Università della Svizzera italiana

Institute of Computing CI

High-Performance Computing Lab **2022**

Student: Claudio Maggioni Discussed with: Gianmarco De Vita

Solution for Project 1

Due date: 12.10.2022 (midnight)

Contents

1. Explaining Memory Hierarchies *(25 Points)*

1.1. Memory Hierarchy Parameters of the Cluster

By invoking `likwid-topology` for the cache topology and `free -g` for the amount of primary memory, the following memory hierarchy parameters are found:

Main memory	62 GB
L3 cache	25 MB per socket
L2 cache	256 kB per core
L1 cache	32 kB per core

All values are reported using base 2 IEC byte units. The cluster has 2 sockets and a total of 20 cores (10 per socket). The cache topology diagram reported by `likwid-topology -g` is shown in Figure 1.

Socket 0:

0	1	2	3	4	5	6	7	8	9
32 kB	32 kB	32 kB	32 kB	32 kB	32 kB	32 kB	32 kB	32 kB	32 kB
256 kB	256 kB	256 kB	256 kB	256 kB	256 kB	256 kB	256 kB	256 kB	256 kB
25 MB									

Socket 1:

10	11	12	13	14	15	16	17	18	19
32 kB	32 kB	32 kB	32 kB	32 kB	32 kB	32 kB	32 kB	32 kB	32 kB
256 kB	256 kB	256 kB	256 kB	256 kB	256 kB	256 kB	256 kB	256 kB	256 kB
25 MB									

Figure 1: Cache topology diagram as outputted by `likwid-topology -g`. Byte sizes all in IEC units.

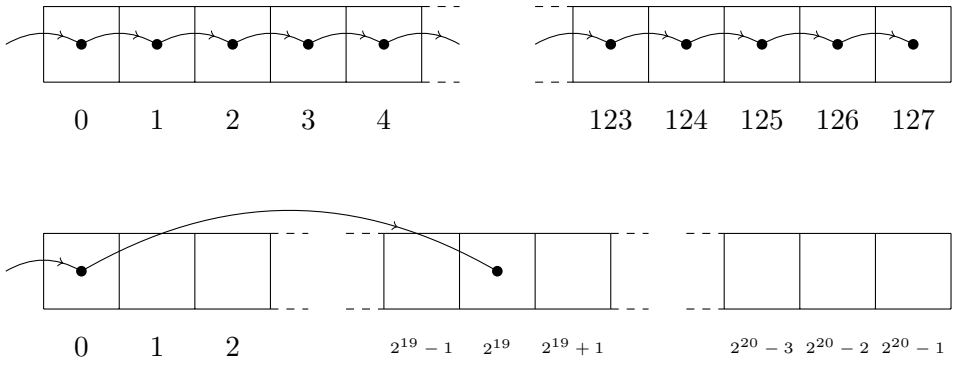

Figure 2: Memory access patterns of `membench.c` for `csize = 128` and `stride = 1` (above) and for `csize` = 2^{20} and `stride` = 2^{19} (below)

1.2. Memory Access Pattern of `membench.c`

The benchmark `membench.c` measures the average time of repeated read and write overations across a set of indices of a stack-allocated array of 32-bit signed integers. The indices vary according to the access pattern used, which in turn is defined by two variables, `csize` and `stride`. `csize` is an upper bound on the index value, i.e. (one more of) the highest index used to access the array in the pattern. `stride` determines the difference between array indexes over access iterations, i.e. a `stride` of 1 will access every array index, a `stride` of 2 will skip every other index, a `stride` of 4 will access one index then skip 3 and so on and so forth. The benchmark stops when the index to access is strictly greater than `csize - stride`.

Therefore, for `csize = 128` and `stride = 1` the array will access all indexes between 0 and 127 sequentially, and for `csize` = 2^{20} and `stride` = 2^{19} the benchmark will access index 0, then index $2^{19} - 1$. The access patterns for these two configurations are shown visually in Figure 2.

