

# Unified Memory Protection with Multi-granular MAC and Integrity Tree for Heterogeneous Processors

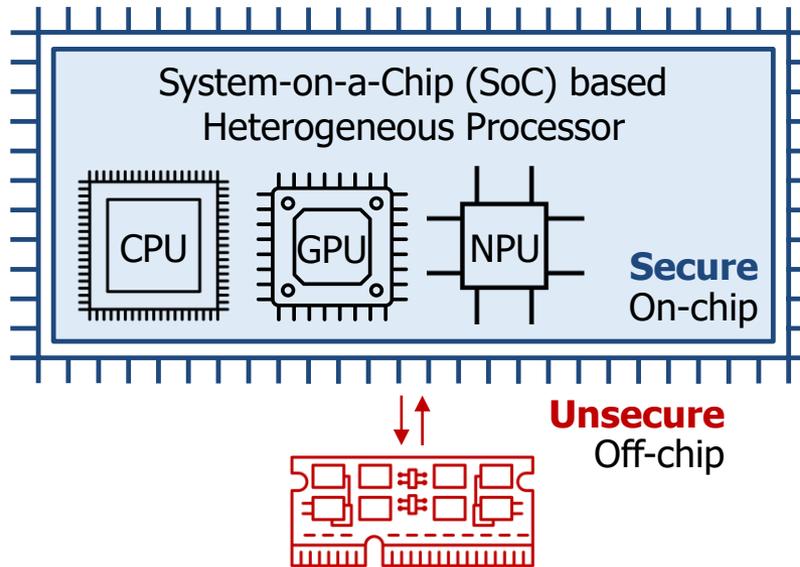
Sunho Lee<sup>1</sup>, Seonjin Na<sup>2</sup>, Jeongwon Choi<sup>1</sup>  
Jinwon Pyo<sup>1</sup>, Jaehyuk Huh<sup>1</sup>

KAIST<sup>1</sup>, Georgia Tech<sup>2</sup>



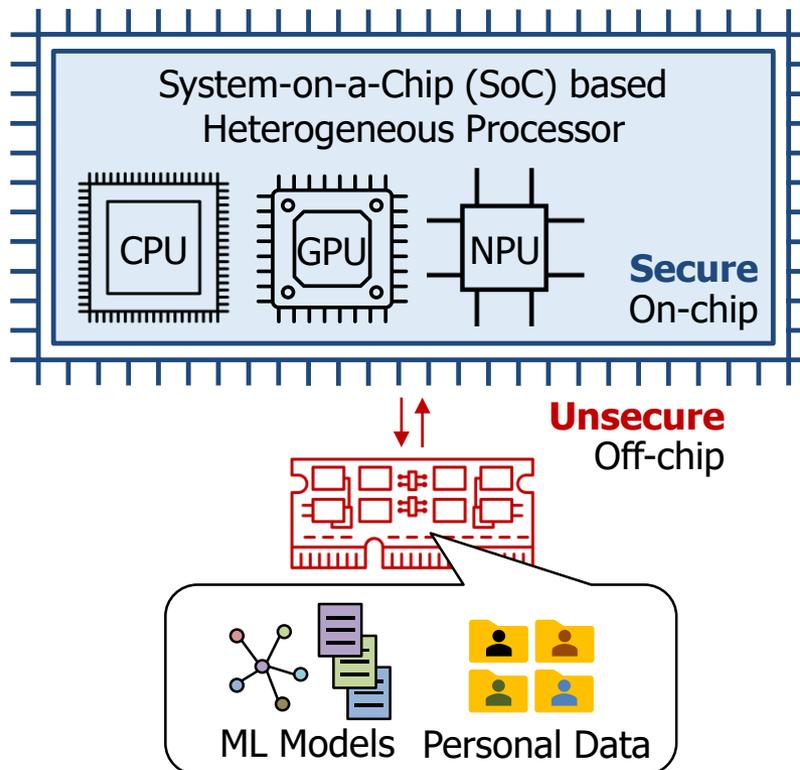
# Secure Heterogeneous Processor

- Heterogeneous processor: SoC with CPU, GPU, NPU



# Secure Heterogeneous Processor

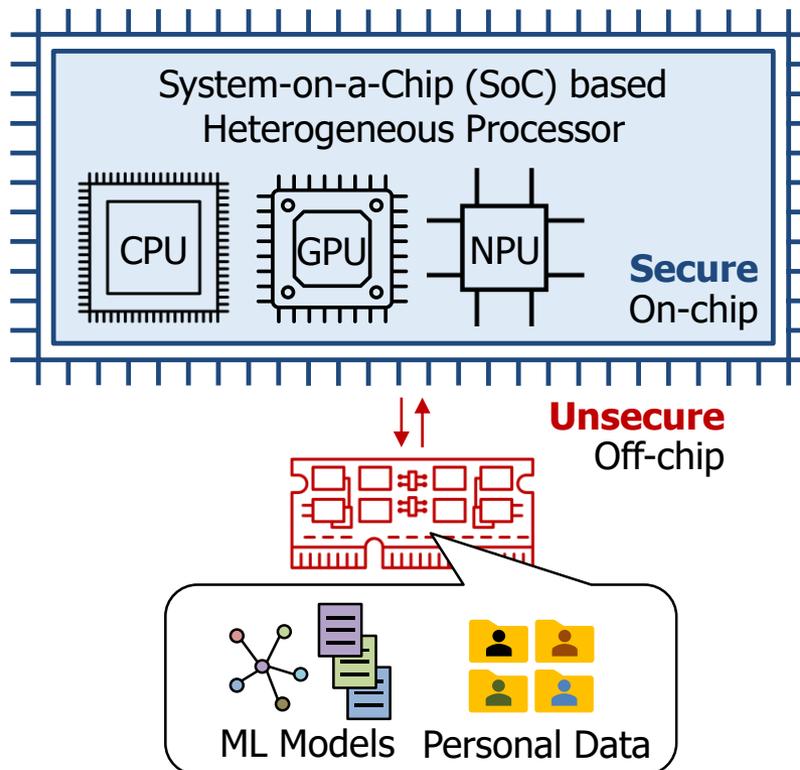
- Heterogeneous processor: SoC with CPU, GPU, NPU
- Data confidentiality & integrity are essential



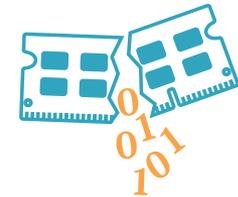
- [1] Lest We Remember: Cold-Boot Attacks on Encryption Keys (USENIX Security 2008)
- [2] RAMBleed: Reading Bits in Memory Without Accessing Them (S&P 2020)
- [3] Direct Memory Attack the Kernel (DEFCON 2016; PCILeech)
- [4] Handbook of Applied Cryptography (Menezes, Alfred J. et. al.)

# Secure Heterogeneous Processor

- Heterogeneous processor: SoC with CPU, GPU, NPU
- Data confidentiality & integrity are essential



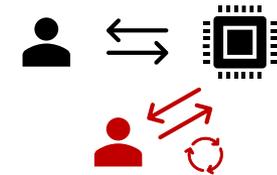
Cold Boot Attack [1]



Rowhammer Attack [2]



DMA Attack [3]



Replay Attack [4]

[1] Lest We Remember: Cold-Boot Attacks on Encryption Keys (USENIX Security 2008)

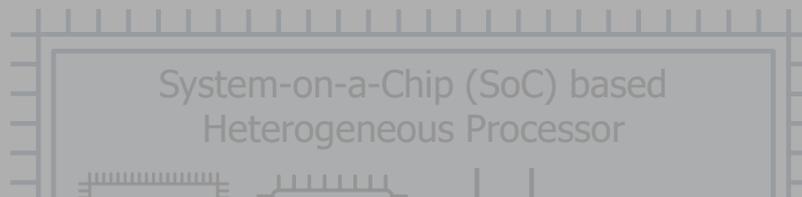
[2] RAMBleed: Reading Bits in Memory Without Accessing Them (S&P 2020)

[3] Direct Memory Attack the Kernel (DEFCON 2016; PCILeech)

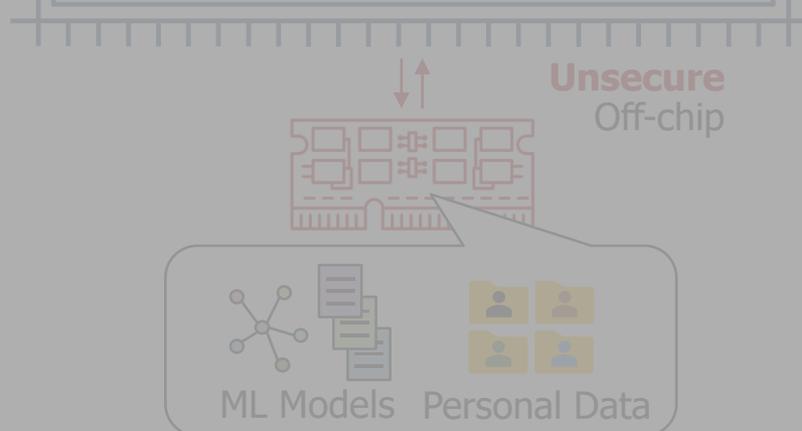
[4] Handbook of Applied Cryptography (Menezes, Alfred J. et. al.)

# Secure Heterogeneous Processor

- Heterogeneous processor: SoC with CPU, GPU, NPU
- Data confidentiality & integrity are essential



## Memory protection is necessary for heterogeneous processors



Cold Boot Attack [1]

Rowhammer Attack [2]



DMA Attack [3]



Replay Attack [4]

[1] Lest We Remember: Cold-Boot Attacks on Encryption Keys (USENIX Security 2008)

[2] RAMBleed: Reading Bits in Memory Without Accessing Them (S&P 2020)

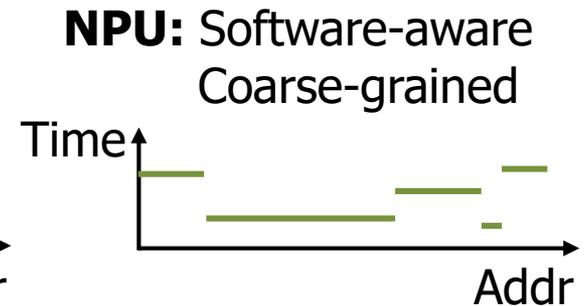
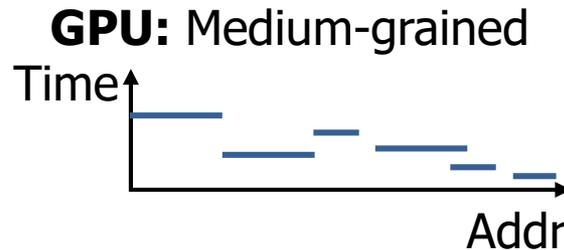
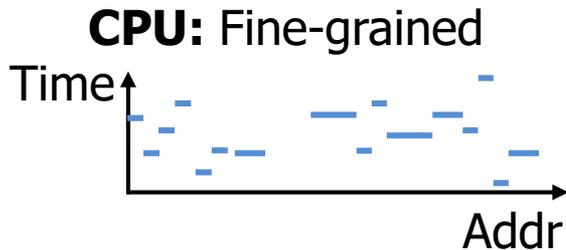
[3] Direct Memory Attack the Kernel (DEFCON 2016; PCIleech)

[4] Handbook of Applied Cryptography (Menezes, Alfred J. et. al.)

