CUDA□□□□

A100□□□□ tensorcore□□□□□□□□□□□□□□□□ https://zhuanlan.zhihu.com/p/620257581

□□ GPGPU□□□□□□ v2.01.pdf

https://zhuanlan.zhihu.com/p/166180054

https://www.tinyedi.com/cuda_learning/#pipeline

https://docs.nvidia.com/cuda/cuda-c-programming-guide/index.html

https://zhuanlan.zhihu.com/p/486224812

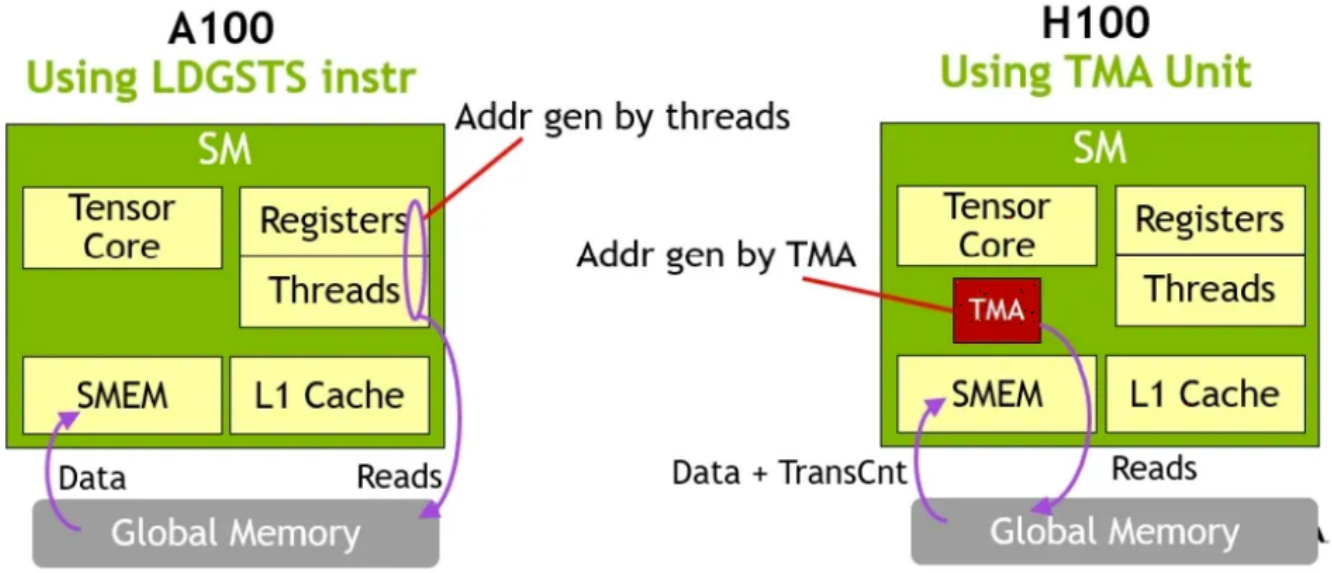

Figure 19. Asynchronous Memory Copy with TMA on H100 vs LDGSTS on A100

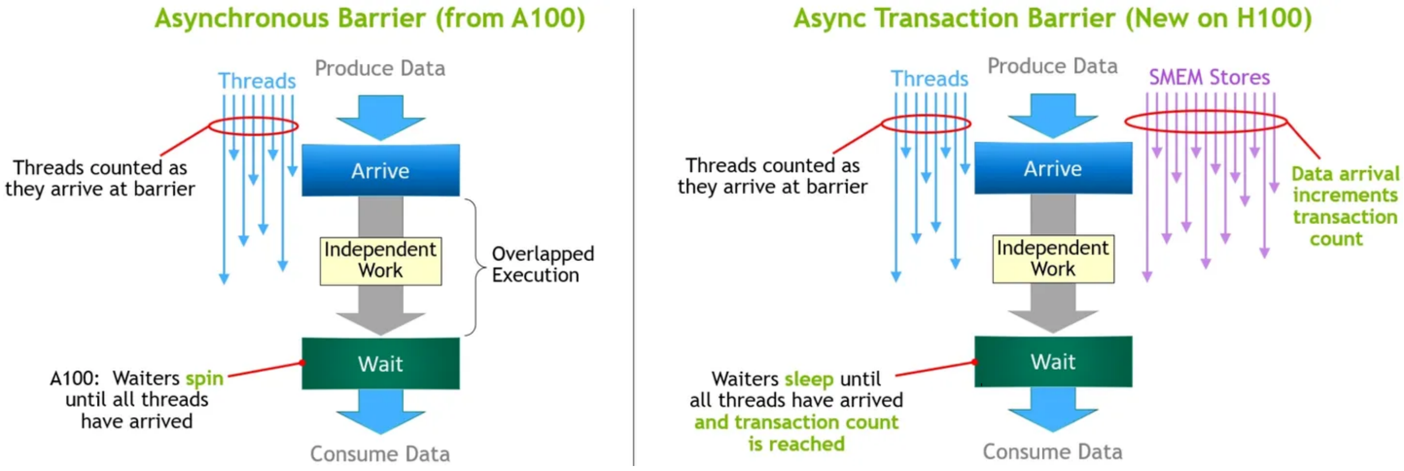

Hopper Hopper Asynchronous Transaction Barrier Asynchronous Transaction Barrier Hopper

Predicate @P0, @!P3 predication PTX guard predicate predicate register P0, P6, PT 4bit predicate 3bit 2^3=8 predicate register P0~P7 P7=PT True 1bit @PT @! PT ……
Predicate conditional branch predicate warp divergence branch latency instruction cache branch divergence mask thread active mask PTX warp vote load %lanemask_* warp mask

Cache control cache line load cacheline cache cache load latency PTX Turing TLB latency latency

Control code reuse read barrier write barrier wait barrier yield hint stall count reuse GPR bank conflict GPR barrier thread yield stall count warp

memory memory bound memory bound latency

[illegible]

SASS[illegible] predicate[illegible]

[illegible] `LDG.E R6, [R2] ;` [illegible] `R[2:3]` [illegible] memory IO[illegible] queue[illegible] conflict[illegible] shared memory[illegible] bank conflict[illegible] register[illegible] bank conflict[illegible] atomic[illegible] conflict[illegible] constant memory[illegible] warp[illegible] conflict[illegible] re-issue[illegible] cache miss[illegible] replay[illegible] dispatch port[illegible]

