

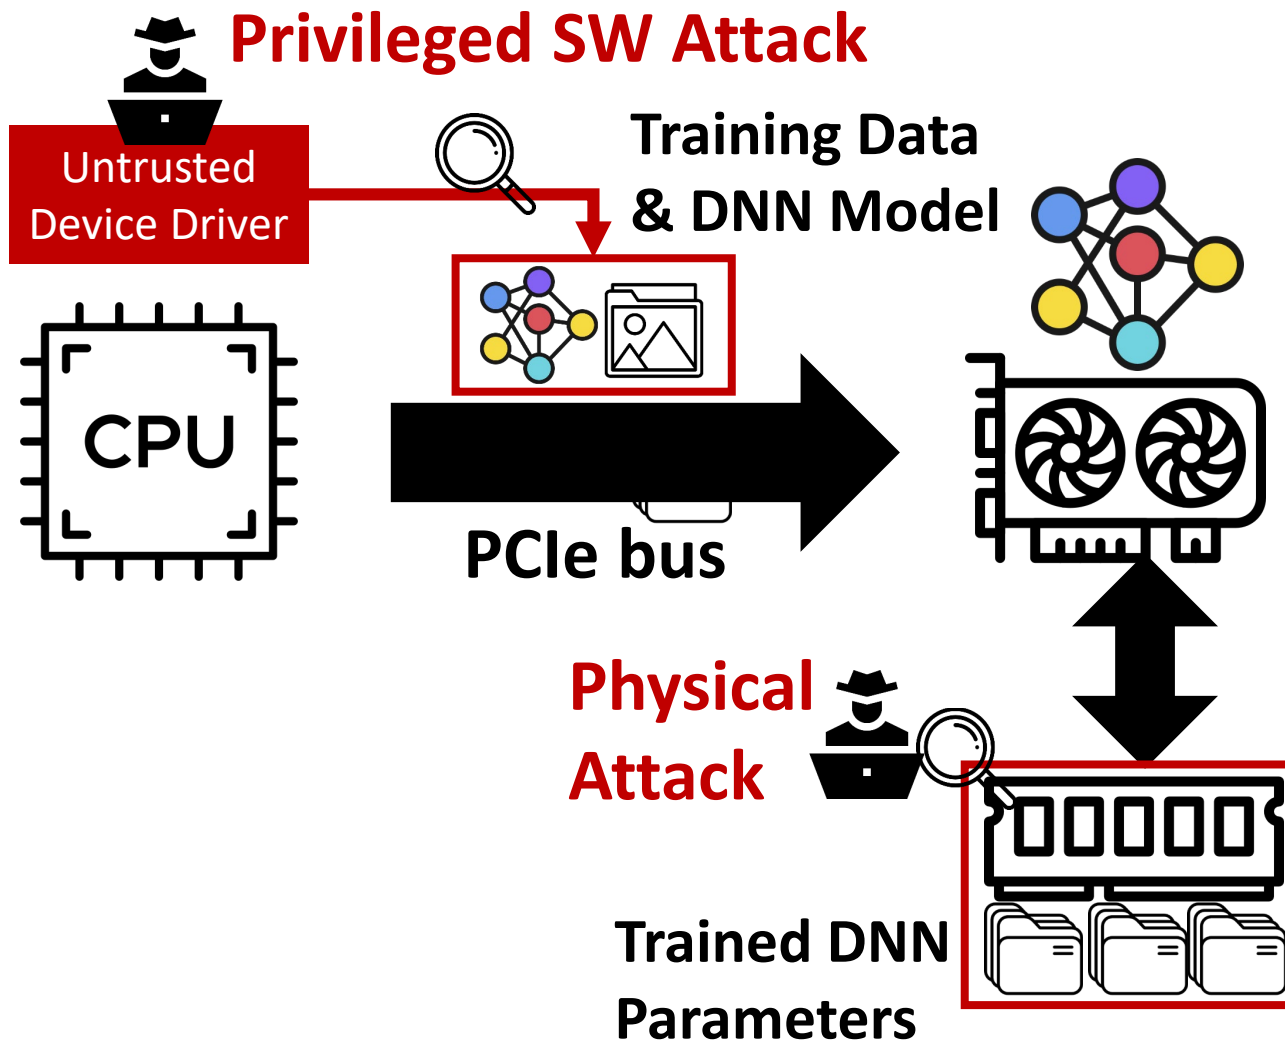
Common Counters: Compressed Encryption Counters for Secure GPU Memory

Seonjin Na, Sunho Lee, Yeonjae Kim, Jongse Park, Jaehyuk Huh

KAIST, School of Computing




Need for Secure GPU Computing



Need for Secure GPU Computing

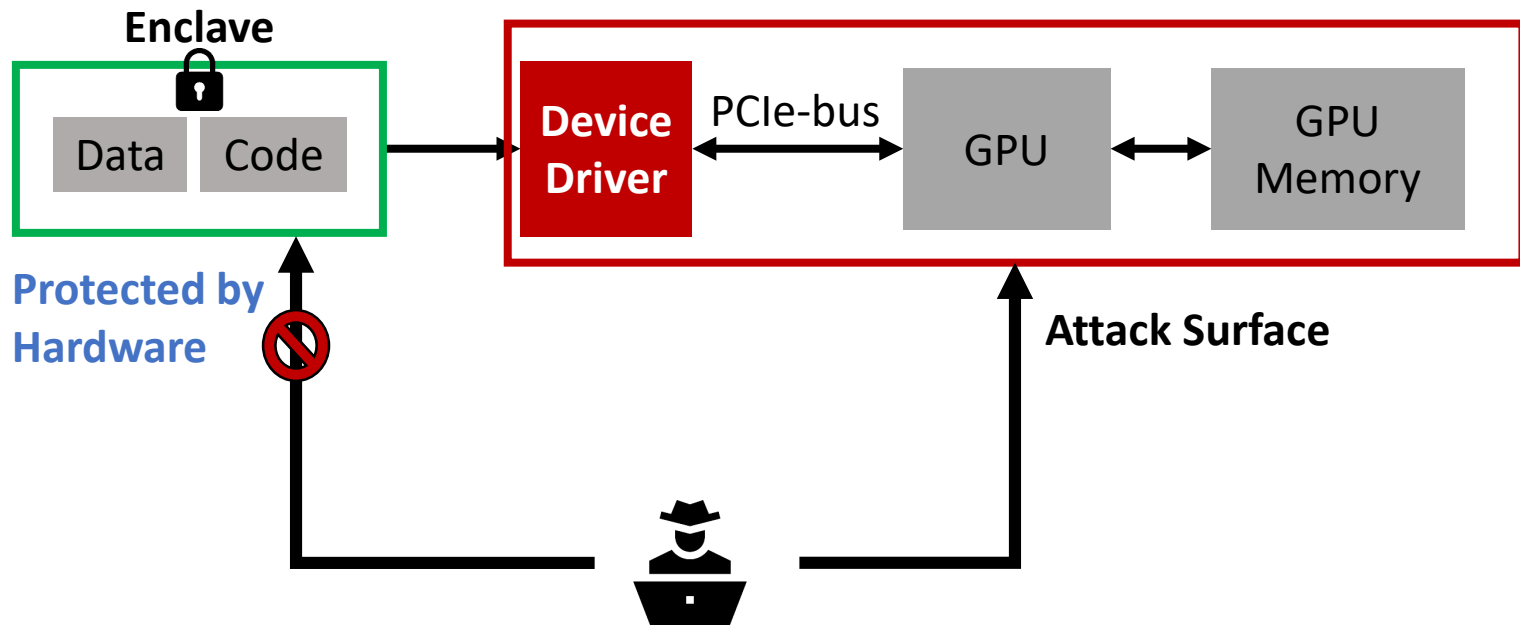


We need to consider Secure GPU computing !



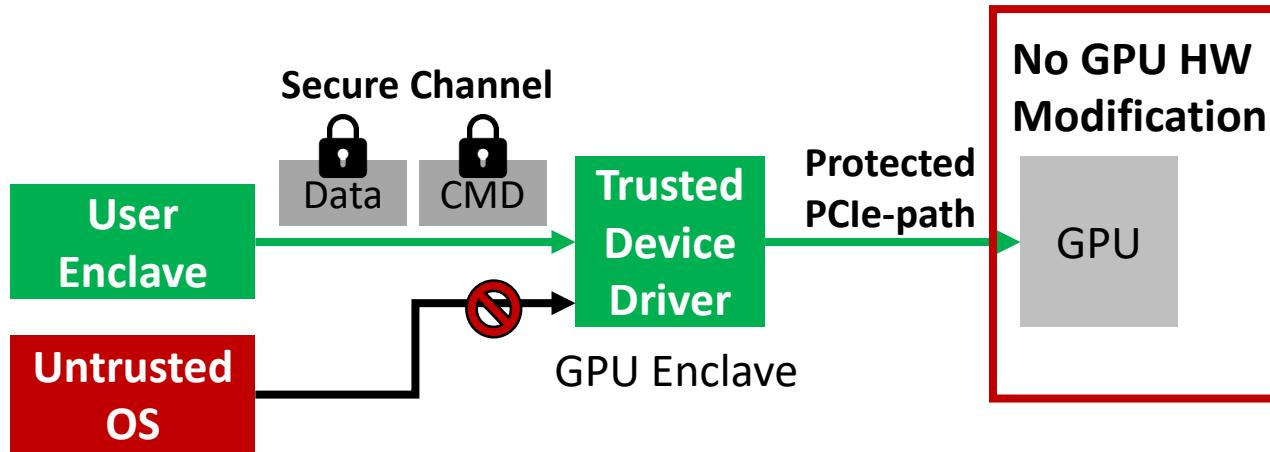
Trusted GPU Computing

- Trusted Execution Environment (TEE)
 - Intel SGX, ARM TrustZone
- Existing TEEs **does not provide TEE on GPUs**



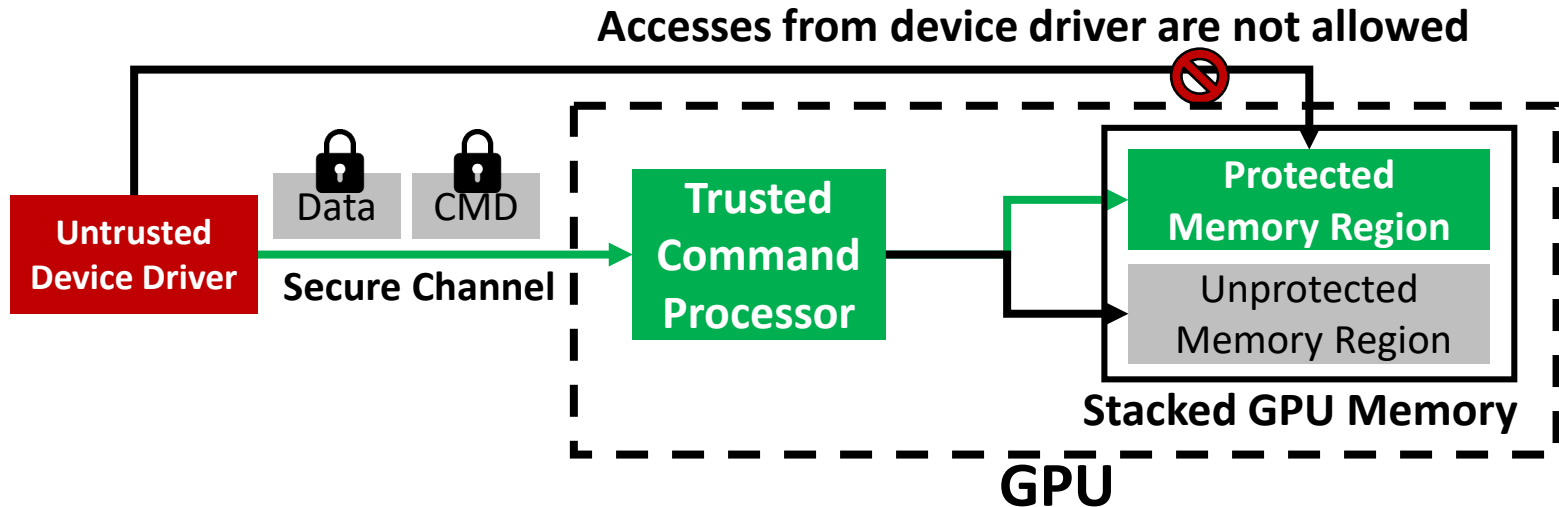
Prior Work : HIX

- **HIX [ASPLOS '19]: Securing I/O Path from CPU to GPU**
 - All device I/O accesses to GPU are controlled by **trusted device driver**



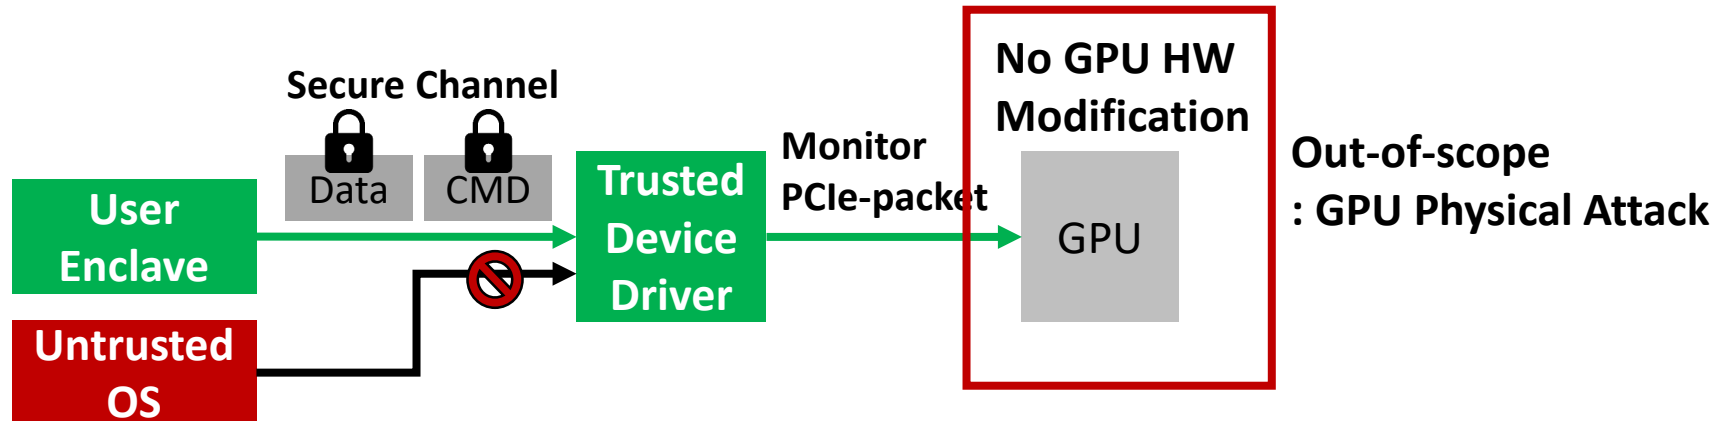
Prior Work : Graviton

- **Graviton [OSDI '18]: Trusted GPU by changing GPU HW**
 - **Trusted Command Processor** handles critical GPU operations instead of driver



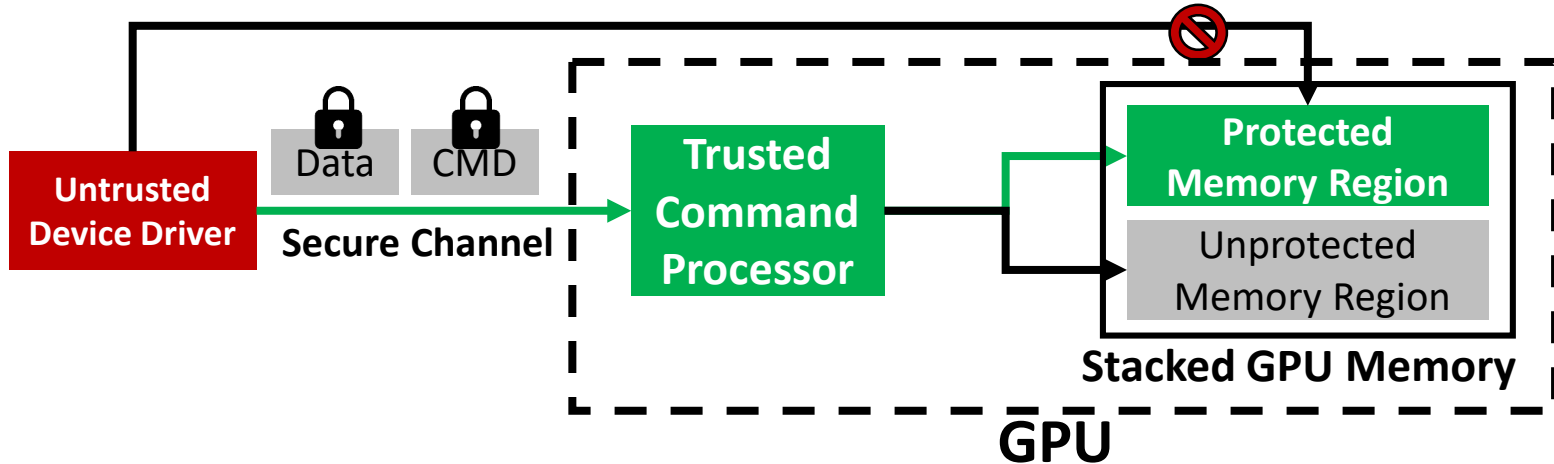
Limitations of Prior Work

- HIX [ASPLOS '19]: Securing I/O Path from CPU to GPU



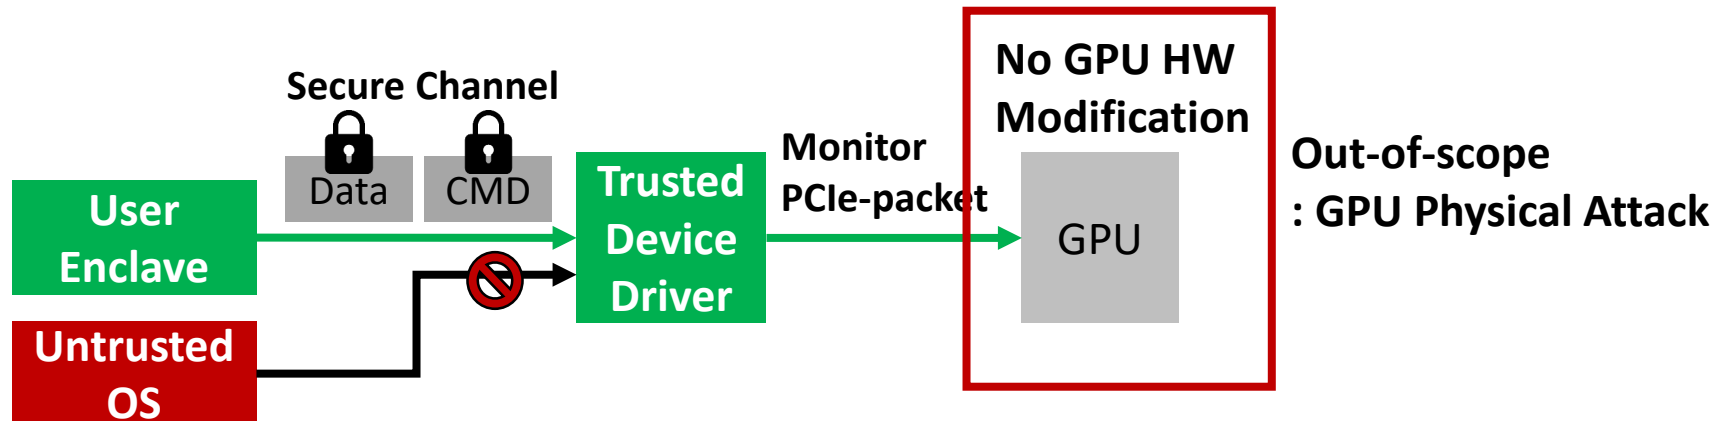
- Graviton [OSDI '18]: Trusted GPU by changing GPU HW