By running the `membench.c` both on my personal laptop and on the cluster, the results shown in Figure 3 are obtained. *csize* values are shown as different data series and labeled by byte size and *stride* values are mapped on the x axis by the byte-equivalent value as well[1]. For `csize = 128` = 512 bytes and `stride = 1` = 4 bytes the mean access time is 0.124 nanoseconds, while for `csize` = 2^{20} = 4MB and for `stride` = 2^{19} = 2MB the mean access time is 1.156 nanoseconds. The first set of parameters performs well thanks to the low *stride* value, thus achiev-

[1]Byte values are a factor of 4 greater than the values used in the code and in Figure 3. This is due to the fact that the array elements used in the benchmark are 32-bit signed integers, which take up 4 bytes each.

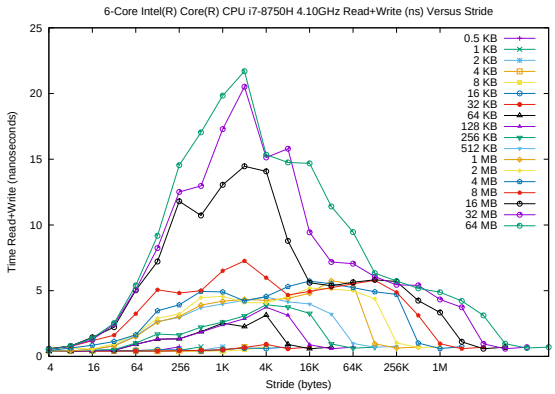

(a) Personal laptop

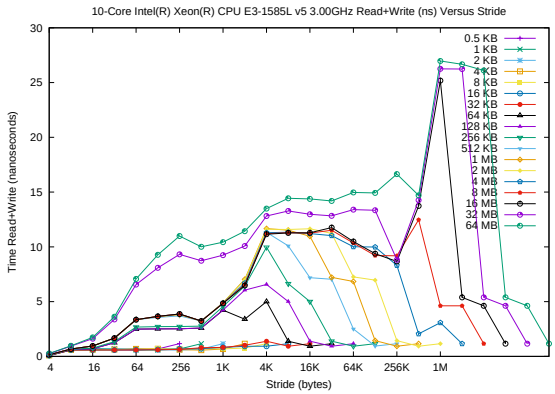

(b) Cluster

Figure 3: Results of the `membench.c` benchmark for both my personal laptop (in Figure 3a) and the cluster (in Figure 3b).

ing very good space locality and maximizing cache hits. However, the second set of parameters achieves good performance as well thanks to the few values accessed with each pass, thus improving the temporal locality of each address accessed. This observation applies for the few last data points in each data series of Figure 3, i.e. for *stride* values close to *csize*.

1.3. Analyzing Benchmark Results

The `membench.c` benchmark results for my personal laptop (Macbook Pro 2018 with a Core i7-8750H CPU) and the cluster are shown in figure 3.

The memory access graph for the cluster's benchmark results shows that temporal locality is best for small array sizes and for small `stride` values. In particular, for array memory sizes of 16MB or lower (`csize` of $4 \cdot 2^{20}$ or lower) and `stride` values of 2048 or lower the mean read+write time is less than 10 nanoseconds. Temporal locality is worst for large sizes and strides, although the largest values of `stride` for each size (like `csize / 2` and `csize / 4`) achieve better mean times for the aforementioned effect of having *stride* values close to *csize*.

The pattern that can be read from the graphs, especially the one for the cluster, shows that the *stride* axis is divided in regions showing memory access time of similar magnitude. The boundary between the first and the second region is a *stride* value of rougly 2KB, while a *stride* of 512KB roughly separates the second and the third region. The difference in performance between regions and the similarity of performance within regions suggest the threshold stride values are related to changes in the use of the cache hierarchy. In particular, the first region may characterize regions where the L1 cache, the fastest non-register memory available, is predominantly used. Then the second region might overlap with a more intense use of the L2 cache and likewise between the third region and the L3 cache.