# Research Goal

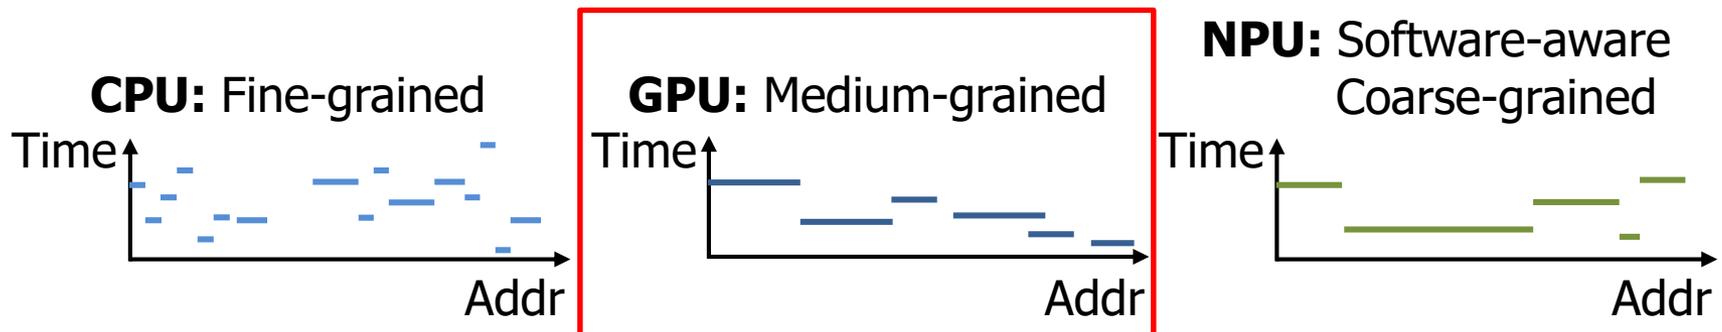
---

- Existing memory protections are tailored to **specific, individual access patterns**



# Research Goal

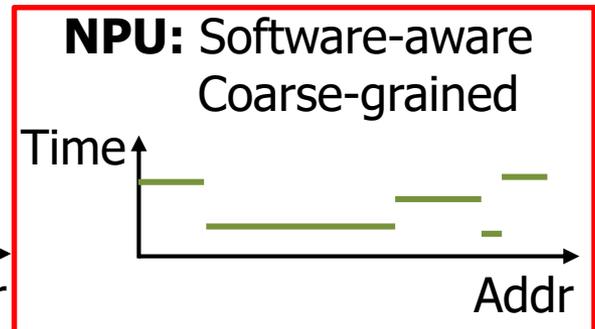
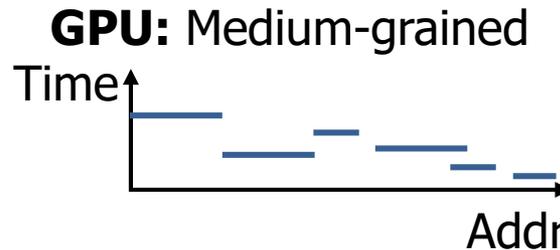
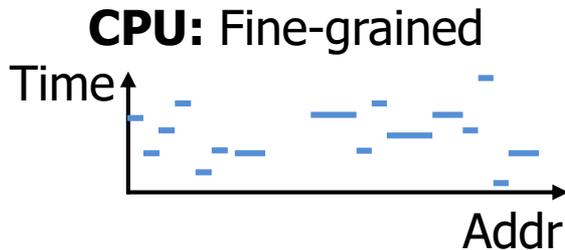
- Existing memory protections are tailored to **specific, individual access patterns**
  - Common Counters [1] → GPU medium-grained pattern



[1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)

# Research Goal

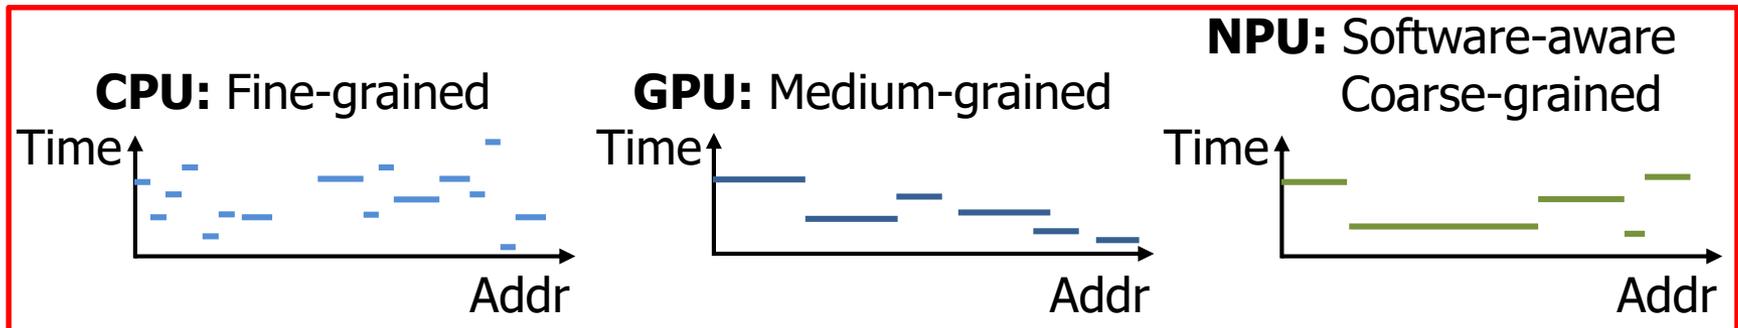
- Existing memory protections are tailored to **specific, individual access patterns**
  - Common Counters [1] → GPU medium-grained pattern
  - Studies of S/W-based counters → NPU S/W-detected pattern



[1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)

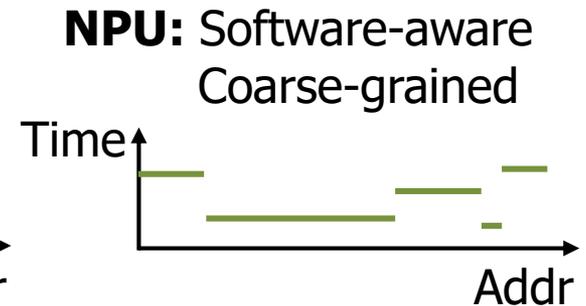
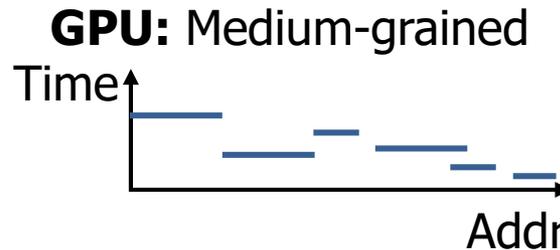
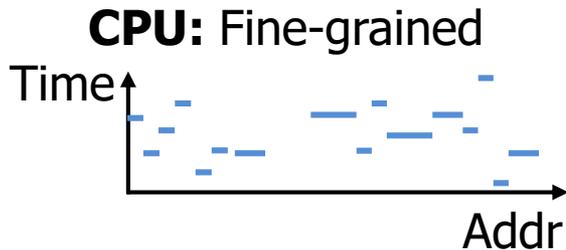
# Research Goal

- Existing memory protections are tailored to **specific, individual access patterns**
  - Common Counters [1] → GPU medium-grained pattern
  - Studies of S/W-based counters → NPU S/W-detected pattern
- Heterogeneous processor → **diverse access pattern**



# Research Goal

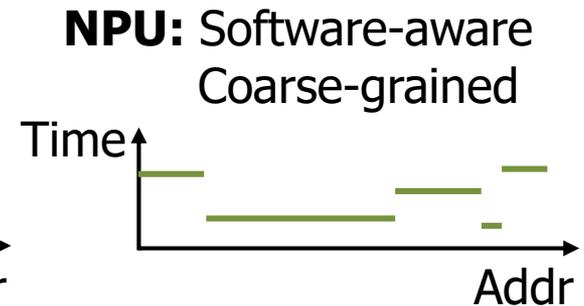
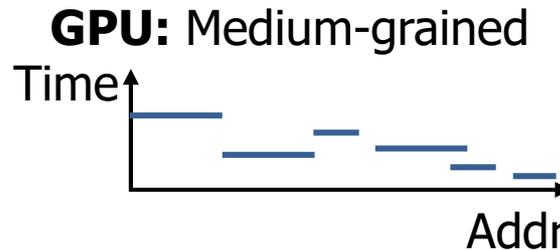
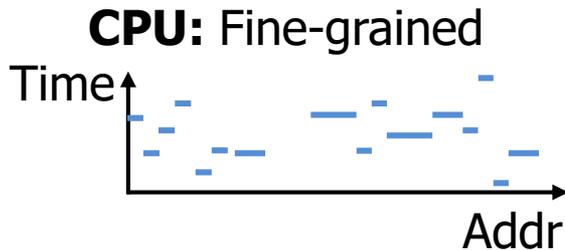
- Existing memory protections are tailored to **specific, individual access patterns**
  - Common Counters [1] → GPU medium-grained pattern
  - Studies of S/W-based counters → NPU S/W-detected pattern
- Heterogeneous processor → **diverse access pattern**
  - A **unified** memory protection for all access patterns



[1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)

# Research Goal

- Existing memory protections are tailored to **specific, individual access patterns**
  - Common Counters [1] → GPU medium-grained pattern
  - Studies of S/W-based counters → NPU S/W-detected pattern
- Heterogeneous processor → **diverse access pattern**
  - A **unified** memory protection for all access patterns
  - Limitation of prior studies: Bypassing **integrity tree** optimization



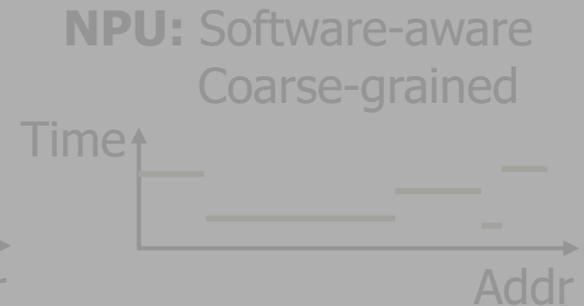
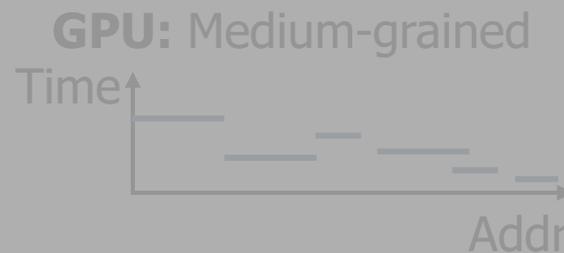
[1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)

# Research Goal

- Existing memory protections are tailored to **specific, individual access patterns**
  - Common Counters [1] → GPU medium-grained pattern
  - Studies of S/W-based counters → NPU S/W-detected pattern

**This study constructs a unified memory protection scheme with integrity tree optimization for heterogeneous processors**

- A **unified** memory protection for all access patterns
- Limitation of prior studies: Bypassing **integrity tree** optimization



[1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)

