



## Quiz navigation

 1  2  3  4

[Finish attempt ...](#)


## Navigation

[Home](#)
[My profile](#)
[Current course](#)
[201510-ITCS-5145-091:ITCS-4145-091-XLSP6201510\\_Com...](#)
[Participants](#)
[Badges](#)
[Week 12](#)
 [Mini-Quiz Week 12](#)
[My courses](#)

## Activities

- [Assignments](#)
- [Forums](#)
- [Questionnaires](#)
- [Quizzes](#)
- [Resources](#)
- [Saba Meeting Events](#)

## Administration

[Quiz administration](#)

- [Edit settings](#)
- [Group overrides](#)
- [User overrides](#)

 [Edit quiz](#)
 [Preview](#)
[Results](#)
 [Locally assigned roles](#)
 [Permissions](#)
 [Check permissions](#)
 [Filters](#)
 [Logs](#)
 [Backup](#)
 [Restore](#)
 [Question bank](#)
[Course administration](#)
[Switch role to...](#)
[My profile settings](#)

## You can preview this quiz, but if this were a real attempt, you would be blocked because:

This quiz is not currently available

### Question 1

Not yet answered

Marked out of 1

[Flag question](#)

[Edit question](#)

In the following, where are **x**, **y**, and **z** stored

```
__global__ mykernel() {
```

```
int x, y, z;
```

```
...
```

```
}
```

Select one:

- a. In GPU (device) registers if possible.
- b. In GPU (device) global memory
- c. In the CPU (host) registers
- d. In GPU (device) cache memory
- e. None of the other answers

### Question 2

Not yet answered

Marked out of 1

[Flag question](#)

[Edit question](#)

In GPUs, would a group of threads access a row of consecutive elements in a 2-dimensional array more efficiently than a column of consecutive elements, each thread accessing one element, according to memory coalescing in general?

Select one:

- a. None of the other answers
- b. It makes no difference
- c. Yes
- d. No

### Question 3

Not yet answered

Marked out of 1

[Flag question](#)

[Edit question](#)

Which CUDA routine is used to synchronize all threads in a multi-block kernel, that is, will wait until all threads in the kernel have completed?

Select one:

- a. CUDABarrier()
- b. None of the other answers
- c. cudaThreadSynchronize()
- d. \_\_threadfence()
- e. \_\_syncthreads()

### Question 4

Not yet answered

Marked out of 1

[Flag question](#)

[Edit question](#)

What does the CUDA qualifier `__device__` indicate when used to prefix an array declaration outside main and any function body?

Select one:

- a. The array will be held in GPU cache memory.
- b. The array will be held in GPU shared memory.
- c. The array will be held in host main memory
- d. The array will be held in GPU global memory.
- e. None of the other answers

[Moodle Docs for this page](#)

You are logged in as [Anthony Wilkinson](#) ([Log out](#))

201510-ITCS-5145-091:ITCS-4145-091-XLSP6201510\_Combined