GMP

Sync And Async

?????

1. □□□□□□□□□□□□□□□
2. □□□□□□□□
3. N to N□□□□□
4. □□□□□□□□□□□
5. □□□□□□□□□□□□□ pulling□□ latency
6. □□□□□□□□□□□□□□
    1. transformer□ Flash-attention□□□□ L1□□□□ fusion□□□□□□ DMA□□□□□□
7. □□□□□□□

??Global??????????????

1. □□□□
2. □□□□□□□□□□□□□□□
3. □□□□□□□□□□□□□□□□□
    1. DMA□ Kernel□□□□□
4. □□□□□□□□□

??????????

1. □□□□□□□□□□ Sync□□□□□□□□□□□□□□□ mailbox□□ N to N□□□□
2. □□□□□□□ producer consumer□□ Pipeline
3. □□□□□□
    1. □□□□□□□□□□□□□□□□□□□
    2. □□□□□□□□□□□□□□□□□□□□□□□□□
4. □□□□□□□□□□□□□□□□□□□□□□□□□□□□
5. □□□□□□□
    1. □□□□□□□□□□□□□□□□□□□□□□□ write mbx□ wait mbx □□□□□□□□□□□ ack□ latency□□□□□□□□□

?????????? mbarrier

1. □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□ M barrier □□□
    1. □□□□□
    2. □□□□□□□□□□□□□□□
2. □□□□□□□□□□□□□□
    1. □□□ □
    2. □□□□□□□ fence
    3. □□□ Core0 Store() Core0 Fence() memory Ack() Core0 Signal(Core1 mbx)

    4. [illegible] Core0 Store_mbarrier() Core1 Wait(mbarrier)
3. [illegible]
    1. [illegible]
4. [illegible]
    1. [illegible] wait [illegible] local [illegible] mbarrier [illegible] counter [illegible] latency [illegible]
    2. [illegible] transaction [illegible]

AI??????2D??

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

?2D Dot?????

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

????

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

GMP

1. [illegible]
    1. [illegible] AI[illegible]
        1. [illegible]
        2. [illegible]
        3. [illegible]
        4. [illegible]
    2. [illegible]
2. [illegible]
    1. [illegible]
        1. [illegible] 2[illegible]
        2. [illegible] 4[illegible]
            1. DRAM[illegible]
            2. [illegible] LUT[illegible]
            3. [illegible]
        3. [illegible]
            1. [illegible] **L0 Cache**[illegible]
    2. [illegible]
        1. MoE
        2. [illegible]
    3. [illegible]
        1. [illegible]
            1. [illegible] credit[illegible] L1[illegible]
            2. [illegible] NoC[illegible]
            3. L2/L3[illegible]
            4. [illegible]
        2. outstanding/[illegible]
        3. [illegible]
        4. launch[illegible]
    4. [illegible]
        1. [illegible]
        2. [illegible] /[illegible]
            1. [illegible] [illegible] [illegible]
        3. [illegible] warp[illegible]
        4. [illegible] [illegible] [illegible]
    5. [illegible] KV Cache ..) [illegible]
        1. [illegible]
        2. [illegible]
    6. [illegible]
        1. **Binary Lut** [illegible] ?
        2. **CIM** [illegible] [illegible] ?
    7. 1D[illegible]

    1. [illegible] sequence[illegible]
    2. [illegible]
    3. [illegible] LD MUL ST[illegible]
       ---> DSA[illegible]
    4. [illegible] L1[illegible] dma[illegible] Mac[illegible] dma[illegible] Mac[illegible] l1
       [illegible]
    5. [illegible]
        1. 1D[illegible] 2D
            1. VLD VST MLD MST VMUL VMUL_reduce MUL_join [illegible] [illegible]
        2. LD/ST[illegible] mbarrier
            1. [illegible] L1[illegible] bank[illegible] bank[illegible] count
               [illegible] count
3. [illegible]
    1. [illegible]
    2. [illegible]
    3. [illegible] edga[illegible]
    4. [illegible]
    5. [illegible]
4. [illegible]
    1. [illegible] Fork[illegible] launch[illegible] join sync
    2. [illegible] [illegible] sync
    3. … …
5. [illegible]
6. [illegible]
    1. load[illegible] fetch
    2. [illegible]
7. ISA
8. Launch
9. Sync

[illegible]

1. [illegible] /[illegible]
2. [illegible]

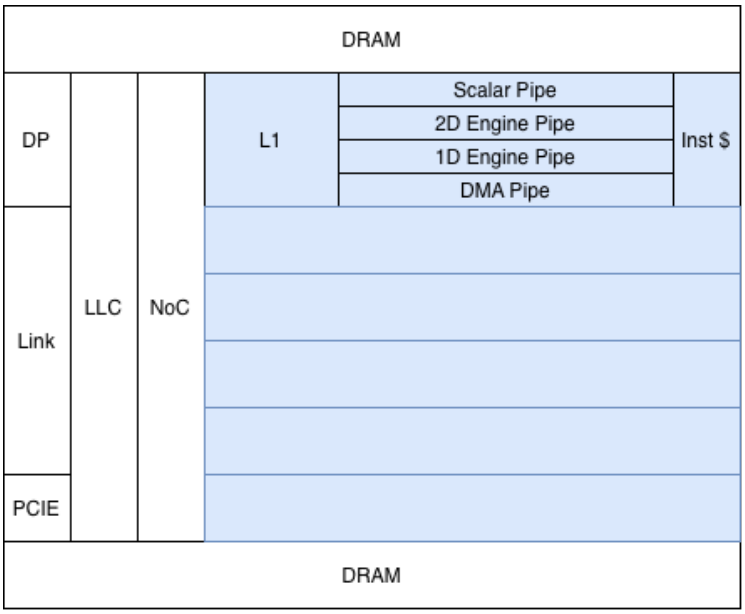
DRAM
DP
L1
Scalar Pipe
2D Engine Pipe
1D Engine Pipe
DMA Pipe
Inst $
Link
LLC
NoC
PCIE
DRAM

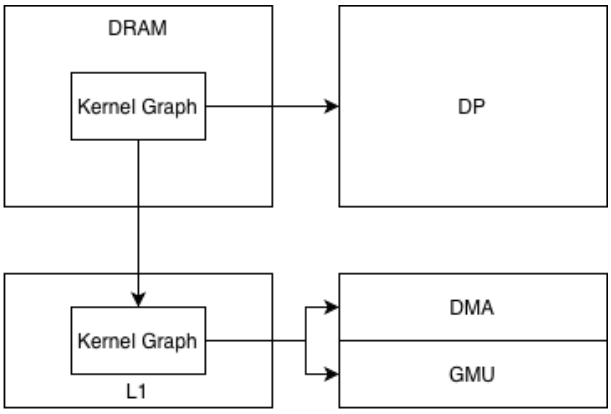
DRAM
Kernel Graph
DP
Kernel Graph
L1
DMA
GMU

