LIBXSMM Samples

CP2K Artificial Benchmark

The first code sample given for LIBXSMM was a performance reproducer exercising the same set of kernels usually generated for CP2K's SMM library. The code sample attempted to model the way "matrix stacks" are processed in CP2K, however there are two different code paths in CP2K: (1) the "main" code path used when processing stacks on the host-side, and (2) a code path targeting offload devices. Beside of the host-sided parallelization via MPI (and perhaps OpenMP), the secondly mentioned code path relies on an additional level of parallelization (which is obviously necessary to drive a potentially highly parallel offload device). Also, the additional level of parallelism is not exactly "nested" in the sense that it participates on sharing the same resources as the host-side. In fact, this "artificial benchmark" (cp2k code sample) is modeling a code path as utilized in the secondly mentioned case (offload device).

Dispatch (Microbenchmark)

This code sample benchmarks the performance of (1) the dispatch mechanism, and (2) the time needed to JIT-generate code for the first time. Both mechanisms are relevant when replacing GEMM calls (see Call Wrapper section of the reference documentation), or in any case of calling LIBXSMM's native GEMM functionality.

Command Line Interface (CLI)

- Optionally takes the number of dispatches/code-generations (default: 10000).
- Optionally takes the number of threads (default: 1).

Measurements (Benchmark)

- Duration of an empty function call (serves as a reference timing).
- Duration to find an already generated kernel (cached/non-cached).
- Duration to JIT-generate a GEMM kernel.

In case of a multi-threaded benchmark, the timings represent a highly contended request (worst case). For thread-scaling, it can be observed that read-only accesses (code dispatch) stay roughly with a constant duration whereas write-accesses (code generation) are serialized and hence the duration scales linearly with the number of threads.

Image Convolution

This code sample aims to provide a simple piece of code, which takes an image and produces a visual result as well. For the convolution, LIBXSMM's DNN-domain is used. This sample code cannot use multiple threads (therefore OMP=0) or JIT code generation since parallelization and JIT-vectorization in the DNN-domain are per multiple images and channels respectively. JIT code is vectorized over image channels according to the native vector-width of the processor hence the sample code falls back to a high-level implementation. The code processes only a single image which consists of a single channel (eventually multiple times as per `nrepeat`).

NOTE: Multicore and JIT code can be only leveraged with code changes and input data that consists of multiple images or channels (a.k.a. "deep neural networks" or "deep learning"). Please note the collection of DNN code samples, which achieves both of which.

The executable can run with the following arguments (all arguments are optional):

> iconv [<filename-in> [<nrepeat> [<kw> [<kh>] [<filename-out>]]]]

For stable timing (benchmark), the key operation (convolution) may be repeated (`nrepeat`). Further, `kw` and `kh` can specify the kernel-size of the convolution. The `filename-in` and `filename-out` name MHD-files (see Meta Image File I/O) used as input and output respectively. The `filename-in` may not exist, and specify the image resolution (`w[xh]` where the file `wxh.mhd` is generated in this case).

To load an image from a familiar format (JPG, PNG, etc.), one may save the raw data using for instance IrfanView and rely on a "header-only" MHD-file (plain text). This may look like:

```
NDims = 2
DimSize = 202 134
ElementType = MET_UCHAR
ElementNumberOfChannels = 1
ElementDataFile = mhd_image.raw
```

In the above case, a single channel (gray-scale) 202x134-image is described with pixel data stored separately (`mhd_image.raw`). Multi-channel images are expected to interleave the pixel data. The pixel type is per `libxsmm_mhd_elemtype` (libxsmm_mhd.h).

NEK Sample Collection

This directory contains kernels taken from Nek{Box,5000}. They aim to represent most of the matrix-matrix workloads.

Please note that the mxm_std.f source code is protected by an (US) GOVERNMENT LICENSE, and under the copyright of the University of Chicago.

stpm

Small tensor-product multiple (stpm) replicates the axhelm kernel, which computes the Laplacian with spectral elements. Usage:

```
./stpm m n k size1 size
```

The elements are m-by-n-by-k, mode picks the LIBXSMM interface used, and size scales the number of spectral elements.

rstr

Restriction operator transforms elements from one size to another. This occurs in multi-grid, the convection operator, and, when the sizes are the same, the local Schwarz solves. Usage:

```
./rstr m n k mm nn kk size1 size
```

The input elements are m-by-n-by-k and the output elements are mm-by-nn-by-kk. When m=mm, n=nn, k=kk, this half of a Schwarz solve.

