

Basic Setup

Follow the instructions for class account setup found on <http://www.ece.lsu.edu/gp/proc.html>. This assignment does not rely upon code in a `hw` directory, instead use the code in file `.../cuda/intro-vtx-transform/vtx-xform-size.cu`.

If the class account has been set up properly, the code can be built from within Emacs by pressing `F9` when visiting any file in the `.../cuda/intro-vtx-transform` directory or when in an Emacs shell buffer (which can be entered using `Alt-x shell Enter`). The code can be built from the command line using the command `make -j 4` (assuming `.../cuda/intro-vtx-transform` is the current directory). Either method runs a makefile that builds all examples in the directory. It builds two versions of each program, one taking the base name of the main file, such as `vtx-xform-size`, and one with the suffix `-debug`, such as `vtx-xform-size-debug`. The versions with the `-debug` suffix are compiled with host optimization turned off, which facilitates debugging. At the moment GPU debugging can only be enabled by editing the makefile, but that should not be a problem for this assignment.

Running `make` on a clean directory will produce a large amount of output, including warnings. The `make` program and input `Makefile` are designed to build executables in a lazy fashion, meaning that they only create a file if it is not present or if its prerequisites have changed. Therefore a second run of `make` will take much less time. (The `Makefile` in the `intro-vtx-xform` directory prints a diagnostic message about the presence of a Phi. When run a second time that's the only output.)

Quickly check whether the build is successful with the command `./vtx-xform-size`. It should produce output ending with a line starting `mxv_sh_ochunk`.

The makefile will compile code for a GPU on the system it was run, favoring the GPU that's not connected to a display. For this assignment multiple types of GPU need to be used. Re-run `make` when moving to a different system. The `Makefile` should automatically detect whether the GPU for which the executable was built matches the GPU on the current system, and re-build if needed.

Using `vtx-xform-size`

The code in `vtx-xform-size.cu` contains several kernels that compute a matrix/vector product using a constant $M \times N$ matrix A applied to S N -component input column vectors, producing S M -component output column vectors. The values of N and M are hard-coded in the file, whereas S is a command-line argument. Vector and matrix elements are of type `Elt_Type`, which is hardcoded to `float`.

The `vtx-xform-size` program takes three command-line arguments. The first indicates how many blocks to launch. If the argument is zero then the number of blocks will be set to the number of multiprocessors (which is the default). If the argument is negative then the number of blocks will be $-aP$, where a is the argument value and P is the number of MPs.

The second argument is the number of threads per block to try to use to launch each kernel. If the argument is omitted 1024 threads are tried. If the argument is omitted or positive, the actual number of threads used in a launch is the minimum of this argument and the kernel's maximum. (For example, when the second argument is 512, but kernel `foo` has a limit of 256 threads, `foo` will be launched with 256 threads.) When the second argument is 0 (zero) or `p` then each kernel will be launched multiple times starting with 4 warps, incrementing by 4 warps until the kernel maximum is reached. When `p` is used additional performance data is shown, which is interesting but it can slow things down.

The third argument specifies the number of *mibions* input vectors to use. One mibion is 2^{20} . This is the premier usage of this term. The default is 1 mibion (1,048,576) vectors. If a_3 is the value of the third argument, the input size will be $a_3 2^{20}$ vectors. The third argument is read as a floating-point number, so “0.5” will result in a 2^{19} vectors input.

Here are some examples: Run with 256 threads per block: `./vtx-xform-size 0 256`. Run with 512 threads per block and twice as many blocks as MPs: `./vtx-xform-size -2 512`. Run with 256 threads per block and 10 blocks: `./vtx-xform-size 10 256`. Run each kernel multiple times using an input size of 1 bibion (2^{30} vectors): `./vtx-xform-size 0 0 1024`.