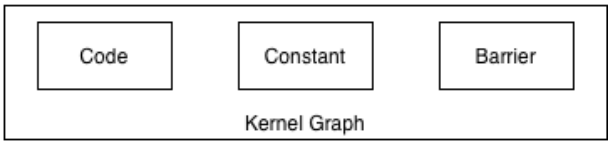
Code
Constant
Barrier
Kernel Graph

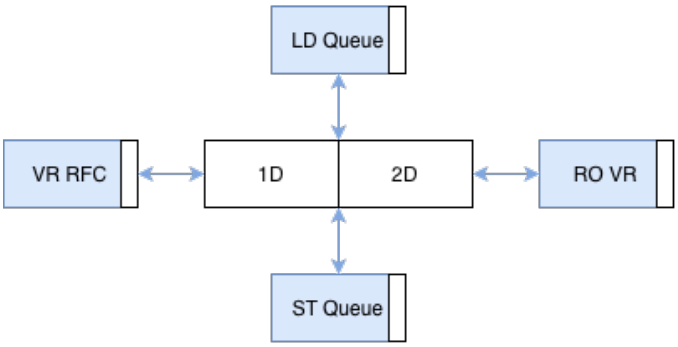
LD Queue
VR RFC
1D
2D
RO VR
ST Queue

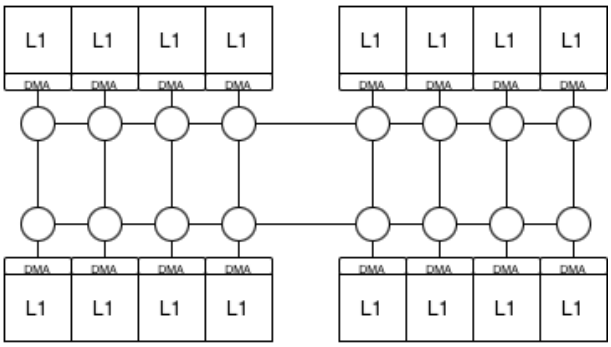
L1
DMA

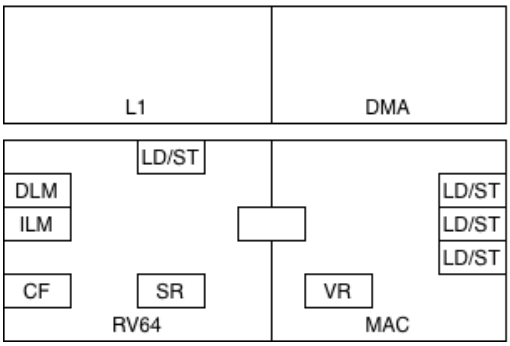
L1
DMA
LD/ST
DLM
ILM
CF
SR
RV64
VR
MAC
LD/ST
LD/ST
LD/ST

??/???

??

1. □□□□□□□□□□□□□□□□□□□□□□□□□□□□
2. □□□□□□ SP-PU-L1-DMA□□□□□□□□□□□□□□□□□□□□□□□□□□□□
   1. □□□□□□□□
   2. □□□□□□□□□
   3. □□□□□□□□□□□□□□□
   4. LD/ST □□□□□ fence□□
3. Launch□ fork
   1. □□□□□□□□□□□□□□□□□□□□□□□□
4. launch/signal/wait:join
   1. launch pu instrution: write_back_id local_id
      1. write_back atomic add/sub
   2. wait instruction: local_id
      1. local_id GE LE counter
   3. wait remote instruction: remote_id
      1. local_id GE LE counter
5. □□□□ &□□
   1. □□□□□□□□□□
   2. □□□□□□□□
      1. □□□□□□ scale
      2. 1D/2D□□
      3. □□ Reduce□□□□□□□□
         1. L3 Atomic
   3. □□□□□□□□□□□□□□□□□□□□□□□□□□□□
   4. □□□□□□□□□
      1. leading tailing latency
         1. LD/ST inflight delay□ DMA Delay
            1. ping-pong
            2. pipeline
         2. kernel□□ latency
            1. kernel□□□
         3. kernel□□ gap
            1. kernel□□□□□□□□□□
            2. □□□□□□ fusion
      2. □□□□□ ready□□□ fence□□□□□
         1. □□□□□□□□
            1. m-barrier
      3. □□□□□□□

MAC

1. Vector计算单元支持乱序和并行
2. 软件编译指定明确的VR寄存器依赖，RO WO属性
3. 自动拆解大指令成Vector指令，并行执行
4. 软件静态排布VR寄存器，生成依赖关系，申请和释放管理
5. Vector指令支持标量对其进行动态配置

Software

1. 使用独立的硬件仿真软件，不依赖硬件仿真运行

ISA

1. 标量
    1. RV64i
2. 向量
    1. VLD VST VMUL VADD REDUCE_ADD REDUCE_MAX REDUCE_MIN VMUL_REDUCE_ADD
    2. MLD MST
3. 张量
    1. GEMM
        1. 左/右数: 128个 =7位 + 3位扩展矩阵 = 10位
        2. 输出: 7位
        3. Opcode
4. fence
    1. L1 cache line 计数
5. VR
    1. 软件管理VR data hazard? VR之间的依赖?
        1. 增加指令的表达信息
    2. 软件分配VR 还是硬件rename ing，解决bank冲突?
    3. 利用VR count? 软件进行管理依赖关系?
        1. LD ST 计算的三类指令之间可以并行，通道内部没有必要并行
            1. 因为硬件资源没有特殊性，不会因为并行而减少气泡
        2. 默认，GEMM指令一定要在前面LD指令之后执行
        3. 默认，ST指令一定要在前面计算指令之后执行
6. L1
    1. 软件管理cache line
    2. **cache line**硬件计数，自动异步等待
    3. 针对**L1 CacheLine**的编程?
        1. 软件指定**Load**到**L1 cache line**的位置和有效长度
        2. 向量指令按照**cache line**的粒度和**mask**来执行指定的计算
        3. 这整个流程都是提前编译好，从**L1-L1**都是提前确定的
        4. 针对不同**MNK**大小需求，可以通过标量指令来快速配置，支持动态性
7. DMA

1. 附带一片ROM空间，可以在Launch Kernel的时候主动加载一片指令编码，和当前kernel绑定的
2. 支持一个cmd端口，接受其他核心的控制命令
3. DMA作为对外的接口，通过配置kernel的linear copy kernel作为launch的起点？
    1. 简化系统复杂度，通用的功能组件，避免开发专用的launch constructor
8. 增加到128个标量寄存器支持RV64扩展?
9. 增加指令流控制
10. 增加配合/加速向量单元的定制指令?通过兼容RiscV-V的指令来实现?？

??_???

??/???

??

1. □□□□□□□□□□□□□□□□□□□□□□□□□□
2. □□□□□□ SP-PU-L1-DMA□□□□□□□□□□□□□□□□□□□□□□□□
    1. □□□□□□□□
    2. □□□□□□□□□
    3. □□□□□□□□□□□□□□
    4. LD/ST □□□□□ fence□□
3. Launch□ fork
    1. □□□□□□□□□□□□□□□□□□□□□□
4. launch/signal/wait:join
    1. launch pu instrution: write_back_id local_id
        1. write_back atomic add/sub
    2. wait instruction: local_id
        1. local_id GE LE counter
    3. wait remote instruction: remote_id
        1. local_id GE LE counter
5. □□□ &□□
    1. □□□□□□□□□□
    2. □□□□□□□□
        1. □□□□□□ scale
        2. 1D/2D□□
        3. □□ Reduce□□□□□□□□
            1. L3 Atomic
    3. □□□□□□□□□□□□□□□□□□□□□□□□□□
    4. □□□□□□□□□
        1. leading tailing latency
            1. LD/ST inflight delay□ DMA Delay
                1. ping-pong
                2. pipeline
            2. kernel□□ latency
                1. kernel□□□
            3. kernel□□ gap
                1. kernel□□□□□□□□□□
                2. □□□□□□ fusion
        2. □□□□□ ready□□□ fence□□□□□

        1. 无缝的同步机制
            1. m-barrier
    3. 不对齐、尾数

MAC

1. Vector计算单元支持乱序和并行
2. 软件编译指定明确的 VR寄存器依赖， RO WO属性
3. 自动拆解大指令成 Vector指令，并行执行
4. 软件静态排布 VR寄存器，生成依赖关系，申请和释放管理
5. Vector指令支持标量对其进行动态配置

Software

1. 使用独立的硬件仿真软件，不依赖硬件仿真运行

ISA

1. 标量
    1. RV64i
2. 向量
    1. VLD VST VMUL VADD REDUCE_ADD REDUCE_MAX REDUCE_MIN VMUL_REDUCE_ADD
    2. MLD MST
3. 张量
    1. GEMM
        1. 左 /右数: 128个 =7位 + 3位扩展矩阵 = 10位
        2. 输出: 7位
        3. Opcode
4. fence
    1. L1 cache line 计数
5. VR
    1. 软件管理 VR data hazard? VR之间的依赖?
        1. 增加指令的表达信息
    2. 软件分配 VR 还是 硬件 rename ing，解决 bank冲突?
    3. 利用 VR count? 软件进行管理依赖关系?
        1. LD ST 计算 的三类指令之间可以并行，通道内部没有必要并行
            1. 因为硬件资源没有特殊性，不会因为并行而减少气泡
        2. 默认， GEMM指令一定要在前面 LD指令之后执行
        3. 默认， ST指令一定要在前面计算指令之后执行
6. L1
    1. 软件管理 cache line
    2. **cache line**硬件计数，自动异步等待
    3. 针对 **L1 CacheLine**的编程?
        1. 软件指定 **Load**到 **L1 cache line**的位置和有效长度

   2. 向量指令按照**cache line**的粒度和**mask**来执行指定的计算
   3. 这整个流程都是提前编译好，从**L1-L1**都是提前确定的
   4. 针对不同MNK大小需求，可以通过标量指令来快速配置，支持动态性
7. DMA
   1. 附带一片ROM空间，可以在Launch Kernel的时候主动加载一片指令编码，和当前kernel绑定的
   2. 支持一个cmd端口，接受其他核心的控制命令
   3. DMA作为对外的接口，通过配置kernel的 linear copy kernel作为launch的起点 **? **
      1. 简化系统复杂度，通用的功能组件，避免开发专用的launch constructor
8. 增加到128个标量寄存器 支持RV64扩展?
9. 增加指令流控制
10. 增加配合/加速向量单元的定制指令? 通过兼容RiscV-V的指令来实现? ?

????

????????

□ [[Pipe]] □ 6 □ §208 □□□□□□□□□□ RTL □□□□□□□□□□ = □□ engine pipe □□□

0. ? DMA ?????

□□□□□ [[DMA]] □□□□ □□□ "engine pipe □□ + sync pool + □□□□ + □□□□ fabric master"□□□□□□□□□

	DMA	□□□□
□□□□	□ → transform → □	□ → □□□□ → □
□□□	□□□□□ L1□ DRAM□	□□□□□ L1□ DRAM□
□□□□	transform buffer	L0 □□□□
cmd_desc □□□	xform_flags, xform_params	compute_op, loop, dtype, accum

□□□□□□□ DMA □□□□□□ [[DMA]] □□□□□□□□□□□□

1. ????

□□	□□	□□	□□
`N_CU`	□□□□□□□□ = □□ engine pipe□	4	□□□□□ §13□
`N_SYNC`	sync counter □	32	
`Q_MAX`	□□□□□□□□□□	16	2 □□
`K_IN / K_OUT / K_SIG`	cmd_desc □□□	4/4/4	
`N_L0_VEC`	□□□ L0 □□□□□	8	
`N_L0_TEN`	□□□ L0 □□□□□	4	
`VEC_LANE`	□□□ lane □	16	INT8 lane
`MAC_M / MAC_K`	2D MAC □□□	8 / 8	systolic
`MAC_ACC_W`	MAC □□□□	32	INT32
`OUT_TYPE_SET`	□□□□□	{8, 16, 32}	round/sat □□
`LOOP_W`	inner-loop □□□□	16	

□□	□□	□□	□□
`STRIDE_DIM / STRIDE_W / SHAPE_W`	□ [[DMA]]	3 / 32 / 24	
□□ `COUNT_W / HANDLE_W / ADDR_W / NOC_W`	□ [[DMA]]		

2. ????

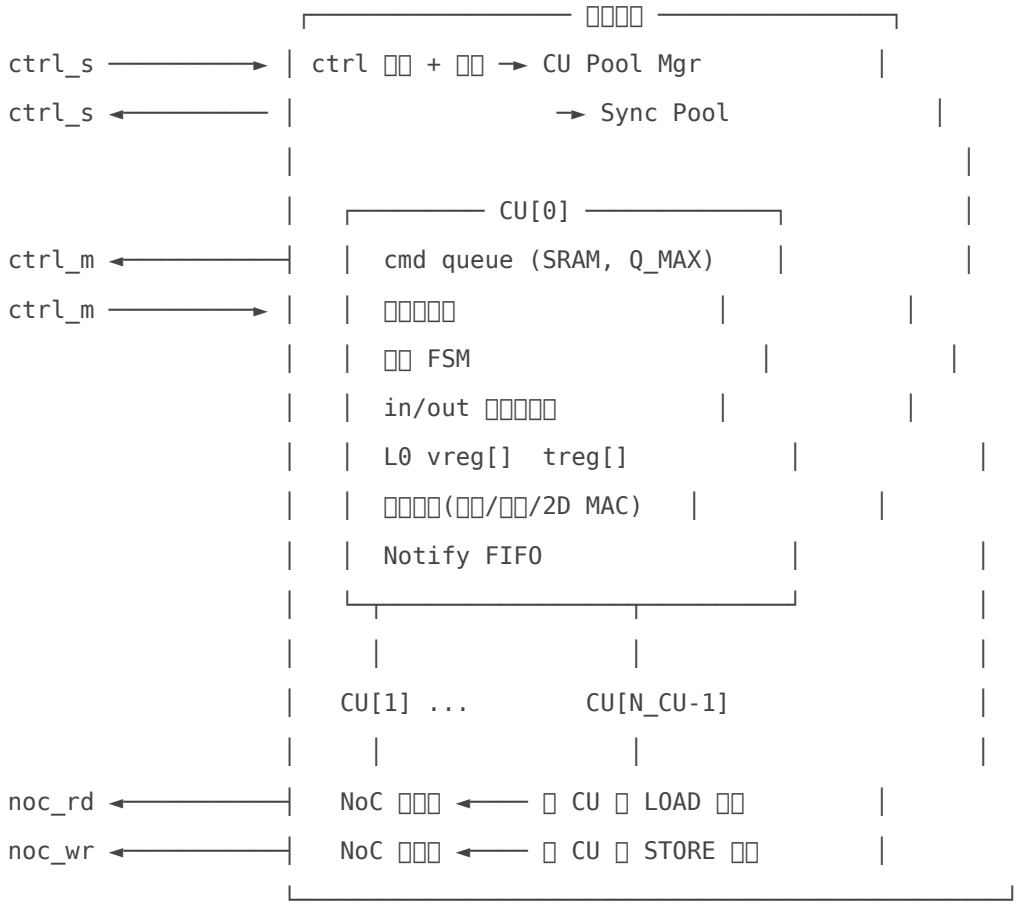

3. ????

□□□□ [[DMA]] §3 □□□□□□□

- `noc_rd_*` / `noc_wr_*` □□□□□□□□□□ L1□ DRAM□□□□ cmd_desc □□ in/out handle.unit_addr □□
- `ctrl_s_*` / `ctrl_m_*` □□□□□□

4. cmd_desc ????

□□ ≤ 1024 bit□□ [[DMA]] §4 □□□□□□ `op + □□□□□` □

□□	□□	□□	□□

0	4	`op`	0=SCALAR_ALU, 1=VEC_ALU, 2=VEC_REDUCE, 3=MAC_2D, 4=ELEMENTWISE_ACT(ReLU/GELU), 5=SOFTMAX_PARTIAL, ...
4	4	`dtype_in`	0=I8, 1=I16, 2=I32, 3=FP8(□□), 4=FP16(□□)
8	4	`dtype_out`	□□
12	2	`accum_mode`	0=□□ , 1=□□□ C += A*B □ , 2=□□ +□□
14	2	`act_kind`	ELEMENTWISE □□□□□□□
16	16	`loop_count`	inner-loop □□□□□□□□ step □□□□
32	48	`src_a_base`	byte_addr □□□□□□□□ pipe □□□ [[L1]] §4.1 / [[DRAM]] □□□□
80	96	`src_a_stride[3]`	
176	72	`src_a_shape[3]`	
248	48	`src_b_base`	MAC_2D / VEC_ALU □
296	96	`src_b_stride[3]`	
368	48	`dst_base`	
416	96	`dst_stride[3]`	
512	72	`dst_shape[3]`	
584	8	`l0_a_reg`	A □□□□□□□ L0 □□□
592	8	`l0_b_reg`	
600	8	`l0_c_reg`	□□□□□
608	4+4+4+4	`n_in / n_out / n_sig / _`	
624	96	`in_list[4]`	□□ {handle:16, delta:8}
720	96	`out_list[4]`	
816	96	`sig_list[4]`	
912	112	reserved/padding	

`src_b_*` □□□□□□□ MAC / VEC_ALU□□□□□□□□□□□□ reduce□□ 0□

5. ??????

