

EXTRA CREDIT HW #8: Implement the 2D SOR algorithm using CUDA on the Tesla C2070 GPU.

(Recall: 2D SOR - on each iteration replace all interior values by the average of their four nearest neighbors)

Simplifying assumption: square array nxn interior (Let's handle n not matching the dimensions of the grid of threads)

a) Maximum threads per block is 1024 (2<sup>10</sup>). If we want to make it 2-dimensional (and square), what would the dimensions of the thread block (DIM by DIM)?

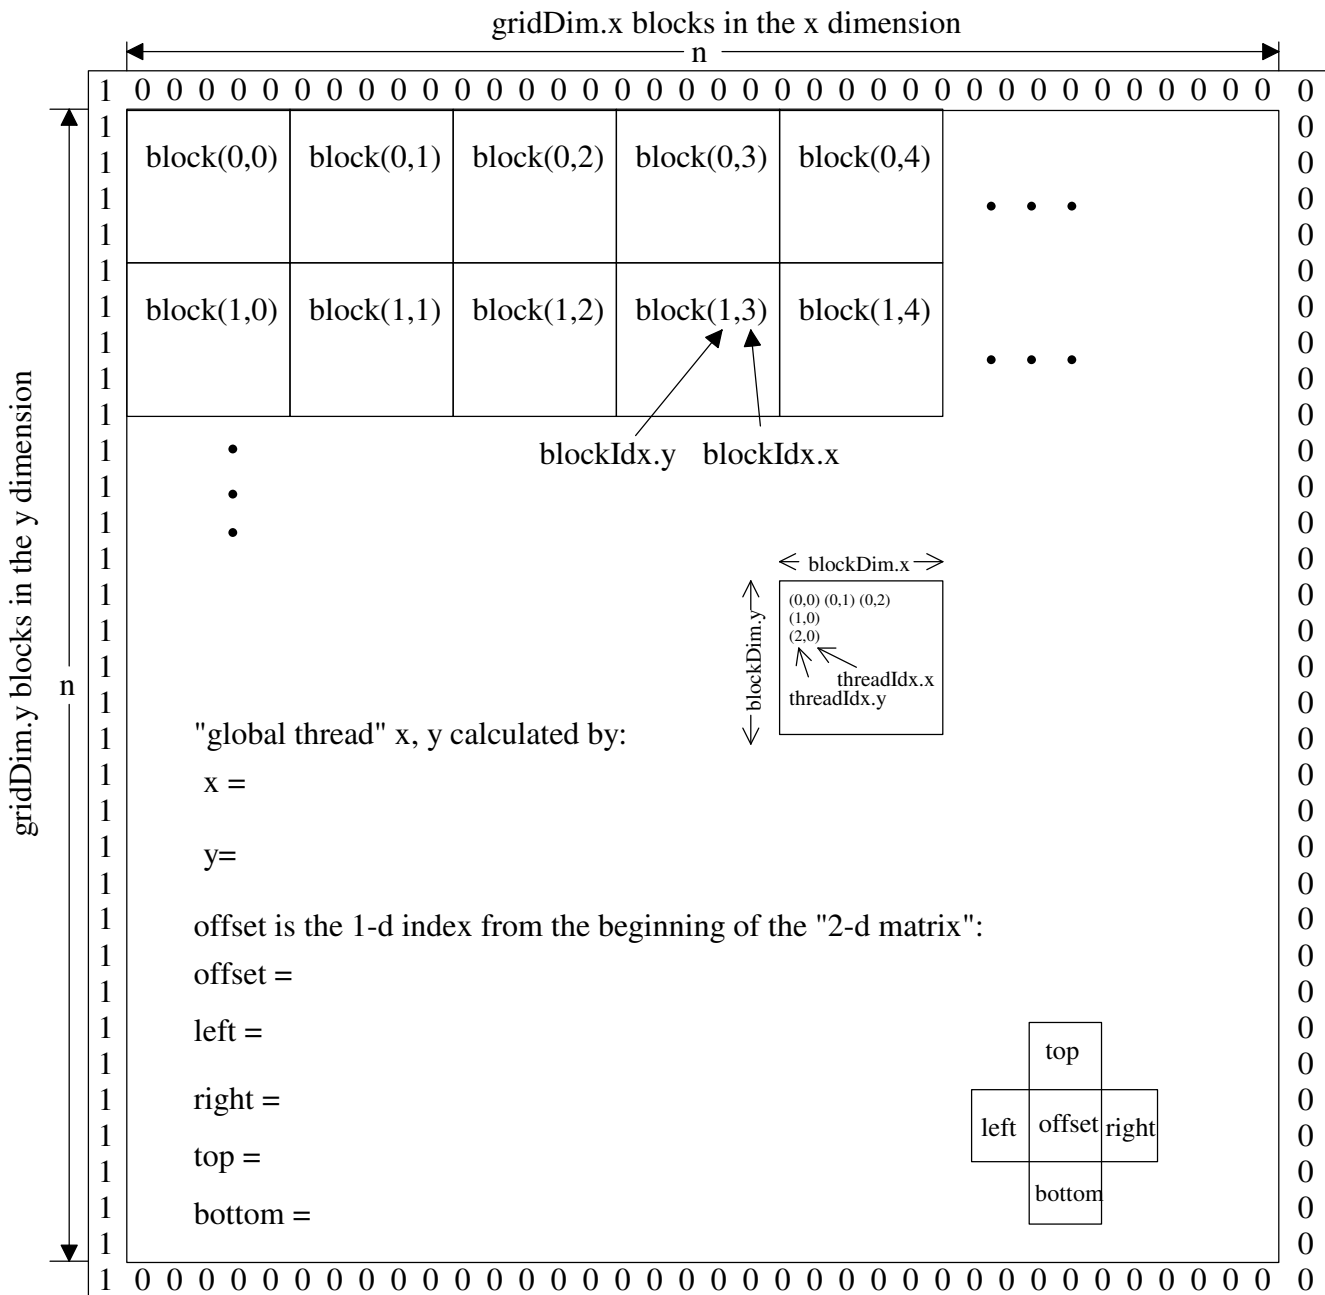
b) If we want to "tile" blocks across the n x n interior of the array, what is the dimension of the grid of blocks?

```
#define DIM
```

```
#define threadsPerBlock
```

```
dim3 threads(          ,          );
```

```
dim3 blocks(          ,          );
```



c) Design the host's algorithm:

d) Design the device's kernel: