Host: DESKTOP-... | Launch: ncu.exe -o ... | Application: ...\Release\kernel.exe | Target Process: 1 | Kernel Selection: Regex: ... | Profiling Mode: Default

Kernel: void ... (1024, 1, 1)x(1024, 1, 1) | Context: 1 | Stream: 7 | Device: NVIDIA GeForce RTX 4090 | CC: 8.9 | Domain: Range: 0 | Process: 1 | Replay Mode: Kernel | Pass Count: 1 | Max Passes: 1 | Max Duration: 0 | Range Mode: Auto | Cache Control: All | Clock Control: None | Apply Rules: Yes | Disable Extra Passes: No | Kernel Profiling: Full | Metric Sets: Default | Apply Rules: Yes | Disable Extra Passes: No | Kernel Profiling: Full | Metric Sets: Default

Time: 07.37 ms | Cycles: 10 430 597 | Regs: 32 | GPU: first 0 GPU (AD) | SM Frequency: 1.41 cycle/nsec | Process: 1 | Process: 1 | Process: 1 | Process: 1

The Speed Of Light breaks several Rules.

GPU Speed Of Light [1]

The high-level overview of the utilization for compute and memory resources of the GPU. For each unit, the Speed Of Light (SOL) reports the achieved percentage of utilization with respect to the theoretical maximum. High level overview of the utilization for compute and memory resources of the GPU presented as a roofline chart.

SM [%]	1.17	Duration [ms]	7.37
Memory [%]	1.17	L1/TEX Cache Throughput [%]	1.17
DRAM Throughput [%]	0.00	L2 Cache Throughput [%]	0.00
Max Bandwidth [%]	1.17	Mem Busy [%]	0.00
Compute (SM) Throughput [%]	1.17	Limiter	Compute (SM)

GPU Utilization

Speed Of Light [%]

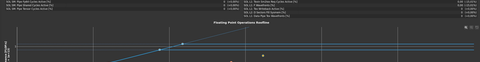

SOL SM Breakdown

SOL SM: Pipe Fma Cycles Active [%]	76.01
SOL SM: Pipe Tensor Cycles Active [%]	0.00
SOL SM: Pipe Fp64 Cycles Active [%]	0.00
SOL SM: Pipe Alu Cycles Active [%]	1.17
SOL SM: Pipe Fp16 Cycles Active [%]	0.00
SOL SM: Pipe Shared Cycles Active [%]	0.00
SOL SM: Pipe Xu Cycles Active [%]	0.00
SOL SM: Pipe Lsu Cycles Active [%]	0.00
SOL SM: Pipe Tex Cycles Active [%]	0.00
SOL SM: Pipe Uniform Cycles Active [%]	0.00
SOL SM: Pipe Vector Cycles Active [%]	0.00
SOL SM: Pipe Tensor Cycles Active [%]	0.00
SOL SM: Pipe Fp64 Cycles Active [%]	0.00
SOL SM: Pipe Fma Cycles Active [%]	76.01
SOL SM: Pipe Alu Cycles Active [%]	1.17
SOL SM: Pipe Fp16 Cycles Active [%]	0.00

SOL Memory Breakdown

SOL L1: Data Pipe Lsu Wavefronts [%]	34.01
SOL L1: Data Pipe Lsu Wavefronts [%]	34.01
SOL L1: Lsu Writeback Active [%]	34.01
SOL L1: Tex Writeback Active [%]	0.00
SOL L1: Tex Data Bank Conflicts [%]	0.00
SOL DRAM: Dram Read Cycles [%]	0.00
SOL DRAM: Dram Write Cycles [%]	0.00
SOL L1: L1 Global Load Hit [%]	0.00
SOL L1: L1 Global Load Miss [%]	0.00
SOL L1: Data Pipe Lsu Active [%]	0.00
SOL L1: L1 Global Store Hit [%]	0.00
SOL L1: L1 Global Store Miss [%]	0.00
SOL L1: L1 Atomic Hit [%]	0.00
SOL L1: L1 Atomic Miss [%]	0.00
SOL L1: L1 Local Load Hit [%]	0.00
SOL L1: L1 Local Load Miss [%]	0.00
SOL L1: L1 Local Store Hit [%]	0.00
SOL L1: L1 Local Store Miss [%]	0.00
SOL L1: Tex Writeback Active [%]	0.00
SOL L1: Data Pipe Tex Wavefronts [%]	0.00
SOL L1: Data Pipe Tex Wavefronts [%]	0.00