□□□ [[DMA]] §5□□ beat `PSEND` □□ 8 × 128 bit□ opcode □□□□

6. CU Pool ???

□□□ [[DMA]] §6□□□□

- `meta[N_CU]` □□□ `cu_class` □□□□□ / □□ / MAC□□□□□□□
- PALLOC □ `params.cu_class` □□□□□□ class □□□□□□□□□□□□□□□□□□□

7. CU ?????×N_CU?

7.1 ????

□□□ [[DMA]] §7.1□

7.2 ?????

□□□ [[DMA]] §7.2□

7.3 ?? FSM

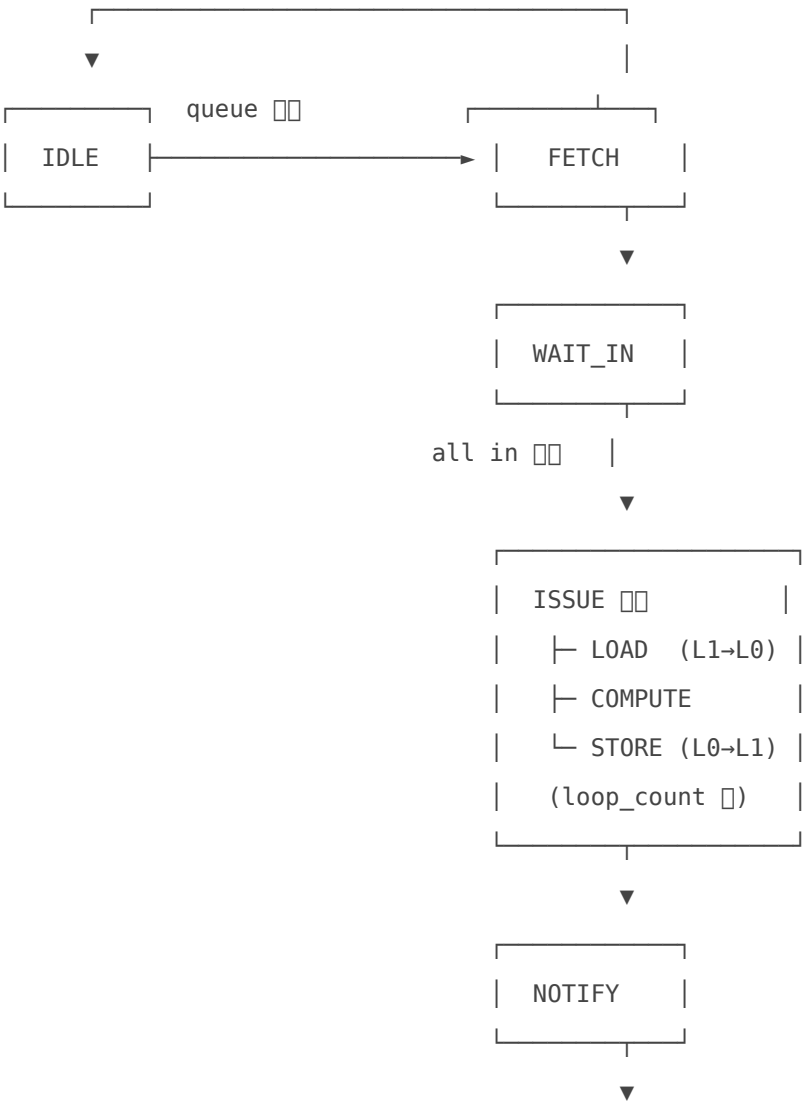

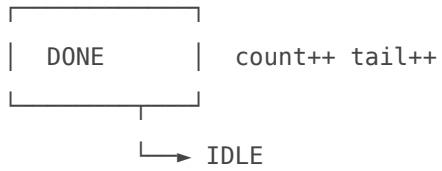

□□	□□	□□
IDLE	1 cycle	
FETCH	1 + (STRIDE_DIM-1) cycle	□ cmd_desc + □□ stride wrap_delta □□ DMA §7.4□
WAIT_IN	□□	□□ COUNT_QUERY in_list
ISSUE	□□□□ §7.6	
NOTIFY	(n_in + n_out + n_sig) cycle	□ Notify FIFO
DONE	1 cycle	

7.4 L0 ????

□	□□
`vreg[N_L0_VEC]`	□□ `VEC_LANE × dtype_in_bits` □□□ vector □□□ / □□
`treg[N_L0_TEN]`	□□ `MAC_M × MAC_K × dtype_in_bits` □□□ tile □□□
□□	vreg: 2R/1W□ treg: 2R/1W□ FMA □□□□
□□	□□□ □ undefined□□□□□□□ LOAD □□
□□	□□□□□□□ CU □□

□□□ vreg □ register file□ small□□ treg □ SRAM macro□□□ treg ≈ 8×8×8 = 512 bit□ N_L0_TEN = 4 □ 2 KiB□□ treg SRAM □□□ 2R/1W □□ MAC 2 □□ 1 □□□

vreg □ treg □□□□ —— □□□□□□□□□□

7.5 ????

□□□□□□□□□□□□□□□□□ cu_class □□□□

7.5.1 ??

- 1 lane□□□ `max(dtype_in, dtype_out)` ≤ 32 bit
- op□ ADD/SUB/MUL/SHIFT/CMP/SELECT/BIT
- latency 1 cycle□ throughput 1/cycle

7.5.2 ??

- `VEC_LANE` □□□ INT8 ALU + 16-bit accumulator option
- op□ lane-wise ADD/SUB/MUL/MAC/CMP□□ lane REDUCE_SUM/MAX/MIN
- latency 1-2 cycle□ throughput 1/cycle
- elementwise activation□ LUT □□ 4 KiB ROM/SRAM□□□ ReLU / GELU / sigmoid □

7.5.3 2D MAC

- `MAC_M × MAC_K` systolic □□
- □□□ INT8 × INT8 → INT16 partial → □□□ MAC_ACC_W □□□
- □□□ A treg (M × K), B treg (K × N□□□□ N = MAC_M)□□□□□□□□□□ cmd_desc □□□
- □□□ C treg (M × N)□□□□□□ treg□ accum_mode□
- pipeline depth = `MAC_K + 2` cycle
- □ cycle □□□□ A□□□ B□ MAC_M = MAC_K = 8 □ 8 cycle □□□□ 9 cycle □□ cycle □□□□□□

□□ / □□ / round□ cmd_desc.dtype_out □ accum_mode □□□□□□□ saturate + truncate / round_to_nearest□□□□□□□□□□□ stage□□□□

7.6 LOAD/COMPUTE/STORE inner-loop ??

□□□ cmd □□ `loop_count` □ micro-op□□□□□□

```
cycle:  N   N+1   N+2   N+3   N+4   N+5
        |    |     |     |     |     |
LOAD    op0  op1   op2   op3   op4   op5
COMPUTE      op0   op1   op2   op3   op4
STORE              op0   op1   op2   op3
```

- L0 □□□□□□□□ buffer□ LOAD □□ COMPUTE □ + □□ STORE □
- □□□ `l0_a_reg / l0_b_reg / l0_c_reg` □□□□ op □□□ reg□□□□ `op_id % N_L0_REG` □ cyclic □□□□□□□□
- leading 2 cycle □ LOAD□ COMPUTE/STORE □□□□□ tailing 2 cycle □ STORE
- inner-loop □□□□□ / □□□□□ stall□□ NoC backpressure□

7.7 in / out ?????

□□□ 3 □□□□□□□ A□ B□ C□□□ [[DMA]] §7.4 □□□□□□□□□□

- □□□□□□□□□□□ L1 / DRAM□□□□□□□
- inner-loop □ `loop_count` □□□□ `shape[3]` □□□□□□
- □□□□□□□□□ NoC □□□ burst □□ = □□ L0 □□□□□ / NOC_W□□□ 1 beat□

7.8 Notify FIFO

□□□ [[DMA]] §7.6□

8. Sync Pipe Pool

□□□ [[DMA]] §8□

9. ??????

9.1 NoC ? / ?

□□□ [[DMA]] §9.1□ NoC □□□□□ L1 / DRAM□□□ L1 □□□□□□□□□ outstanding □□□□□ DRAM □□ DRAM □□□□□□□

9.2 ctrl_m ??

□□□ [[DMA]] §9.2□

9.3 ctrl_s ????

□□□ [[DMA]] §9.3□□□□ PALLOC □ `params.cu_class` □□□□□ cu_class □□□□

10. ????????

