

Tools/likwid/example_bench_Ice_Lake

Example: likwid-bench on Intel Xeon Ice Lake

- List available micro benchmarks

```
likwid-bench -a | \
    grep -e stream_avx -e stream_mem_avx
```

```
stream_avx - Double-precision stream triad  $A(i) = B(i)*c + C(i)$ , optimized for AVX
stream_avx512 - Double-precision stream triad  $A(i) = B(i)*c + C(i)$ , optimized for AVX-
stream_avx512_fma - Double-precision stream triad  $A(i) = B(i)*c + C(i)$ , optimized for AVX
stream_avx_fma - Double-precision stream triad  $A(i) = B(i)*c + C(i)$ , optimized for AVX
stream_mem_avx - Double-precision stream triad  $A(i) = B(i)*c + C(i)$ , uses AVX and non-t
stream_mem_avx512 - Double-precision stream triad  $A(i) = B(i)*c + C(i)$ , uses AVX-
stream_mem_avx_fma - Double-precision stream triad  $A(i) = B(i)*c + C(i)$ , optimized for
```

- List properties of test

```
likwid-bench -l stream_mem_avx_fma
```

```
Name: stream_mem_avx_fma
```

```
Description: Double-precision stream triad  $A(i) = B(i)*c + C(i)$ , optimized for AVX FMAs
```

```
Number of streams: 3
```

```
Loop stride: 16
```

```
Data Type: Double precision float
```

```
Flops per element: 2
```

```
Bytes per element: 24
```

```
Load bytes per element: 16
```

```
Store bytes per element: 8
```

```
Load Ops: 2
```

```
Store Ops: 1
```

```
Constant instructions: 17
```

```
Loop instructions: 15
```

```
Loop micro Ops (OPs): 22
```

- List available thread domains

```
likwid-bench -p
```

```

Number of Domains 9
Domain 0:
    Tag N: 0 76 1 77 2 78 3 79 4 80 5 81 6 82 7 83 8 84 9 85 10 86 11 87 12 88 13 8
Domain 1:
    Tag S0: 0 76 1 77 2 78 3 79 4 80 5 81 6 82 7 83 8 84 9 85 10 86 11 87 12 88 13
Domain 2:
    Tag S1: 38 114 39 115 40 116 41 117 42 118 43 119 44 120 45 121 46 122 47 123 4
Domain 3:
    Tag D0: 0 76 1 77 2 78 3 79 4 80 5 81 6 82 7 83 8 84 9 85 10 86 11 87 12 88 13
Domain 4:
    Tag D1: 38 114 39 115 40 116 41 117 42 118 43 119 44 120 45 121 46 122 47 123 4
Domain 5:
    Tag C0: 0 76 1 77 2 78 3 79 4 80 5 81 6 82 7 83 8 84 9 85 10 86 11 87 12 88 13
Domain 6:
    Tag C1: 38 114 39 115 40 116 41 117 42 118 43 119 44 120 45 121 46 122 47 123 4
Domain 7:
    Tag M0: 0 76 1 77 2 78 3 79 4 80 5 81 6 82 7 83 8 84 9 85 10 86 11 87 12 88 13
Domain 8:
    Tag M1: 38 114 39 115 40 116 41 117 42 118 43 119 44 120 45 121 46 122 47 123 4

```

- Run micro benchmark `stream_mem_avx_fma` on memory domain 0 to 2 with 38 threads each skipping hyperthreads

```

likwid-bench \
    -t stream_mem_avx_fma \
    -i 40 \
    -w M0:16GB:38:1:2 \
    -w M1:16GB:38:1:2

```

```

Warning: Sanitizing vector length to a multiple of the loop stride 16 and thread count
Allocate: Process running on hwthread 0 (Domain M0) - Vector length 666666528/533333222
Allocate: Process running on hwthread 0 (Domain M0) - Vector length 666666528/533333222
Allocate: Process running on hwthread 0 (Domain M0) - Vector length 666666528/533333222
Initialization: First thread in domain initializes the whole stream
Warning: Sanitizing vector length to a multiple of the loop stride 16 and thread count
Allocate: Process running on hwthread 38 (Domain M1) - Vector length 666666528/533333222
Allocate: Process running on hwthread 38 (Domain M1) - Vector length 666666528/533333222
Allocate: Process running on hwthread 38 (Domain M1) - Vector length 666666528/533333222
Initialization: First thread in domain initializes the whole stream

