AI??????2D??

1. [illegible] Dot[illegible] 9[illegible]
2. [illegible]
3. L1/L2/L0 [illegible] reshape[illegible] swizzel
4. [illegible] layout [illegible] NHWC[illegible] [illegible] BPI BPK FF
5. [illegible]
6. L0 L1 Fusion[illegible]
7. [illegible] mapping [illegible] fusion
8. feature*weight vs weightT*featueT

?2D Dot?????

1. [illegible] 1D[illegible] 2D[illegible] fusion[illegible]
    1. [illegible] L0[illegible] broadcast[illegible] L1 [illegible] 2D[illegible]
    2. [illegible] load/store[illegible]
    3. [illegible] **2D**[illegible]
    4. [illegible] /[illegible] /[illegible]
2. cuda[illegible] simt[illegible] warp[illegible] thread [illegible] 2D[illegible] PPA[illegible]
    1. [illegible] thread[illegible]
    2. [illegible] 2D[illegible]
    3. [illegible]
    4. [illegible]
    5. [illegible] sm subcore[illegible] block group[illegible]
3. XXU4.0[illegible] thread[illegible] footprint order[illegible] thread[illegible]
    1. [illegible] L1load[illegible] load[illegible] L0[illegible] subcore[illegible]
4. 2D-1D-2D[illegible] Fusion
    1. [illegible]
        1. 8bit int8 [illegible] 64[illegible] 512bit[illegible] 32bit[illegible] 16[illegible]
    2. [illegible]
    3. layout[illegible] NCHW NHWC
    4. [illegible] stride[illegible] 1
5. [illegible] 2D[illegible] 1D[illegible] 2D[illegible]
6. [illegible] L1[illegible] latency[illegible] L1[illegible] L1_base x latency[illegible] L0 [illegible] L0_base[illegible] L1_base[illegible]
7. **swizzle[illegible] renaming[illegible] L0[illegible] bank conflict[illegible] renaming[illegible]**

????

1. 算力和存储器（SRAM）都不能做到单体很大
    1. 很大的单体有利于软件使用方便
    2. 多个小的存储器实例（bank）组成一个大的抽象存储器
2. 物理单元之间很难方便做到cycle by cycle的同步，必须使用额外的同步逻辑
3. 硬件电路虽然可以并行，理论算力很高，但是latency是必然存在的

NV Blackwell

从PTX、SASS可以看到AI面对的问题和很多解决方法都是共同的。

https://docs.nvidia.com/cuda/parallel-thread-execution/index.html#tensorcore-5th-generation-family-instructions https://docs.nvidia.com/cuda/cuda-binary-utilities/index.html#blackwell-instruction-set

1. L0的memory变化
    1. 显式的手动管理tensor memory，是不是2D/1D更像是独立的一个NPU/DSA单元?
    2. Tensor计算直接使用专门的Tensor memory，做memory to memory的Tensor计算，且支持软件动态分配
    3. 1D可以直接load/store这块memory
    4. 这样让L0可以动态，目的估计是破除2D/1D每写一种L0 fusion就要写一份code示例，Fusion算子可以隔离开发
2. 更大的矩阵形状128/256，降低2D带宽需求
3. 支持TPC内2个SM之间的2D协同计算，看起来是2个core支持的读数据复用
4. 支持OCP-MX的micro-scaling
5. 标量thread issue Tensor计算指令，不再是SIMT style的多个thread
6. 直接支持卷积指令，支持weight-stationary GEMM，带上mask bit表达哪些是padding的0

Revision #8
Created 2025-01-11 09:46:28 UTC by Colin
Updated 2026-04-29 07:33:45 UTC by Colin