

Tools/likwid/example_perfctr_stream

Contents

Example likwid-perfctr performance group MEM on benchmark stream	1
Example likwid-perfctr performance group UPI on benchmark stream	6

Example likwid-perfctr performance group MEM on benchmark stream

- Prepare environment

```
module purge
module add compiler/intel/2022
```
- Build stream benchmark with Intel compiler

```
icc -std=c11 -Ofast -xHost -ipo -qopenmp \
    stream.c -o stream
```
- Set up OpenMP environment

```
export OMP_NUM_THREADS=76
export KMP_AFFINITY="verbose,granularity=core,respect,scatter"
```
- List available performance groups

```
likwid-perfctr -a
```

```
...
      MEM  Memory bandwidth in MBytes/s
      MEM_DP  Overview of arithmetic and main memory performance
      MEM_FREERUN  Memory bandwidth in MBytes/s
      MEM_SP  Overview of arithmetic and main memory performance
      ...
```
- Get detailed information on performance groups

```
likwid-perfctr -H --group MEM
```

Group MEM:

Formulas:

Memory read bandwidth [MBytes/s] = 1.0E-06*(SUM(CAS_COUNT_RD))*64.0/runtime

Memory read data volume [GBytes] = 1.0E-09*(SUM(CAS_COUNT_RD))*64.0

Memory write bandwidth [MBytes/s] = 1.0E-06*(SUM(CAS_COUNT_WR))*64.0/runtime

Memory write data volume [GBytes] = 1.0E-09*(SUM(CAS_COUNT_WR))*64.0

Memory bandwidth [MBytes/s] = 1.0E-06*(SUM(CAS_COUNT_RD)+SUM(CAS_COUNT_WR))*64.0/runtime

Memory data volume [GBytes] = 1.0E-09*(SUM(CAS_COUNT_RD)+SUM(CAS_COUNT_WR))*64.0

-

Profiling group to measure memory bandwidth drawn by all cores of a socket.

Since this group is based on Uncore events it is only possible to measure on a per socket base. Some of the counters may not be available on your system.

Also outputs total data volume transferred from main memory.

- Measure performance group MEM for benchmark stream on CPU Hyper-thread 0 to 151

```
likwid-perfctr \  
    --group MEM \  
    -c 0-151 \  
    ./stream -n 1000000000
```

STREAM version \$Revision: 5.10 \$

This system uses 8 bytes per array element.

Array size = 999999944 (elements)

Memory per array = 7629.4 MiB (= 7.5 GiB).

Total memory required = 22888.2 MiB (= 22.4 GiB).

Each kernel will be executed 10 times.

The *best* time for each kernel (excluding the first iteration)
will be used to compute the reported bandwidth.

OpenMP version (yyymm): 201611

Number of Threads requested = 76

Number of Threads counted = 76

Your clock granularity appears to be 1000 ticks per microseconds.

Each test below will take on the order of 50701 microseconds.

(= 50701851 clock ticks)

Increase the size of the arrays if this shows that
you are not getting at least 20 clock ticks per test.

WARNING -- The above is only a rough guideline.

For best results, please be sure you know the

precision of your system timer.

Function	Best Rate MB/s	Med time	Min time	Max time
Copy:	315666.4	0.050794	0.050686	0.052201
Scale:	313917.3	0.051067	0.050969	0.052823
Add:	319821.6	0.075599	0.075042	0.077726
Triad:	319194.3	0.076074	0.075189	0.077786