# Counter-mode Memory Protection

---

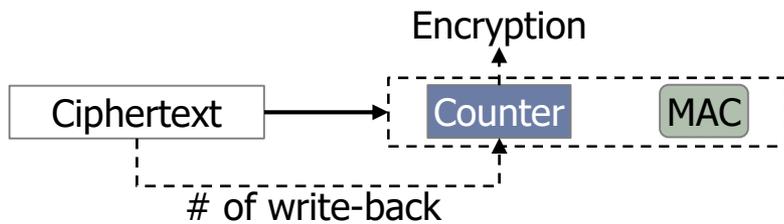
# Counter-mode Memory Protection

---



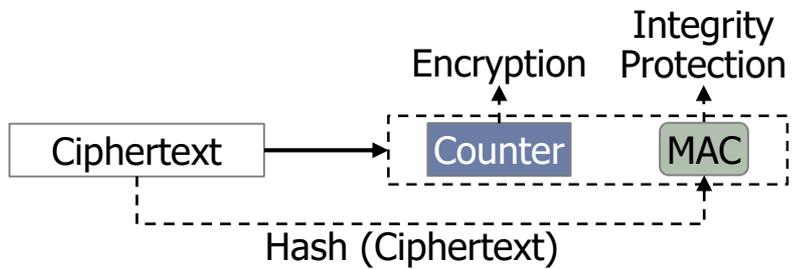
# Counter-mode Memory Protection

---



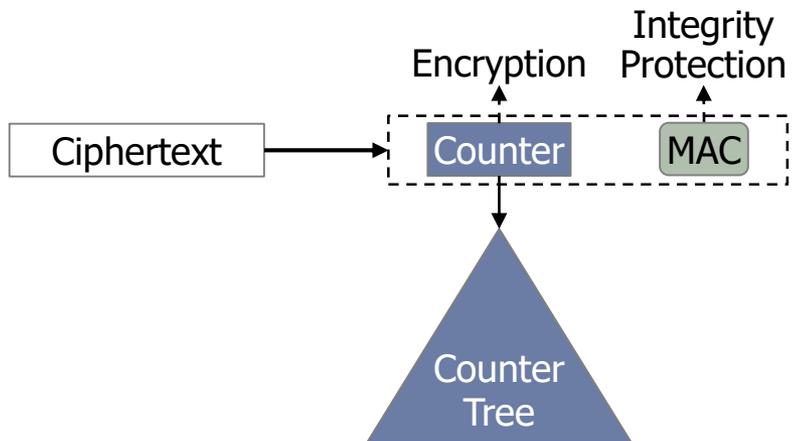
# Counter-mode Memory Protection

---



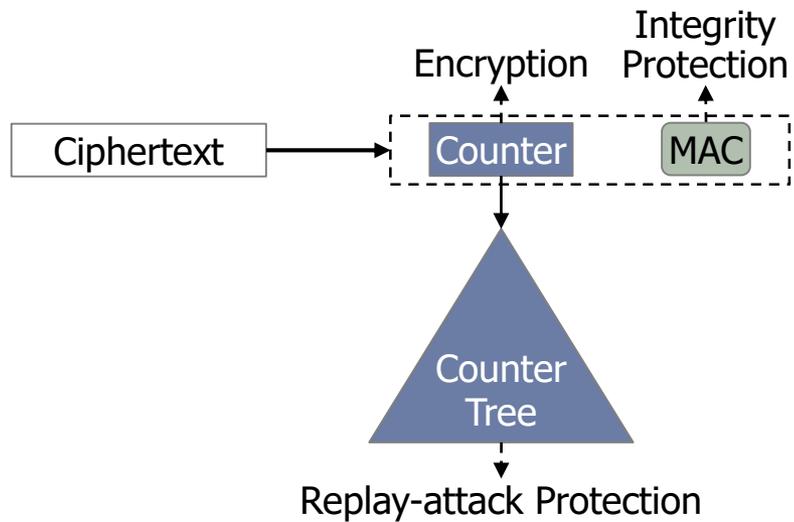
# Counter-mode Memory Protection

---



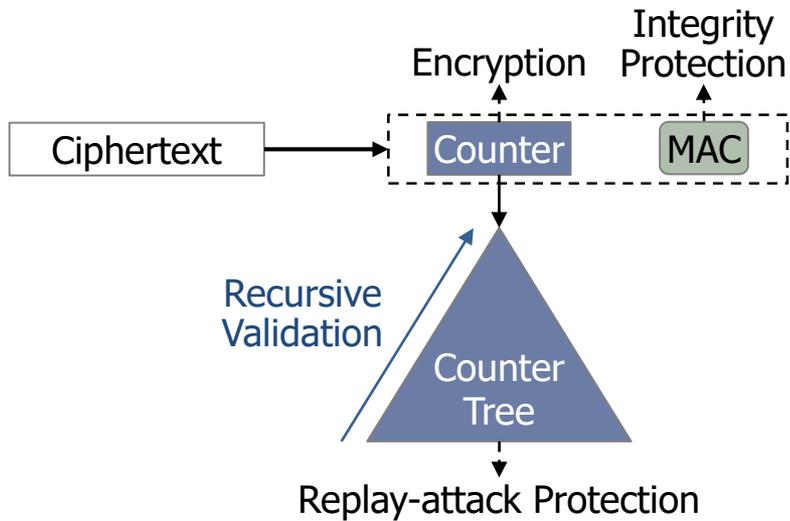
# Counter-mode Memory Protection

---



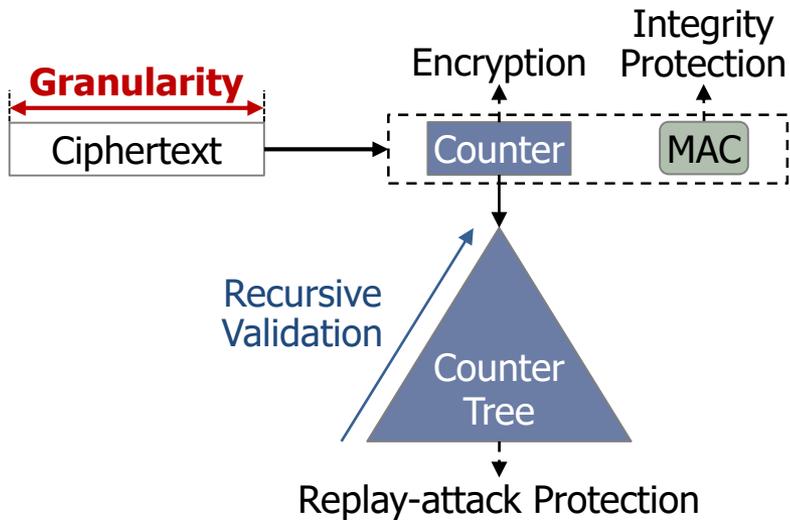
# Counter-mode Memory Protection

---



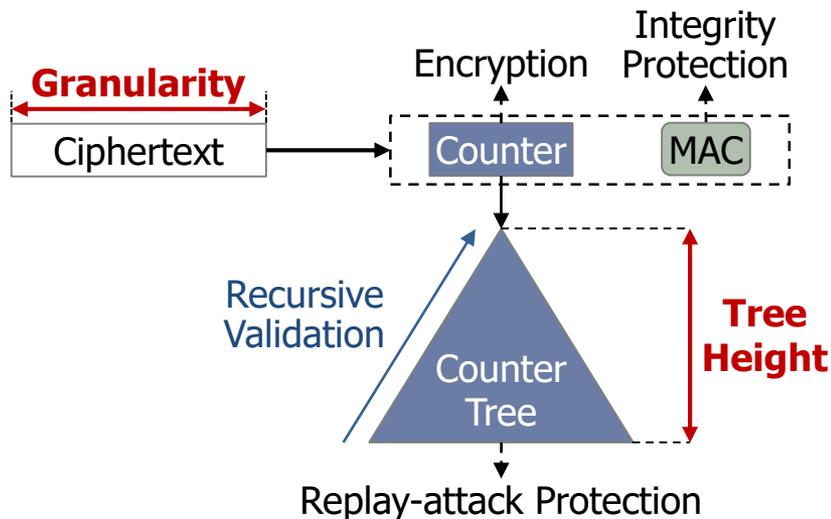
# Counter-mode Memory Protection

- Critical factors of memory protection
  - Amount of counters and MACs: Granularity



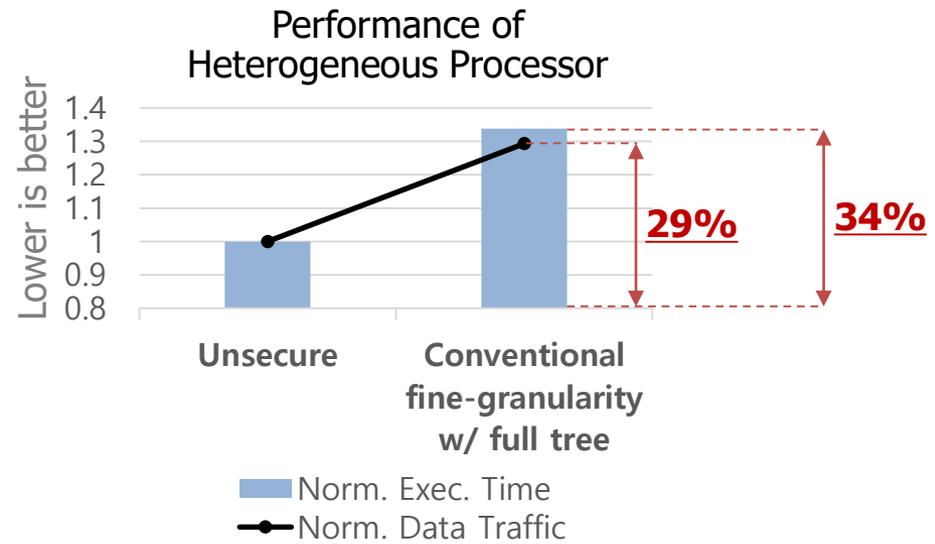
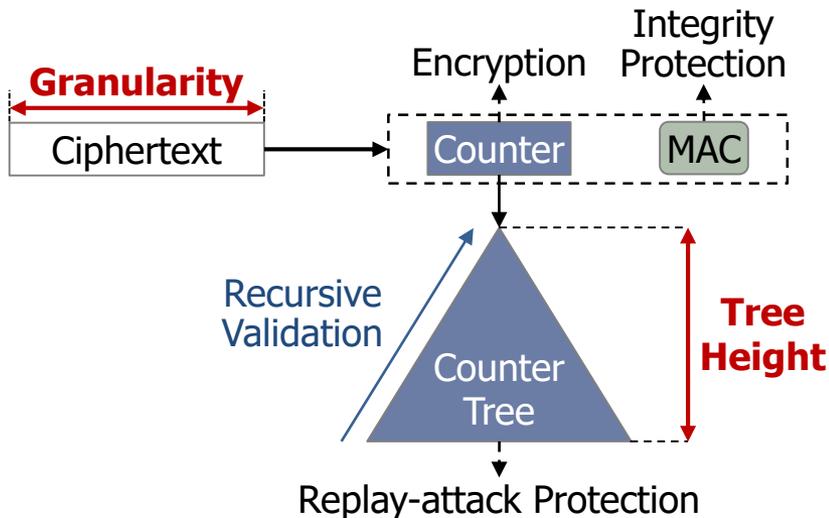
# Counter-mode Memory Protection

- Critical factors of memory protection
  - Amount of counters and MACs: Granularity
  - Overhead of recursive validation: Height of integrity tree



# Counter-mode Memory Protection

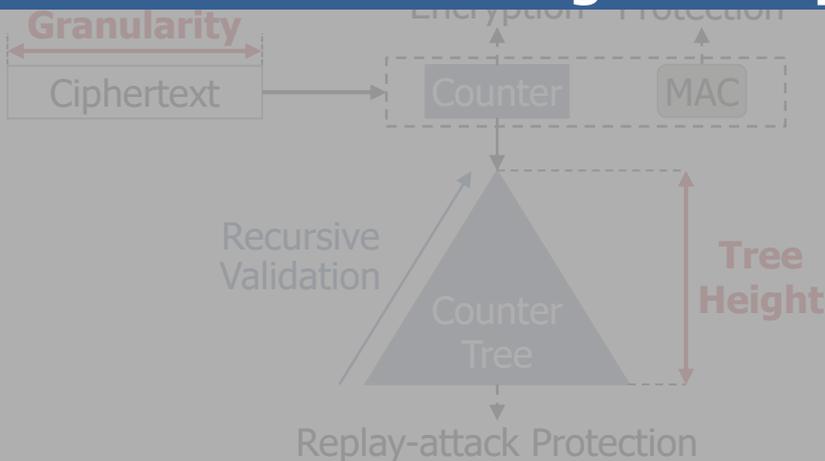
- Critical factors of memory protection
  - Amount of counters and MACs: Granularity
  - Overhead of recursive validation: Height of integrity tree
- 34% delay with 29% data traffic increment



# Counter-mode Memory Protection

- Critical factors of memory protection
  - Amount of counters and MACs: Granularity
  - Overhead of recursive validation: Height of integrity tree
- 34% delay with 29% data traffic increment