Program Output

A run of `vtx-xform-size` produces the following output:

A stray message about a CUDA API call.

```
[koppel@dmk-laptop intro-vtx-transform]$ ./vtx-xform-size
Call of cuDeviceGetCount, cbid 4, serial 1
```

Ignore it (the Call message).

Information about each GPU connected to the system, followed by a line showing the chosen GPU.

```
GPU 0: Quadro M2200 @ 1.04 GHz WITH 4010 MiB GLOBAL MEM
GPU 0: L2: 1024 kiB MEM<->L2: 88.1 GB/s
GPU 0: CC: 5.2 MP: 8 CC/MP: 128 DP/MP: 4 TH/BL: 1024
GPU 0: SHARED: 49152 B/BL 98304 B/MP CONST: 65536 B # REGS: 65536
GPU 0: PEAK: 1061 SP GFLOPS 33 DP GFLOPS COMP/COMM: 48.2 SP 3.0 DP
Using GPU 0
```

Most lab computers have two GPUs, please pay attention to the GPU that is actually being used, GPU 0 above.

Most fields are self-explanatory. For example, L2 is the size of the level-2 cache and CC indicates that the device is of compute capability 5.2 (a cheap Maxwell). The MEM<->L2 field shows the off-chip bandwidth. MP indicates the number of multiprocessors, also called streaming multiprocessors (SM’s). CC/MP indicates the number of CUDA cores (single-precision functional units) per MP, DP/MP indicates the number of double-precision functional units per MP, and TH/BL is the maximum number of threads per block.

The amount of shared memory available is shown per block (B/BL) and per MP, this does not indicate whether any particular kernel is using that much shared memory or could use that much. The same line shows the amount of constant memory, and the number of registers available.

The next line, PEAK, shows FP operation bandwidth in which a fused multiply-add is counted as one operation. The COMP/COMM line gives the computation to communication ratio in floating-point operations per floating-point element transfers. (The information above was collected in part using the runtime library’s `cudaGetDeviceProperties` function.)

The program will next print information about each kernel:

CUDA Kernel Resource Usage:

For `mxv_g_only`:

0 shared, 1088 const, 0 loc, 32 regs; 1024 max threads per block.

For `mxv_i_lbuf`:

0 shared, 1088 const, 0 loc, 32 regs; 1024 max threads per block.

For `mxv_o_lbuf`:

0 shared, 1088 const, 0 loc, 32 regs; 1024 max threads per block.

For mxv_o_per_thd:
 0 shared, 1088 const, 0 loc, 31 regs; 1024 max threads per block.
 For mxv_vec_ld:
 0 shared, 1088 const, 64 loc, 40 regs; 1024 max threads per block.
 For mxv_vls:
 16384 shared, 1088 const, 0 loc, 48 regs; 1024 max threads per block.
 For mxv_sh:
 36864 shared, 1088 const, 0 loc, 32 regs; 1024 max threads per block.
 For mxv_sh_ochunk:
 4096 shared, 1088 const, 0 loc, 54 regs; 1024 max threads per block.

The max threads per block shown above is based on the kernel and reflects register usage. Though it does not happen above, it is possible that a kernel can be limited to less than 1024 threads per block because it uses more than 64 registers (on a CC 5.2 device).

Next, the program provides information on the input size and launch configuration.

Matrix size: 16 x 16. Vectors: 1048576. 8 blocks of 1024 thds.

The input size can be changed using command-line arguments, that is explained further below.

The program can either launch each kernel once, with a particular configuration (number of blocks and number of threads per block), or it can launch each kernel multiple times, each with a different block size. Without arguments it runs each kernel once and prints one line per kernel.