Scratch Memory Allocation (Microbenchmark)

This code sample aims to benchmark the performance of the scratch memory allocation. This facility is a viable option to satisfy the need for temporary memory when using the DNN domain of LIBXSMM (small convolutions). Although any kind of readable/writable buffer can be bound to a convolution handle, LIBXSMM's `libxsmm_aligned_scratch` features a thread-safe linear allocator mechanism which can help to lower allocation overhead.

SMM Sample Collection

This collection of code samples exercises different memory streaming cases when performing the matrix multiplication $C_{mxn} = alpha \cdot A_{mxk} \cdot B_{kxn} + beta \cdot C_{mxn}$: (1) streaming the matrices A, B, and C which is usually referred as batched matrix multiplication, (2) streaming the inputs A and B but accumulating C within cache, (3) streaming the A and C matrices while B is kept in cache, (4) streaming the B and C matrices while A is kept in cache, and (4) not streaming any of the operands but repeating the very same multiplication until the requested number of matrix multiplications has been completed.

Beside of measuring the duration of a test case, the performance is presented in GFLOPS/s. As an alternative metric, the memory bandwidth is given (the artificial "cached" case omits to present the cache-memory bandwidth). The "pseudo-performance" given in FLOPS/cycle is an artificial scoring, it not only uses a non-standard formula for calculating the FLOPS ($2 * M * N * K - M * N$ rather than $2 * M * N * K$) but also relies on pseudo clock cycles:

```
$ ./specialized.sh 0
m=32 n=32 k=32 size=87381 memory=2048.0 MB (DP)

Batched (A,B,C)...
        pseudo-perf.: 10.7 FLOPS/cycle
        performance: 23.9 GFLOPS/s
        bandwidth: 11.1 GB/s
        duration: 239 ms
Finished
```

There are two sub collections of samples codes: (1) a collection of C++ code samples showing either BLAS, Compiler-generated code (inlined code), LIBXSMM/dispatched, LIBXSMM/specialized functions to carry out the multiplication, and (2) a Fortran sample code showing BLAS versus LIBXSMM including some result validation.

C/C++ Code Samples: Command Line Interface (CLI)

- Takes an optional number (1st arg.) to select the streaming-case (0...8)

- Optionally takes the M, N, and K parameter of the GEMM in this order
- If only M is supplied, the N and K "inherit" the M-value
- Example I (A,B,C): ./specialized.sh 0 16 8 9
- Example II (A,B): ./specialized.sh 6 16

Fortran Code Sample: Command Line Interface (CLI)

- Optionally takes the M, N, and K parameter of the GEMM in this order
- Optional problem size (in MB) of the workload; M/N/K must have been supplied
- Optional total problem size (in MB) implying the number of repeated run
- If only M is supplied, the N and K are "inheriting" the M-value
- Shows the performance of each of the streaming cases
- Example I: ./smm.sh 16 8 9 1024 16384
- Example II: ./smm.sh 16

SPECFEM Sample

This sample contains a dummy example from a spectral-element stiffness kernel taken from SPECFEM3D_GLOBE.

It is based on a 4th-order, spectral-element stiffness kernel for simulations of elastic wave propagation through the Earth. Matrix sizes used are (25,5), (5,25) and (5,5) determined by different cut-planes through a three dimensional (5,5,5)-element with a total of 125 GLL points.

Usage Step-by-Step

This example needs the LIBXSMM library to be built with static kernels, using MNK="5 25" (for matrix size (5,25), (25,5) and (5,5)).