```

```

-----
LIKWID MICRO BENCHMARK
Test: stream_mem_avx_fma
-----

```

```

Using 2 work groups
Using 76 threads
-----

```

Running without Marker API. Activate Marker API with `-m` on commandline.

```

Group: 0 Thread 5 Global Thread 5 running on hwthread 5 - Vector length 17543856 Offset
Group: 0 Thread 1 Global Thread 1 running on hwthread 1 - Vector length 17543856 Offset
Group: 0 Thread 3 Global Thread 3 running on hwthread 3 - Vector length 17543856 Offset
Group: 0 Thread 4 Global Thread 4 running on hwthread 4 - Vector length 17543856 Offset
Group: 0 Thread 2 Global Thread 2 running on hwthread 2 - Vector length 17543856 Offset
Group: 0 Thread 11 Global Thread 11 running on hwthread 11 - Vector length 17543856 Off
Group: 0 Thread 9 Global Thread 9 running on hwthread 9 - Vector length 17543856 Offset
Group: 0 Thread 14 Global Thread 14 running on hwthread 14 - Vector length 17543856 Off
Group: 0 Thread 21 Global Thread 21 running on hwthread 21 - Vector length 17543856 Off
Group: 0 Thread 35 Global Thread 35 running on hwthread 35 - Vector length 17543856 Off
Group: 0 Thread 22 Global Thread 22 running on hwthread 22 - Vector length 17543856 Off
Group: 0 Thread 12 Global Thread 12 running on hwthread 12 - Vector length 17543856 Off
Group: 0 Thread 20 Global Thread 20 running on hwthread 20 - Vector length 17543856 Off
Group: 1 Thread 1 Global Thread 39 running on hwthread 39 - Vector length 17543856 Offs
Group: 0 Thread 34 Global Thread 34 running on hwthread 34 - Vector length 17543856 Off
Group: 1 Thread 2 Global Thread 40 running on hwthread 40 - Vector length 17543856 Offs
Group: 0 Thread 13 Global Thread 13 running on hwthread 13 - Vector length 17543856 Off
Group: 1 Thread 3 Global Thread 41 running on hwthread 41 - Vector length 17543856 Offs
Group: 0 Thread 18 Global Thread 18 running on hwthread 18 - Vector length 17543856 Off
Group: 1 Thread 5 Global Thread 43 running on hwthread 43 - Vector length 17543856 Offs
Group: 0 Thread 16 Global Thread 16 running on hwthread 16 - Vector length 17543856 Off
Group: 0 Thread 6 Global Thread 6 running on hwthread 6 - Vector length 17543856 Offset
Group: 1 Thread 25 Global Thread 63 running on hwthread 63 - Vector length 17543856 Off
Group: 0 Thread 23 Global Thread 23 running on hwthread 23 - Vector length 17543856 Off
Group: 0 Thread 15 Global Thread 15 running on hwthread 15 - Vector length 17543856 Off
Group: 1 Thread 13 Global Thread 51 running on hwthread 51 - Vector length 17543856 Off
Group: 0 Thread 29 Global Thread 29 running on hwthread 29 - Vector length 17543856 Off
Group: 0 Thread 24 Global Thread 24 running on hwthread 24 - Vector length 17543856 Off
Group: 0 Thread 27 Global Thread 27 running on hwthread 27 - Vector length 17543856 Off
Group: 0 Thread 30 Global Thread 30 running on hwthread 30 - Vector length 17543856 Off
Group: 0 Thread 28 Global Thread 28 running on hwthread 28 - Vector length 17543856 Off
Group: 1 Thread 23 Global Thread 61 running on hwthread 61 - Vector length 17543856 Off
Group: 0 Thread 8 Global Thread 8 running on hwthread 8 - Vector length 17543856 Offset
Group: 0 Thread 36 Global Thread 36 running on hwthread 36 - Vector length 17543856 Off
Group: 1 Thread 36 Global Thread 74 running on hwthread 74 - Vector length 17543856 Off
Group: 0 Thread 31 Global Thread 31 running on hwthread 31 - Vector length 17543856 Off
Group: 0 Thread 32 Global Thread 32 running on hwthread 32 - Vector length 17543856 Off
Group: 0 Thread 33 Global Thread 33 running on hwthread 33 - Vector length 17543856 Off
Group: 0 Thread 19 Global Thread 19 running on hwthread 19 - Vector length 17543856 Off
Group: 1 Thread 10 Global Thread 48 running on hwthread 48 - Vector length 17543856 Off
Group: 1 Thread 17 Global Thread 55 running on hwthread 55 - Vector length 17543856 Off
Group: 1 Thread 4 Global Thread 42 running on hwthread 42 - Vector length 17543856 Offs
Group: 1 Thread 6 Global Thread 44 running on hwthread 44 - Vector length 17543856 Offs
Group: 1 Thread 7 Global Thread 45 running on hwthread 45 - Vector length 17543856 Offs