Launching with 8 blocks of up to 1024 threads.

mxv_g_only	32 wp	6118 μ s	43.873 GF	21.937 GB/s	1.21 I/F	12.5%
mxv_i_lbuf	32 wp	6836 μ s	39.266 GF	19.633 GB/s	1.22 I/F	12.5%
mxv_o_lbuf	32 wp	6905 μ s	38.873 GF	19.437 GB/s	1.22 I/F	12.5%
mxv_o_per_thd	32 wp	4602 μ s	58.331 GF	29.166 GB/s	3.44 I/F	12.5%
mxv_vec_ld	32 wp	4144 μ s	64.776 GF	32.388 GB/s	1.22 I/F	50.0%
mxv_vls	32 wp	2606 μ s	103.017 GF	51.509 GB/s	1.50 I/F	100.0%
mxv_sh	32 wp	4084 μ s	65.733 GF	32.867 GB/s	4.14 I/F	100.0%
mxv_sh_ochunk	32 wp	2598 μ s	103.337 GF	51.669 GB/s	2.00 I/F	100.0%

The μ s values are the execution time, the GF values show the floating point throughput (assuming NM FP operations per input vector), and the GB/s value is the off-chip data throughput, assuming that $4(N + M)B$ crosses the chip boundary for each element.

The I/F value shows the number of instructions executed per expected multiply-add. A value of 2.00 means that $2NM$ instructions are executed for each input vector. Presumably one of those instructions is a multiply-add. The I/F value is computed by actually measuring the number of instructions executed using event counters built in to the GPU, and dividing that count by the number by SNM , the number of multiply-add instructions expected for S inputs. A value of exactly 1 for I/F would be suspiciously low because that would mean there could be no load instructions to read the input vector nor stores to write the output vector. Including these would give a ratio of $\frac{MN+M+N}{MN} = 1 + \frac{1}{M} + \frac{1}{N}$ or 1.125 for $M = N = 8$.

The last column, showing percents, shows read request utilization.

When run with a 0 as the second argument, such as `./vtx-xform-size -1 0`, the program, launches each kernel multiple times, starting with 4 warps per block, up to 32 warps per block. Run time and other information will be shown for each launch. An excerpt for one kernel appears below:

```
Kernel mxv_o_per_thd:
wp ac   t/ $\mu$ s FP  $\theta$  GB/s Bandwidth Util-----
  4  4  24269   11    6 ***
  8  8  13178   20   10 *****
 12 12   8574   31   16 *****
 16 16   6741   40   20 *****
 20 20   6059   44   22 *****
 24 24   5198   52   26 *****
 28 28   4935   54   27 *****
 32 32   4439   60   30 *****
```

The `wp` column shows the number of warps per block that the kernel was launched with. The `ac` column shows the number of warps assigned to an MP (which is the product of the number of warps per block and the number of active blocks per MP). The number in the `ac` column is computed by an NVIDIA API using information about the kernel and the GPU. In the example above the `wp` and `ac` numbers are the same because the number of blocks is the same as the number of MPs and so there is no way to have more than one block per MP.

The `t/ μ s` column shows the measured execution time. The `FP θ` column shows the FP throughput based on the measured execution time and the assumption that *SMP* floating point operations were performed. The number under `GB/s` is the minimum off-chip bandwidth, computed by dividing $4S(M + N)$ by the measured execution time. The stars in last column show bandwidth utilization based on the `GB/s` number. If the stars extend to the maximum length (shown by the hyphens to the right of `Bandwidth Util` in the column heading) then off-chip bandwidth is being saturated. Note that this number is computed using measured time and an ideal amount of data crossing the chip boundary.

When run using a `p` instead of a `0`, `vtx-xform-size` collects hardware utilization data related to load and store instructions. A sample is shown below.

```
Kernel mxv_o_per_thd:
wp ac   t/ $\mu$ s FP  $\theta$  Ld eff L2r $\theta$  L2w $\theta$  DRr $\theta$  DRw $\theta$  GB/s Bandwidth Util-----
  4  4  24520   11 12.5% 23.7  2.7  4.6  2.7   5 *
  8  8  12458   22 12.5% 45.0  5.4  7.3  5.4  11 ***
 12 12   8820   30 12.5% 62.9  7.6  9.6  7.6  15 ****
 16 16   8167   33 12.5% 67.4  8.2  9.8  8.2  16 ****
 20 20   6189   43 12.5% 88.7 10.8 12.8 10.8  22 *****
 24 24   6336   42 12.5% 86.4 10.6 12.3 10.6  21 *****
 28 28   5065   53 12.5% 107.8 13.2 15.1 13.2  26 *****
 32 32   5364   50 12.5% 101.8 12.5 14.1 12.5  25 *****
```