Floating Point Operations Roofline

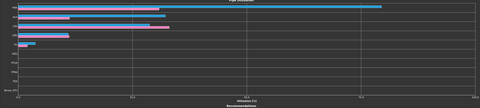

Arithmetic Intensity [FLOP/byte] Recommendations

Softwares | Warning | Compute is more heavily utilized than Memory: look at the [Source] report section to see what the compute pipelines are spending their time doing

Roofline Analysis | The ratio of peak float (fp32) to double (fp64) performance on this device is 1:1. The kernel achieved 0% of this device's fp32 peak performance and 0% of its fp64 peak performance.

Compute Workload Analysis

Detailed analysis of the compute resources of the streaming multiprocessors (SM), including the achieved instructions per clock (IPC) and the utilization of each available pipeline. Pipelines with very high utilization might limit the overall performance.

Executed Ipc Active	0.01	Issued Ipc Active	0.01
Executed Ipc Elapsed	0.01	Issue Slots Busy	0.29 %
Issued Ipc Elapsed	0.01	Issue Slots Busy	0.29 %

Pipe Utilization

Recommendations

High Pipe Utilization | FMA is the highest-utilized pipeline (76.0%). It accepts 32-bit floating-point (FADD, FMUL, FMAD, ...) and integer (IMAD, IMADC) operations. The pipeline is well utilized and might become a bottleneck if more work is added.

Memory Workload Analysis [2]

Detailed analysis of the memory resources of the GPU. Memory can become a limiting factor for the overall kernel performance when fully utilizing the involved hardware units (Mem Busy), exhausting the available communication bandwidth between those units (Max Bandwidth), or by reaching the maximum throughput of issuing memory instructions (Mem Pipes Busy). Detailed chart of the memory units. Detailed tables with data for each memory unit.

Memory Throughput	1.17 %	Mem Busy	0.00 %
Max Bandwidth	1.17 %	L1/TEX Hit Rate	0.00 %
L2 Hit Rate	0.00 %	Mem Pipes Busy	0.00 %

Memory Chart

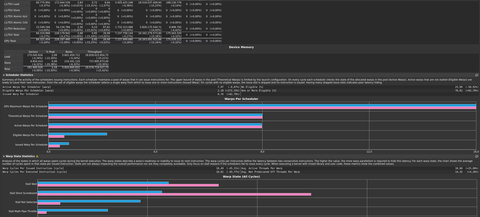

Shared Memory

	Instructions	Requests	Wavefronts	% Peak	Bank Conflicts
Shared Load	4,513,332.00	4,513,332.00	4,513,332.00	0.00	1,440,270.00
Shared Store	139,722	139,722	139,722	0.00	1,536
Shared Atom	1,048,576	1,048,576	1,048,576	0.00	1,048,576
Shared Atom Cas	4,194,304.00	4,194,304.00	4,194,304.00	0.00	4,194,304.00
Other					2,490,878.00
Total	9,895,934.00	9,895,934.00	9,895,934.00	0.00	8,675,284.00