```

```

Group: 0 Thread 7 Global Thread 7 running on hwthread 7 - Vector length 17543856 Offset
Group: 1 Thread 0 Global Thread 38 running on hwthread 38 - Vector length 17543856 Offs
Group: 0 Thread 26 Global Thread 26 running on hwthread 26 - Vector length 17543856 Off
Group: 0 Thread 17 Global Thread 17 running on hwthread 17 - Vector length 17543856 Off
Group: 1 Thread 19 Global Thread 57 running on hwthread 57 - Vector length 17543856 Off
Group: 1 Thread 14 Global Thread 52 running on hwthread 52 - Vector length 17543856 Off
Group: 1 Thread 11 Global Thread 49 running on hwthread 49 - Vector length 17543856 Off
Group: 1 Thread 12 Global Thread 50 running on hwthread 50 - Vector length 17543856 Off
Group: 1 Thread 18 Global Thread 56 running on hwthread 56 - Vector length 17543856 Off
Group: 1 Thread 16 Global Thread 54 running on hwthread 54 - Vector length 17543856 Off
Group: 1 Thread 9 Global Thread 47 running on hwthread 47 - Vector length 17543856 Offs
Group: 1 Thread 15 Global Thread 53 running on hwthread 53 - Vector length 17543856 Off
Group: 1 Thread 8 Global Thread 46 running on hwthread 46 - Vector length 17543856 Offs
Group: 0 Thread 25 Global Thread 25 running on hwthread 25 - Vector length 17543856 Off
Group: 0 Thread 0 Global Thread 0 running on hwthread 0 - Vector length 17543856 Offset
Group: 1 Thread 20 Global Thread 58 running on hwthread 58 - Vector length 17543856 Off
Group: 1 Thread 28 Global Thread 66 running on hwthread 66 - Vector length 17543856 Off
Group: 1 Thread 29 Global Thread 67 running on hwthread 67 - Vector length 17543856 Off
Group: 1 Thread 26 Global Thread 64 running on hwthread 64 - Vector length 17543856 Off
Group: 0 Thread 10 Global Thread 10 running on hwthread 10 - Vector length 17543856 Off
Group: 1 Thread 31 Global Thread 69 running on hwthread 69 - Vector length 17543856 Off
Group: 1 Thread 22 Global Thread 60 running on hwthread 60 - Vector length 17543856 Off
Group: 1 Thread 27 Global Thread 65 running on hwthread 65 - Vector length 17543856 Off
Group: 1 Thread 30 Global Thread 68 running on hwthread 68 - Vector length 17543856 Off
Group: 1 Thread 21 Global Thread 59 running on hwthread 59 - Vector length 17543856 Off
Group: 1 Thread 24 Global Thread 62 running on hwthread 62 - Vector length 17543856 Off
Group: 1 Thread 33 Global Thread 71 running on hwthread 71 - Vector length 17543856 Off
Group: 1 Thread 32 Global Thread 70 running on hwthread 70 - Vector length 17543856 Off
Group: 1 Thread 37 Global Thread 75 running on hwthread 75 - Vector length 17543856 Off
Group: 1 Thread 34 Global Thread 72 running on hwthread 72 - Vector length 17543856 Off
Group: 1 Thread 35 Global Thread 73 running on hwthread 73 - Vector length 17543856 Off
Group: 0 Thread 37 Global Thread 37 running on hwthread 37 - Vector length 17543856 Off

```

```

Cycles:                9639007356
CPU Clock:             2394359220
Cycle Clock:           2394359220
Time:                  4.025715e+00 sec
Iterations:            3040
Iterations per thread: 40
Inner loop executions: 1096491
Size (Byte):           31999993344
Size per thread:       421052544
Number of Flops:       106666644480
MFlops/s:              26496.32
Data volume (Byte):    1279999733760

```

```

MByte/s:          317955.89
Cycles per update: 0.180731
Cycles per cacheline: 1.445851
Loads per update:  2
Stores per update:  1
Load bytes per element: 16
Store bytes per elem.: 8
Load/store ratio:   2.00
Instructions:       49999989617
UOPs:              73333318080