The `Ld eff` column indicates the average percentage of each load request that is used. The 12.5% value shown above is what one would expect for scattered accesses to 4-byte values. The `L2r θ` column shows the throughput from the L2 caches to the MPs due to load instructions, in `GB/s`. The `L2w θ` column shows the throughput from the MPs to the L2 caches due to store instructions, also in units of `GB/s`. The `DRr θ` and `DRw θ` columns show the throughput from DRAM (off chip) to the L2 cache and the L2 cache to DRAM, respectively, in `GB/s`. These four throughput values are based on request sizes, not on how much of those requests are actually needed. This data is

collected using event counters.

Problem 1: In class we noted that kernels `mxv_i_1buf` and `mxv_o_1buf` suffer from request under-utilization but noted that their performance was actually better than what one would expect if only $\frac{1}{8}$ of off-chip bandwidth carried useful data. We attributed this better-than-expected performance to the the level 2 (L2) cache.

Indicate whether data from runs using the commands described above provides evidence that the L2 cache is helping. Do so for a CC 3.5 device (there are two in the lab) and a CC 6.X device. Use the event counter data (the data provided when the `p` option is present) to justify your answer. In your solution indicate the type of GPU, show the performance data, and explain what the L2 cache is doing based on the data.

Based on the data indicate what might be limiting performance.

Problem 2: The number of threads needed in a kernel launch can be determined using a technique called *GPU interval analysis*. Interval analysis starts with an *interval* which is the time during the execution of one iteration of a well-chosen loop. The loop is well chosen if the execution time of the program is some multiple of the time for an execution of the loop.

For the `mxv_o_1buf` kernel the interval will be one iteration of the `h` loop.

The first step in GPU interval analysis is to compute the latency for a thread executing one iteration of the loop while assuming no resource contention due to other threads. Compute this latency so that if L_i is the latency of one iteration, then the time needed for x iterations of the loop by a thread is xL_i .

For this assignment estimate the interval latency by accounting for *instruction issue rate* and *exposed* operation latency. The instruction issue rate will be used to compute the earliest time that the instruction can be issued based on instruction throughput and assuming one warp.

Consider Maxwell- and Pascal-generation devices. Most have four warp schedulers per MP, 128 single-precision functional units per MP, and 64 load/store (LDST) units per MP, or 32 SP FP units per scheduler and 16 LDST units per scheduler. Assume that at most one instruction per cycle can be issued. Then for a thread SP FP instructions, such as `FFMA`, will be issued at a rate of one per cycle. (Remember that we are assuming only one warp of threads for purposes of issue.) For memory instructions (loads and stores) instructions will be issued every two cycles. See the code excerpt and pipeline diagram below, where `FFMA` instructions advance at one per cycle but `LDG` advance at half that rate.

Let L_F denote the latency of a floating-point operation, and L_M denote the minimum latency of a load instruction that misses the L2 cache. For Maxwell- and Pascal-generation devices $L_F = 6$ cyc and a typical value is $L_M = 400$ cyc. Note that L_F is a fixed duration, meaning that once an instruction is dispatched to the functional unit a dependent instruction can be dispatched L_F cycles later and expect to find its source value ready. In contrast L_M holds for a load to uncached data on a lightly loaded system. The actual latency would be higher if there were contention for, say, off-chip bandwidth, or lower if the load hit the L2 cache. For interval analysis contention is usually assumed to be zero, but whether an accesses hits the L2 cache is part of the analysis.