Accesses from device driver are not allowed

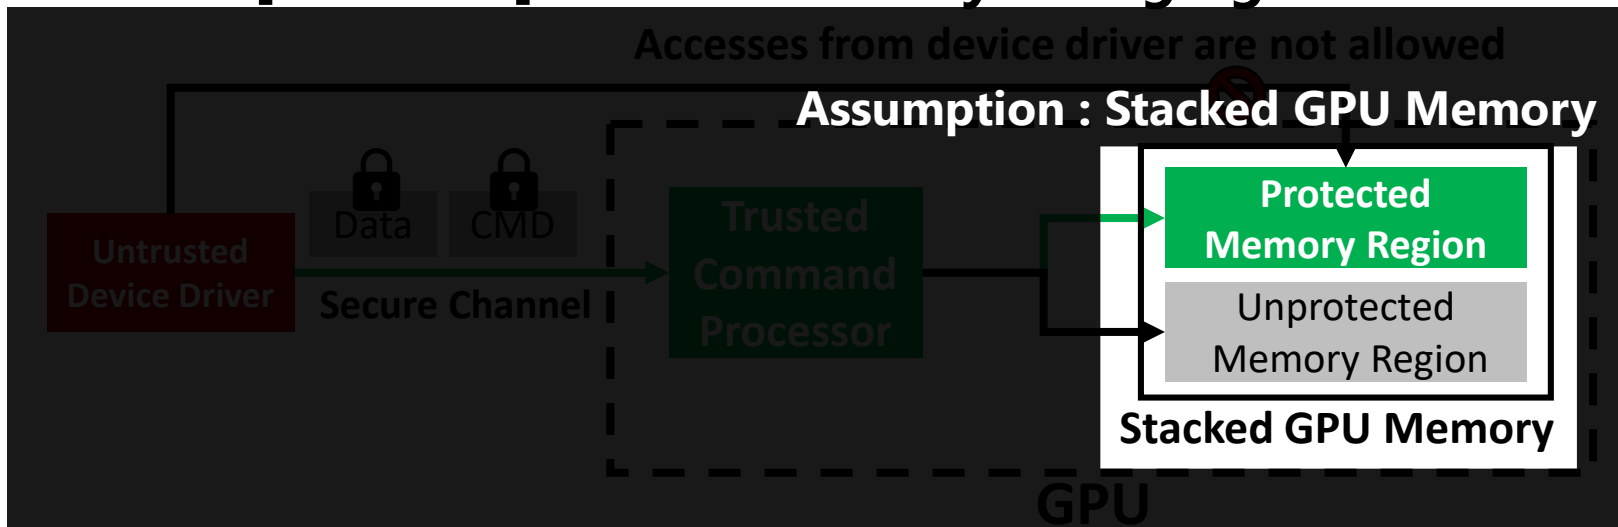


Limitations of Prior Work

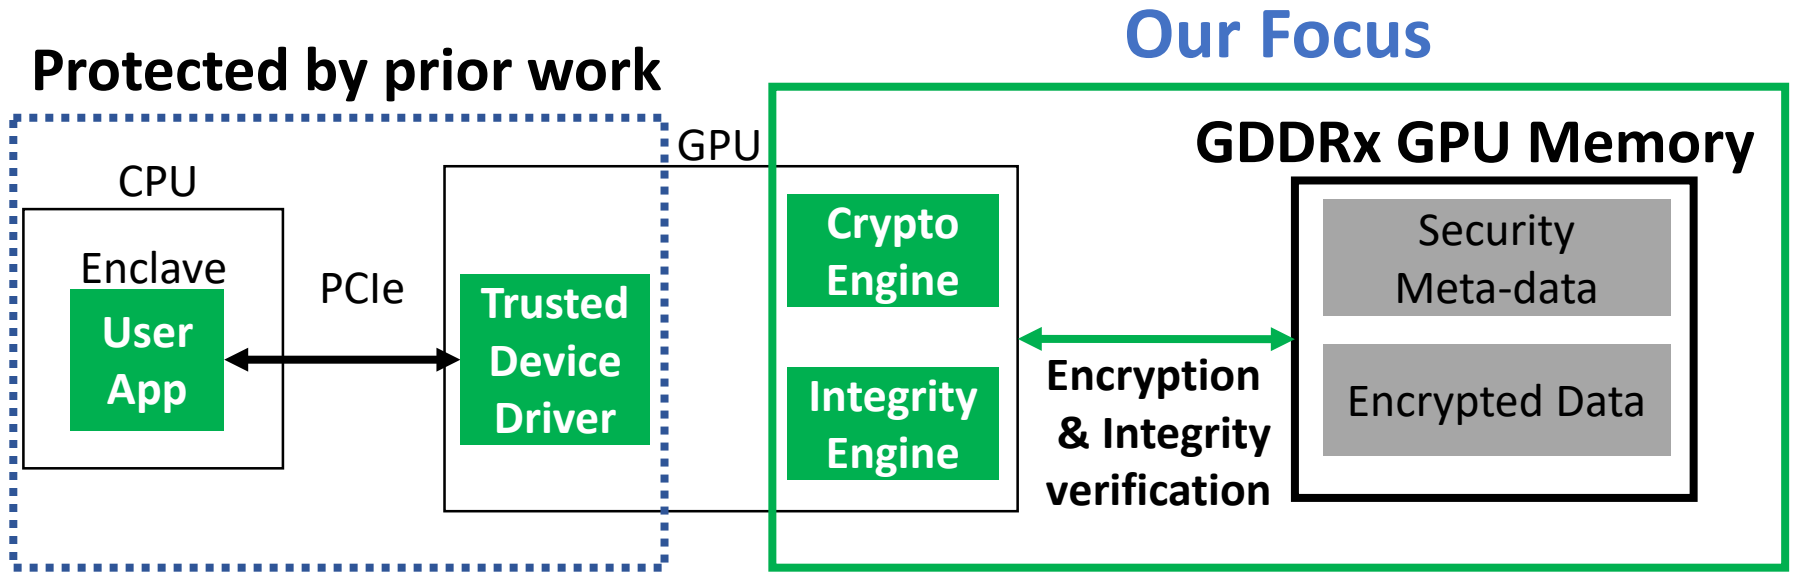
- HIX [ASPLOS '19]: Securing I/O Path from CPU to GPU



- Graviton [OSDI '18]: Trusted GPU by changing GPU HW



Goal: Secure GPU Memory



• Main Contributions

- Provide secure GPU memory **with low performance overheads**
- Exploit **unique memory update behavior** of common GPU applications
- Reduce the average performance overhead to **2.9 %**

Threat Model & Assumptions

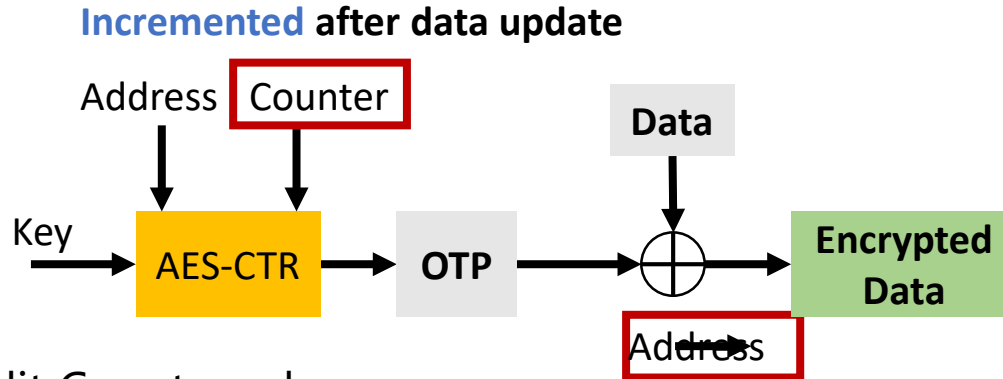
- Threat Model
 - Attackers can fully control **operating system/hypervisor**
 - Attackers can **physically access the whole system**
- Trusted Computing Base (TCB)
 - GPU processor & GPU software running on the GPU
 - CPU chip & user application in an CPU Enclave
- Out of Scope
 - Denial of Service(DoS) attacks
 - Side-channel attacks

Outline

- Introduction
- **Background & Motivation**
- Common Counter
 - Main Idea
 - Additional Metadata
 - Common Counter Mechanism
- Evaluation

Background : Securing Memory


- Memory Encryption
 - Counter mode encryption



- Split Counter scheme

Counter = Major | Minor

Minor Counters(7-bit for each)



Cache block :128B
→ 128 minor counters

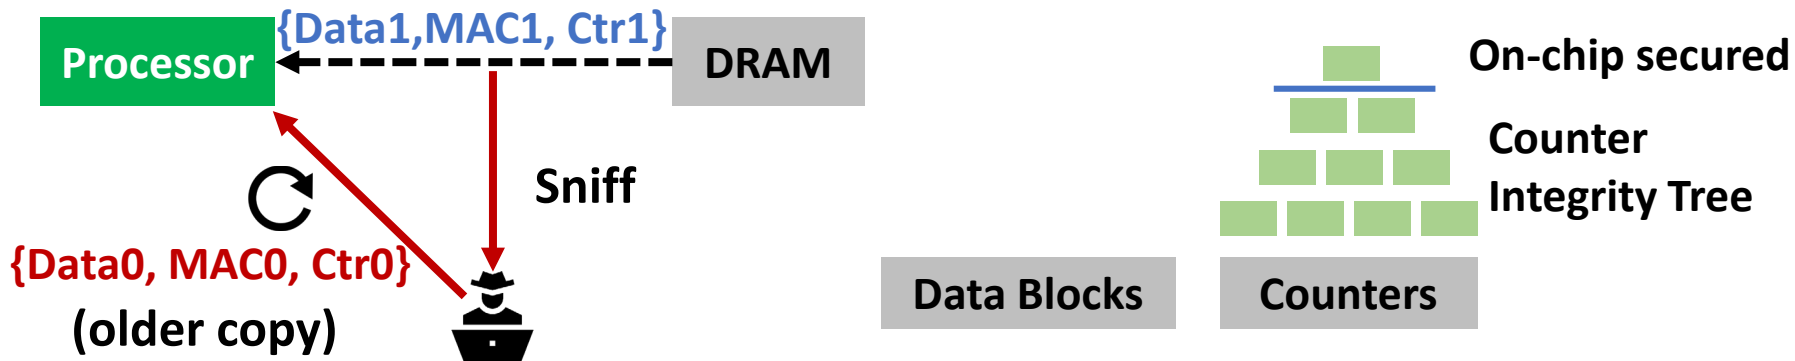
Background : Securing Memory

- Memory Integrity Verification
 - Message Authentication Code (MAC)