L1/TEX Cache

	Instructions	Requests	Wavefronts	Sectors	Sectors/Req	Hit Rate	Bytes	Lines Misses to L2 Pipe Lsu	Returns to SM	Pipe Lsu
Local Load	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0.00%	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)
Global Load	32,032,000 (+0.00%)	32,032,000 (+0.00%)	40,120,931 (+0.00%)	40,120,931 (+0.00%)	1.25 (+0.00%)	20.11%	64,064,000 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)
Surface Load	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0.00%	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)
Texture Load	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0.00%	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)
Global Store	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0.00%	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)
Local Store	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0.00%	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)
Surface Store	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0.00%	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)
Global Atomic	1,048,576 (+0.00%)	1,048,576 (+0.00%)	1,048,576 (+0.00%)	1,048,576 (+0.00%)	1.00 (+0.00%)	0.00%	4,194,304 (+0.00%)	1,048,576 (+0.00%)	1,048,576 (+0.00%)	1,048,576 (+0.00%)
Global Atomic Cas	4,194,304 (+0.00%)	4,194,304 (+0.00%)	4,194,304 (+0.00%)	4,194,304 (+0.00%)	1.00 (+0.00%)	0.00%	16,777,216 (+0.00%)	4,194,304 (+0.00%)	4,194,304 (+0.00%)	4,194,304 (+0.00%)
Surface Atomic	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0.00%	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)
Reduction	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0.00%	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)
Global Reduction	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0.00%	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)
Surface Reduction	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0.00%	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)
Global Atomic Store	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0.00%	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)
Global Atomic Cas	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0.00%	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)
Surface Atomic Cas	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0.00%	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)
Loads	36,080,576 (+0.00%)	36,080,576 (+0.00%)	44,168,507 (+0.00%)	44,168,507 (+0.00%)	1.22 (+0.00%)	18.26%	68,258,304 (+0.00%)	1,048,576 (+0.00%)	1,048,576 (+0.00%)	1,048,576 (+0.00%)
Stores	4,194,304 (+0.00%)	4,194,304 (+0.00%)	4,194,304 (+0.00%)	4,194,304 (+0.00%)	1.00 (+0.00%)	0.00%	16,777,216 (+0.00%)	4,194,304 (+0.00%)	4,194,304 (+0.00%)	4,194,304 (+0.00%)
Total	40,274,880 (+0.00%)	40,274,880 (+0.00%)	48,362,811 (+0.00%)	48,362,811 (+0.00%)	1.20 (+0.00%)	16.74%	85,035,520 (+0.00%)	5,242,880 (+0.00%)	5,242,880 (+0.00%)	5,242,880 (+0.00%)

L2 Cache

	Requests	Sectors	Sectors/Req	Hit Rate	Bytes	Throughput	Lts Returns to Sm	Returns to Sm Hit	Returns to Sm Miss
L2 Cmp Load	5,242,880 (+0.00%)	5,242,880 (+0.00%)	1.00 (+0.00%)	0.00%	20,971,520 (+0.00%)	2.71 GB/s	5,242,880 (+0.00%)	0 (+0.00%)	5,242,880 (+0.00%)
L2 Rop Atomic Load	1,048,576 (+0.00%)	1,048,576 (+0.00%)	1.00 (+0.00%)	0.00%	4,194,304 (+0.00%)	542.29 MB/s	1,048,576 (+0.00%)	0 (+0.00%)	1,048,576 (+0.00%)
L2 Rop Atomic Cas	4,194,304 (+0.00%)	4,194,304 (+0.00%)	1.00 (+0.00%)	0.00%	16,777,216 (+0.00%)	2.17 GB/s	4,194,304 (+0.00%)	0 (+0.00%)	4,194,304 (+0.00%)
L2 Rop Reduction	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)	0.00%	0 (+0.00%)	0 B/s	0 (+0.00%)	0 (+0.00%)	0 (+0.00%)
L2 Rop Store	4,194,304 (+0.00%)	4,194,304 (+0.00%)	1.00 (+0.00%)	0.00%	16,777,216 (+0.00%)	2.17 GB/s	4,194,304 (+0.00%)	0 (+0.00%)	4,194,304 (+0.00%)
Sys Total	14,680,064 (+0.00%)	14,680,064 (+0.00%)	1.00 (+0.00%)	0.00%	58,720,256 (+0.00%)	7.59 GB/s	14,680,064 (+0.00%)	0 (+0.00%)	14,680,064 (+0.00%)

Device Memory

	Device	% Peak	Throughput	% Peak
Load	1,048,576 (+0.00%)	1,048,576 (+0.00%)	542.29 MB/s	0.00%
Store	4,194,304 (+0.00%)	4,194,304 (+0.00%)	2.17 GB/s	0.00%
Total	5,242,880 (+0.00%)	5,242,880 (+0.00%)	2.71 GB/s	0.00%