In the code fragment below F0 depends on L0, and so stalls for 394 cycles. Instruction F8 depends on F0, but since the issue time is larger than latency there is no stall.

```

L0:    LDG.E R16, [R2];
L1:    LDG.E R12, [R14];
L2:    LDG.E R11, [R14+0x4];
...
F0:    FFMA R13, R16, c[0x3][0x4], RZ; // Issue: t = 0, avail t = 6
F1:    FFMA R17, R16, c[0x3][0x24], RZ; // Issue: t = 1, avail t = 7
F2:    FFMA R18, R16, c[0x3][0x44], RZ; // ...
F3:    FFMA R19, R16, c[0x3][0x64], RZ;
F4:    FFMA R21, R16, c[0x3][0x84], RZ;
F5:    FFMA R22, R16, c[0x3][0xa4], RZ;
F6:    FFMA R14, R16, c[0x3][0xc4], RZ;
F7:    FFMA R15, R16, c[0x3][0xe4], RZ;
F8:    FFMA R13, R12, c[0x3][0x8], R13; // Issue:
...
S0:    STG.E [R4], R2;
S1:    STG.E [R4+0x4], R14;
S2:    STG.E [R4+0x8], R13;
...

```

```

Time 0  1  2  3  4  5  6  ... 400 401 402 403 404 405 406 407 408
wp0: [L0  ] [L1  ] [L2  ]          F0 F1 F2 F3 F4 F5 F6 F7 F8

```

Consider a modified version of the `mxv_o_1buf` kernel in which the entire input vector is read before performing arithmetic:

```

for ( int h=start; h<stop; h += inc )
{
    Elt_Type vout[M]; for ( int r=0; r<M; r++ ) vout[r] = 0;
    Elt_Type vin[N]; for ( int c=0; c<N; c++ ) vin[c] = d_app.d_in[ h * N + c ];

    for ( int c=0; c<N; c++ )
        for ( int r=0; r<M; r++ ) vout[r] += d_app.matrix[r][c] * vin[c];

    for ( int r=0; r<M; r++ ) d_app.d_out[ h * M + r ] = vout[ r ];
}

```

Accounting only for the load, multiply/add, and store instructions we get $L_\ell(K_0) = L_M + MN + M$ if $N \geq L_F$ and $L_M \geq N/2$, where K_0 is a shorthand name for the modified kernel and its assumed code. For values $M = N = 16$ and the latencies above we get $L_\ell(K_0) = 400 + 256 + 16 = 672$ cyc.

Once the latency for an interval is computed, it is possible to bound the number of threads that can be executed in parallel by using resource constraints. For example, let d_ℓ denote the amount of off-chip data read or written by a thread during one iteration and let Θ_M indicate the off-chip bandwidth. Let P indicate the number of threads in a launch. The amount of data read during the interval is Pd_ℓ , the data throughput would be $\frac{Pd_\ell}{L_\ell}$, solving $\frac{Pd_\ell}{L_\ell} = \Theta_M$ yields $P = \Theta_M \frac{L_\ell}{d_\ell}$, the number of threads needed to saturate off-chip bandwidth.

For our matrix/vector kernels, $d_\ell = 4(M + N)B$. Then $P = \Theta_M \frac{(L_M + MN + M) \text{cyc}}{4(M + N)B}$. For $N = M$ (square matrices) we have $P = \Theta_M \frac{(L_M + N^2 + N) \text{cyc}}{8NB} = \frac{\Theta_M}{8} \left(\frac{L_M}{N} + N + 1 \right) \frac{\text{cyc}}{B}$ for $N \geq L_F$. Notice that this expression has a minimum at $N = \sqrt{L_M}$, or 20 for our default global memory latency. As N

shrinks below 20 more and more threads are needed because each thread is loading less data but the iteration time, L_ι , can't fall below L_M . When N grows above 20 the computation time, N^2 , grows quadratically while data grows linearly, so more threads are needed. Note that Θ_M is often given in GB/s while the latency is given in cycles. Cycles can be converted to seconds by dividing by the clock frequency, often denoted ϕ . In that case $L_\iota(K_0) = (L_M + MN + M) \text{ cyc} = \frac{1}{\phi}(L_M + MN + M) \text{ s}$.