```

- Loop over number of cores used

```

NUM_DOMAINS=2
for ((NUM_CORES=1; NUM_CORES <= 76; NUM_CORES++))
do
    # Distribute cores in round robin mode
    declare -i -a CORES_PER_DOMAIN=()
    for ((COUNT=0; COUNT < NUM_CORES; COUNT++))
    do
        let CORES_PER_DOMAIN[$((COUNT % NUM_DOMAINS))]++
    done
    COMMAND=(
        likwid-bench
            -t stream_mem_avx_fma
            -i 100
    )
    for ((DOMAIN=0; DOMAIN<NUM_DOMAINS; DOMAIN++))
    do
        if [[ ${CORES_PER_DOMAIN[${DOMAIN}]} -gt 0 ]]; then
            COMMAND+=( -w M${DOMAIN}:16GB:${CORES_PER_DOMAIN[${DOMAIN}]}:1:2 )
        fi
    done
    echo "${COMMAND[@]}"
    "${COMMAND[@]}"
done |
grep -e "^likwid-bench" -e "Using .* threads" -e "MByte/s:"

likwid-bench -t stream_mem_avx_fma -i 100 -w M0:16GB:1:1:2
likwid-bench -t stream_mem_avx_fma -i 100 -w M0:16GB:1:1:2 -w M1:16GB:1:1:2
likwid-bench -t stream_mem_avx_fma -i 100 -w M0:16GB:2:1:2 -w M1:16GB:1:1:2
likwid-bench -t stream_mem_avx_fma -i 100 -w M0:16GB:2:1:2 -w M1:16GB:2:1:2
...
likwid-bench -t stream_mem_avx_fma -i 100 -w M0:16GB:38:1:2 -w M1:16GB:37:1:2
likwid-bench -t stream_mem_avx_fma -i 100 -w M0:16GB:38:1:2 -w M1:16GB:38:1:2

```

#Cores	stream_mem_avx_fma(MByte/s)	% of Max
1	18932.74	6%
2	37606.15	12%
4	71658.90	22%
8	128337.42	40%
12	178572.55	55%
16	221021.84	68%
32	303688.83	94%
50	323435.61	100%
64	320327.16	99%
76	316941.00	98%

=> 1 core can only get about 6% of peak memory bandwidth => 12 cores are needed for 1/2 of peak memory bandwidth => Peak bandwidth is reached with 50 cores, after which the bandwidth slowly drops

- Loop over memory size used (L1 (32 kB), L2 (256 kB), L3 (25 MB))

```

for BENCHMARK in stream_mem_avx_fma stream_avx_fma
do
    echo "${BENCHMARK}"

    for (( MEM_SIZE=2; MEM_SIZE <= 16*1024*1024; MEM_SIZE*=2 ))
    do
        likwid-bench \
            -t ${BENCHMARK} \
            -i 100 \
            -w M0:${MEM_SIZE}KB:38:1:2 \
            -w M1:${MEM_SIZE}KB:38:1:2
    done 2>/dev/null |
        grep -e 'Size (Byte):' -e 'MByte/s:'
done

```

MEM_SIZE(Byte)	stream_mem_avx_fma(MByte/s)	stream_avx_fma(MByte/s)	stream_mem_avx_fma(MByte/s)
29184	77263.25	171384.46	
58368	137188.23	405930.71	
116736	238430.74	775374.46	
233472	269033.89	1249811.39	
466128	274972.44	1928920.73	
1021440	266466.69	2478547.23	
2042880	286880.57	3080760.17	
4085760	665429.01	3708201.59	
8171520	684543.73	4277030.97	
16372224	689832.27	5382984.20	

MEM_SIZE(Byte)	stream_mem_avx_fma(MByte/s)	stream_avx_fma(MByte/s)	stream_mem_avx_fma(MByte/s)
32744448	480203.78	6254640.60	
65518080	405960.67	5563724.35	
131065344	640713.77	895231.74	
262130688	607510.76	546060.93	
524261376	405670.12	362572.25	
1048551936	345566.80	332700.65	
2097133056	329452.87	315844.12	
4194295296	322903.94	308891.59	
8388590592	319732.19	310216.19	
16777210368	319405.49	308877.21	
33554420736	318836.42	308883.68	

=> **stream_avx_fma** does not use streaming stores. All stores go to the cache first. As long as you stay in the cache **stream_avx_fma** is much faster as **stream_mem_avx_fma**

=> **stream_mem_avx_fma** use streaming stores. All stores go directly to main memory. When you leave the cache **stream_mem_avx_fma** is faster than **stream_avx_fma**

=> Intel Xeon Ice Lake *SpecI2M* optimization: Use streaming stores when memory subsystem is heavily loaded (see: HotChips 2020: New 3rd Gen Intel Xeon Scalable Processor)