10.1 CU ? pending wait CAM

□ [[DMA]] §10.1□

10.2 CU ? ctrl_m

□ [[DMA]] §10.2□

10.3 CU ? NoC????????

- LOAD□ CU □ noc_rd_req □□□□□□ L1 / DRAM□ master□□□□□ §7.7 □□□ resp □□□ req_id □□□ CU □ L0 □□
- STORE□ CU □ noc_wr_req + wr_data □□□□□□□□ ack□ fire-and-forget□□ NOTIFY □□ COUNT_DELTA □ CU □□ FSM □ STORE □□□□□□□□□□□□□□□□□□□□□□ NoC □ id □□□□□□□□

10.4 ? cmd_desc ? in ? out ???? pipe

□□ [[Pipe]] §289□□□□□□□□□□□□□

1. WAIT_IN □□□□□ pipe □ count ≥ in □□
2. LOAD □□□□□□□□□□□□□
3. COMPUTE
4. STORE □□□□□□□
5. NOTIFY □□□□□ pipe □□ `COUNT_DELTA(-in_delta)` □□□ `COUNT_DELTA(+out_delta)`

□□□□□□□□□□□□□□□ In-place □□□□□□□□□ NoC □ id □□□□□□

11. ?????

□□ / □□	□□□
□ [[DMA]] §11 □□□	
L0 vreg / treg	□□□ □ undefined□

□□ / □□	□□□
□□□□ pipeline □□□	□□
MAC □□□	0

L0 □□□□□□□□ LOAD □ COMPUTE □□□ bug□□□□□□□□□□□□□□

12. ??????

□ [[DMA]] §12□□□

□□	□□	□□
`dbg_cu_state[N_CU]`	N_CU × 3	□ CU FSM □□
`dbg_mac_active[N_CU]`	N_CU	MAC □□□□
`dbg_inner_loop_idx[N_CU]`	N_CU × LOOP_W	□□ micro-op □□
`dbg_l0_read_hazard[N_CU]`	N_CU	LOAD/COMPUTE □ reg □□□□□□□ 0□
`dbg_pipeline_stall_cycles[N_CU]`	N_CU × 16	NoC backpressure □□ stall

13. ??????

1. `N_CU` □□□□□□□ 1 □ 2D MAC □□ + 2 □□□□□ + 1 □□□□□ vs □□□
2. `VEC_LANE / MAC_M / MAC_K` □□□□ —— □□□□□□
3. `dtype` □□□□□ INT4 / FP8 / FP16 / BF16□
4. MAC □□□□□ systolic 1D vs 2D□ weight-stationary vs output-stationary□
5. L0 □□□□□□□□□□□ FMA □ 3R+1W □□□□ 2R+1W□
6. inner-loop □□□□□ 3 □ vs 4 □□□□□□ stall □□□
7. □□ round / saturate □□□□□□
8. □□□□□□□□ LUT □□□□ per-CU □□

14. ??? Pipe ??

□□□□□	□□□□□
engine pipe □□	§7.1 □□ + §7.2 count + §7.3 FSM
sync pipe □□	§8
□□ A	§6 + §8 □□□□ bitmap
□□ B	§3 ctrl vs noc □□□□
□□ D	§8 sync □□
□□ 3□□ pipe □□□□	§7.3 FSM □ channel □□□□
□□ 6□□□□□□□	§5 / §6 / §7.4 L0 □□□□□□□
□□ 7□□□□□□□□□	§6 alloc □ cu_class / queue_depth □□

□□□□□	□□□□□
□□ 8□□□ pipe□	engine □ §7□ sync □ §8
§289□ in/out □ pipe□	§10.4 NOTIFY □□

DynamicGraphMultiProcessor??

Dynamic Graph Multi Processor ??

??

1. Etched□□□ GPU□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□ TFLOPS□□□□□
2. □□□□□ +□□□□□□□□□ +□□□□□□□□□□□□□□□□
3. Etched□□□□□ H100□ 800□□□□□□□□□ 3.3% □□□□□□□□□□□□□□□□□□□□□□□□□ Tranformer□ Sohu□□ FLOPS □□□□□□□ 90%□ GPU□□□ 30%□
4. □□ Scaling Law□□□□□□□□□ Scale Up□□□□□□□□□□□□□□□□□□□□ Scaling Law□□□□□□□□□□ Scale Down□□□□□□□□□□□□□□□□□

????

- □□ AI□□□□□□□□□□
- □□□□□□□□□□□□□□ Int8□□□□□□□□□□□□ pattern□□□□□□□□
- □□□□□□□□□□□□□□□□□□□□□□□□□
- □□□□□□□□□□□□□□□ 2D□□□ => □□□□□□□□
- GPU□□□□□□□□□□ DRAM□□□□□□□□□
- □□□□□□□□□□ NoC□ Cache□□□□□□□□□□□□□□□□□□□□□□□ Fork/Join □□□□□□□□□□□□□□□□□□□□□□□□□ NoC□ flow control□□□
- □□□□□□□□□□□□□□□□□ NoC/Cache□□□□□□□□□□
- □□□□□□□□□□□□ consistency corhenrency□ □□□□ fence+sync □□□□□□□□□□□□□□
  1. GMP□□□□□□□□□□ fork/join□□□□□□□□
  2. □□□□□□□□□□□□□□□□□□ fence□□□
  3. □□□□□□□□□□□ fence□□
- □□□□□□□□□□□□□□□□□□□□□□ IP□□□□□□□□□□□□□□□ cpu □□□□□□□□□□□□□□□□□□□□□□ IP□□□□□□□ IP□□□□□□□ cache □□□□□□□□ cache_hint□□□□
- □□□□□□□□□□□□□□□□□
  - □□□□□□□□□□□ DSA□
  - □□ SOC□□□□□□□
  - □ NOC□□□□□□□□□□□□
- □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□
  - □□□□□□□□□□□□□□□□□□□□□□□□□□□□ x□□□□

    - □□□□□□□□□
- DSA□□□□ GPU□□□□

??/??

1. □□□□□□□□
    1. □□□□ graph□□□□□□□□ DRAM□□□□
    2. □□ LLC□□ DRAM□□□□□ LLC□□□□□ L3->LLC->DMA->L1->L0->MALU
2. □□□
    1. □□ fork/join□□□□□□□□□□□□□□□□□□□□ L1/L0□□□□ DMA□□□□□□
    2. □□□□□□□□□□□□□□□□□□
3. □□
    1. Root □□ & □□□□ & □□□□□
    2. □□□□ NoC□ Cache
    3. Fork Join
    4. □□□□ CPU□□□□□□□ RISCV□□□□□□□□□□□□□□□□□□□□ IP □□□□□□□□□□ ForkJoin□
4. □□
    1. NPU□□□□ NPU Kernel□ 1D 2D□□□□□
    2. DSA Confige
    3. Atomic Reduce□□□□□□□□□□□□□□□□□□
    4. LD/ST+Relu□□□□

??

1. NPU□□□□□□□□□□□ LD/ST□□□□□□□□□□□ L2/L1 <-> L2/L1
2. □□ MLI□□□□□□□□ HLI□□□□□□□□□
3. □□□□□□□ □□□□□□□□□□□□
4. □□□□□□□□□□□□□ MLI□□□
5. □□□□□ Leading/Tainling □□□ scheduler□□□□□□□□□□□□□□□□□□□□□ pipeline□ NPU□□□□□□□
6. □□□□□□□□□□□□□□□□□□ debug□□□□□□□□□□□ /□□ /□□□□

???ASIC??

1. MLI□□□□□□□□□□□□□□□□□□□□□□□
2. scheduler□□ NPU□□□□□□
3. □□ latency
4. □□□□□ HLI

??