**Significant overhead caused by the conventional 64B-granular protection with a full integrity tree**

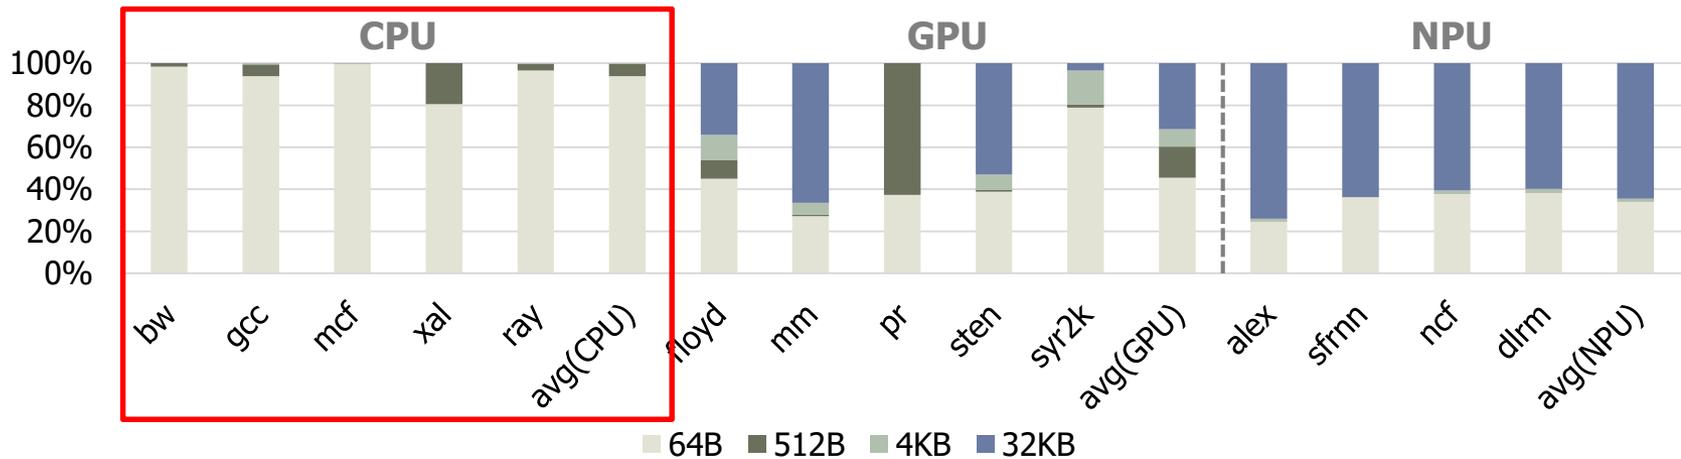


# Diverse Access Granularity

---

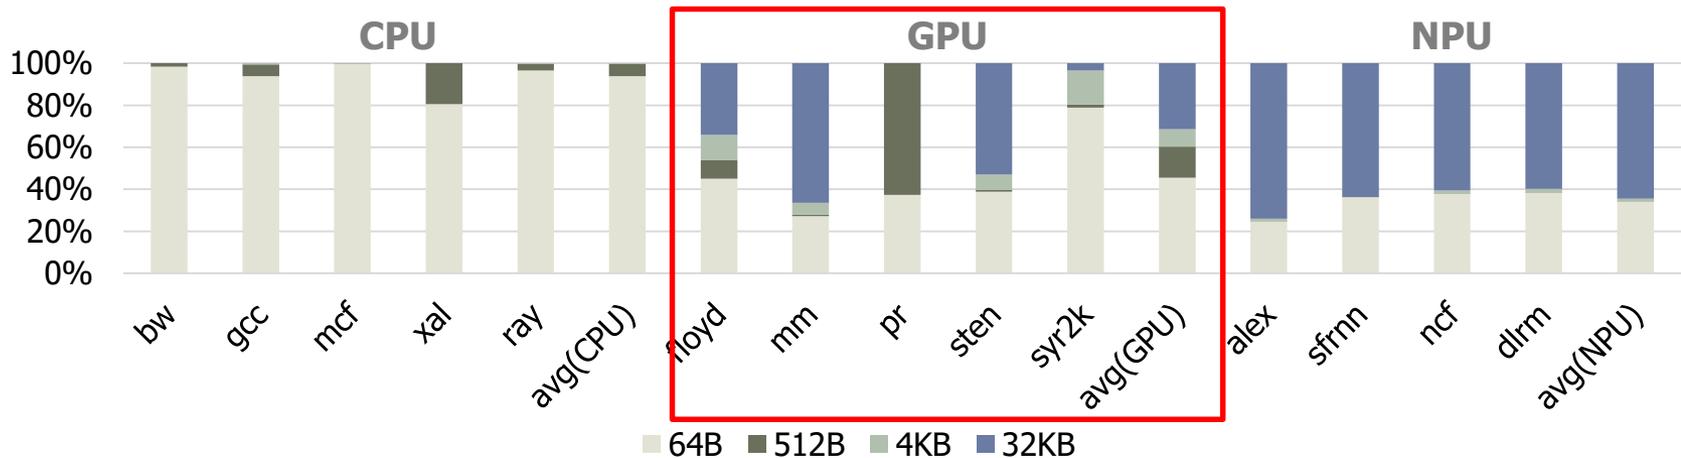
# Diverse Access Granularity

- Major access chunks (consecutive access blocks)
  - Fine-grained (64B): CPU



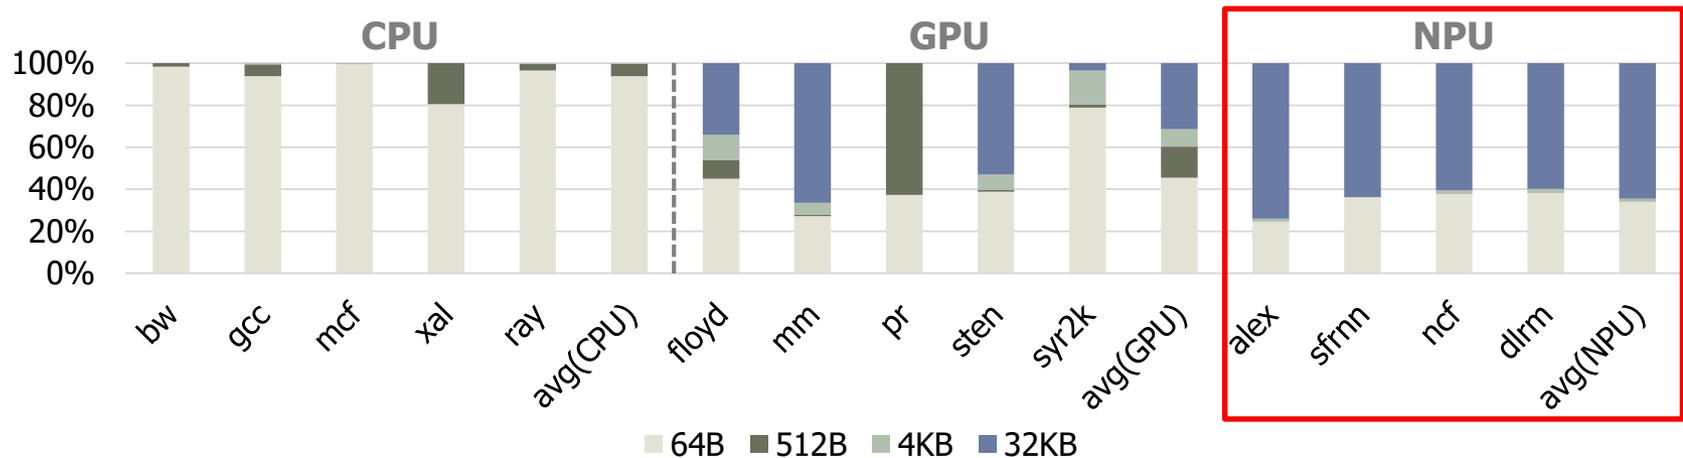
# Diverse Access Granularity

- Major access chunks (consecutive access blocks)
  - Fine-grained (64B): CPU
  - Medium-grained (512B, 4KB): GPU



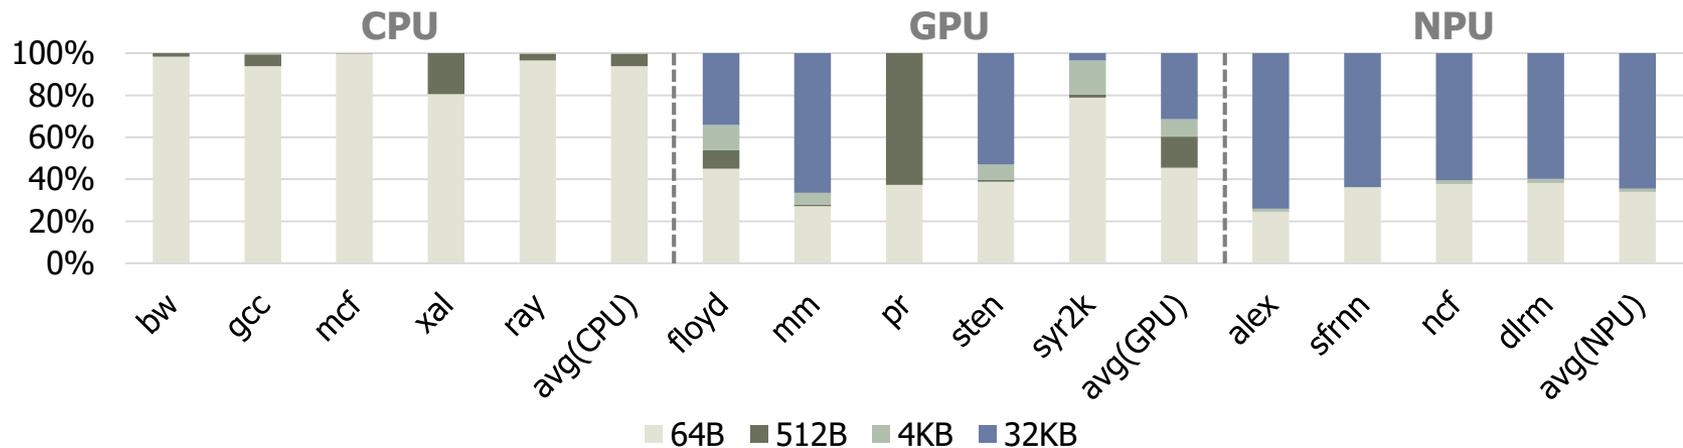
# Diverse Access Granularity

- Major access chunks (consecutive access blocks)
  - Fine-grained (64B): CPU
  - Medium-grained (512B, 4KB): GPU
  - Coarse-grained (32KB): NPU



# Diverse Access Granularity

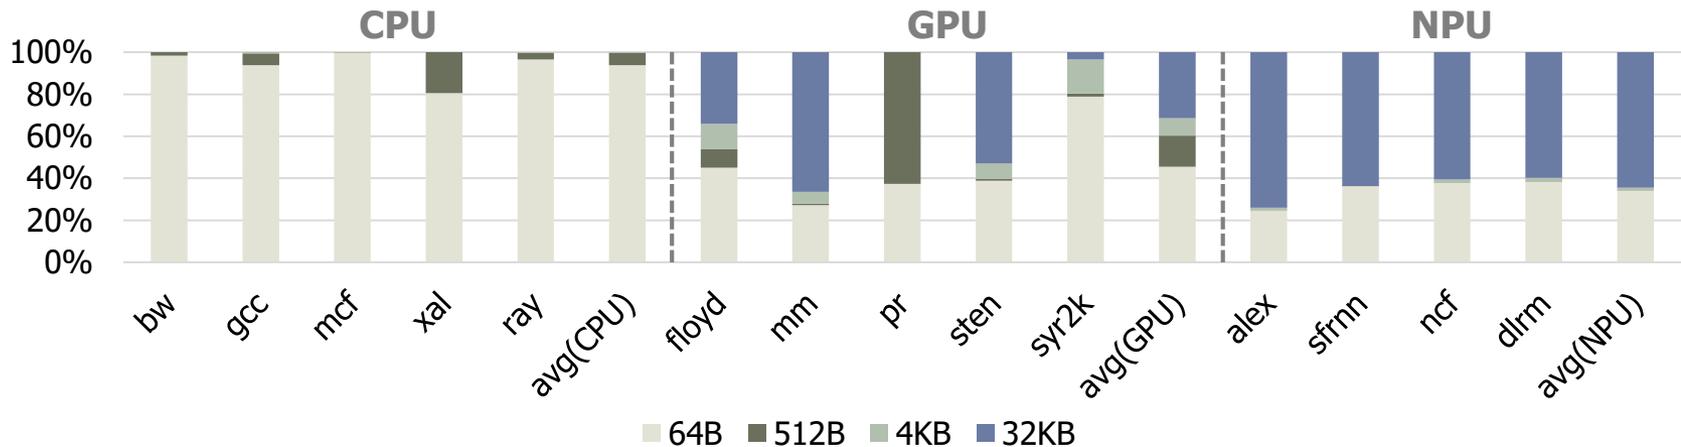
- Major access chunks (consecutive access blocks)
  - Fine-grained (64B): CPU
  - Medium-grained (512B, 4KB): GPU
  - Coarse-grained (32KB): NPU



Matching security granularity to access granularity

# Diverse Access Granularity

- Major access chunks (consecutive access blocks)
  - Fine-grained (64B): CPU
  - Medium-grained (512B, 4KB): GPU
  - Coarse-grained (32KB): NPU



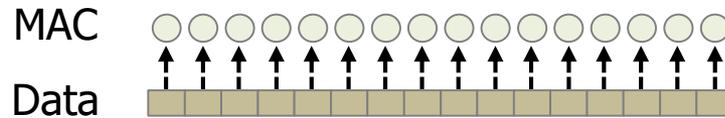
**Matching security granularity to access granularity**  
→ Requirement: Multi-granularity for MACs and counters