1. In LIBXSMM root directory, compile the library with:

- general default compilation:

```
make MNK="5 25" ALPHA=1 BETA=0
```

additional compilation examples are:

- compilation using only single precision version & aggressive optimization:

```
make MNK="5 25" ALPHA=1 BETA=0 PRECISION=1 OPT=3
```

- for Sandy Bridge CPUs:

```
make MNK="5 25" ALPHA=1 BETA=0 PRECISION=1 OPT=3 AVX=1
```

- for Haswell CPUs:

```
make MNK="5 25" ALPHA=1 BETA=0 PRECISION=1 OPT=3 AVX=2
```

- for Knights Corner (KNC) (and thereby creating a Sandy Bridge version):

```
make MNK="5 25" ALPHA=1 BETA=0 PRECISION=1 OPT=3 AVX=1 \
OFFLOAD=1 KNC=1
```

- installing libraries into a sub-directory workstation/:

```
make MNK="5 25" ALPHA=1 BETA=0 PRECISION=1 OPT=3 AVX=1 \
OFFLOAD=1 KNC=1 \
PREFIX=workstation/ install-minimal
```

1. Compile this example code by typing:

- for default CPU host:

```
cd sample/specfem
make
```

- for Knights Corner (KNC):

```
cd sample/specfem
make KNC=1
```

- additionally, adding some specific Fortran compiler flags, for example:

```
cd sample/specfem
make FCFLAGS="-O3 -fopenmp" [...]
```

Note that steps 1 & 2 could be shortened:

- by specifying a "specfem" make target in the LIBXSMM root directory:

```
make MNK="5 25" ALPHA=1 BETA=0 PRECISION=1 OPT=3 AVX=1 specfem
```

- for Knights Corner, this would need two steps:

```
make MNK="5 25" ALPHA=1 BETA=0 PRECISION=1 OPT=3 AVX=1 OFFLOAD=1 KNC=1
make OPT=3 specfem_mic
```

Run the performance test:

- for default CPU host:

```
./specfem.sh
```

- for Knights Corner (KNC):

```
./specfem.sh -mic
```

Results

Using Intel Compiler suite: icpc 15.0.2, icc 15.0.2, and ifort 15.0.2

Sandy Bridge - Intel(R) Xeon(R) CPU E5-2670 0 @ 2.60GHz

- library compilation by (root directory):

```
make MNK="5 25" ALPHA=1 BETA=0 PRECISION=1 OPT=3 AVX=1
```

- single threaded example run:

```
cd sample/specfem
make; OMP_NUM_THREADS=1 ./specfem.sh
```

Output:

```
===================================================================
average over             15 repetitions
 timing with Deville loops    =    0.1269
 timing with unrolled loops   =    0.1737 / speedup =   -36.87 %
 timing with LIBXSMM dispatch =    0.1697 / speedup =   -33.77 %
 timing with LIBXSMM prefetch =    0.1611 / speedup =   -26.98 %
 timing with LIBXSMM static   =    0.1392 / speedup =    -9.70 %
===================================================================
```

Haswell - Intel(R) Xeon(R) CPU E5-2680 v3 @ 2.50GHz

- library compilation by (root directory):

```
make MNK="5 25" ALPHA=1 BETA=0 PRECISION=1 OPT=3 AVX=2
```

- single threaded example run:

```
cd sample/specfem
make; OMP_NUM_THREADS=1 ./specfem.sh
```

Output:

```
===================================================================
average over             15 repetitions
 timing with Deville loops    =    0.1028
 timing with unrolled loops   =    0.1385 / speedup =   -34.73 %
 timing with LIBXSMM dispatch =    0.1408 / speedup =   -37.02 %
 timing with LIBXSMM prefetch =    0.1327 / speedup =   -29.07 %
 timing with LIBXSMM static   =    0.1151 / speedup =   -11.93 %
===================================================================
```

- multi-threaded example run:

```
cd sample/specfem
make OPT=3; OMP_NUM_THREADS=24 ./specfem.sh
```

Output:

```
OpenMP information:
  number of threads =          24

[...]


===================================================================
average over             15 repetitions
 timing with Deville loops    =    0.0064
 timing with unrolled loops   =    0.0349 / speedup =   -446.71 %
 timing with LIBXSMM dispatch =    0.0082 / speedup =    -28.34 %
 timing with LIBXSMM prefetch =    0.0076 / speedup =    -19.59 %
 timing with LIBXSMM static   =    0.0068 / speedup =     -5.78 %
===================================================================
```

Knights Corner - Intel Xeon Phi B1PRQ-5110P/5120D

- library compilation by (root directory):

```
make MNK="5 25" ALPHA=1 BETA=0 PRECISION=1 OPT=3 OFFLOAD=1 KNC=1
```