[illegible] CUDA[illegible] Out of order[illegible]

Register Reuse Cache : [illegible] hint[illegible] cache [illegible] read

Wait Dependency Barrier : [illegible] 6bit[illegible] bit[illegible] dependency barrier[illegible] 6[illegible] dependency barrier[illegible] barrier[illegible] Read[illegible] Write[illegible] wait dependency barrier[illegible] SASS[illegible] `DEPBAR.LE SB0, 0x0, {2,1} ;`

Read dependency barrier : [illegible] 3bit[illegible] 6[illegible] barrier[illegible] 0~5[illegible] barrier 1-6 [illegible] barrier[illegible] 0b111[illegible] Read dependency barrier [illegible] hold[illegible] GPR [illegible] GPR[illegible] **Read dependency barrier**[illegible] **memory**[illegible] `F2I/I2F` [illegible]

Write dependency barrier : [illegible] read dependency[illegible] 3bit[illegible] barrier[illegible] Read[illegible] Write[illegible] dependency barrier[illegible] wait[illegible] 6[illegible] Write dependency barrier [illegible] GPR[illegible] predicate[illegible] **barrier** [illegible] **data race**[illegible] latency[illegible] latency [illegible] stall cycle[illegible] race[illegible]

Yield hint flag : [illegible] 1bit[illegible] Yield[illegible] cycle[illegible] warp[illegible] yield[illegible] bit[illegible] stall count[illegible] bit[illegible] 0[illegible] stall[illegible] cycle[illegible] >16[illegible] warp [illegible]

Stall count : [illegible] 4bit[illegible] stall[illegible] cycle[illegible]