1. □ GM□□□□□□□□
    1. scheduler□□□□□□□□□□□□
2. □□□
    1. □□□□□□□□□□□□ <8bit□ DMA□□□□□□□ 8bit□□□□□□□□□□□ bit □□□

3. Launch/Sync/DataFlow
    1. [illegible] Launch [illegible] Sync [illegible] Fork/Join [illegible]
        1. [illegible] **Fork** [illegible] **Join** [illegible] **A fork B** [illegible] **C join B**
        2. fork [illegible] launch [illegible] join [illegible] wait
        3. [illegible] fork [illegible] join [illegible]
        4. [illegible] fork [illegible]
    2. [illegible] Launch [illegible] Sync
    3. [illegible] NoC [illegible] Launch [illegible]
    4. [illegible] Graph Kernel
    5. [illegible] **latency** [illegible]
4. [illegible] launch [illegible]
    1. [illegible] -> [illegible] -> [illegible] in flight [illegible] pipeline [illegible]
    2. [illegible]
    3. [illegible]
    4. [illegible]
    5. [illegible] SRAM [illegible]
        1. [illegible] SRAM [illegible] L1 [illegible] L1 [illegible] L2 [illegible] LLC
        2. [illegible] L1 [illegible] cache [illegible] GEMM [illegible]
        3. [illegible]
        4. [illegible] L1 [illegible] bank [illegible] thread [illegible] Engine [illegible] ALU [illegible] RegRead [illegible] MALU [illegible] ld Cal St [illegible]
    6. [illegible] graph [illegible] launch [illegible]
    7. [illegible] DSA [illegible] PPA [illegible] gpu [illegible] gs [illegible] simt [illegible] simi [illegible] data [illegible] structure [illegible] hazard [illegible]
    8. [illegible] DSA [illegible] fence [illegible] TensorALU [illegible] LD [illegible] ST [illegible] latency [illegible] standby [illegible]

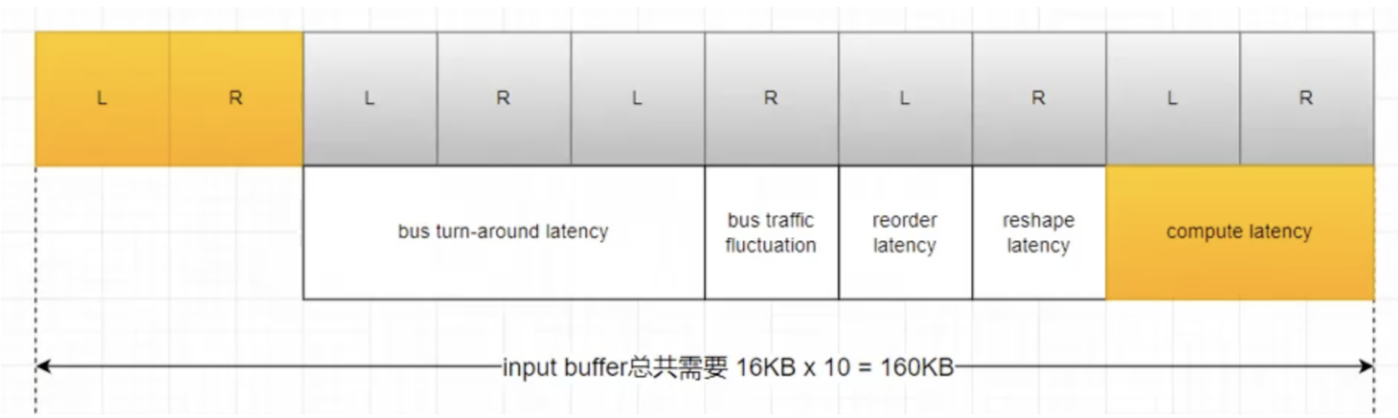

??

1. [illegible]
    1. [illegible]
    2. [illegible] pipeline[illegible]
    3. [illegible] L2L1[illegible]
2. [illegible]
    1. [illegible] graph[illegible]
    2. [illegible] DMA[illegible] /[illegible]
    3. DMA[illegible] NPU[illegible] Scheduler[illegible] Inst SRAM[illegible]
    4. [illegible] scheduler[illegible] fork[illegible] join [illegible]
    5. [illegible] fork[illegible] join [illegible]
3. VMM[illegible] VLD VST[illegible] pipeline
4. kernel[illegible]
5. [illegible] IP[illegible]
    1. HWSync[illegible] RISCV CPU[illegible]
    2. [illegible] IP[illegible] mailbox[illegible] trigger[illegible] mailbox[illegible] pipeline [illegible] forkjoin[illegible]
6. [illegible] fence[illegible] leading tailing[illegible]
    1. [illegible] NPU[illegible] leading[illegible] tailing[illegible] scheduler [illegible]