- multi-threaded example run:

```
cd sample/specfem
make FCFLAGS="-O3 -fopenmp -warn" OPT=3 KNC=1; ./specfem.sh -mic
```

Output:

```
OpenMP information:
  number of threads =         236

[...]


===================================================================
average over             15 repetitions
 timing with Deville loops    =    0.0164
 timing with unrolled loops   =    0.6982 / speedup = -4162.10 %
 timing with LIBXSMM dispatch =    0.0170 / speedup =    -3.89 %
 timing with LIBXSMM static   =    0.0149 / speedup =     9.22 %
===================================================================
```

Matrix Transpose (TCOPY)

Overview

This code sample aims to benchmark the performance of matrix transposes. The C/C++ and FORTRAN sample code differ slightly with the C/C++ code sample offering a richer set of command line options as well as build settings available inside of the translation unit.

The available command line options of the sample code may be reviewed by looking into the source code. Generally, the idea is to support the following:

> transpose [<kind> [<m> [<n> [<ldi> [<ldo>]]]]]
> transposef [<m> [<n> [<ldi> [<ldo>]]]]

Above, m and n specify the matrix shape, and ldi the leading dimension of the matrix. The argument ldo allows to specify an output dimension, which may differ from ldi. The transpose kind shall be either out-of-place (o) or in-place (i).

Running the C sample code may look like:

```
$ ./transpose.sh o 20000
m=20000 n=20000 ldi=20000 ldo=20000 size=3052MB (double, out-of-place)
        bandwidth: 18.8 GB/s
        duration: 159 ms
```

Instead of executing a wrapper script, one may affinitize the multi-threaded execution manually (OpenMP runtime). In case of an executable built using the Intel Compiler this may look like:

```
LIBXSMM_VERBOSE=2 KMP_AFFINITY=balanced,granularity=fine,1 \
./transpose o 20000
m=20000 n=20000 ldi=20000 ldo=20000 size=3052MB (double, out-of-place)
        bandwidth: 21.1 GB/s
        duration: 141 ms

Registry: 20 MB (gemm=0 mcopy=0 tcopy=1)
```

In the above case one can see from the verbose output (`LIBXSMM_VERBOSE=2`) that one kernel (tcopy) served transposing the entire matrix. To avoid duplicating JIT-kernels under contention (code registry), one may also consider `LIBXSMM_TRYLOCK=1`, which is available per API-call as well.

OpenTuner

To tune the tile sizes ("block sizes") internal to LIBXSMM's transpose routine, the OpenTuner extensible framework for program autotuning can be used. A tuning script (`transpose_opentuner.py`) is provided, which accepts a range of matrix sizes as command line arguments.

```
./transpose_opentuner.py &lt;begin&gt; &lt;end&gt; &lt;nruns&gt;
```

To start a tuning experiment for a new set of arguments, it is highly recommended to start from scratch. Otherwise the population of previously generated tuning results is fetched from a database and used to tune an eventually unrelated range of matrix shapes.

```
rm -rf opentuner.db
```

The script tunes matrices with randomized shape according to the specified range. The optimizer not only maximizes the performance but also minimizes the value of $M * N$ (which also helps to prune duplicated results due to an additional preference).

```
rm -rf opentuner.db
./transpose_opentuner.py --no-dups 1 1024 1000

rm -rf opentuner.db
./transpose_opentuner.py --no-dups 1024 2048 100

rm -rf opentuner.db
./transpose_opentuner.py --no-dups 2048 3072 20

rm -rf opentuner.db
./transpose_opentuner.py --no-dups 3072 4096 20

rm -rf opentuner.db
./transpose_opentuner.py --no-dups 4096 5120 16

rm -rf opentuner.db
./transpose_opentuner.py --no-dups 5120 6144 12

rm -rf opentuner.db
./transpose_opentuner.py --no-dups 6144 7168 8

rm -rf opentuner.db
./transpose_opentuner.py --no-dups 7168 8192 6
```

The tuning script uses the environment variables `LIBXSMM_TRANS_M` and `LIBXSMM_TRANS_N`, which are internal to LIBXSMM. These variables are used to request a specific tiling-scheme inside of LIBXSMM's `libxsmm_otrans_omp` routine.