2. Optimize Square Matrix-Matrix Multiplication *(60 Points)*

The file `matmult/dgemm-blocked.c` contains a C implementation of the blocked matrix multiplication algorithm presented in the project. A pseudocode listing of the implementation is provided in Figure 4.

In order to achieve a correct and fast execution, my implementation:

- Handles the edge cases related to the "remainders" in the matrix block division, i.e. when the division between the size of the matrix and the block size yields a remainder. Assuming only squared matrices are multiplied through the algorithm (as in the test suite provided) the

```
INPUT: A (n by n), B (n by n), n
OUTPUT: C (n by n)

s := 26 # block dimension

A_row := <matrix A converted in row major form>
C_temp := <empty s by s matrix>

for i := 0 to n by s:
    i_next := min(i + s, n)

    for j := 0 to n by s:
        j_next := min(j + s, n)

        <set all cells in C_temp to 0>

        for k := 0 to n by s:
            k_next := min(k + s, n)

            # Perform naive matrix multiplication, incrementing cells of C_temp
            # with each multiplication result
            naivemm(A_row[i, k][i_next, k_next], B[k, j][k_next, j_next],
                    C_temp[0, 0][i_next - i, j_next - j])
        end for

        C[i, j][i_next, j_next] = C_temp[0, 0][i_next - i, j_next - j]
    end for
end for
```

Figure 4: Pseudocode listing of my blocked matrix multiplication implementation. Matrix indices start from 0 (i.e. row 0 and column 0 denotes the top-left-most cell in a matrix). `M[a, b][c, d]` denotes a rectangular region of the matrix M whose top-left-most cell is the cell in M at row a and column b and whose bottom-right-most cell is the cell in M at row $c - 1$ and column $d - 1$.

block division could yield rectangular matrix blocks located in the last rows and columns of each matrix, and the bottom-right corner of the matrix will be contained in a square matrix block of the size of the remainder. The result of this process is shown in Figure 5;

- Converts matrix A into row major format. As shown in Figure 6, by having A in row major format and B in column major format, iterations across matrix block in the inner most loop of the algorithm (the one calling *naivemm*) cache hits are maximised by achieving space locality between the blocks used. This achieved approximately an increase of performance of two percentage points in terms of CPU utilization (i.e. from a baseline of 4% to 6%),
- Caches the result of each innermost iteration into a temporary matrix of block size before storing it into matrix C. This achieves better space locality when *naivemm* needs to store values in matrix C. The block size temporary matrix has virtually no stride and thus cache hits are maximised. The copy operation is implemented with bulk copy `memcpy` calls. This optimization achieves an extra half of a percentage point in terms of CPU utilization (i.e. from the 6% discussed above to a final 6.5%).

The chosen matrix block size for running the benchmark on the cluster is:

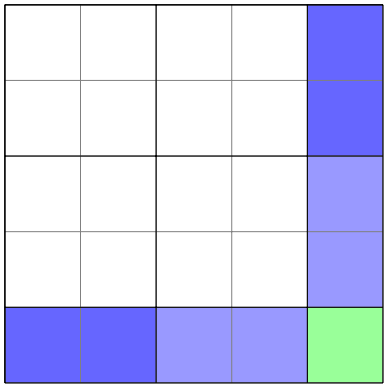

Figure 5: Result of the block division process of a square matrix of size 5 using a block size of 2. The 2-by-1 and 1-by-2 rectangular remainders are shown in blue and the square matrix of remainder size (i.e. 1) is shown in green.

