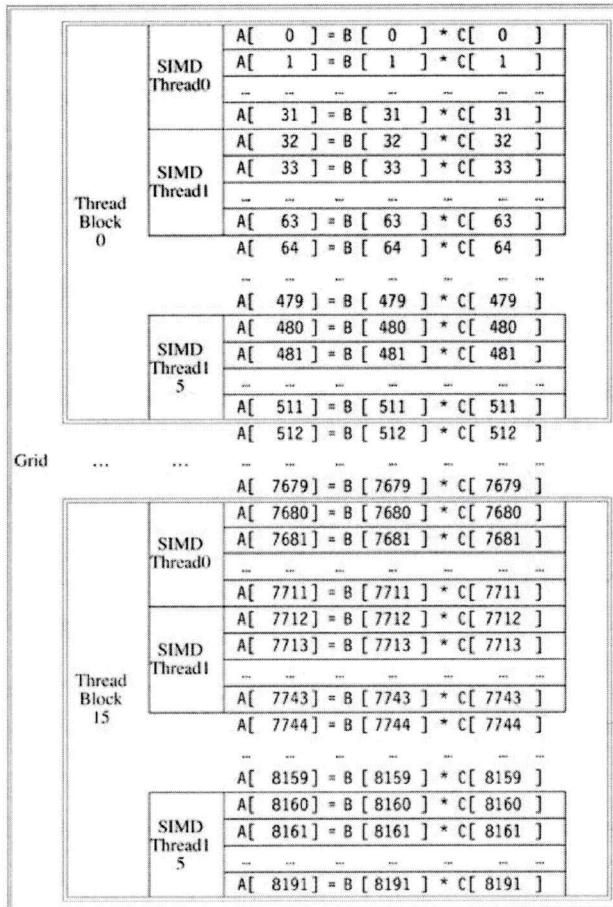


8192 Element Vector * Vector:

$$A[] = B[] * C[]$$

From Hennessy & Patterson, Figure 4.13



CUDA:

gridDim = 16; blockDim = 512 (total 8192 threads); warp size = 32

OpenCL:

global_size = 8192; local_size = 512; wavefront size = 32

H&P: "SIMD Thread *i*" means warp (wavefront) *i*.

For Both:

Each thread block (work group) assigned in its entirety to one SM. This will be of major importance later when we discuss use of shared memory (local memory).