Wrapped DGEMM

This code sample is calling DGEMM and there is no dependency on the LIBXSMM API as it only relies on LAPACK/BLAS interface. Two variants are linked when building the source code: (1) code which is dynamically linked against LAPACK/BLAS, (2) code which is linked using `--wrap`=*symbol* as possible when using a GNU GCC compatible tool chain. For more information, see the Call Wrapper section of the reference documentation.

The code will execute in three flavors when running `dgemm-test.sh`: (1) code variant which is dynamically linked against the originally supplied LAPACK/BLAS library, (2) code variant which is linked using the wrapper mechanism of the GNU GCC tool chain, and (3) the first code but using the LD_PRELOAD mechanism (available under Linux).

Command Line Interface (CLI)

- Optionally takes the number of repeated DGEMM calls
- Shows the performance of the workload (wall time)

XGEMM: Tiled GEMM Routines

Overview

This sample code calls the `libxsmm_?gemm_omp` routines provided by the LIBXSMM extension library (`libxsmmext`). These routines are meant for big(ger) xGEMM routines, and thereby provide an OpenMP-based parallelization.

The driver program (`xgemm.c`) currently accepts all typical GEMM arguments (except for the transposition specifier): `m`, `n`, `k`, `lda`, `ldb`, `ldc`, `alpha`, and `beta`. All arguments are optional (or will inherit defaults from previously specified arguments). Matrix transposition as part of the `libxsmm_?gemm_omp` routines will become available in an upcoming release of LIBXSMM. Please also note that unsupported Alpha or Beta values will cause a fall back to the related BLAS routine. The single-precision matrix multiplications require to change the `REAL_TYPE` in `xgemm.c`.

```
./xgemm.sh 2000
```

OpenTuner

To tune the tile sizes ("block sizes") internal to LIBXSMM, the OpenTuner extensible framework for program autotuning can be used. A tuning script (`xgemm_opentuner.py`) is provided, which optionally accepts a list of grouped parameters as command line arguments. The syntax of the arguments is per LIBXSMM's `MNK` build-option, and expands to "triplets" specifying the matrix shapes. For instance, four matrix multiplications of square-matrices can be benchmarked and tuned using the following command.

```
./xgemm_opentuner.py 1024,1280,1536,1792
```

To start a tuning experiment for a new set of arguments, it is highly recommended to start from scratch. Otherwise the population of previously generated tuning results is fetched from a database and used to tune an unrelated range of matrix shapes.

```
rm -rf opentuner.db
```

The script tunes the geometric mean of the performance for each of the requested triplets. However, the optimizer not only maximizes the performance but also minimizes the value of $M * N * K$ (which also helps to prune duplicated results due to an additional preference). As a limitation of the current implementation, the multiplication kernels are not accompanied by copy-kernels (and not accompanied by transpose kernels). This negatively impacts performance on power-of-two matrix shapes (POT) due to trashing the LLC. However, it has been found, that tuning for POT shapes likely achieves superior performance when compared to tuning for non-POT shapes of the same range.

```
rm -rf opentuner.db
./xgemm_opentuner.py --no-dups 192,256,320,512,768

rm -rf opentuner.db
./xgemm_opentuner.py --no-dups 1024,1280,1536,1792

rm -rf opentuner.db
./xgemm_opentuner.py --no-dups 2048,2304,2560,2816

rm -rf opentuner.db
./xgemm_opentuner.py --no-dups 3072,3328,3584,3840

rm -rf opentuner.db
./xgemm_opentuner.py --no-dups 4096,4416,4736

rm -rf opentuner.db
./xgemm_opentuner.py --no-dups 5120,5440,5760

rm -rf opentuner.db
./xgemm_opentuner.py --no-dups 6144,6464,6784

rm -rf opentuner.db
./xgemm_opentuner.py --no-dups 7168,7488,7808
```

Above, the series of matrix multiplications from 192-8K is separately tuned in eight ranges. The tuning script uses the environment variables `LIBXSMM_TGEMM_M`, `LIBXSMM_TGEMM_N`, and `LIBXSMM_TGEMM_K` which are internal to LIBXSMM. These variables are used to request a specific tiling-scheme within LIBXSMM's `libxsmm_?gemm_omp` routines.