Intensive Computation

Prof. A. Massini

April 17, 2023

Midterm test

Student's Name

_____

***Matricola* number**

_____

Exercise 1 (5 points)	
Exercise 2 (3 points)	
Exercise 3 (3 points)	
Exercise 4 (5 points)	
Exercise 5 (4 points)	
Exercise 6 (5 points)	
Question (4 points)	
Exercise 7 (3 points)	
Total (32 points)	

Exercise 1 (5 points) - GPU & CUDA

Technical specifications	Compute capability (version)									
	1.0	**1.1**	**1.2**	**1.3**	**2.x**	**3.0**	**3.5**	**3.7**	**5.0**	**5.2**
Maximum dimensionality of grid of thread blocks	2				3					
Maximum x-dimension of a grid of thread blocks	65535					2^{31}-1				
Maximum y-, or z-dimension of a grid of thread blocks	65535									
Maximum dimensionality of thread block	3									
Maximum x- or y-dimension of a block	512				1024					
Maximum z-dimension of a block	64									
Maximum number of threads per block	512				1024					
Warp size	32									
Maximum number of resident blocks per multiprocessor	8					16			32	
Maximum number of resident warps per multiprocessor	24		32		48	64				
Maximum number of resident threads per multiprocessor	768		1024		1536	2048				
Technical specifications	**1.0**	**1.1**	**1.2**	**1.3**	**2.x**	**3.0**	**3.5**	**3.7**	**5.0**	**5.2**
	Compute capability (version)									

Consider a tridimensional matrix of size 1200x1200x1200. You would like to assign one thread to each matrix element.

a) How would you select the **3D grid** dimensions and **3D cubic block** dimensions of your kernel to **minimize the number of idle threads** on a device having compute capability 3.0?

b) How would you select the 3D block dimensions of your kernel if you want tridimensional blocks of maximum size and you do not need blocks are cubic, on a device having compute capability 3.0?

Exercise 2 (3 points) - GPU & CUDA

A CUDA device's SM (Streaming Multiprocessor) can take up to 2048 threads and up to 8 thread blocks. Which of the following block configuration would result in the most number of threads in the SM?
(A) 128 threads per block
(B) 256 threads per block
(C) 512 threads per block
(D) 1024 threads per block

Give a comment for each answer.

Exercise 3 (3 points) – Pipeline

Consider an architecture where each instruction (unpipelined) takes 98 ns. Consider the pipeline implementation of instructions taking 105 ns, using 7 pipe stages.
i) Compute the time required to execute 80 instructions without and with pipeline.
ii) Compute the speedup of the pipelined solution with respect to the unpipelined one (for 80 instructions).

Exercise 4 (5 points) - Number representation

- Represent the natural number range [0; 790] using the residue number system, considering:
  - the conventional choice consisting of 3 moduli S1={2^n -1; 2^n ; 2^n +1}
  - a moduli set consisting of 3 moduli at your choice S2.
- Give an estimation of the representational efficiency in both cases.
- Represent A= 25 in **both residue systems** S1 and S2 defined and in the **mixed radix representation** associated.

Exercise 5 (4 points) – Number representation

a) Given the values A= 00 11 10 01 11 and B = 00 10 10 11 01 in the signed RB (Redundant Binary) representation, convert A and B in decimal.

b) Show the execution of operation A-B using the look-up table for addition. Verify the correctness of the result.

Exercise 6 (6 points)

a) Consider the sparse matrix here below, whose pattern is shown on the right.

	1	2	3	4	5	6	7	8	9	10	11	12
1	16.9	19.1	0	12.8	1.25	0	0	0	0	0	0	0
2	19.1	25.5	13.2	21.9	0	25.5	0	0	0	0	0	0
3	9.9	13.2	9	0	9.5	1.13	13.9	0	0	0	0	0
4	0	0	0	18.4	12.9	8.2	4.5	0	2.7	0	0	0
5	0	0	0	12.9	1.1	6.1	0	1.2	3.9	0	0	0
6	0	0	0	0	6.1	4.6	0	2.7	3.9	0	0	0
7	0	0	0	0	0	0	0	0	0	0	0	0
8	0	0	0	0	0	0	0	0	0	0	0	0
9	0	0	0	0	0	0	0	0	0	0	0	0
10	0	0	0	0	0	0	0	16.3	0	12.8	0	9.5
11	0	0	0	0	0	0	12.5	24.9	0	16.3	22.9	0
12	0	0	0	0	0	0	18.4	22.5	25.5	17.7	25.5	13.9

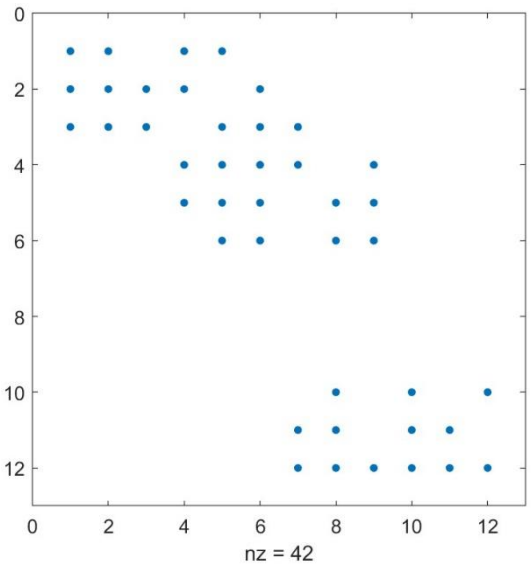

Specify which arrays you need for the following compressed representations and how many bytes they occupy in memory.

BSR

Skyline

b) Explain how arrays change after the insertion of element $m_{8,8}$=1.3 and what is the new memory occupation.

BSR

Skyline

c) Explain what operations must be executed for deleting element $m_{4,6}$

BSR

Skyline

Question (4 points)

Briefly describe the main characteristics of the vector architectures.