# Multi-granular MAC

---

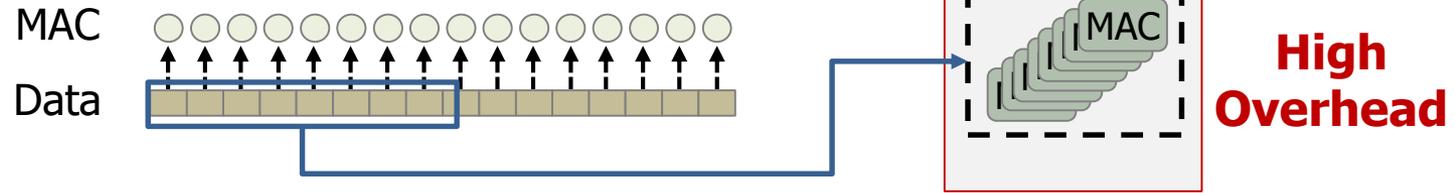


**Conventional  
Fine-granular  
MAC**

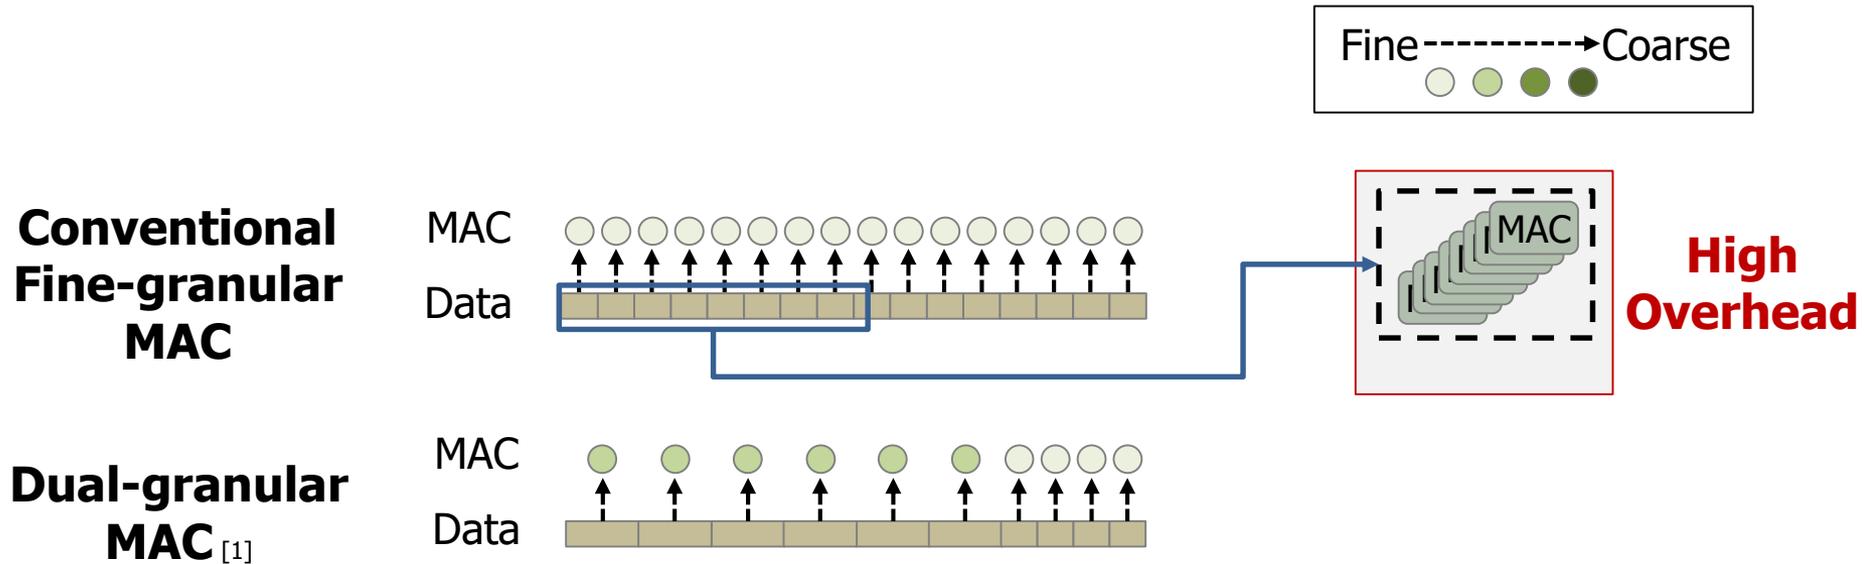


# Multi-granular MAC

## Conventional Fine-granular MAC

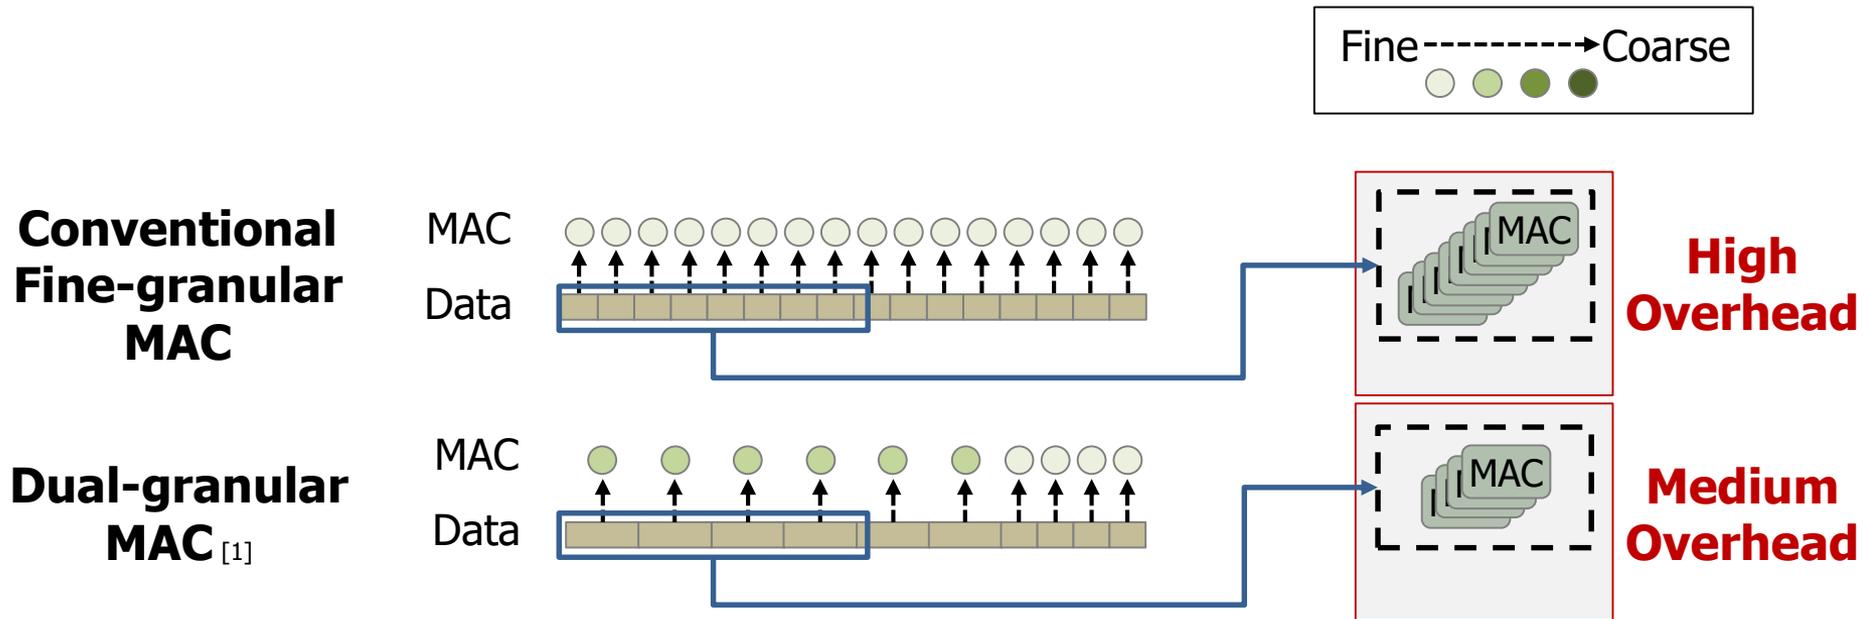


# Multi-granular MAC



[1] Adaptive Security Support for Heterogeneous Memory on GPUs (HPCA 2022)

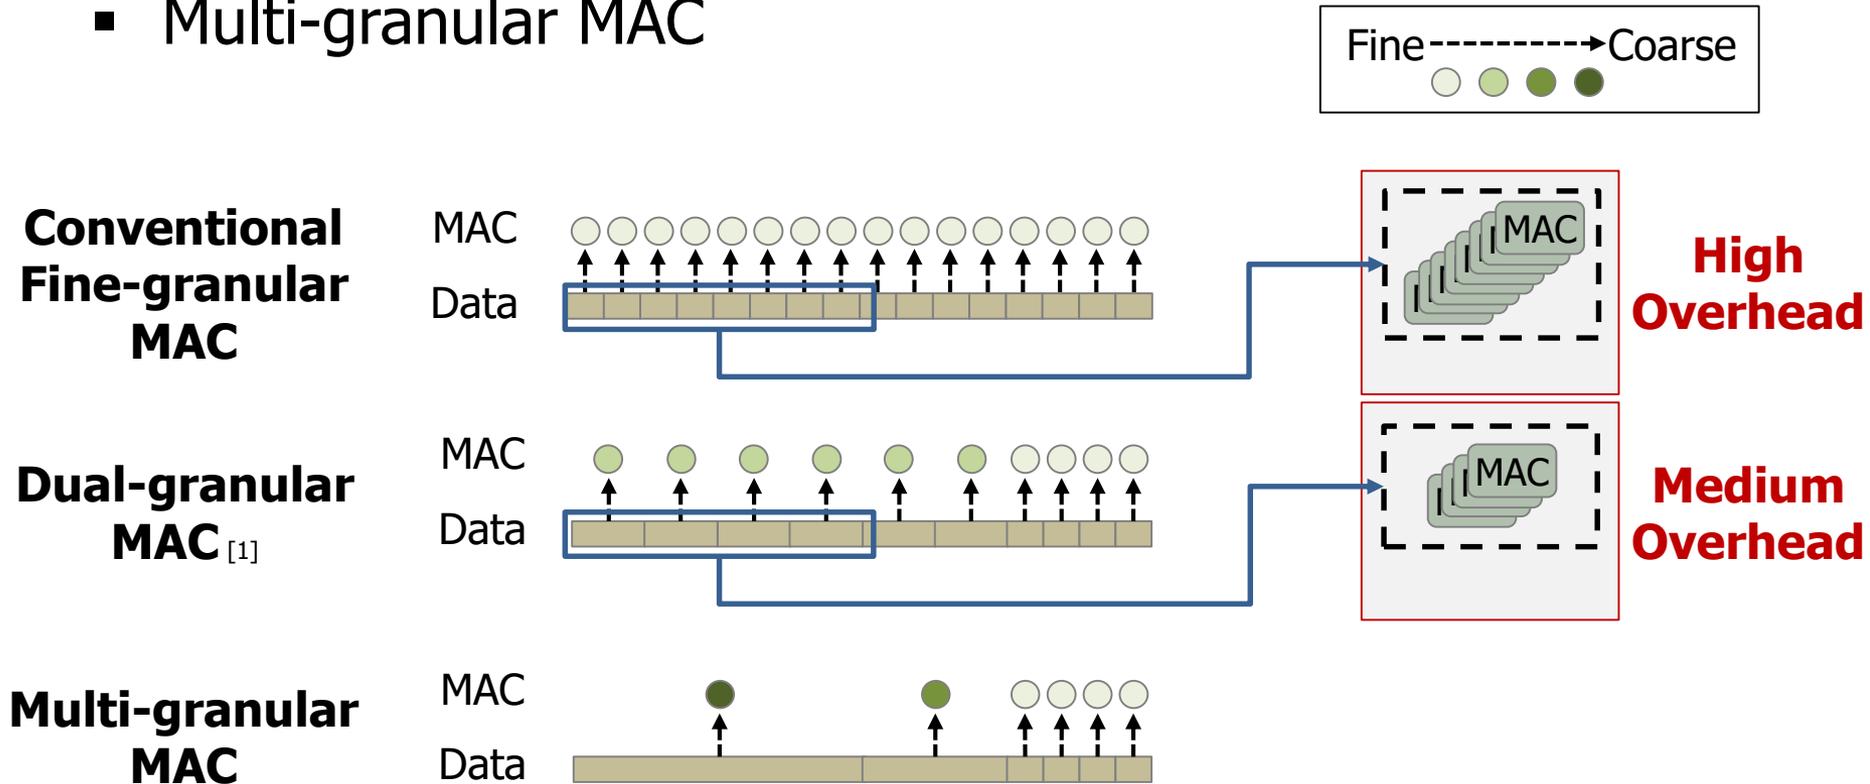
# Multi-granular MAC



[1] Adaptive Security Support for Heterogeneous Memory on GPUs (HPCA 2022)