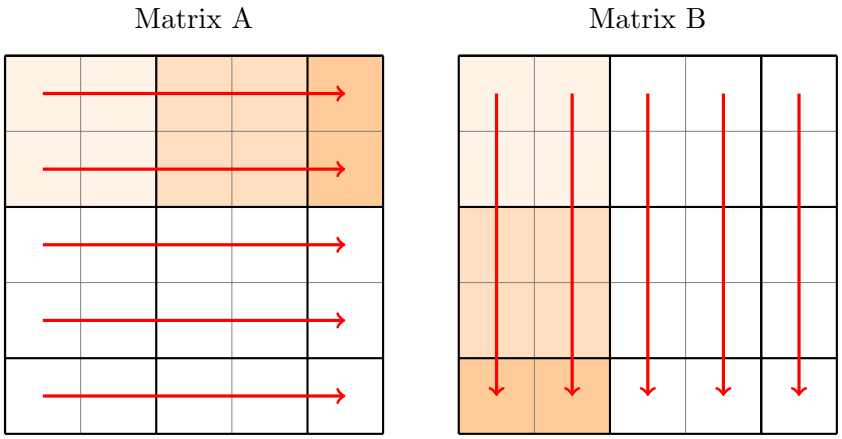

Figure 6: Inner most loop iteration of the blocked GEMM algorithm across matrices A and B. The red lines represent the "majorness" of each matrix (A is converted by the algorithm in row-major form, while B is given and used in column-major form). The shades of orange represent the blocks used in each iteration.

$$s = 26$$

as shown in the pseudocode. This value has been obtained by running an empirical binary search on the value using the benchmark as a metric, i.e. by running `./run_matrixmult.sh` several times with different values. For square blocks (i.e. the worst case) the total size for the matrix A and B sub-block and the `C_temp` temporary matrix block for C is:

$$\text{Bytes} = \text{cellSize} * s^2 * 3 = 8 * 26^2 * 3 = 16224$$

given that a double-precision floating point number, the data type used for matrix cells in the scope of this project, is 8 bytes long. The obtained total bytes size is fairly close to the L1 cache size of the processor used in the cluster (32Kb = 32768 bytes), which is expected given that the algorithm needs to exploit fast memory as much as possible. The reason the empirically best value results in a theoretical cache allocation that is only half of the complete L1 cache size is due to some real-life factors. For example, cache misses tipically result in aligned page loads which may load unnecessary data.

A potential way to exploit the different cache levels is to apply the blocked matrix algorithm iteratively multiple times. For example, OpenBLAS implements DGEMM by having two levels of matrix blocks to better exploit the L2 and L3 caches found on most processors.