- Counter Integrity Tree

Replay Attack

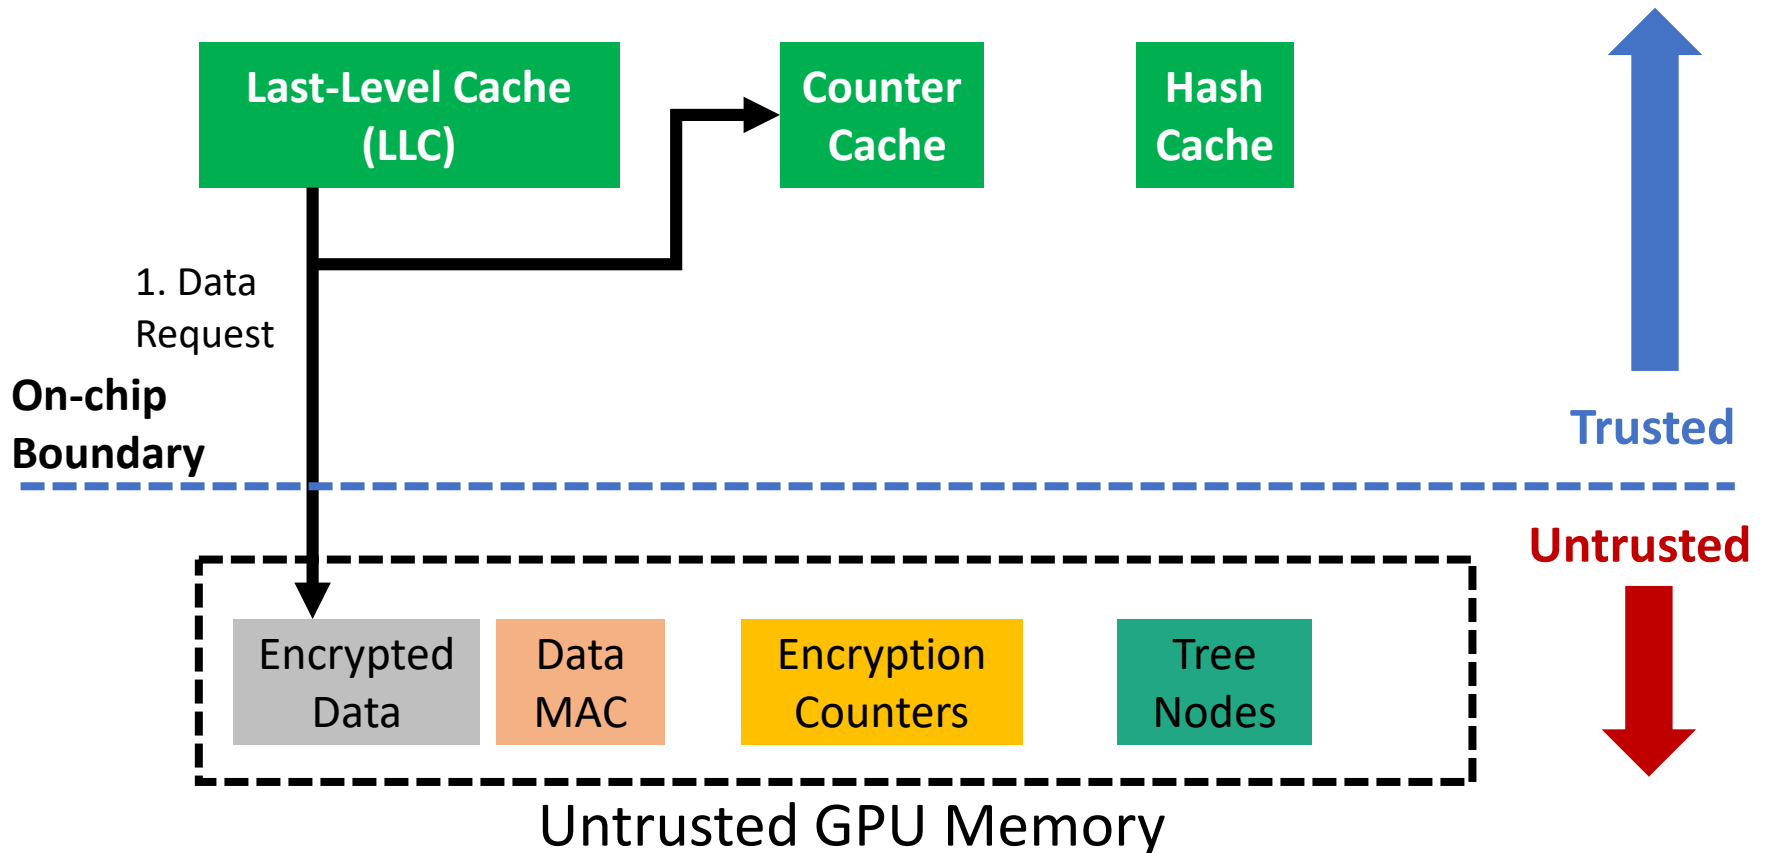


Our baseline : **SC-128**

128-ary (Split Counter + Counter integrity tree)