Solution Validates: avg error less than 1.000000e-13 on all three arrays

Group 1: MEM

Event	Counter	HWThread 0	HWThread 1	HWThread 2	HWThread 3
INSTR_RETIRED_ANY	FIXC0	411100	1184249672	149965269	1104712055
CPU_CLK_UNHALTED_CORE	FIXC1	2985530	9439091228	2038977682	9433147786
CPU_CLK_UNHALTED_REF	FIXC2	2263680	7190013696	1536292800	7184417760
TOPDOWN_SLOTS	FIXC3	10657535	47172282470	5156016260	47153531225
CAS_COUNT_RD	MBOX0C0	565487546	0	0	0
CAS_COUNT_WR	MBOX0C1	394832802	0	0	0
CAS_COUNT_RD	MBOX1C0	565673329	0	0	0
CAS_COUNT_WR	MBOX1C1	394772844	0	0	0
CAS_COUNT_RD	MBOX2C0	565520834	0	0	0
CAS_COUNT_WR	MBOX2C1	394661003	0	0	0
CAS_COUNT_RD	MBOX3C0	565668266	0	0	0
CAS_COUNT_WR	MBOX3C1	394751578	0	0	0
CAS_COUNT_RD	MBOX4C0	565524973	0	0	0
CAS_COUNT_WR	MBOX4C1	394643153	0	0	0
CAS_COUNT_RD	MBOX5C0	565691821	0	0	0
CAS_COUNT_WR	MBOX5C1	394667249	0	0	0
CAS_COUNT_RD	MBOX6C0	565545031	0	0	0
CAS_COUNT_WR	MBOX6C1	394475734	0	0	0
CAS_COUNT_RD	MBOX7C0	565684926	0	0	0
CAS_COUNT_WR	MBOX7C1	394625582	0	0	0

Event	Counter	Sum	Min	Max	Avg
INSTR_RETIRED_ANY STAT	FIXC0	86285205997	4585	1279573502	5.676658
CPU_CLK_UNHALTED_CORE STAT	FIXC1	720357096927	17183	9503596926	4.739191
CPU_CLK_UNHALTED_REF STAT	FIXC2	548464184352	12576	7235373600	3.608317
TOPDOWN_SLOTS STAT	FIXC3	3586852836495	64205	47515962200	2.359772
CAS_COUNT_RD STAT	MBOX0C0	1131151698	0	565664152	7.441787

CAS_COUNT_WR STAT	MBOX0C1	790129862	0	395297060	5.198223
CAS_COUNT_RD STAT	MBOX1C0	1131368409	0	565695080	7.443213
CAS_COUNT_WR STAT	MBOX1C1	790131145	0	395358301	5.198231
CAS_COUNT_RD STAT	MBOX2C0	1131223082	0	565702248	7.442257
CAS_COUNT_WR STAT	MBOX2C1	789790486	0	395129483	5.195990
CAS_COUNT_RD STAT	MBOX3C0	1131410934	0	565742668	7.443493
CAS_COUNT_WR STAT	MBOX3C1	789953933	0	395202355	5.197065
CAS_COUNT_RD STAT	MBOX4C0	1131095868	0	565570895	7.441420
CAS_COUNT_WR STAT	MBOX4C1	789684268	0	395041115	5.195291
CAS_COUNT_RD STAT	MBOX5C0	1131636477	0	565944656	7.444977
CAS_COUNT_WR STAT	MBOX5C1	789969836	0	395302587	5.197170
CAS_COUNT_RD STAT	MBOX6C0	1131135233	0	565590202	7.441679
CAS_COUNT_WR STAT	MBOX6C1	789535002	0	395059268	5.194309
CAS_COUNT_RD STAT	MBOX7C0	1131559078	0	565874152	7.444468
CAS_COUNT_WR STAT	MBOX7C1	789845999	0	395220417	5.196355

Metric	HWThread 0	HWThread 1	HWThread 2	HWThread
Runtime (RDTSC) [s]	3.1710	3.1710	3.1710	3.1710
Runtime unhaltd [s]	0.0012	3.9422	0.8516	3.939
Clock [MHz]	3157.8904	3143.3382	3177.8188	3143.805
CPI	7.2623	7.9705	13.5963	8.539
Memory read bandwidth [MBytes/s]	91322.3211	0	0	
Memory read data volume [GBytes]	289.5870	0	0	
Memory write bandwidth [MBytes/s]	63725.2563	0	0	
Memory write data volume [GBytes]	202.0755	0	0	
Memory bandwidth [MBytes/s]	155047.5774	0	0	
Memory data volume [GBytes]	491.6625	0	0	

