

1

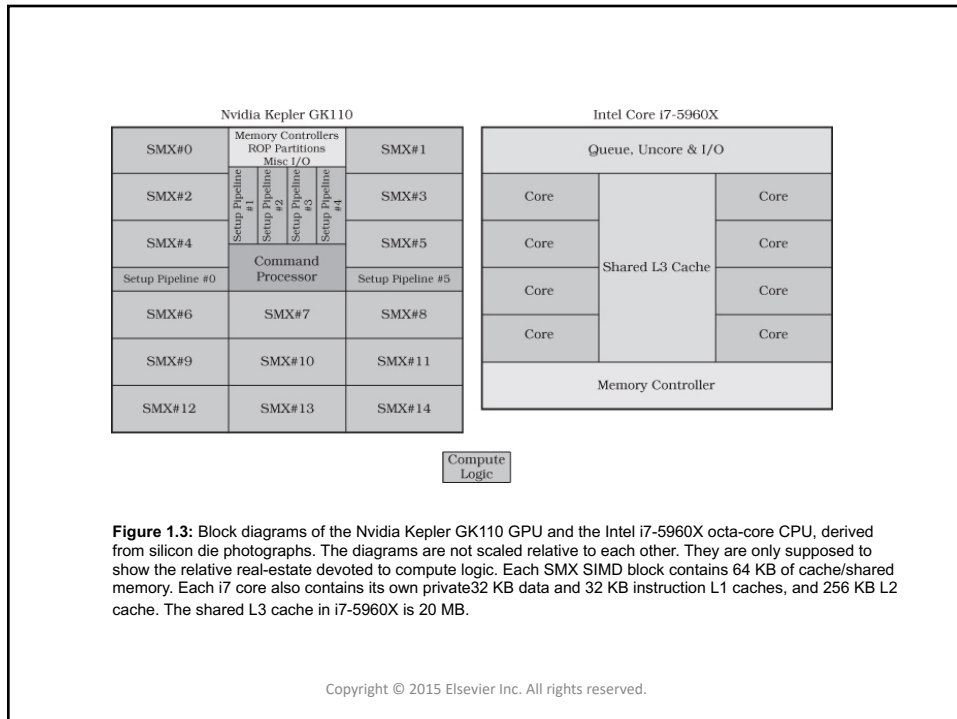


Figure 1.3: Block diagrams of the Nvidia Kepler GK110 GPU and the Intel i7-5960X octa-core CPU, derived from silicon die photographs. The diagrams are not scaled relative to each other. They are only supposed to show the relative real-estate devoted to compute logic. Each SMX SIMD block contains 64 KB of cache/shared memory. Each i7 core also contains its own private 32 KB data and 32 KB instruction L1 caches, and 256 KB L2 cache. The shared L3 cache in i7-5960X is 20 MB.

Copyright © 2015 Elsevier Inc. All rights reserved.

2