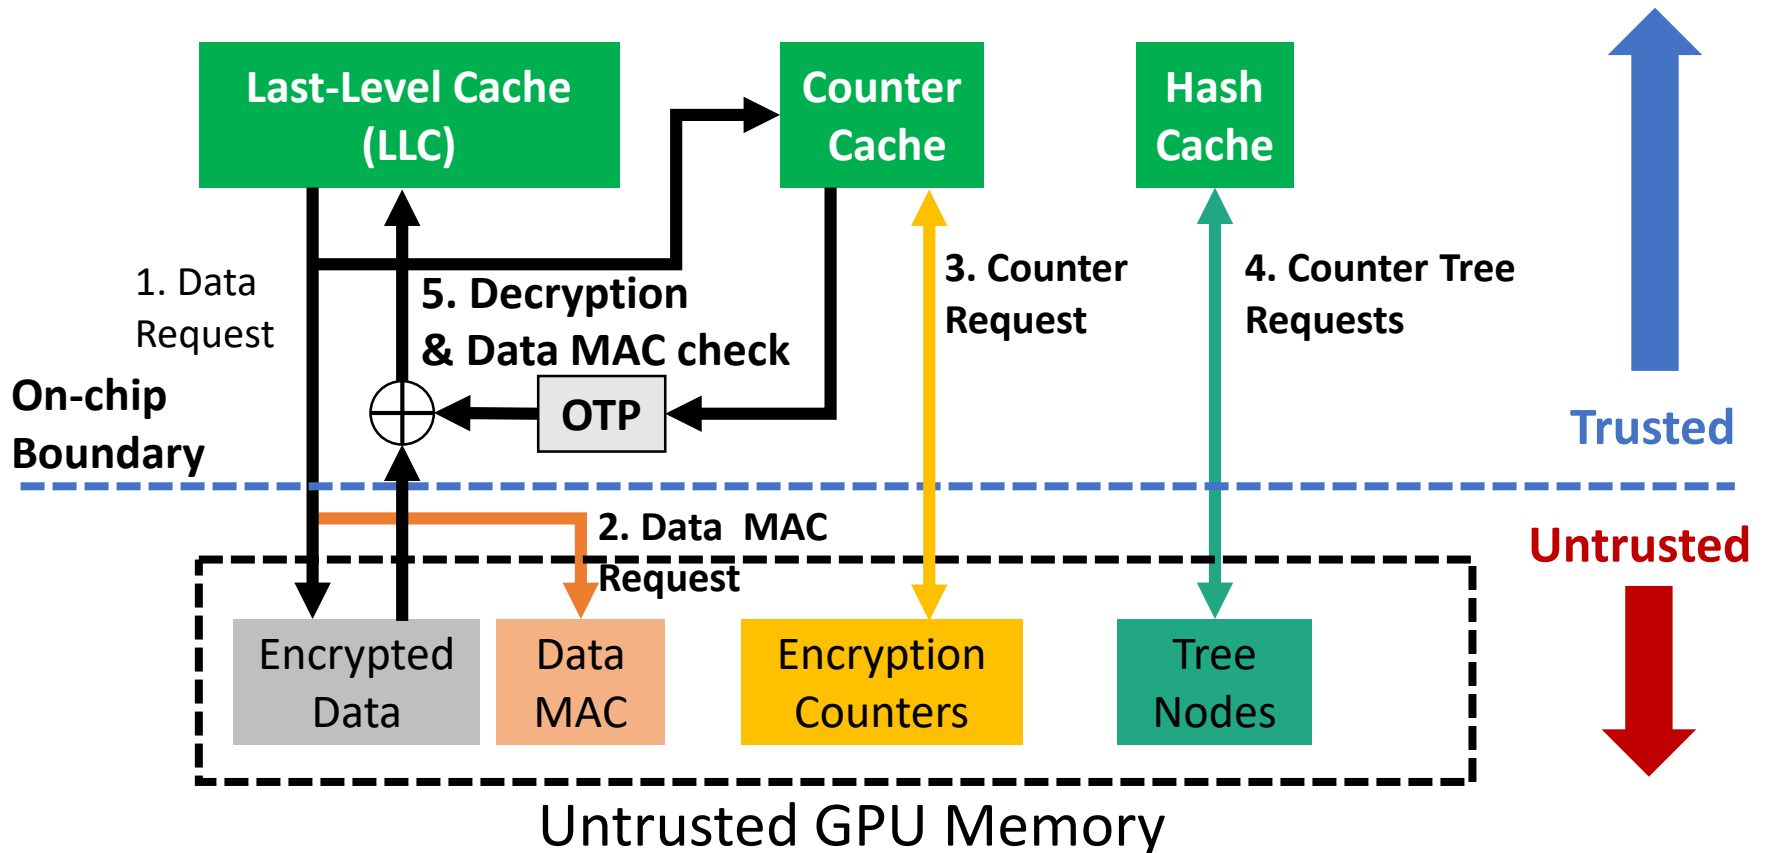
Problem : Performance Overhead

- Secure memory require **additional meta-data requests**



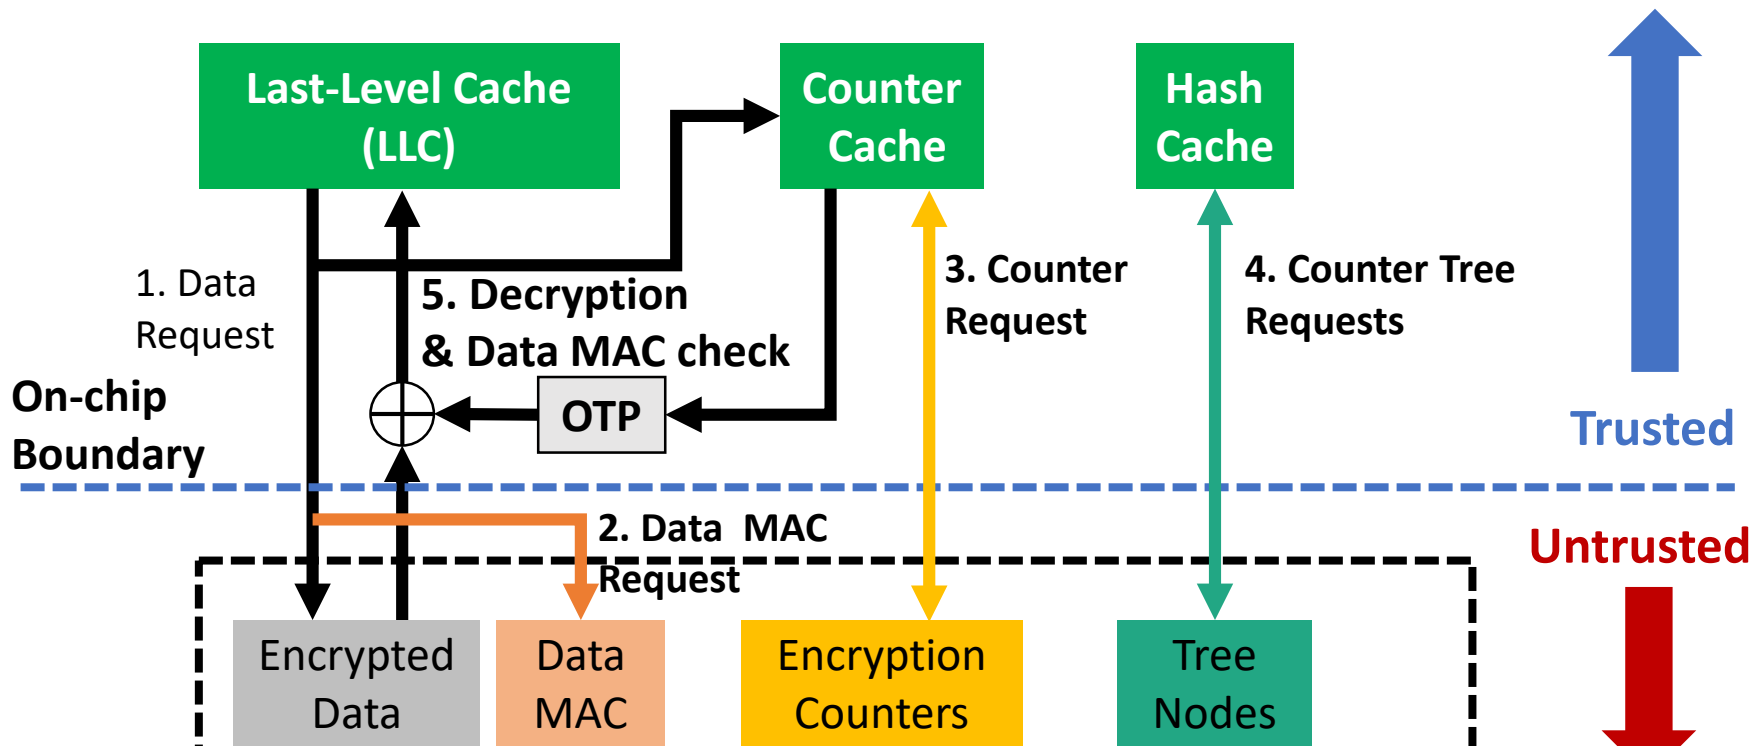
Problem : Performance Overhead

- Secure memory require **additional meta-data requests**



Problem : Performance Overhead

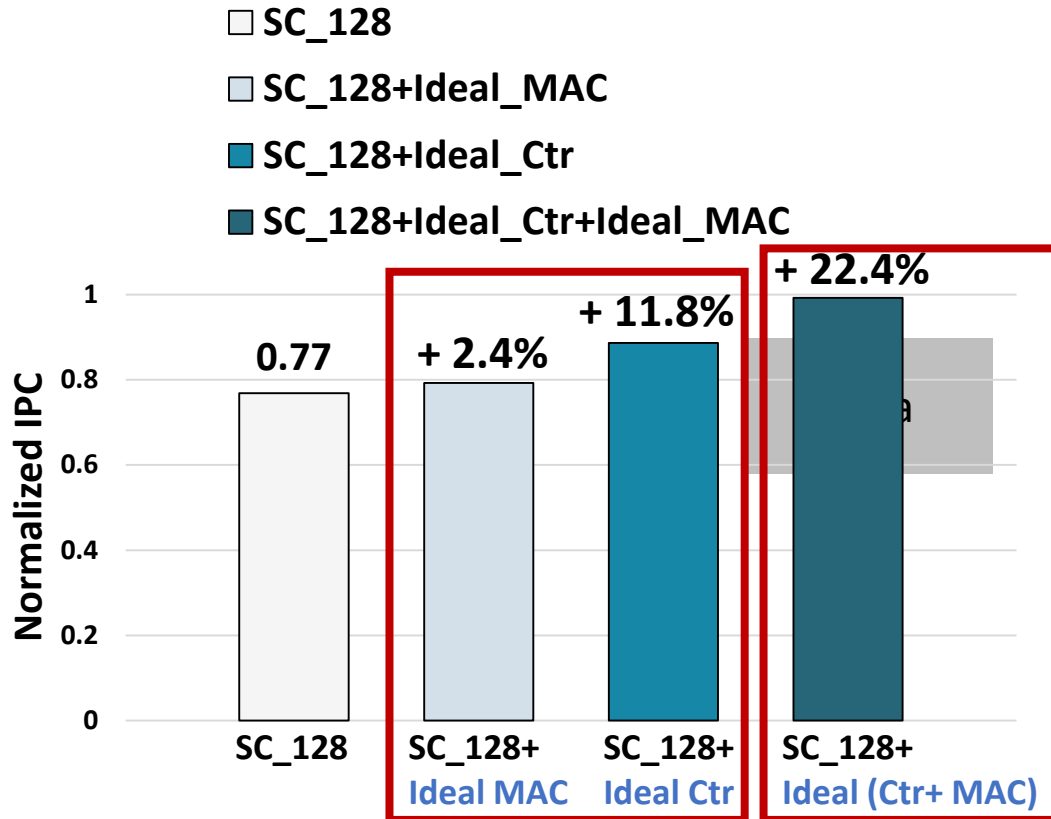
- Secure memory require **additional meta-data requests**




Secure memory **adds decryption latency**
and **increases memory bandwidth**

Performance Breakdown Analysis

- GPU memory protection overhead result for GPU benchmark suites



For data MAC overhead




With ECC memory,
Data & MAC can be provided by
1 memory access using [Synergy\[1\]](#)


[1] :SYNERGY: Rethinking Secure-Memory Design for Error-Correcting Memories, HPCA'18

Performance Breakdown Analysis

- GPU memory protection overhead result for GPU benchmark suites



For data MAC overhead



With ECC memory,
Data & MAC can be provided by
1 memory access using [Synergy\[1\]](#)

Counter mode encryption is one of the key bottlenecks


[1] :SYNERGY: Rethinking Secure-Memory Design for Error-Correcting Memories, HPCA'18

Uniformly Updated Segments

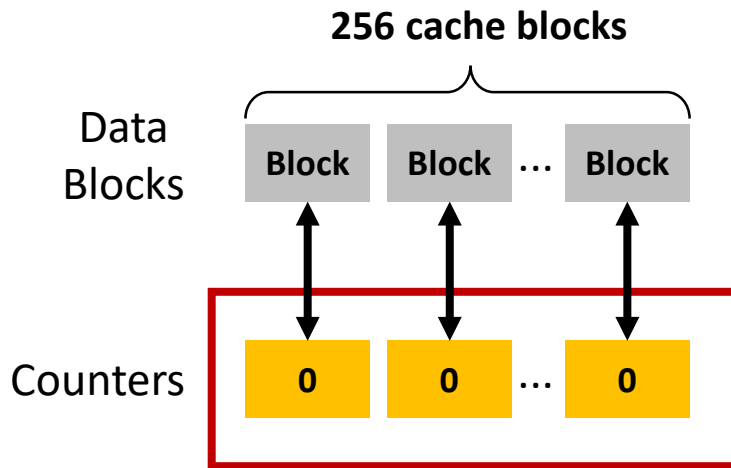
- Memory segment: Contiguous memory region