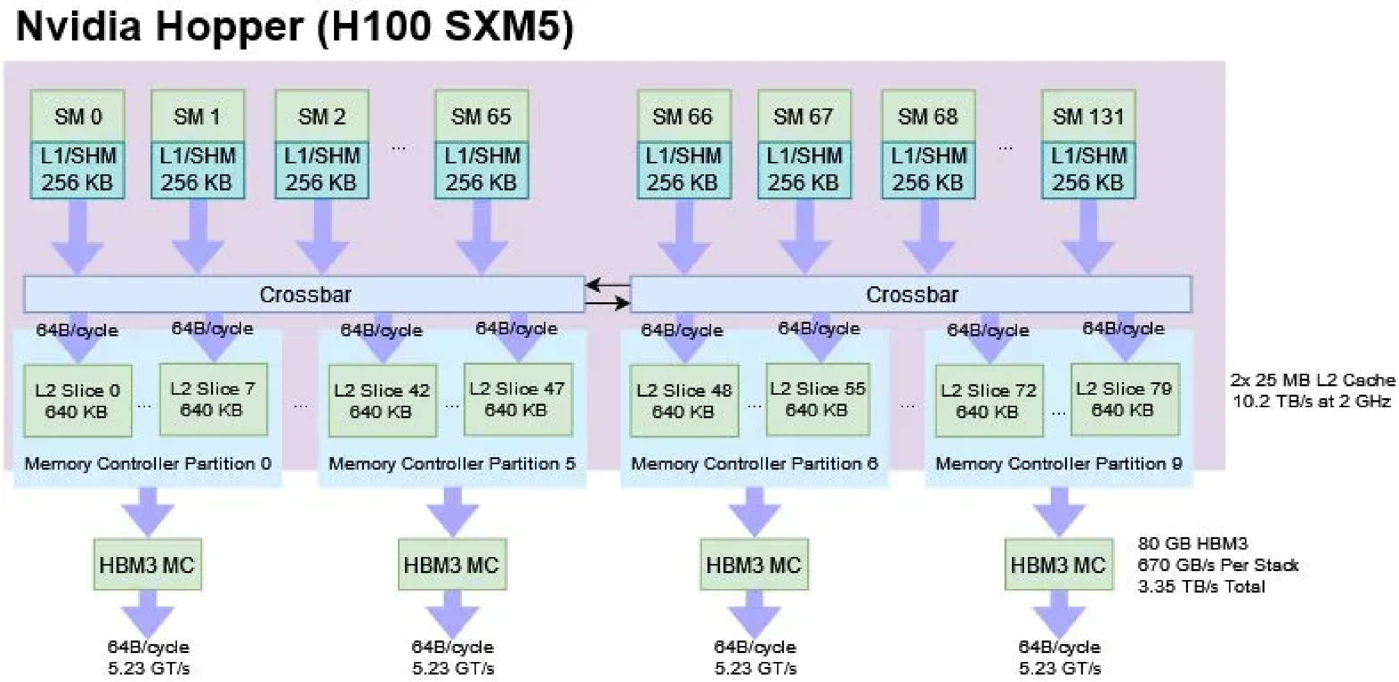

CUDA[illegible]

- [illegible]
- [illegible]
- [illegible] Adjust kernel launch configuration to maximize device utilization.
- [illegible]
- [illegible] Minimize redundant accesses to global memory whenever possible.
- [illegible] Avoid long sequences of diverged execution by threads within the same warp.

TMA

TMA [illegible] SM [illegible] 1D [illegible] 5D [illegible] [illegible] [illegible] add/min/max [illegible] / [illegible]

memcpy_async[illegible] API[illegible] , [illegible] group,barrier,pipeline.

- group: [illegible] `cooperative_groups::wait(group)` [illegible] ,[illegible] `group.sync()`, [illegible] group[illegible] . `cooperative_groups::wait` [illegible] ,[illegible] `group.sync()` [illegible] .
- barrier: [illegible] `barrier.arrive_and_wait()` [illegible] ,[illegible] ,[illegible] barrier[illegible] count,[illegible] count[illegible] async_memcpy[illegible]
- pipeline: [illegible] queue

Revision #1
Created 11 January 2025 09:46:27 by Colin

Updated 12 January 2025 06:33:19 by Colin