Metric	Sum	Min	Max	
Runtime (RDTSC) [s] STAT	481.9920	3.1710	3.1710	
Runtime unhaltd [s] STAT	300.8552	7.176428e-06	3.9691	
Clock [MHz] STAT	480262.0290	3092.9626	3271.8067	3
CPI STAT	1392.3753	0.5205	47.9745	
Memory read bandwidth [MBytes/s] STAT	182664.5690	0	91342.2479	1
Memory read data volume [GBytes] STAT	579.2372	0	289.6502	
Memory write bandwidth [MBytes/s] STAT	127534.8890	0	63809.6327	
Memory write data volume [GBytes] STAT	404.4186	0	202.3431	
Memory bandwidth [MBytes/s] STAT	310199.4580	0	155151.8806	2
Memory data volume [GBytes] STAT	983.6558	0	491.9933	

- All memory related performance counters are only accounted on first CPU core on the socket

- Validity check

```

Socket 0:
Memory read bandwidth: 91322.3211 MBytes/s
Memory write bandwidth: 63725.2563 MBytes/s
+
-----
155047.5774 MBytes/s
Memory bandwidth: 155047.5774 MBytes/s

Memory write data volume socket 0: 202.0755 GBytes
Memory write data volume socket 1: 202.3431 GBytes
+
-----
404.4186 GBytes
Memory write data volume [GBytes] STAT: 404.4186 GBytes

Memory read data volume socket 0: 289.5870 GBytes
Memory read data volume socket 1: 289.6502 GBytes
+
-----
579.2372 GBytes
Memory read data volume [GBytes] STAT: 579.2372 GBytes

#Elements/vec = 1,000,000,000
#Bytes/Element = 8
#Bytes/vec = 1,000,000,000 * 8 = 8,000,000,000
#GBytes/vec = 8,000,000,000 / 1,000,000,000 = 8
#Num repetition = 10
#Num worm up rounds = 1
#Num iterations = Num repetition + Num worm up rounds = 11

Copy: 1 vec. read, 1 vec. write
Scale: 1 vec. read, 1 vec. write
Add: 2 vec. read, 1 vec. write
Triad: 2 vec. read, 1 vec. write

Write:
-----
* vec initialization: 3 vec. write * 8 GBytes/vec = 24 GB
* timer initialization: 1 vec. write * 8 GBytes/vec = 8 GB
* iteration: 4 vec. write * 11 iterations * 8 GBytes/vec = 352 GB
-> Sum: 384 GB
~ 404.4186 Memory write data volume [GBytes] STAT

Read:
-----
* timer initialization: 1 vec. read * 8 GBytes/vec = 8 GB

```

```

* iteration:          6 vec. read * 11 iterations * 8 GBytes/vec = 528 GB
* validation:         3 vec. read          * 8 GBytes/vec = 24 GB
-> Sum: 560 GB
~ 579.2372 Memory read data volume [GBytes] STAT

```

Example likwid-perfctr performance group UPI on benchmark stream

- Prepare environment

```

module purge
module add compiler/intel/2022

```

- Build stream benchmark with Intel compiler

```

icc -std=c11 -Ofast -xHost -ipo -qopenmp \
    stream.c -o stream

```

- Set up OpenMP environment

```

export OMP_NUM_THREADS=76
export KMP_AFFINITY="verbose,granularity=core,respect,scatter"

```

- List available performance groups

```

likwid-perfctr -a

...
                UPI  UPI data traffic
...

```

- Get detailed information on performance groups

```

likwid-perfctr -H --group UPI

```

Group UPI:

Formulas:

Received data bandwidth [MByte/s] = $1.0E-06 * (\text{SUM}(\text{RXL_FLITS_ALL_DATA}) / 9.0) * 64.0 / \text{runtime}$

Received data volume [GByte] = $1.0E-09 * (\text{SUM}(\text{RXL_FLITS_ALL_DATA}) / 9.0) * 64.0$

Sent data bandwidth [MByte/s] = $1.0E-06 * (\text{SUM}(\text{TXL_FLITS_ALL_DATA}) / 9.0) * 64.0 / \text{time}$