Example granularity: 32KB


$$32\text{KB} / 128\text{B} = 256 \text{ Data cache blocks}$$



- Uniformly updated segment: **Read-only** + **uniformly written**



1. Read-only

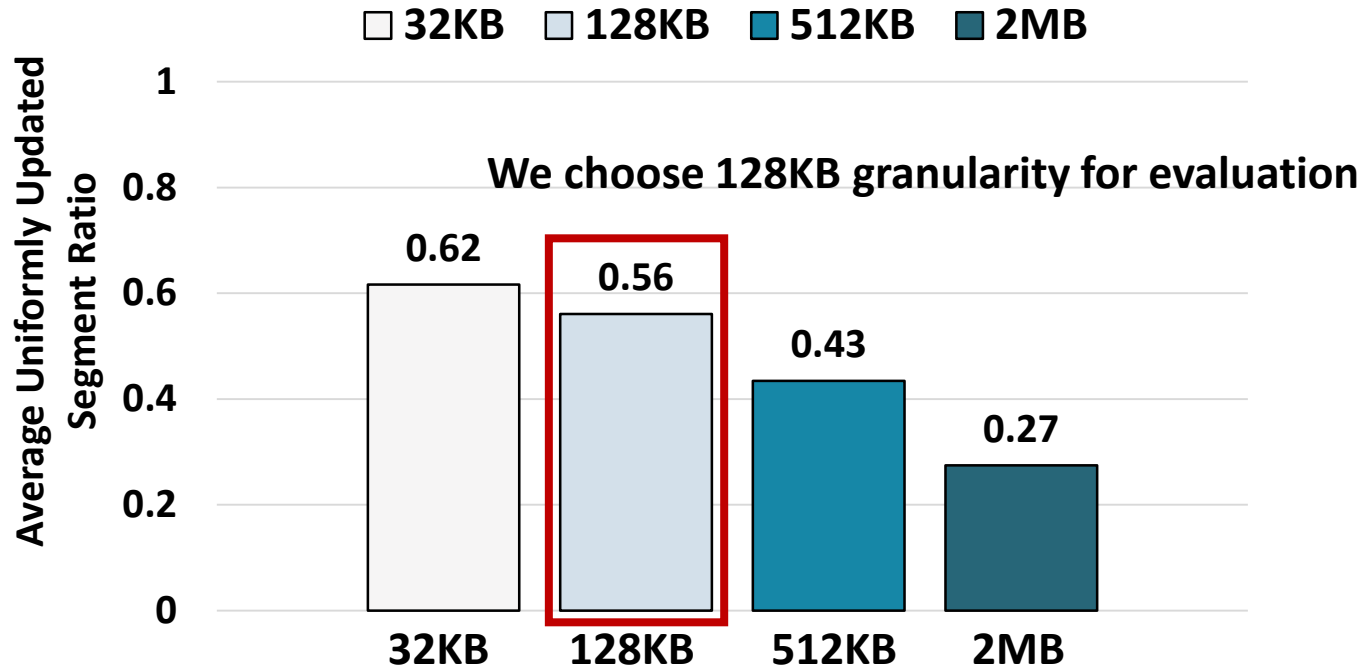


2. Uniformly written

Observation : GPU SW Write Patterns

- Analyze memory read/write behavior by using NVBit [MICRO '19]

Result of GPU Benchmark Suite




Observation 1: GPU programs tend to uniformly update memory
Observation 2: The number of distinct counter values is small

Outline

- Introduction
- Background & Motivation
- **Common Counter**
 - **Main Idea**
 - **Additional Metadata**
 - **Common Counter Mechanism**
- Evaluation