# Multi-granular MAC

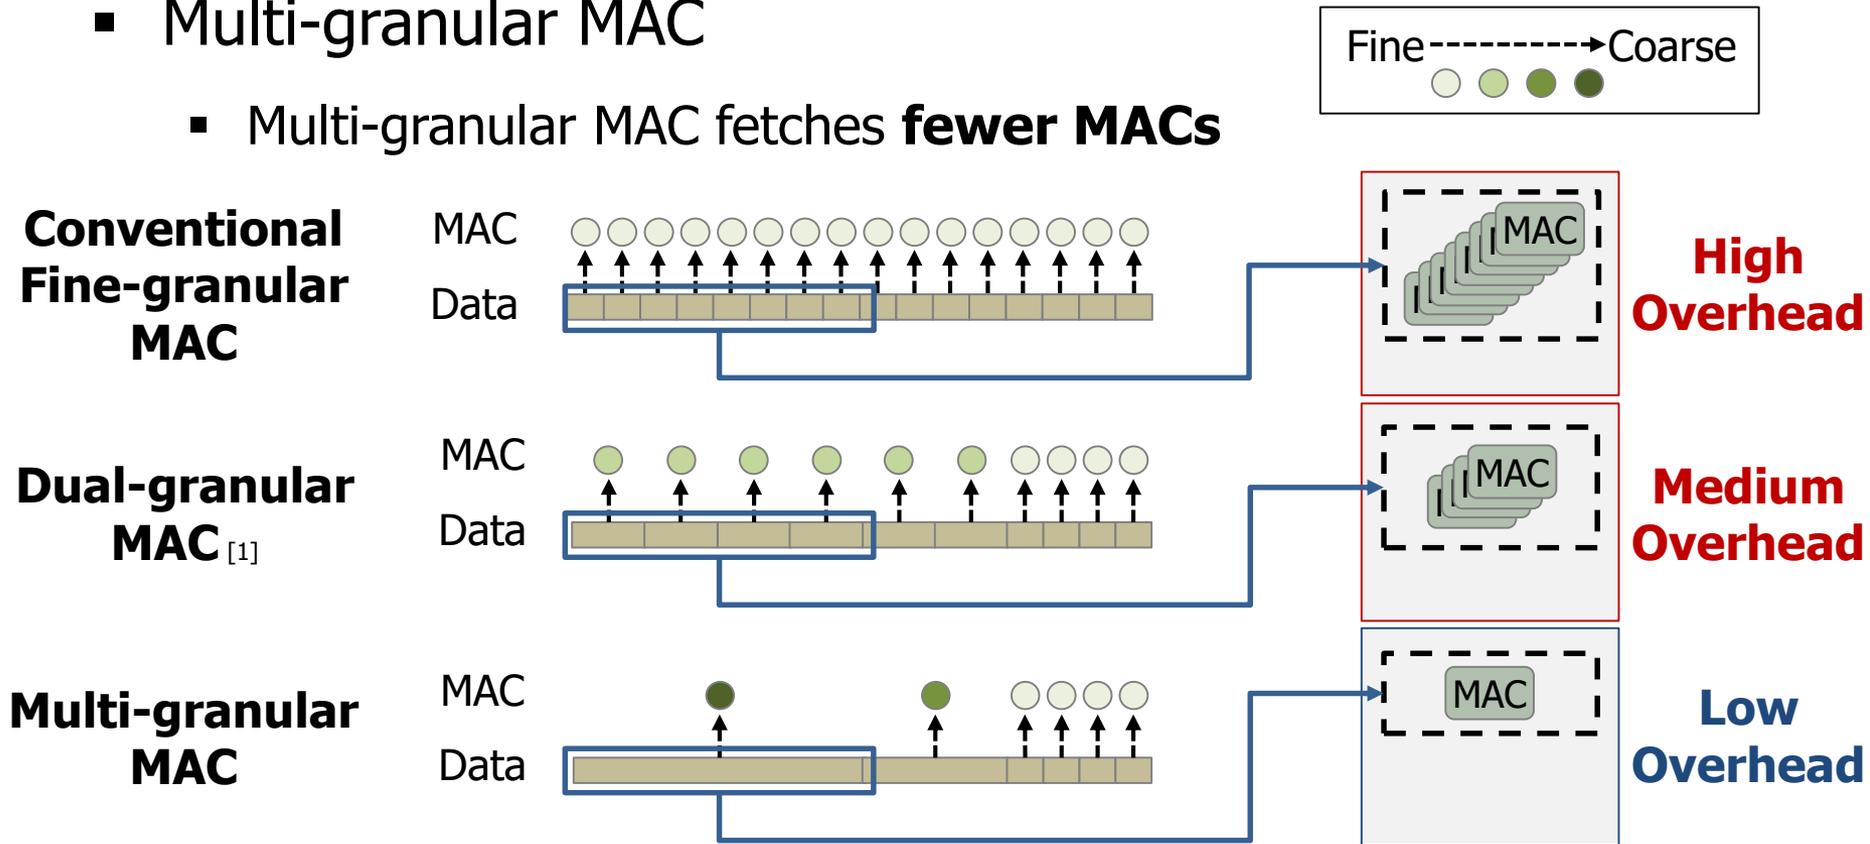
- Multi-granular MAC



[1] Adaptive Security Support for Heterogeneous Memory on GPUs (HPCA 2022)

# Multi-granular MAC

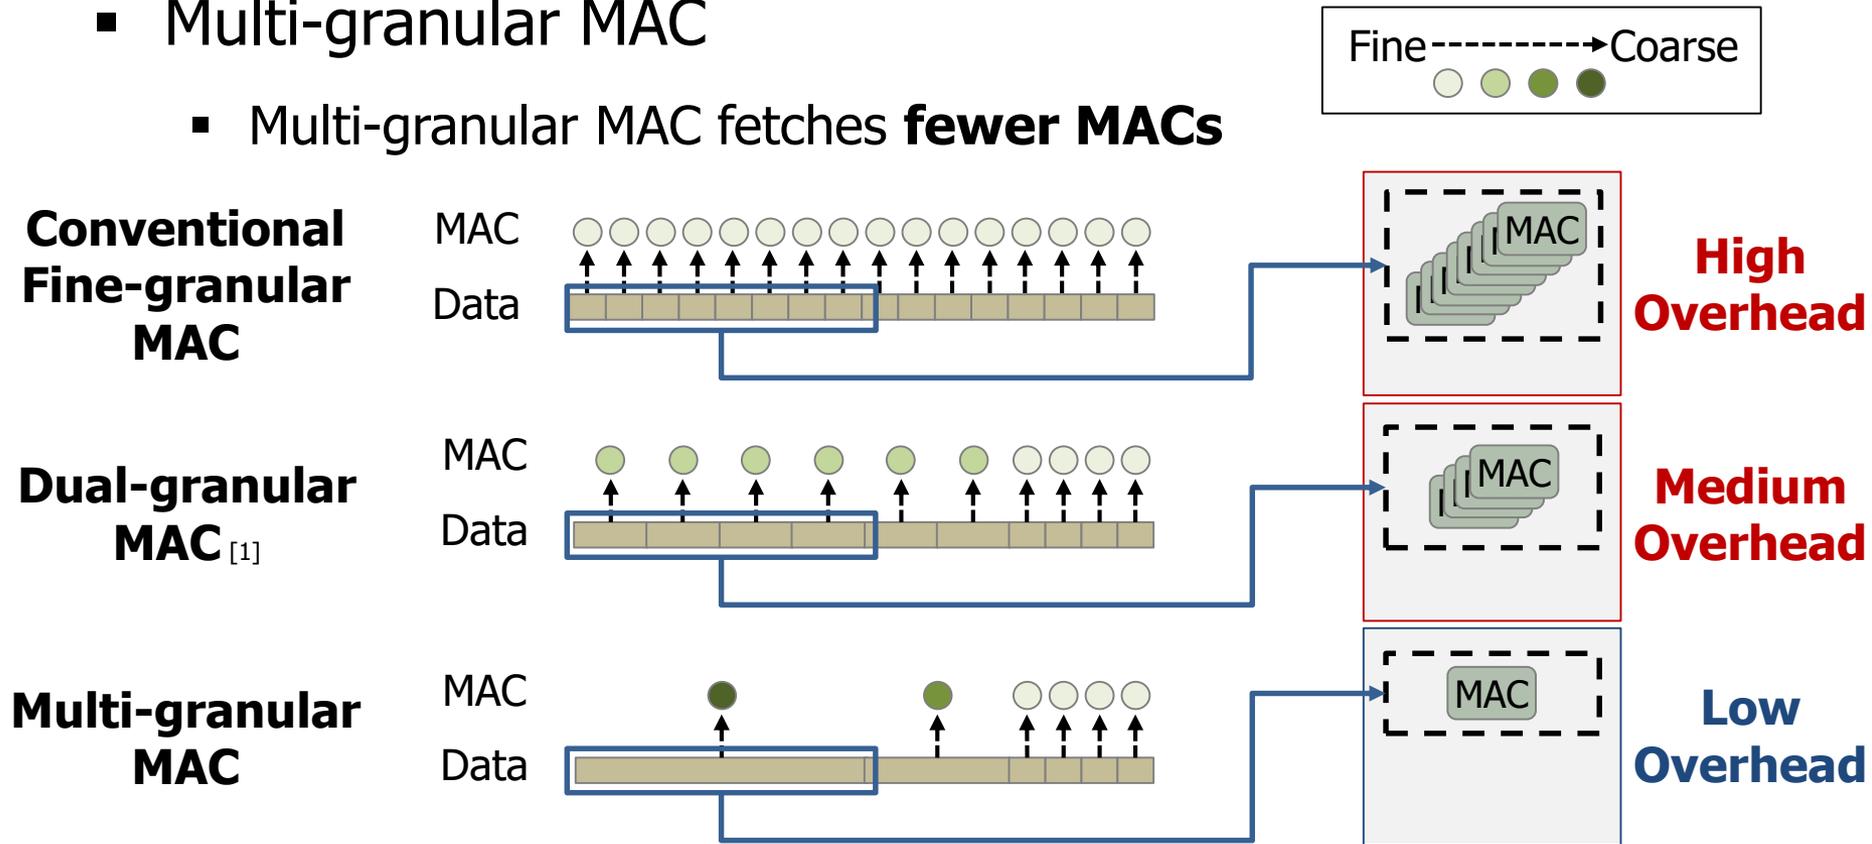
- Multi-granular MAC
  - Multi-granular MAC fetches **fewer MACs**



[1] Adaptive Security Support for Heterogeneous Memory on GPUs (HPCA 2022)

