

## **ECE408/CS483 Practice Exam #2, Fall 2012**

- You are allowed to use any notes, books, papers, and other reference material as you desire. No electronic assistance other than a calculator is permitted.
- No interactions with humans other than course staff are allowed.
- This exam is designed to take 80 minutes to complete.
- This exam is based on lectures as well as lab MPs/projects.
- The questions are randomly selected from the topics we covered this semester.
- You can write down the reasoning behind your answers for possible partial credit.
- **Good luck!**

### Question 1 (50 points): Short Answer

Answer each of the following questions in as few words as you can. Your answer will be graded based on completeness, correctness, and conciseness.

1. (7 points) Give two reasons for adding extra "padding" elements to arrays allocated in GPU global memory?

2. (7 points) Give two potential disadvantages associated with increasing the amount of work done in each CUDA thread, such as loop unrolling techniques, using fewer threads in total?

3. (8 points) The following CUDA code is intended to perform a sum of all elements of the **partialSum** array. First, provide the missing operation, and second estimate the fraction of the iterations of the **for** loop that will have branch divergence.

```
__shared__ float partialSum[];
unsigned int tid = threadIdx.x;

for (unsigned int stride = blockDim.x;
     stride > 0;
     stride = stride >> 1)
{
    _____ ;; missing operation
    if (tid < stride)
        partialSum[tid] += partialSum[tid+stride];
}
```

4. (7 points) For the following code, estimate the average run time in cycles of a thread. Assume for simplicity that the `__sin()` function operates on radians, and that all operations including the `__sin()` take 1 cycle.

```
if (__sin( (float) threadIdx.x) > 0.5)
{
    :
    A operations
    :
}
else
{
    :
    B operations
    :
}
```

5. (7 points) Provide two advantages of a programmer-managed memory, such as the CUDA shared memory, over a hardware managed cache?

6. (7 points) Assuming capacity were not an issue for registers or shared memory, give one case where it would be valuable to use shared memory instead of registers to hold values fetched from global memory?

7. (7 points) Assume the following code will execute all in the same warp of a CUDA thread block. Write a piece of CUDA code in which the threads within the warp reverse an array `list` of length 32 in shared memory. Make your code as efficient in time and space as possible.

```
__shared__ float list[32]
```

## Question 2 (25 points): CUDA Basics

The following code computes 1024 dot products, each of which is calculated from a pair of 256-element sub-vectors. Assume that the code is executed on G80. Use the code to answer the following questions.

```
1  #define VECTOR_N 1024
2  #define ELEMENT_N 256
3  const int DATA_N      = VECTOR_N * ELEMENT_N;
4  const int DATA_SZ     = DATA_N * sizeof(float);
5  const int RESULT_SZ    = VECTOR_N * sizeof(float);
6  ...
7  float *d_A, *d_B, *d_C;
8  ...
9  cudaMalloc((void **)&d_A, DATA_SZ);
10 cudaMalloc((void **)&d_B, DATA_SZ);
11 cudaMalloc((void **)&d_C, RESULT_SZ);
12 ...
13 scalarProd<<<VECTOR_N, ELEMENT_N>>>(d_C, d_A, d_B, ELEMENT_N);
14 __global__ void
15 scalarProd(float *d_C, float *d_A, float *d_B, int ElementN)
16 {
17     __shared__ float accumResult[ELEMENT_N];
18     //Current vectors bases
19     float *A = d_A + ElementN * blockIdx.x;
20     float *B = d_B + ElementN * blockIdx.x;
21     int tx = threadIdx.x;
22
23     accumResult[tx] = A[tx] * B[tx];
24
25     for(int stride = ElementN / 2; stride > 0; stride >>= 1)
26     {
27         __syncthreads();
28         if(tx < stride)
29             accumResult[tx] += accumResult[stride + tx];
30     }
31     d_C[blockIdx.x] = accumResult[0];
32 }
```

1. (2 pts) How many threads are there in total?
2. (2 pts) How many threads are there in a Warp?
3. (2 pts) How many threads are there in a Block?
4. (2 pts) How many global memory loads and stores are done for each thread?
5. (2 pts) How many accesses to shared memory are done for each block? Count one access

from all threads as *blockdim.x* accesses.

6. (3 pts) List the source code lines, if any, that cause shared memory bank conflicts.

7. (3 pts) How many iterations of the *for* loop (line 23) will have branch divergence? Show your derivation.

8. (3 pts) What is the largest and smallest ratios of floating point arithmetic to global memory access in each thread?

9. (3 pts) Identify an opportunity to significantly reduce the bandwidth requirement on the global memory. How would you achieve this? How many accesses can you eliminate?

10. (3 pts) How many registers will be needed to accommodate variable A in each block?

**Question 3 (25 points).** The following kernel is executed on a large matrix, which is tiled into submatrices. To manipulate tiles, a novice CUDA programmer has written the following device kernel to transpose each tile in the matrix. The tiles are of size BLOCK\_SIZE by BLOCK\_SIZE, and each of the dimensions of matrix A is known to be a multiple of BLOCK\_SIZE. The kernel invocation and code are shown below. BLOCK\_SIZE is known at compile-time, but could be set anywhere from 1 to 20.

```
dim3 blockDim(BLOCK_SIZE, BLOCK_SIZE);
dim3 gridDim(A_width/blockDim.x, A_height/blockDim.y);
BlockTranspose<<<gridDim, blockDim>>>(A, A_width, A_height);

__global__ void
BlockTranspose(float* A_elements, int A_width, int A_height)
{
    __shared__ float blockA[BLOCK_SIZE][BLOCK_SIZE];

    int baseIdx = blockDim.x * BLOCK_SIZE + threadIdx.x;
    baseIdx += (blockDim.y * BLOCK_SIZE + threadIdx.y) * A_width;

    blockA[threadIdx.y][threadIdx.x] = A_elements[baseIdx];

    A_elements[baseIdx] = blockA[threadIdx.x][threadIdx.y];
}
```

1. (7 points) Out of the possible range of values for BLOCK\_SIZE, for what values of BLOCK\_SIZE will this kernel function correctly when executing on a current device?
  
2. (6 points) If the code does not execute correctly for all BLOCK\_SIZE values, suggest a fix to the code to make it work for all BLOCK\_SIZE values.
  
3. (6 points) Out of the possible range of values for BLOCK\_SIZE, for what values of BLOCK\_SIZE will the kernel completely avoid uncoalesced accesses to global memory?
  
4. (6 points) For what value of BLOCK\_SIZE (from 1 to 20) will the kernel experience the most severe shared-memory conflicts?