Scheduler Statistics

Summary of the activity of the schedulers issuing instructions. Each scheduler can issue one instruction every cycle. The upper bound of issued instructions per cycle (IPC) is 1. Every cycle in which a scheduler cannot issue an instruction is counted. The higher this value, the more work is available per scheduler to hide latencies between issued instructions. Note that all available warps might be waiting for memory operations to complete, which can also completely prevent issuing. For more details on warp states, see the Warp State Statistics section.

One or More Eligible	0.01	Issued Warp Per Scheduler	0.01
No Eligible	99.99	Active Warps Per Scheduler	1.96
Eligible Warps Per Scheduler	0.01		

Warps Per Scheduler

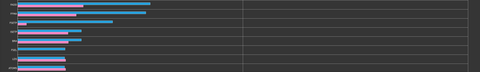

Warp State Statistics [3]

Analysis of the states in which all warps spent cycles during their kernel execution. The warp states describe a warp's readiness or inability to issue its next instruction. The sampling interval is identical to the one used by the Scheduler Statistics section. On average, each of those cycles is divided into issuing an instruction, stalling, or not having an eligible warp to issue it. For each warp state, the chart shows the average number of cycles spent in that state per issued instruction. Stalled warps are not eligible to issue their next instruction. Each cycle with no eligible warp is counted in the Scheduler Statistics section. To see what is stalling warps, inspect the Source Counters tables or page. Reducing the number of cycles a warp is stalled or selected may increase the eligible warp count, which can help to hide latencies and thereby improve overall throughput.

Avg Active Threads Per Warp	32	Avg Not Predicated Off Threads Per Warp	31.98
Avg Stalled Barrier Per Instruction Issued	0.00	Avg Stalled Membar Per Instruction Issued	0.00
Avg Stalled Tex Throttle Per Instruction Issued	0.00	Avg Stalled Lg Throttle Per Instruction Issued	0.00
Avg Stalled Mio Throttle Per Instruction Issued	0.00	Avg Stalled Short Scoreboard Per Instruction Issued	0.00
Avg Stalled Wait Per Instruction Issued	0.00	Avg Stalled Sleeping Per Instruction Issued	0.00
Avg Stalled Misc Per Instruction Issued	0.00	Avg Stalled Drain Per Instruction Issued	0.00
Avg Stalled Dispatch Stall Per Instruction Issued	0.00	Avg Stalled Imc Miss Per Instruction Issued	0.00
Avg Stalled Long Scoreboard Per Instruction Issued	0.00	Avg Stalled Math Pipe Throttle Per Instruction Issued	0.00
Avg Stalled Not Selected Per Instruction Issued	0.00	Avg Stalled No Instruction Per Instruction Issued	0.00
Avg Stalled Branch Resolving Per Instruction Issued	0.00	Avg Stalled Tex Per Instruction Issued	0.00
Avg Stalled Barrier Per Warp	0.00	Avg Stalled Membar Per Warp	0.00
Avg Stalled Tex Throttle Per Warp	0.00	Avg Stalled Lg Throttle Per Warp	0.00
Avg Stalled Mio Throttle Per Warp	0.00	Avg Stalled Short Scoreboard Per Warp	0.00
Avg Stalled Wait Per Warp	0.00	Avg Stalled Sleeping Per Warp	0.00
Avg Stalled Misc Per Warp	0.00	Avg Stalled Drain Per Warp	0.00
Avg Stalled Dispatch Stall Per Warp	0.00	Avg Stalled Imc Miss Per Warp	0.00
Avg Stalled Long Scoreboard Per Warp	0.00	Avg Stalled Math Pipe Throttle Per Warp	0.00
Avg Stalled Not Selected Per Warp	0.00	Avg Stalled No Instruction Per Warp	0.00
Avg Stalled Branch Resolving Per Warp	0.00	Avg Stalled Tex Per Warp	0.00

Warp State (All Cycles)

Cycles per Instruction

Recommendations

Thread Divergence | Warning | On average, each executed branch instruction only causes 0 threads to diverge (=execute a subset of threads in the warp) out of the 32 possible threads in a warp. This value is computed as the number of active threads in each executed branch instruction, averaged over all executed branch instructions. Lower values indicate frequent divergence of threads, causing inefficient execution. Consider if algorithms can be changed to avoid divergence, for example by using different input data for neighboring threads. Typical sources for divergence are data-dependent conditional statements.