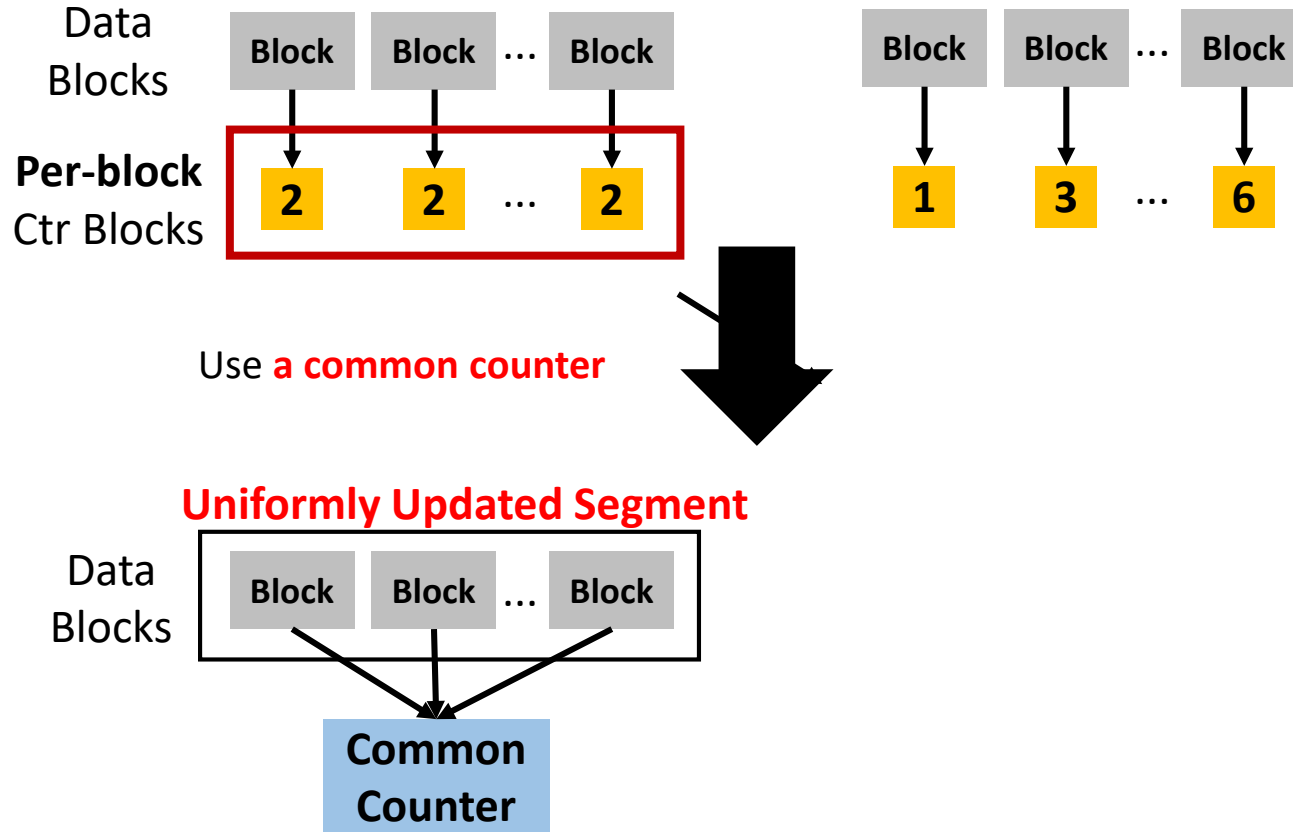
Common Counter : Main Idea

- Use **coarse-grained counters** for uniformly updated segments



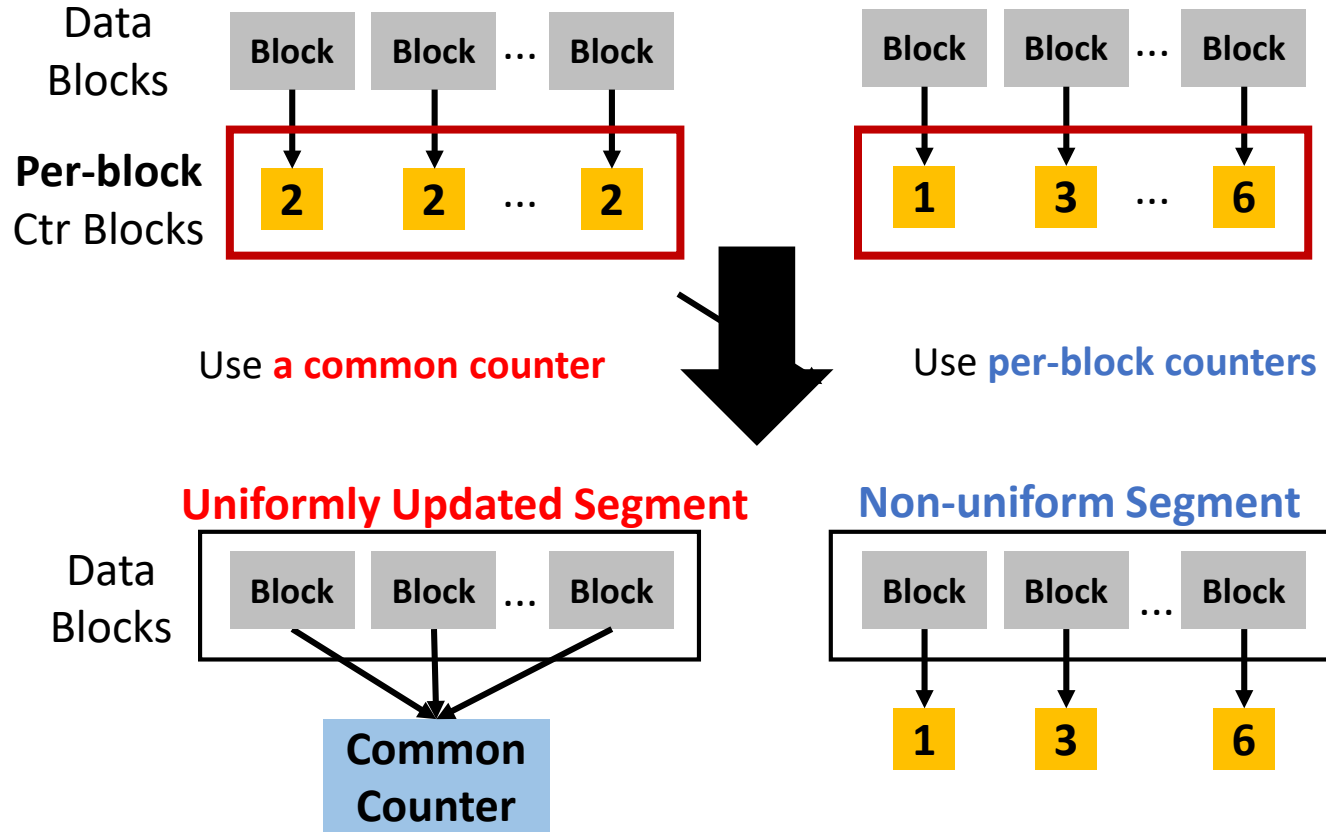
Common Counter : Main Idea

- Use **coarse-grained counters** for uniformly updated segments



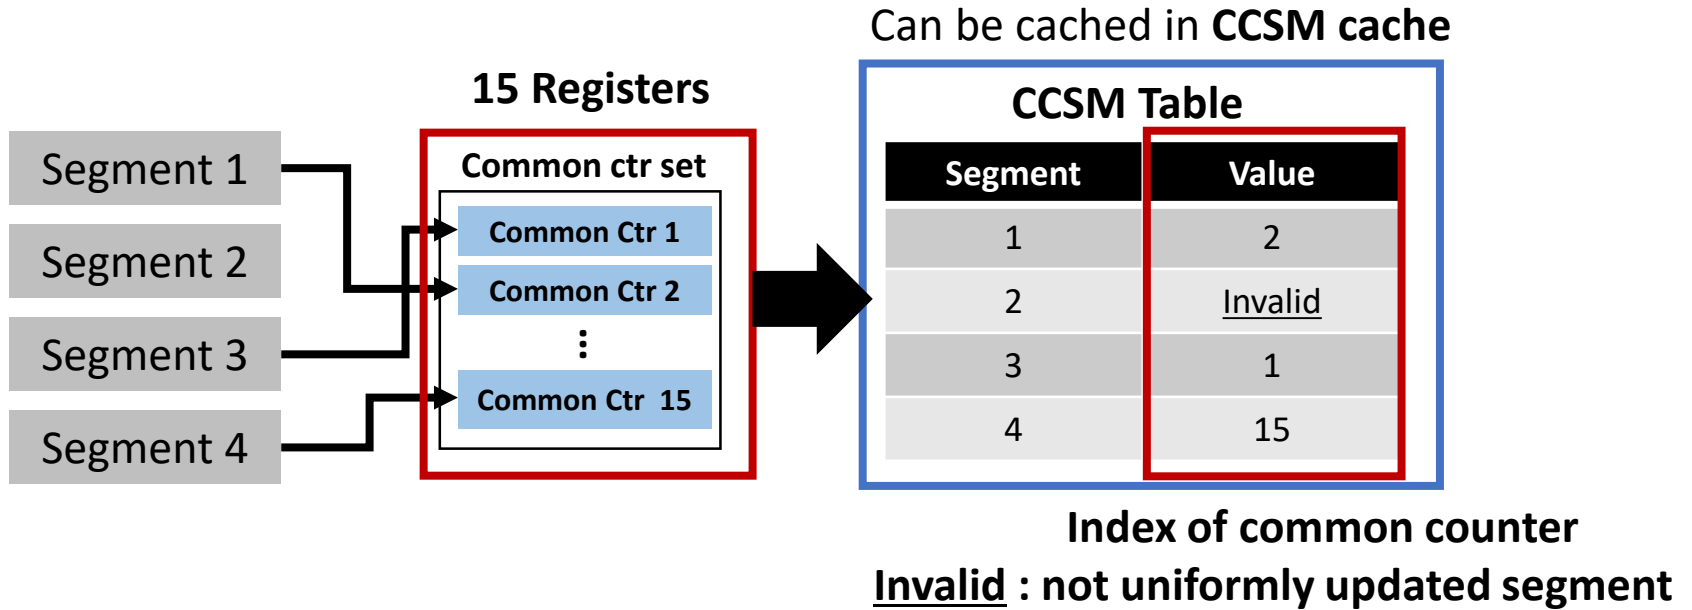
Common Counter : Main Idea

- Use **coarse-grained counters** for uniformly updated segments



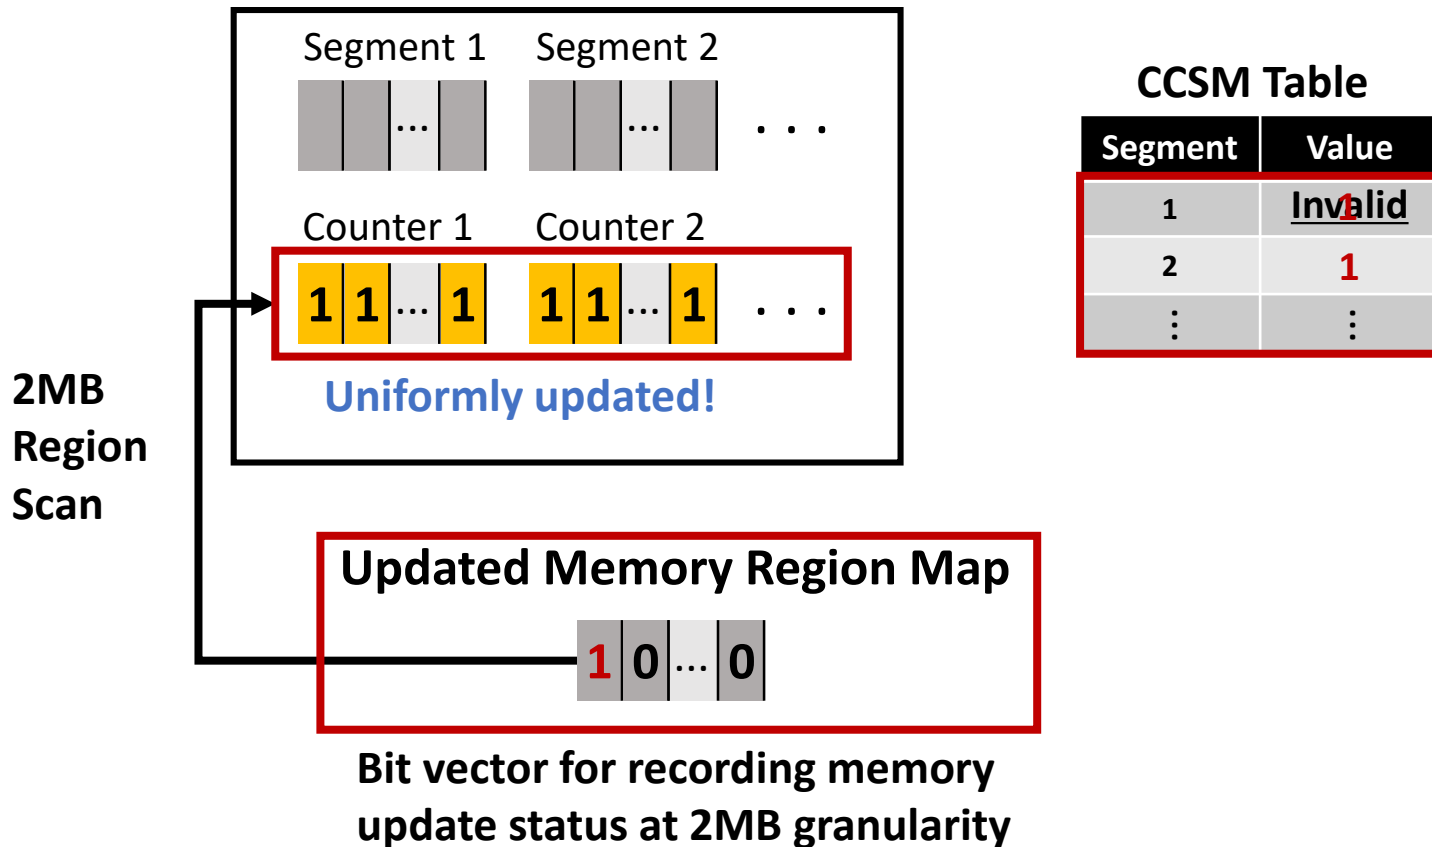
Finding Uniformly Updated Segments

- Common Counter Status Map (CCSM)
 - Check whether a memory segment uses a common counter or not