# Multi-granular MAC

- Multi-granular MAC
  - Multi-granular MAC fetches **fewer MACs**

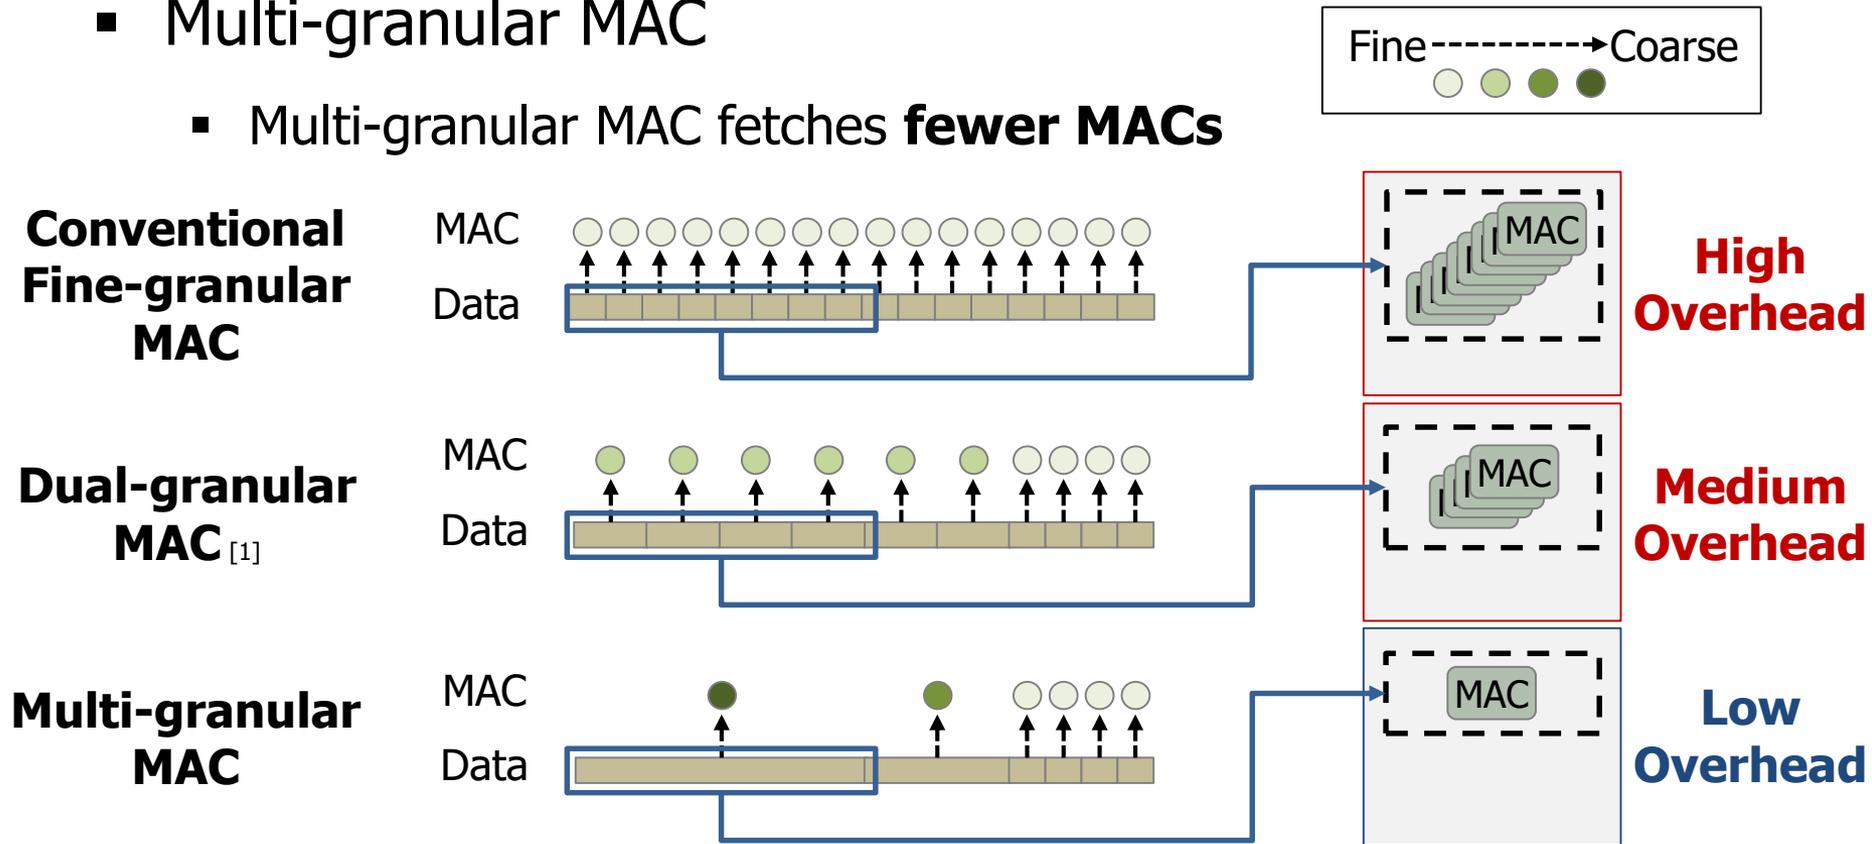


**Multi-granular MAC → Only managing the security granularity**

[1] Adaptive Security Support for Heterogeneous Memory on GPUs (HPCA 2022)

# Multi-granular MAC

- Multi-granular MAC
  - Multi-granular MAC fetches **fewer MACs**



**Multi-granular MAC → Only managing the security granularity  
What about multi-granular counters?**

[1] Adaptive Security Support for Heterogeneous Memory on GPUs (HPCA 2022)

# Prior Multi-granular Counter

---

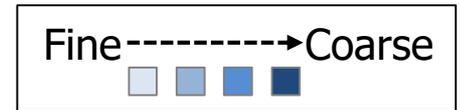
# Prior Multi-granular Counter

---

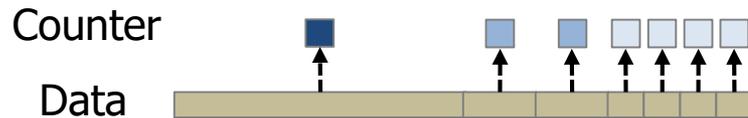
- Due to **integrity tree**, prior studies bypass integrity tree under specific conditions.

# Prior Multi-granular Counter

- Due to **integrity tree**, prior studies bypass integrity tree under specific conditions.



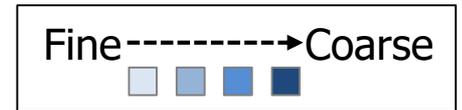
## Common Counters [1]



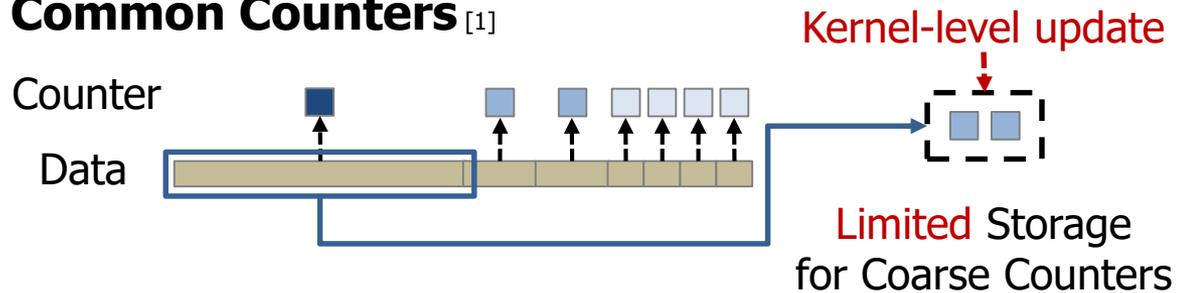
[1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)

# Prior Multi-granular Counter

- Due to **integrity tree**, prior studies bypass integrity tree under specific conditions.



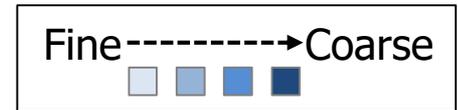
## Common Counters [1]



[1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)

# Prior Multi-granular Counter

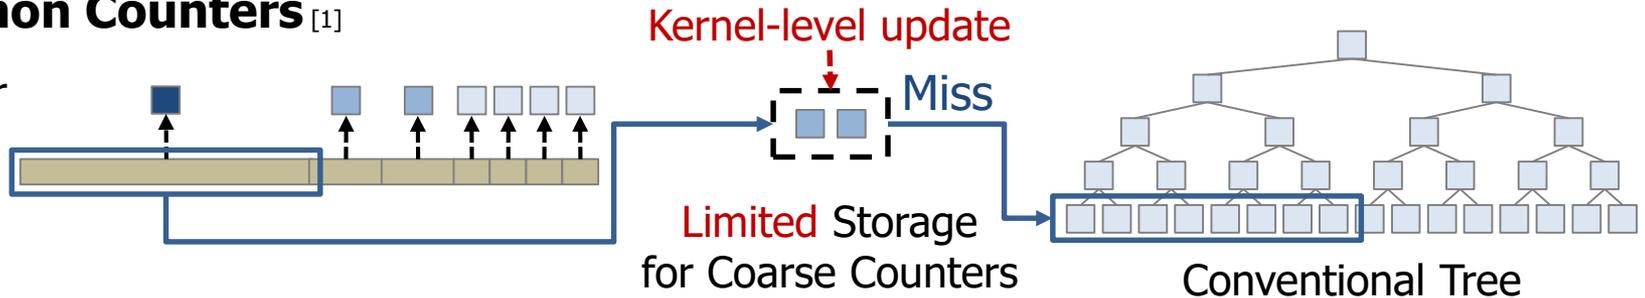
- Due to **integrity tree**, prior studies bypass integrity tree under specific conditions.



## Common Counters [1]

Counter

Data



[1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)

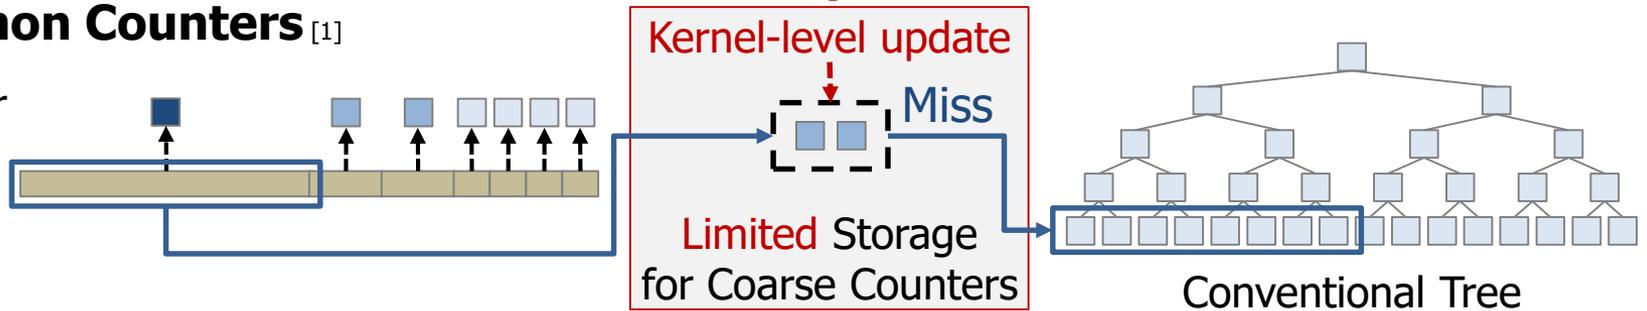
# Prior Multi-granular Counter

- Due to **integrity tree**, prior studies bypass integrity tree under specific conditions.

## Common Counters [1]

Counter

Data



[1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)

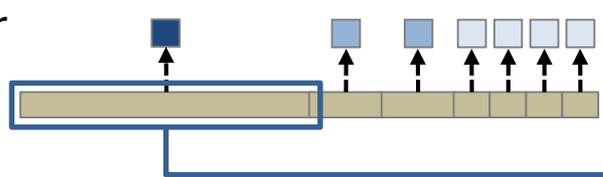
# Prior Multi-granular Counter

- Due to **integrity tree**, prior studies bypass integrity tree under specific conditions.

## Common Counters [1]

Counter

Data

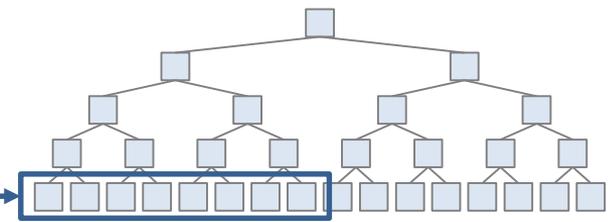
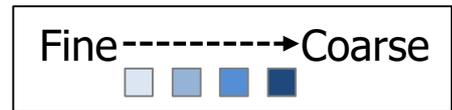


**GPU-specific**

Kernel-level update

Miss

Limited Storage  
for Coarse Counters



Conventional Tree

## S/W-managed Studies [2-5]

Counter

Data



[1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)

[2] TNPU: Supporting Trusted Execution with Tree-less Integrity Protection for Neural Processing Unit (HPCA 2022)

[3] MGX: Near-zero Overhead Memory Protection for Data-intensive Accelerators (ISCA 2022)

[4] GuardNN: Secure Accelerator Architecture for Privacy-preserving Deep Learning (DAC 2022)

[5] SoftVN: Efficient Memory Protection via Software-provided Version Numbers (ISCA 2022)

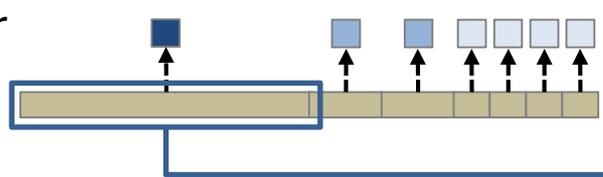
# Prior Multi-granular Counter

- Due to **integrity tree**, prior studies bypass integrity tree under specific conditions.

## Common Counters [1]

Counter

Data

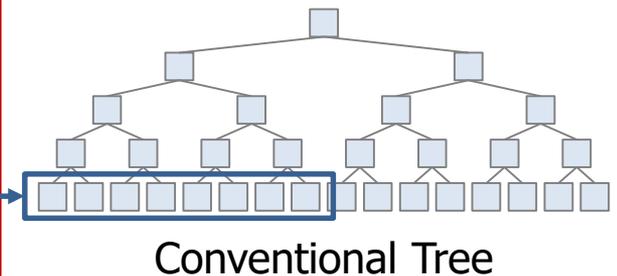
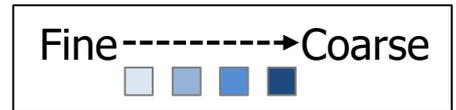


**GPU-specific**

Kernel-level update

Miss

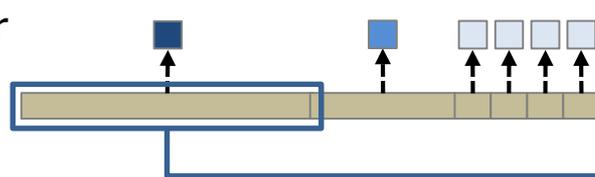
Limited Storage  
for Coarse Counters



## S/W-managed Studies [2-5]

Counter

Data



S/W-managed  
Storage

for Coarse Counters

[1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)

[2] TNPU: Supporting Trusted Execution with Tree-less Integrity Protection for Neural Processing Unit (HPCA 2022)

[3] MGX: Near-zero Overhead Memory Protection for Data-intensive Accelerators (ISCA 2022)

[4] GuardNN: Secure Accelerator Architecture for Privacy-preserving Deep Learning (DAC 2022)

[5] SoftVN: Efficient Memory Protection via Software-provided Version Numbers (ISCA 2022)

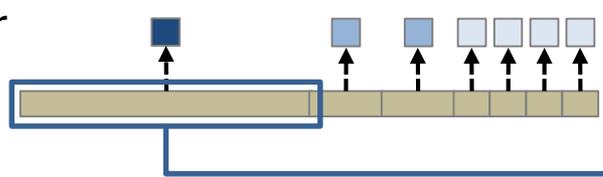
# Prior Multi-granular Counter

- Due to **integrity tree**, prior studies bypass integrity tree under specific conditions.