GPU Stall Barrier | Warp stall analysis for 'Barrier' issues | Apply

GPU Stall Branch Resolving | Warp stall analysis for 'Branch Resolving' issues | Apply

GPU Stall Dispatch Stall | Warp stall analysis for 'Dispatch Stall' issues | Apply

GPU Stall Drain | Warp stall analysis for 'Drain' issues | Apply

GPU Stall IMC Miss | Warp stall analysis for 'Immediate constant cache (IMC) issues | Apply

GPU Stall LG Throttle | Warp stall analysis for 'LG Throttle' issues | Apply

GPU Stall Long Scoreboard | Warp stall analysis for 'Long Scoreboard' issues | Apply

GPU Stall Math Pipe Throttle | Warp stall analysis for 'Math Pipe Throttle' issues | Apply

GPU Stall Membar | Warp stall analysis for 'Membar' issues | Apply

GPU Stall MIO Throttle | Warp stall analysis for 'MIO Throttle' issues | Apply

GPU Stall Misc | Warp stall analysis for 'Misc' issues | Apply

GPU Stall No Instruction | Warp stall analysis for 'No Instruction' issues | Apply

GPU Stall Not Selected | Warp stall analysis for 'Not Selected' issues | Apply

GPU Stall Short Scoreboard | Warp stall analysis for 'Short Scoreboard' issues | Apply

GPU Stall Sleeping | Warp stall analysis for 'Sleeping' issues | Apply

GPU Stall Tex Throttle | Warp stall analysis for 'Tex Throttle' issues | Apply

GPU Stall Wait | Warp stall analysis for 'Wait' issues | Apply

Instruction Statistics

Overview of the executed and issued instructions. Issued instructions are greater-equal than executed instructions as they also include instructions that had to be issued multiple times to execute (e.g. memory operations with bank conflicts). Executed instructions are all instructions that finished execution. Predicated-off instructions are counted as executed, because they require an issue slot, but do not modify any state. Ratio of executed to issued instructions (Issue Efficiency) is 100% if no instruction had to be issued multiple times.

Executed Instructions	9,895,934	Issued Instructions	9,895,934
Issue Efficiency	100.00%		

Executed Instruction Mix

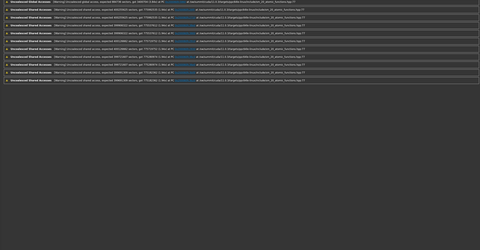

Instruction Statistics

Warp Level Instruction Mix

	Instructions	Thread Instructions	Predicated Off Thread Instructions
Control Flow	1,048,576	33,554,432	0
Integer	1,048,576	33,554,432	0
Floating Point	7,798,784	249,561,088	0
Conversion	0	0	0
Load	0	0	0
Store	0	0	0
Uniform	0	0	0
Inter Thread Communication	0	0	0
Miscellaneous	0	0	0
Total	9,895,934	316,670,208	0

Thread Instruction Mix

Warp Launch Statistics

Block Size	1,024	Grid Size	1,024
Threads	1,048,576	Waves Per SM	0.68
Launch Occupancy Limit	Block Size		

Occupancy | Warning | This launch is theoretical occupancy is limited by the number of required registers

Theoretical Occupancy

Occupancy is the ratio of the number of active warps per multiprocessor to the maximum number of possible active warps. Another way to view occupancy is the percentage of the hardware's ability to process warps that is actively in use. Higher occupancy does not always result in higher performance, however, low occupancy always reduces the ability to hide latencies, resulting in overall performance loss. If occupancy is limited by the number of required registers, consider reducing the number of registers needed by each thread, or by changing the block size to better accommodate the register count.