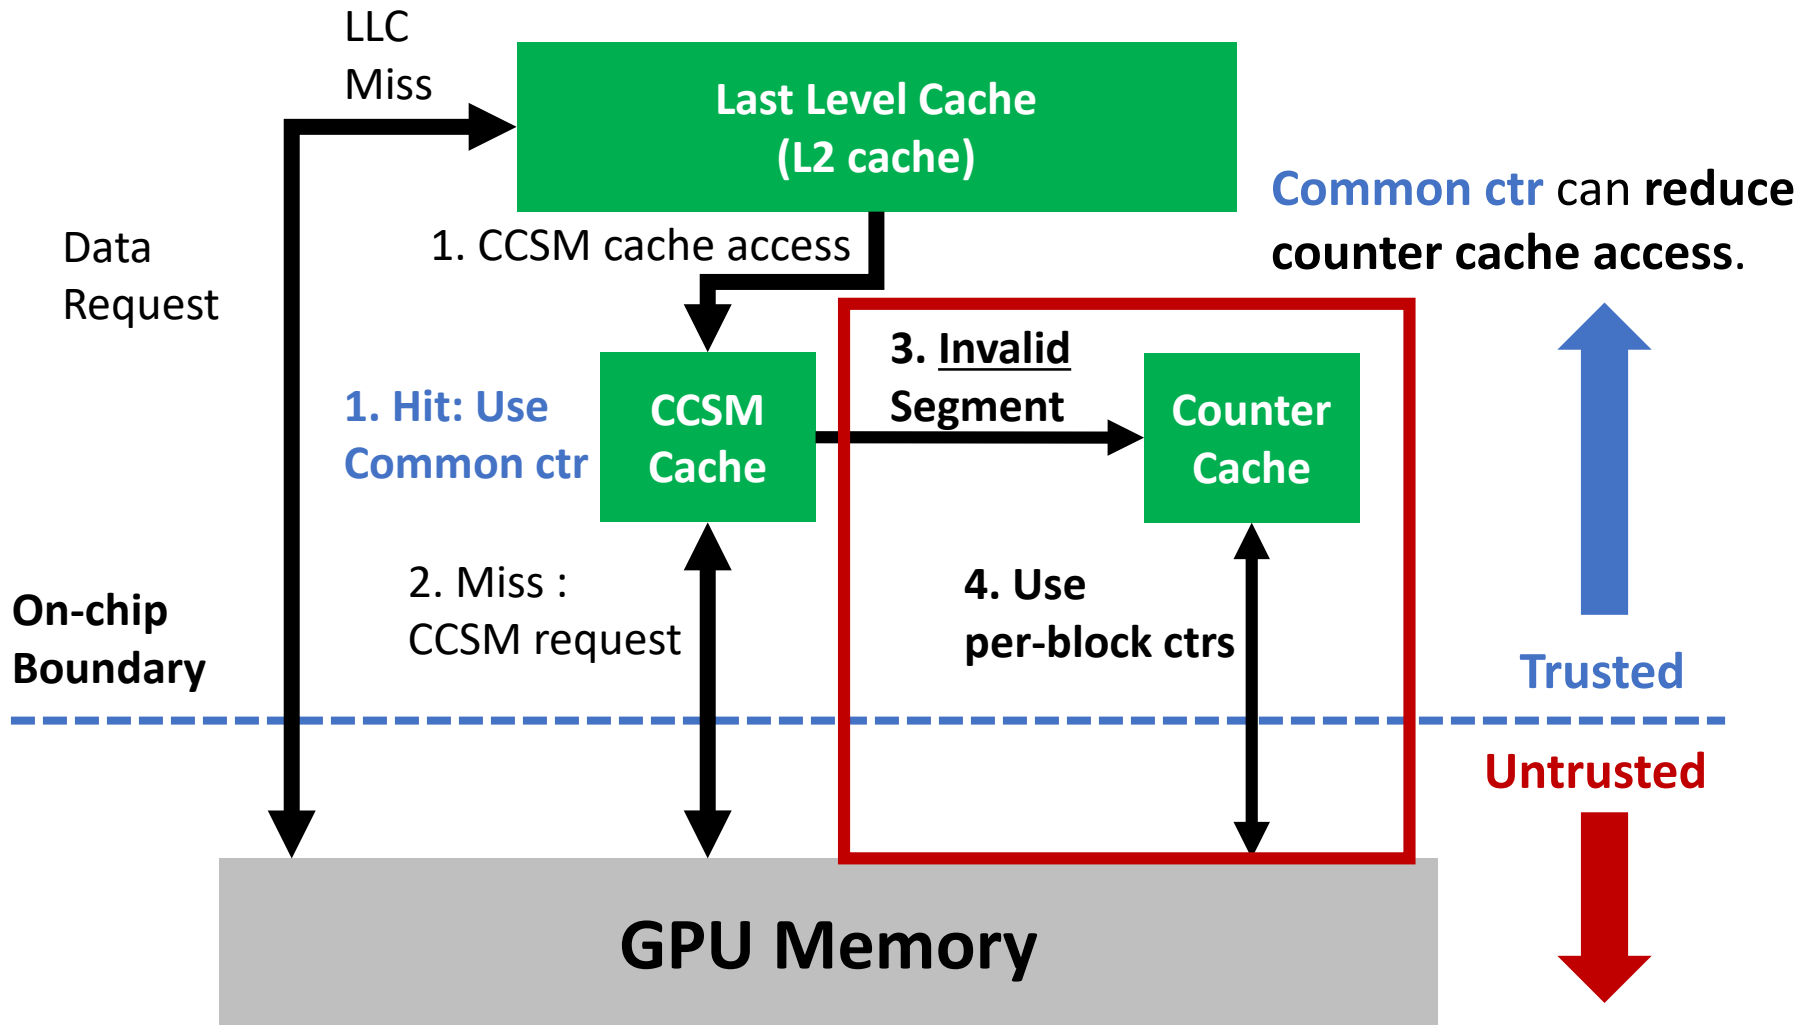


Updating CCSM Table

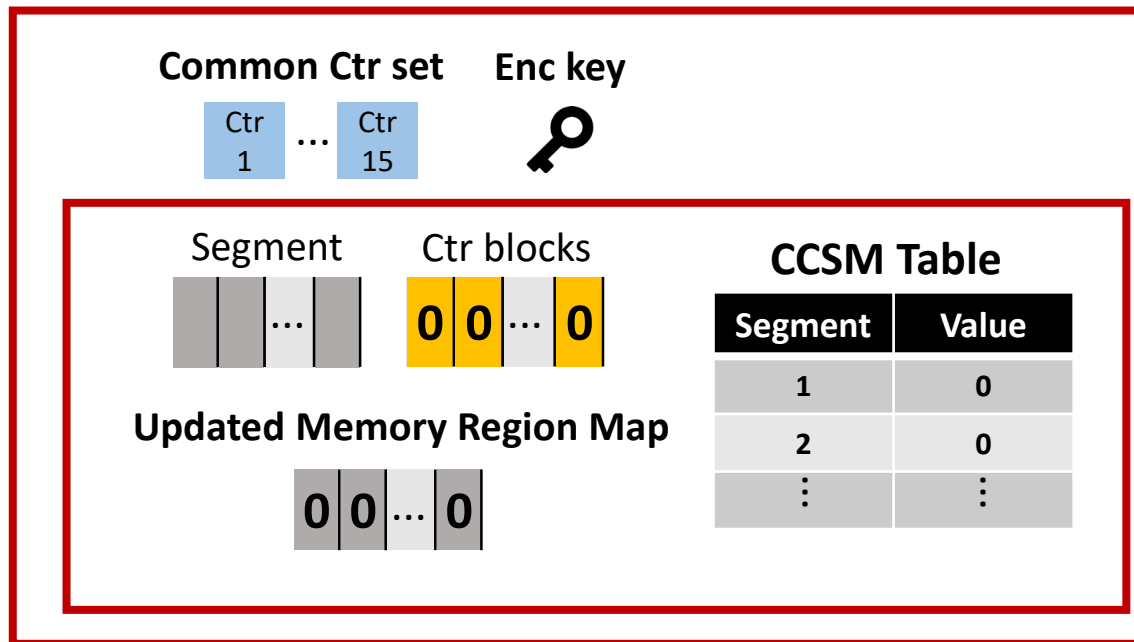
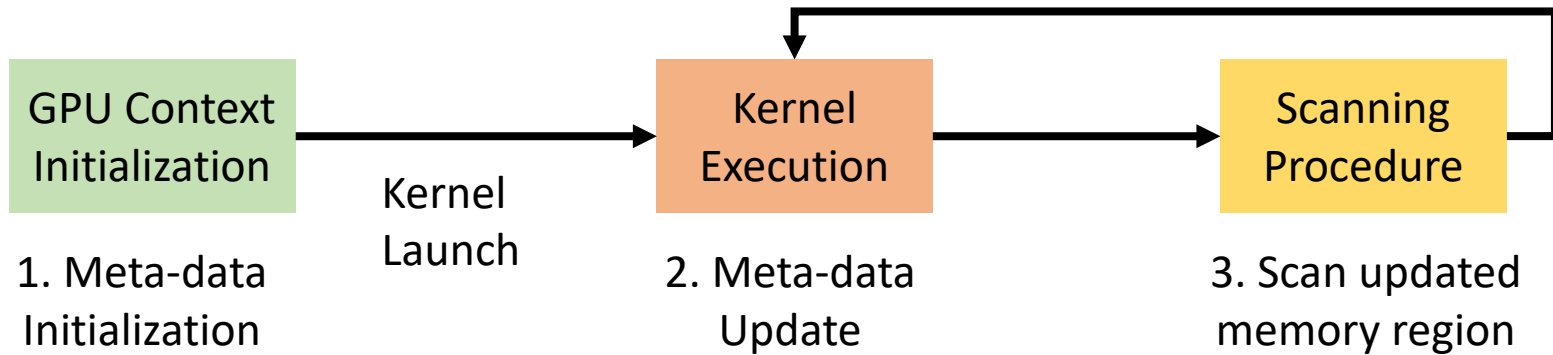
- Initialized at application launch
- Scanning Procedure
 - When? After a kernel is completed



LLC Miss Handling with Common Counters



GPU Execution with Common Counter



Outline

- Introduction
- Background & Motivation
- Common Counter
 - Main Idea
 - Additional Metadata
 - Common Counter Mechanism
- **Evaluation**


Methodology

- Simulator: GPGPU-Sim
- Workloads: ISPASS, Rodinia, Polybench, Pannotia
- System configuration: Models NVIDIA TITAN X Pascal GPU

GPU Core Configuration	
System overview	28 SMs, 64 warps per SM
Shader core	1,417 MHz, GTO Scheduler
Cache & Memory Configuration	
L1 cache	48 KB
L2 cache	3 MB
DRAM	GDDR5X 1,251 MHz, 12 GB
Counter cache, Hash cache	16 KB
CCSM cache	1 KB
Segment size	128 KB
Number of common ctrs	15

Performance: Separate MACs


- Performance overhead analysis (Baseline: Non-secure GPU)
 - SC-128: **128-arity** split counter
 - Morphable Counter [MICRO '18]: **256-arity** split counter
 - Common_Ctr: Implemented on top of SC-128 (**128-arity**)



Performance: Synergy In-line MACs

- **Common counter** reduces the performance degradation to **2.9%**

color, mis,lib,bfs : As kernel runs, # of requests served by **common counters** decreases.



[1] :SYNERGY: Rethinking Secure-Memory Design for Error-Correcting Memories, HPCA'18

More Results in the Paper

- Uniformly updated ratios of real-world GPU Applications
- Hardware area/energy cost for common counter mechanism
- Ratios of LLC misses served by common counters
- Scanning Overheads
- Counter cache sensitivity experiments

Please Refer to our paper for more details!

Conclusion

- **Result**

- **Common Counter** reduces the performance degradation to 2.9%

- **Problem**

- Memory encryption is one of the critical bottlenecks for secure GPU memory

- **Key Observation**

- GPU programs tend to uniformly update memory
- The number of distinct common counters is small

- **Our Approach**

- **Common Counter** provides compressed representation of per-block counters