies:
2.301, 2.3, 2.2, 2.1, 2.0, 1.9, 1.8, 1.7, 1.6, 1.5, 1.4, 1.3, 1.2
$ likwid-setFrequencies -p
Current frequencies:
CPU 0: governor performance frequency 2.301 GHz
CPU 1: governor performance frequency 2.301 GHz
CPU 2: governor performance frequency 2.301 GHz
CPU 3: governor performance frequency 2.301 GHz
[...]
$ likwid-setFrequencies -f 2.3
$
```



Turbo mode



- **Useful only if you know what you are looking for**
 - PM bears potential of acquiring massive amounts of data for nothing!
- **Resource-based metrics are most useful**
 - Cache lines transferred, work executed, loads/stores, cycles
 - Instructions, CPI, cache misses may be misleading
- **Caveat: Processor work != user work**
 - Waiting time in libraries (OpenMP, MPI) may incur lots of instructions
 - → distorted application characteristic
- **Another very useful application of PM: validating performance models!**
 - Roofline is data centric → measure data volume through memory hierarchy