Graph Multiprocessor = GM

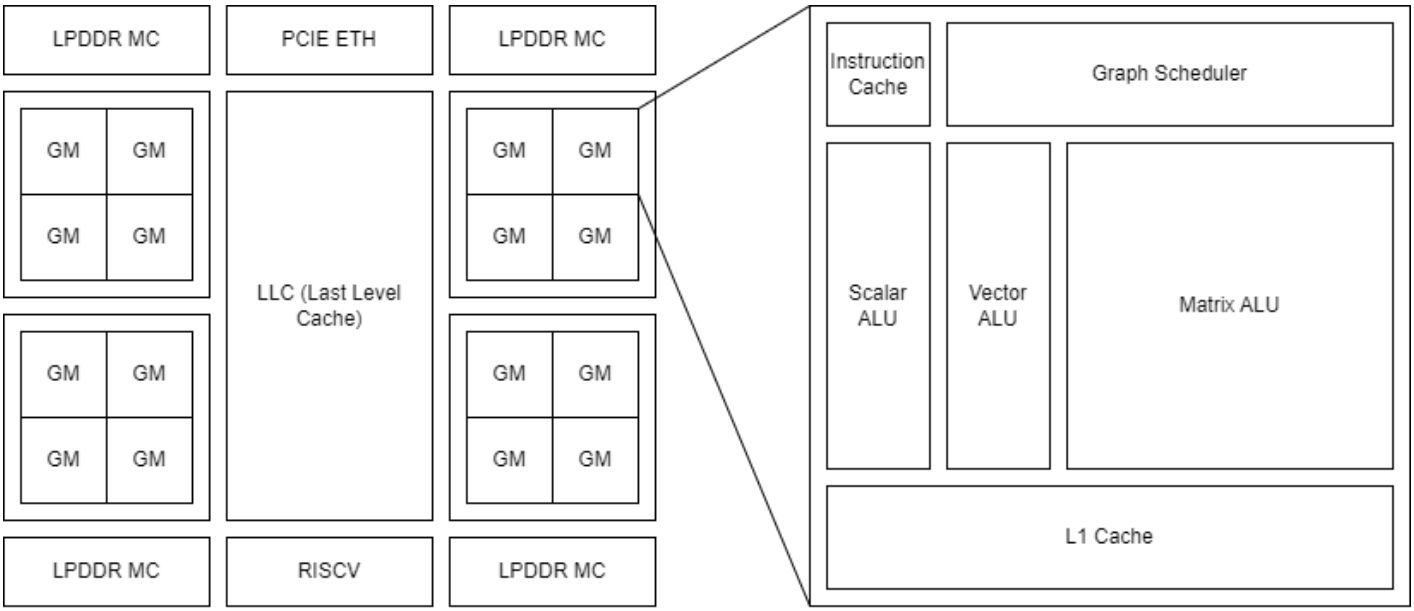

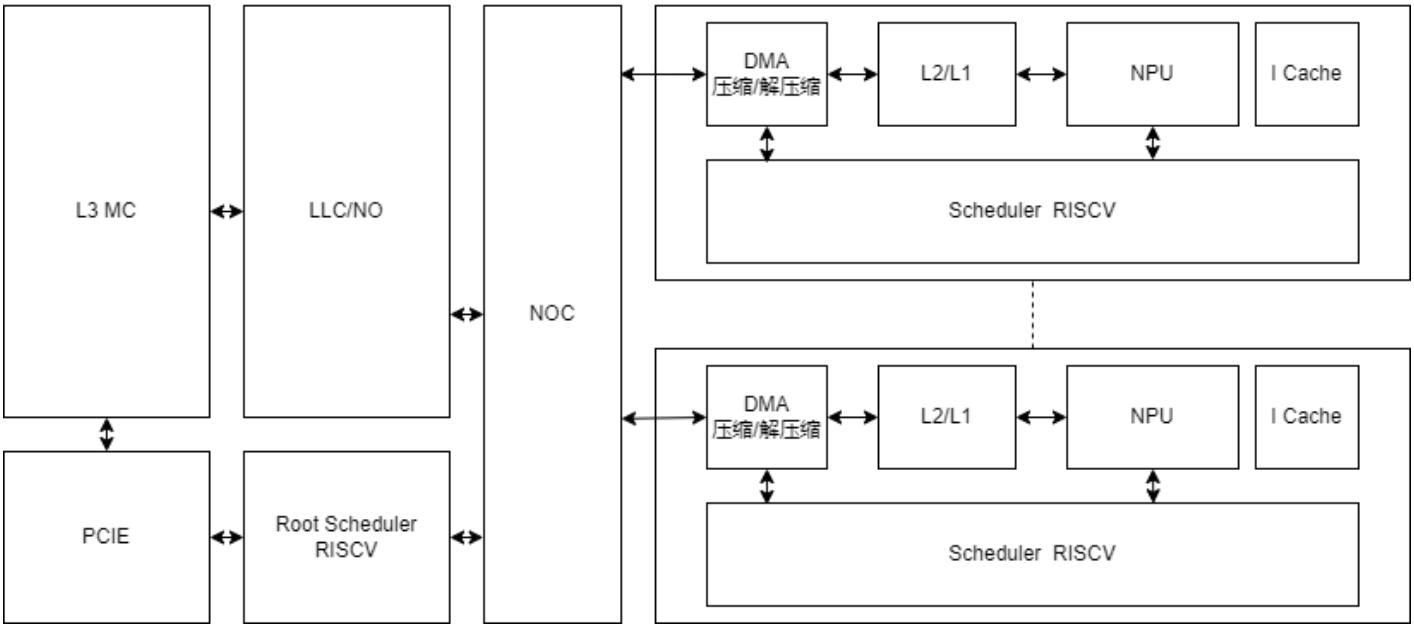

1. 同一片代码多个并行传递参数问题
    1. fork支持参数的计算，批量 fork自动累加传递的参数，类似于 for功能的批量 launch，
    2. 一个 kernel最小的 fork出来的执行单位表达类似于一个 gemm指令
2. 架构
    1. 32bit指令宽度
    2. 标量 +向量 +2D混合指令

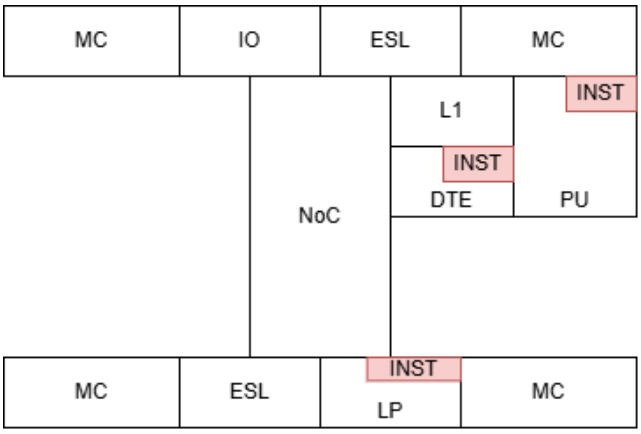

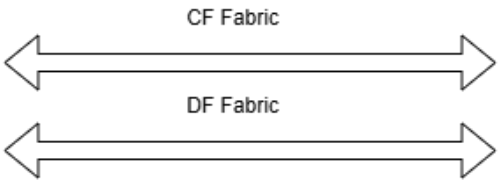

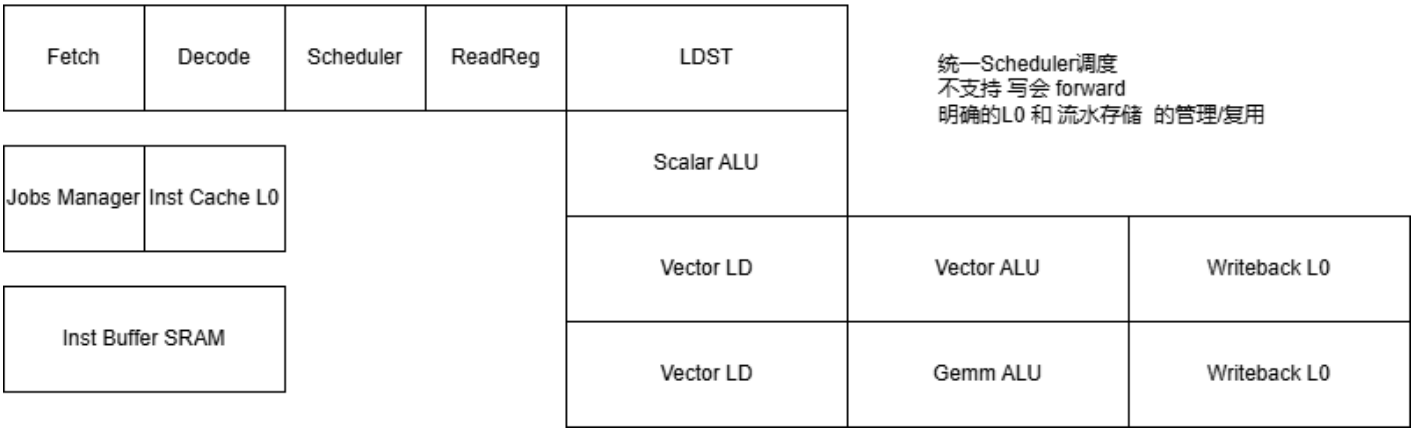

????

1. □□□□□□□□□□□□□□□□□□□□
   1. □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□
   2. □□□ launch/fork□ join
2. □□□□□
   1. SIMT□□□□□□□□ thread□□□□□□□□□□□□□□□□□□□□ thread
   2. **GMP**□□□ **MGMT**□ **Multi Graph Multi Thread**□□□□□□□□□ **Graph**□ **Graph** □□□□□□□□□□□□□□□□□□ **Launch**□□□□□□□□□□□□□□□□
   3. □□□ fork/join□□□□□□□ graph□ launch□□□
   4. □□□□□□□□□□□□□□□□ free□□ run□□□□□□□□□□□□□□□ graph□ processor□□□□□□□□ graph□□□□□□□□□□□□□□□□□□□□
   5. □□ fork□□□ subgraph□□□□□□□□□□□□□□□□ PU□□□ graph□ fork□□□ subgraph□□□ hw thread□□□
   6. □□□□□□□□□□□□ graph□□□□□
   7. fork□ join□□ “□□□□ ”□ fork□□□□□ graph□□□□□□ graph□□□□□□□ graph □□□□□□□□□□□□□□□□□□□□□□ join□□□□□□□□ graph□□□□ join□□□ barrier□
   8. **counter** : wait □□□□□□□□□□□□ counter□□ signal□□□□□ counter□ □□□□□□□□□□□□□□□□
      1. □□□ pull□□□□□□□□ wait□□□□□□ wait signal

    2. [illegible] launch[illegible]
    3. [illegible] launch[illegible] launch[illegible]
    4. [illegible] launch[illegible] sub graph[illegible] graph group[illegible]
       “signal[illegible] ”[illegible]
  9. [illegible] delay[illegible]
    1. [illegible] data hzd
10. [illegible] stall[illegible] [illegible] [illegible] stall[illegible]

??

1. Fork Join
  1. fork[illegible]
    1. [illegible]
    2. [illegible] join[illegible]
    3. [illegible] index
    4. [illegible] launch[illegible] ID[illegible] ID[illegible] IP[illegible] memory mapping
  2. fork[illegible] load[illegible]
  3. join[illegible]
    1. launch[illegible] ID
2. [illegible] VR L0.5 [illegible] L0.5[illegible]
3. [illegible] VR L0[illegible] 4[illegible]
4. [illegible] TR L0.5 , [illegible] L0.5[illegible]
5. DTE[illegible] transpose pad slice deslice
6. PU[illegible]
7. [illegible]
  1.

Width (bits)	4	6	3	3	1	
Meaning	Reuse flags	Wait barrier mask	Read barrier index	Write barrier index	Yield flag	St cycl

  2. Reuse flags [illegible] 4[illegible] register[illegible] cache
  3. 6[illegible] barrier[illegible] thread[illegible] cuda
     [illegible] 6[illegible]
8. async_group [illegible]
  1. async copy bulk [illegible]
9. mbarrier[illegible] fence
   [illegible]
10. [illegible] credit[illegible] valid/ready
   [illegible]
  1. [illegible]
  2. [illegible]
  3. [illegible]

????