Sent data volume [GByte] = $1.0E-09 * (\text{SUM}(\text{TXL_FLITS_ALL_DATA}) / 9.0) * 64.0$

Total data bandwidth [MByte/s] = $1.0E-06 * ((\text{SUM}(\text{RXL_FLITS_ALL_DATA}) + \text{SUM}(\text{TXL_FLITS_ALL_DATA})) / 9.0) * 64.0 / \text{time}$

Total data volume [GByte] = $1.0E-09 * ((\text{SUM}(\text{RXL_FLITS_ALL_DATA}) + \text{SUM}(\text{TXL_FLITS_ALL_DATA})) / 9.0) * 64.0$

--

This group measures the data traffic on the UPI (socket interconnect). The group measures all filled data slots (9 slots per 64 Byte data transfer), that's why the count needs to be divided by 9. These 9 data chunks are not transferred in a single flit but there is one flit for the header and three flits for the data. The metrics show higher values as expected because the events count also different transfers which include data.

- Measure performance group UPI for benchmark **stream** on CPU Hyper-thread 0 to 151 with locally allocated memory

```
likwid-perfctr --group UPI -c 0-151 \
    numactl --localalloc \
    ./stream -n 1000000000
```

...

Function	Best Rate MB/s	Med time	Min time	Max time
Copy:	315404.1	0.050860	0.050729	0.052256
Scale:	313710.9	0.051093	0.051002	0.052351
Add:	319827.1	0.075225	0.075041	0.077298
Triad:	317870.4	0.075667	0.075502	0.077569

...

Metric	Sum	Min	Max	Avg
Runtime (RDTSC) [s] STAT	489.0904	3.2177	3.2177	3.2
Runtime unhaltd [s] STAT	300.4339	0	3.9619	1.9
Clock [MHz] STAT	476654.1148	3094.0504	3257.4571	3135.8
CPI STAT	2147.5606	0.4915	202.9889	14.1
Received data bandwidth [MByte/s] STAT	321.8206	0	181.1360	2.1
Received data volume [GByte] STAT	1.0355	0	0.5828	0.0
Sent data bandwidth [MByte/s] STAT	321.8082	0	181.1440	2.1
Sent data volume [GByte] STAT	1.0355	0	0.5829	0.0
Total data bandwidth [MByte/s] STAT	643.6286	0	321.8285	4.2
Total data volume [GByte] STAT	2.0710	0	1.0355	0.0

– Total UPI data volume and bandwidth are very low => low inter socket traffic

- Measure performance group UPI for benchmark **stream** on CPU Hyper-thread 0 to 151 with all allocated memory in NUMA domain 0

```
likwid-perfctr --group UPI -c 0-151 \
    numactl --membind=0 \
    ./stream -n 1000000000
```

...

Function	Best Rate MB/s	Med time	Min time	Max time
Copy:	112951.9	0.141838	0.141653	0.141931
Scale:	112253.1	0.142650	0.142535	0.142793
Add:	114055.6	0.210627	0.210424	0.210888
Triad:	114081.4	0.210572	0.210376	0.210740

...

Metric	Sum	Min	Max	
Runtime (RDTSC) [s] STAT	1333.6328	8.7739	8.7739	
Runtime unhaltd [s] STAT	834.0245	2.164660e-06	10.9996	
Clock [MHz] STAT	481297.2958	3093.2066	3247.5981	31
CPI STAT	3052.8227	0.4719	244.3819	
Received data bandwidth [MByte/s] STAT	56071.5598	0	33598.6098	3
Received data volume [GByte] STAT	491.9660	0	294.7907	
Sent data bandwidth [MByte/s] STAT	56071.5652	0	33598.5985	3
Sent data volume [GByte] STAT	491.9661	0	294.7906	
Total data bandwidth [MByte/s] STAT	112143.1250	0	56071.5765	7
Total data volume [GByte] STAT	983.9321	0	491.9662	

- Total UPI data volume and bandwidth are very high => high inter socket traffic
- Memory bandwidth halved