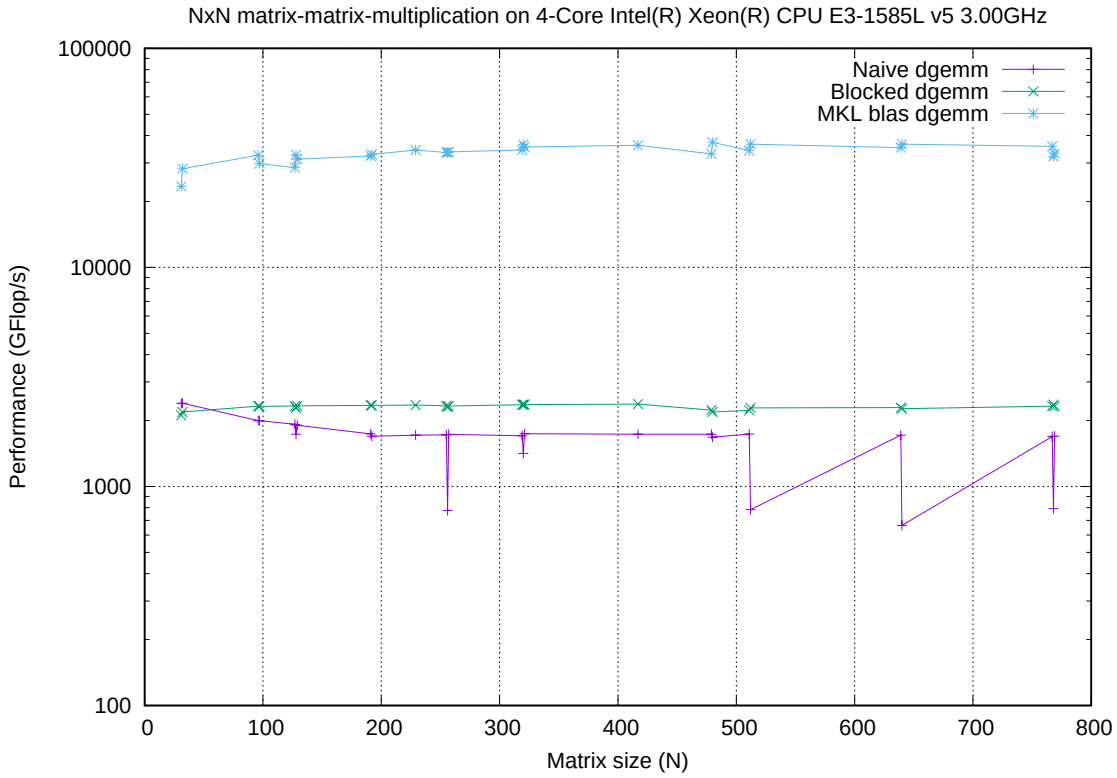

Figure 7: GFlop/s per matrix size of the matrix multiplication benchmark for the naive, blocked, and BLAS implementations. The Y-axis is log-scaled.

The results of the matrix multiplication benchmark for the naive, blocked, and BLAS implementations are shown in Figure 7 as a graph of GFlop/s over matrix size or in Figure 8 as a table. The blocked implementation achieves on average 50% more FLOPS than the naive implementation thanks to the optimisations in space and temporal cache locality described. However, the blocked implementation achives less than a tenth of FLOPS compared to Intel MKL BLAS based one due to the microarchitecture optimization the latter one is able to exploit.

I was unable to run this benchmark suite on my personal machine due to Intel MKL installation issues that prevented the code to compile.

Size	Naive		Blocked		BLAS	
	MFLOPS	% CPU	MFLOPS	% CPU	MFLOPS	% CPU
31	2393.33	6.50	2112.63	5.74	23449.20	63.72
32	2400.13	6.52	2187.44	5.94	28198.90	76.63
96	1998.74	5.43	2325.39	6.32	32542.30	88.43
97	1996.01	5.42	2322.81	6.31	29801.30	80.98
127	1923.81	5.23	2330.30	6.33	28557.80	77.60
128	1731.98	4.71	2282.93	6.20	32643.30	88.70
129	1903.31	5.17	2334.25	6.34	31198.20	84.78
191	1736.78	4.72	2345.91	6.37	32247.30	87.63
192	1694.44	4.60	2345.38	6.37	32830.60	89.21
229	1715.10	4.66	2351.01	6.39	34360.90	93.37
255	1720.39	4.67	2335.21	6.35	33477.70	90.97
256	777.65	2.11	2306.48	6.27	33473.90	90.96
257	1729.27	4.70	2330.68	6.33	33686.50	91.54
319	1704.80	4.63	2360.03	6.41	34335.20	93.30
320	1414.84	3.84	2364.53	6.43	36438.10	99.02
321	1741.30	4.73	2366.38	6.43	35433.70	96.29
417	1733.00	4.71	2378.34	6.46	36133.70	98.19
479	1731.17	4.70	2233.05	6.07	32951.40	89.54
480	1678.77	4.56	2187.87	5.95	37260.00	101.25
511	1733.60	4.71	2224.61	6.05	34128.00	92.74
512	782.96	2.13	2284.85	6.21	36526.40	99.26
639	1714.42	4.66	2292.78	6.23	35249.20	95.79
640	663.42	1.80	2264.70	6.15	36538.70	99.29
767	1690.82	4.59	2324.83	6.32	35718.50	97.06
768	792.04	2.15	2363.92	6.42	32116.80	87.27
769	1696.95	4.61	2321.31	6.31	33033.90	89.77

Figure 8: MFlop/s and CPU utlisation per matrix size of the matrix multiplication benchmark for the naive, blocked, and BLAS implementations.