1. □□□□　　die□□
    1. 4T int8□□
    2. □□　256MB DDR3 DRAM 1866 16bit □ 3.7GB/s□□
    3. □□□□□□□□□□□□□□□□□□□□□□□□□□□□
2. □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□
    1. □□□□□□□□□　　die□□□
    2. □□□□□□□□□□□□□□□□□□□□
    3. die to die □□□□□□□□□□□□□□□□□　　serdes□□

ISA

ctrl

1. 1. Fork(code_index)
    2. launch
    3. Join()

Scalar

1. scalar □□

Vector

Tensor

???????

□□□ `design/` □□□□□□ `.md` □□□□□□□□ [[LLC]] + `LLC.pipeline.html` □

0. ??

- □□□□□□ □□□ **RTL / [[logix]]** □□□□□□□ □□□ `.md` □□□□ RTL□□□□□□
- □□□□□□□ `.md` □□ `L1.md` □ `DMA.md` □ `LLC.md` □□□□□□□□□□
- □□□□ □□ □□□□□□□□□□□□□□□□
- □□□□□□□□□ HTML □□□□□□□□ `<unit>.pipeline.html` □□□□□ git LFS □□□□□□□□

1. ?????

□ `.md` □□□□□□□□□□□□□□ `#### N. □□` □□

□□	□□	□□
§1	□□□□	□□ / □□ / □□□□□□□
§2	□□□□	ASCII □□□ mermaid □□□□□□□□□□□□□□
§3	□□□□	□□ / □□ / □□□□□□□□□□□□□□□□□□□□ □
§4	□□□□	□□□□□□□□□□□□
§5	□□□□□□□□	opcode □□□□□□□□□□□□□□□□□□□□ □
§6 - §M	□□□□□□□□	□□□□□□□□□□□□□□□□□□ / □□□□□□□□□□
§M+1	□□□	□□□□□□□□ **§2**
§M+2 -	Pin / Lock / Bypass □□□□□□□□□□□□□□□□□□ □□□□□□□□□ DGMP □□	

□□□□□ H4 `#### N.` □□□□ H5 `##### N.M` □□□□□ H6 `###### N.M.K` □□□□□□

2. ?????

□□□□□□□□□□ "□□□ "□□□□□□□□□□ N □□□□□□□□□□□□ LLC.md □ §11□□

2.1 ??

- §a-§b
- SRAM 1RW vs 1R1W
- `<unit>.pipeline.html`
- Mermaid / WaveDrom / YAML

2.2 §N.1 ????

`<name>`	`<arb / sram_port / ff_rw / cam / table / comb>`	§x.y

—— SRAM FF

2.3 §N.2 ????

mermaid `flowchart LR` stage

2.4 §N.3 ~ §N.X ????

H6 ###### §N.X.Y

1. mermaid `flowchart LR` stage +
2. WaveDrom JSON cycle stage / latch / `edge`
3. **latch** stage→stage latch +
4. ×

2.5 §N.{?? 2} ??????????

- × / /
- **stall** stall
- **bank**

2.6 §N.{?? 1} ??? schema

YAML logix / RTL schema

```
unit: <name>
clock: <main_clk>

resources:
```

```
  <name>: { kind: ..., ports/slots/width/banks: ... }

paths:
  <path_name>:
    stages:
      - { id: <S0/S1/...>, in: [...], uses: [...], out: [...] }

arbitration:
  - { resource: <res>, contenders: [...], rule: "<text>" }

stall_conditions:
  - { src: <res.cond>, effect: "<text>" }
```

□□□□□□□□□□□□□□□□□□□□□□ `<unit>.pipeline.yaml` □□□

3. ??????

□□	□□	□
□□□	S0 / S1 / S2 / S3 / S4	□□□
□□□□□	S3' / S4'□□□□	□□□□ S2 □□□ mshr
□□□□□	F0 / F1 / F2	□□□ fill□
□□ FSM	W0 / W1 / W2 / W3	□□□□ walker
□□□□□	B0 / B1	□□ bypass
□□	S2'□□□□□□□□□□ + □□□	□□□ master □□□ S2

□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□

4. ????????

- □□□□ □ `§N` □ `§N.M` □ `§N.M.K`
- □□□ □ `[[□□□]] §N` □ Obsidian wiki link □□□□□□□□□ `.md` □□
- □□□□□□□□□□□□□□□ "□□□□ [[L1]] §3.2 □ NoC slave □□□ "□□□□□□□□□

5. ???????

□□	□□	□□
Mermaid `flowchart`	□□□□□□□ + □□□	□□ GitHub / Obsidian / VS Code □□□□
WaveDrom JSON	□□□□□□□ + □□□□ + □□□□	□□ wavedrom.com/editor.html□ Obsidian + WaveDrom □□□

□□	□□	□□
□□□ YAML	□□ schema□ resources / paths / stages / arbitration□	□□

□□ □□□□ □□□□□□□□□□□□□□□□□□□□ YAML □□□□ ——Mermaid / WaveDrom □□□□□□□□□□□□□□□□□□□□□□□□□□□□

□□□□□ drawio□□□ diff□□□□□□□□□□□□ PNG/SVG □□□□□ + □□□□□□□□ Word/PPT□□□□□

6. ??? HTML???????

□□□□□□ `<unit>.pipeline.html`□□□ `.md` □□□

□□□□□□□□□□ CSS / JS □□□□□□□□□□□□□□□□□□□□ `<unit>.pipeline.yaml` □□□□□□□□□□□□ `<script>` □ JS □□□□□ schema □□□□□□

□□□□□□□ `LLC.pipeline.html` □□□□

1. □□□□□□ □□□ cycle□□□□□□ stage □□□ cycle □□
2. □□□□□□ □□□ cycle□□□□□□□□□ / □□□□□□□
3. □□□□□□ □□□ cycle□□□□□□□□□□□□□□□□
4. □□□□□ □□□□□□□ stage □□□□□□□□□□ / □□ / □□ / □□ latch

□ 3 / □ 4 □□□□□□□□□□□□□□ stage / □□□□□□□□□□□□□□

7. ????

1. □□□□ □□□ `.md`□□□ RTL□ `.md` □□□□□□ RTL □□□□□□□□
2. □□□□ □□□□□□□□□ "□□□ "□□□□□□□□□□□□□□ "□□□ "□□□□ "□□□□□ "□□
3. □□□□□□ □ 1RW vs 1R1W□ SRAM □□□□□□□□□□□□□□□□□□□□□□□□
4. □□□ □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□
5. □□ **vs** □□ □□□□□□□□□□□□□□□□□□□□□□□□ stage □□ `comb` □□□
6. □□ **LLC** □□□□ □ cacheline / set / way / pinned way / bypass / fill / writeback / mshr / walker□□□□□□□□□□□□□□

8. ?????????

□□□□□□□□□□□□□

1. □□□□ □□□□□□□□□□□□□□□□ ECC / prefetcher □□□□□□
2. □□□□□ □□□□□□□□ retire □□□ FSM □□□□□□□□ floorplan / □□□□□

□□□ "□□□□□ "□□□□□□□□□ ——□□□□□□□□□□□□□□

9. ?????

□□	□□
□□□□□	`design/LLC.md`
□□□□□□□	`design/LLC.md` §11
□□□ HTML □□	`design/LLC.pipeline.html`
□□□□□□□□□□	`design/Pipe.md` □ `design/□□_□□□.md`

□□□□□□□□□ `LLC.md` □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□ §11 □□□□□□

10. ????

□□□□□□□

- □ §1 - §M □□□□□□□□□□□□□□□□□□□□□
- □ □□□□□□□□□□□□□□□ / □□□□ / □□□□□□□ / □□□□□□ / □□□ YAML
- □ □□□□□□□□
- □ □□□□□□ §3
- □ □□□□□□ `[[...]] §N` □□
- □ □□□□□□□□□
- □ □□□□□□□□□□□□□□□□□□□□□□□□
- □ □□□□□ HTML□□□□□□□□□□
- □ □□□□□□□□□□□□□□ LLC.md / Pipe.md / GMP.md□