Another limiter is instruction throughput. For that one must compute the time needed during an interval to issue the instructions. Let Θ_{IF} indicate the bandwidth of FP instructions and Θ_{IM} indicate the bandwidth of load and store instructions on a multiprocessor. Let n_{IF} and n_{IM} indicate the number of instructions of the respective type executed by a thread in the interval. Then the time for Q threads to issue for one interval is $Q(\frac{n_{\text{IF}}}{\Theta_{\text{IF}}} + \frac{n_{\text{IM}}}{\Theta_{\text{IM}}})$ and based on this instruction issue limit $Q = L_\iota / (\frac{n_{\text{IF}}}{\Theta_{\text{IF}}} + \frac{n_{\text{IM}}}{\Theta_{\text{IM}}})$. Note that Q is the number of active threads on a multiprocessor, which is some multiple of the block size.

For K_0 we have $n_{\text{IF}} = MN$ and $n_{\text{IM}} = M + N$. Substituting these values and L_ι gives

$$\begin{aligned} Q(K_0) &= L_\iota / \left(\frac{n_{\text{IF}}}{\Theta_{\text{IF}}} + \frac{n_{\text{IM}}}{\Theta_{\text{IM}}} \right) \\ &= (L_M + MN + M) / \left(\frac{MN}{\Theta_{\text{IF}}} + \frac{M + N}{\Theta_{\text{IM}}} \right) \\ &= (L_M + N^2 + N) / \left(\frac{N^2}{\Theta_{\text{IF}}} + \frac{2N}{\Theta_{\text{IM}}} \right) \\ &= (L_M/N + N + 1) / \left(\frac{N}{\Theta_{\text{IF}}} + \frac{2}{\Theta_{\text{IM}}} \right) \\ &= (L_M/N^2 + 1 + 1/N) / \left(\frac{1}{\Theta_{\text{IF}}} + \frac{2}{N\Theta_{\text{IM}}} \right) \end{aligned}$$

For large N the interval time is dominated by the N^2 instructions so the number of threads is based on FP instruction bandwidth Θ_{IF} . When N is smaller load latency determines the thread count.

(a) Does the interval analysis above explain the behavior of the `mxv_o_1buf` kernel obtained on a CC 6.1 device? If not, provide a possible explanation and suggest a way of modifying the analysis.

(b) Using interval analysis as described above, determine the minimum number of threads needed to saturate off-chip bandwidth and instruction issue for kernel `mxv_o_1buf` assuming that the order of machine (SASS) instructions follows the order of instructions in the CUDA code. That is, don't assume the entire input vector will be loaded before starting on the multiply-adds. In your analysis assume that the level 2 cache is effective. Use $L_M = 400 \text{ cyc}$ for the latency of a load that misses and $L_2 = 200 \text{ cyc}$ for the latency of a load that hits the level 2 cache. Also use $L_F = 6 \text{ cyc}$ for instruction latency of non-memory instructions.

Show the number of threads needed as expressions, and compute actual values for one of the Pascal CC 6.1 GPUs.

One advantage of following the order of instructions in the CUDA code is that $N - 1$ fewer registers are needed. Is this worth it?

(c) Perform the interval analysis on kernel `mxv_o_per_thd`. When working out the analysis don't forget that each thread is computing only one component of the output vector and that neighbors are using the same input vector. Take that into account when computing d_ι .

Show the number of threads needed as expressions, and compute actual values for one of the Pascal CC 6.1 GPUs.