Block Limit SM	1	Block Limit Registers	1
Block Limit Shared Mem	32	Block Limit Warps	1
Theoretical Active Warps per SM	32	Theoretical Occupancy	100 %
Achieved Occupancy [%]	96.25	Achieved Active Warps Per SM	30.80

Impact of Varying Register Count Per Thread

Impact of Varying Block Size

Impact of Varying Shared Memory Usage Per Block

Occupancy Theoretical Per Block

Recommendations

Source Counters [4]

Branch Instructions Ratio	0.11	Branch Instructions	1,048,576
Branch Efficiency	100.00%	Avg Divergent Branches	0

Sampling Data (All)

Sampling Data (Not Issued)

Most Instructions Executed

	Value
[redacted]	262,144.00
[redacted]	262,144.00
[redacted]	262,144.00
[redacted]	262,144.00
[redacted]	262,144.00

Recommendations

Uncoalesced Global Accesses | Warning | Uncoalesced global access, expected 10149842 sectors, get 40120931 (1.32) at ptr [redacted]

Uncoalesced Global Accesses | Warning | Uncoalesced global access, expected 10149842 sectors, get 40120931 (1.32) at ptr [redacted]

Uncoalesced Global Accesses | Warning | Uncoalesced global access, expected 10149842 sectors, get 40120931 (1.32) at ptr [redacted]

Uncoalesced Global Accesses | Warning | Uncoalesced global access, expected 10149842 sectors, get 40120931 (1.32) at ptr [redacted]

Uncoalesced Global Accesses | Warning | Uncoalesced global access, expected 10149842 sectors, get 40120931 (1.32) at ptr [redacted]

Uncoalesced Global Accesses | Warning | Uncoalesced global access, expected 10149842 sectors, get 40120931 (1.32) at ptr [redacted]

Uncoalesced Global Accesses | Warning | Uncoalesced global access, expected 10149842 sectors, get 40120931 (1.32) at ptr [redacted]

Uncoalesced Global Accesses | Warning | Uncoalesced global access, expected 10149842 sectors, get 40120931 (1.32) at ptr [redacted]

Uncoalesced Global Accesses | Warning | Uncoalesced global access, expected 10149842 sectors, get 40120931 (1.32) at ptr [redacted]

Uncoalesced Global Accesses | Warning | Uncoalesced global access, expected 10149842 sectors, get 40120931 (1.32) at ptr [redacted]

Uncoalesced Global Accesses | Warning | Uncoalesced global access, expected 10149842 sectors, get 40120931 (1.32) at ptr [redacted]

Uncoalesced Global Accesses | Warning | Uncoalesced global access, expected 10149842 sectors, get 40120931 (1.32) at ptr [redacted]

Uncoalesced Global Accesses | Warning | Uncoalesced global access, expected 10149842 sectors, get 40120931 (1.32) at ptr [redacted]

Uncoalesced Global Accesses | Warning | Uncoalesced global access, expected 10149842 sectors, get 40120931 (1.32) at ptr [redacted]

Uncoalesced Global Accesses | Warning | Uncoalesced global access, expected 10149842 sectors, get 40120931 (1.32) at ptr [redacted]

Uncoalesced Global Accesses | Warning | Uncoalesced global access, expected 10149842 sectors, get 40120931 (1.32) at ptr [redacted]

Uncoalesced Global Accesses | Warning | Uncoalesced global access, expected 10149842 sectors, get 40120931 (1.32) at ptr [redacted]

Uncoalesced Global Accesses | Warning | Uncoalesced global access, expected 10149842 sectors, get 40120931 (1.32) at ptr [redacted]

Uncoalesced Global Accesses | Warning | Uncoalesced global access, expected 10149842 sectors, get 40120931 (1.32) at ptr [redacted]