## Common Counters [1]

Counter

Data



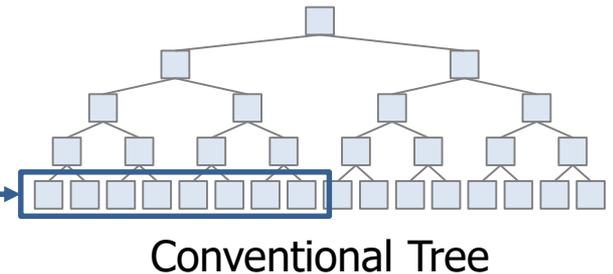
**GPU-specific**

Kernel-level update

Miss

Limited Storage  
for Coarse Counters

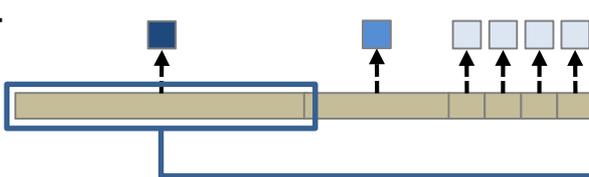
Fine----->Coarse



## S/W-managed Studies [2-5]

Counter

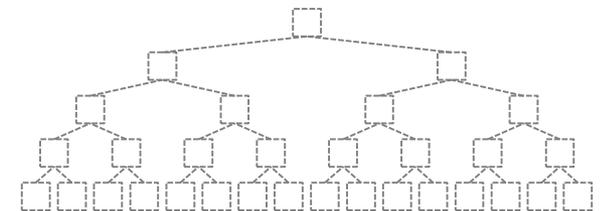
Data



S/W-managed

Storage

for Coarse Counters



[1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)

[2] TNPU: Supporting Trusted Execution with Tree-less Integrity Protection for Neural Processing Unit (HPCA 2022)

[3] MGX: Near-zero Overhead Memory Protection for Data-intensive Accelerators (ISCA 2022)

[4] GuardNN: Secure Accelerator Architecture for Privacy-preserving Deep Learning (DAC 2022)

[5] SoftVN: Efficient Memory Protection via Software-provided Version Numbers (ISCA 2022)

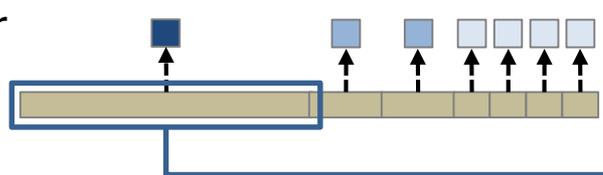
# Prior Multi-granular Counter

- Due to **integrity tree**, prior studies bypass integrity tree under specific conditions.

## Common Counters [1]

Counter

Data



## GPU-specific

Kernel-level update

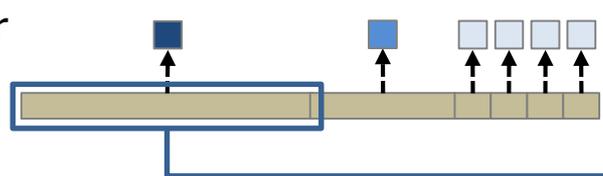
Miss

Limited Storage  
for Coarse Counters

## S/W-managed Studies [2-5]

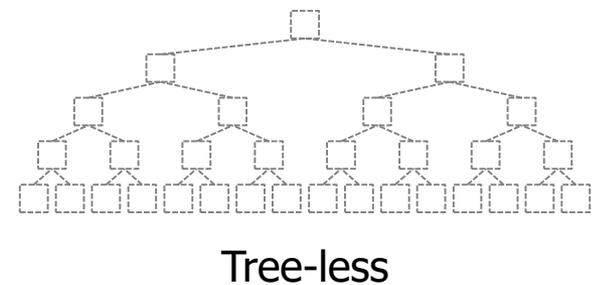
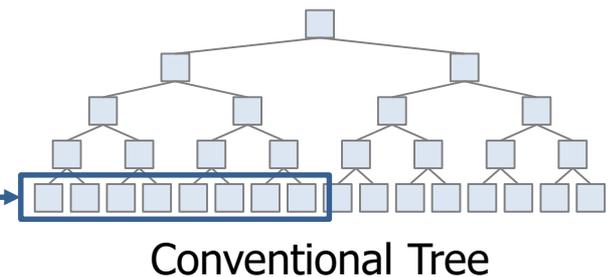
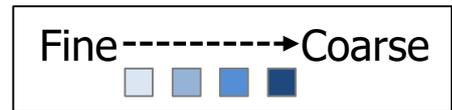
Counter

Data



## ML-specific

S/W-managed  
Storage  
for Coarse Counters



[1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)

[2] TNPU: Supporting Trusted Execution with Tree-less Integrity Protection for Neural Processing Unit (HPCA 2022)

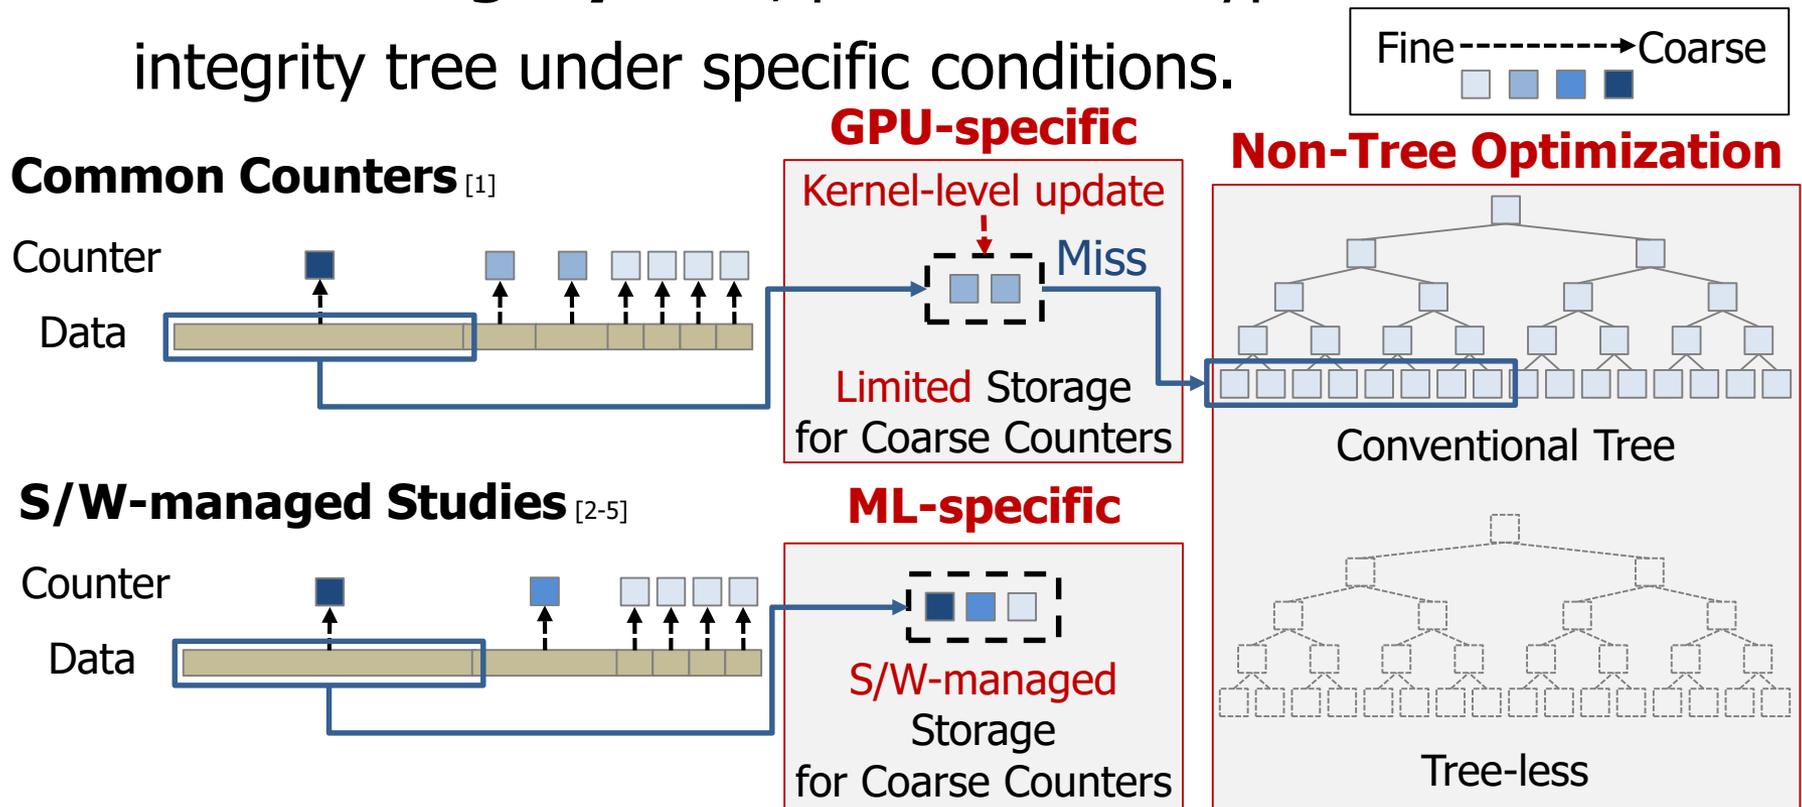
[3] MGX: Near-zero Overhead Memory Protection for Data-intensive Accelerators (ISCA 2022)

[4] GuardNN: Secure Accelerator Architecture for Privacy-preserving Deep Learning (DAC 2022)

[5] SoftVN: Efficient Memory Protection via Software-provided Version Numbers (ISCA 2022)

# Prior Multi-granular Counter

- Due to **integrity tree**, prior studies bypass integrity tree under specific conditions.



[1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)

[2] TNPU: Supporting Trusted Execution with Tree-less Integrity Protection for Neural Processing Unit (HPCA 2022)

[3] MGX: Near-zero Overhead Memory Protection for Data-intensive Accelerators (ISCA 2022)

[4] GuardNN: Secure Accelerator Architecture for Privacy-preserving Deep Learning (DAC 2022)

[5] SoftVN: Efficient Memory Protection via Software-provided Version Numbers (ISCA 2022)

# Prior Multi-granular Counter

- Due to **integrity tree**, prior studies bypass integrity tree under specific conditions.

## Common Counters [1]

Counter

Data



## GPU-specific

Kernel-level update

Miss



Fine -----> Coarse



## Non-Tree Optimization

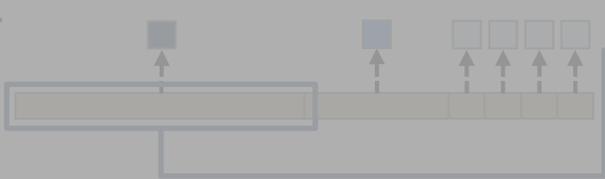


**Multi-granular counter integrity tree is necessary**

## S/W-managed Studies [2-5]

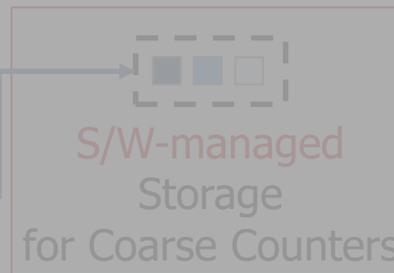
Counter

Data



## ML-specific

S/W-managed  
Storage  
for Coarse Counters



[1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)

[2] TNPU: Supporting Trusted Execution with Tree-less Integrity Protection for Neural Processing Unit (HPCA 2022)

[3] MGX: Near-zero Overhead Memory Protection for Data-intensive Accelerators (ISCA 2022)

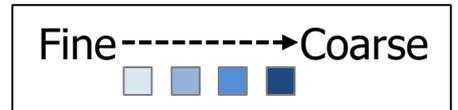
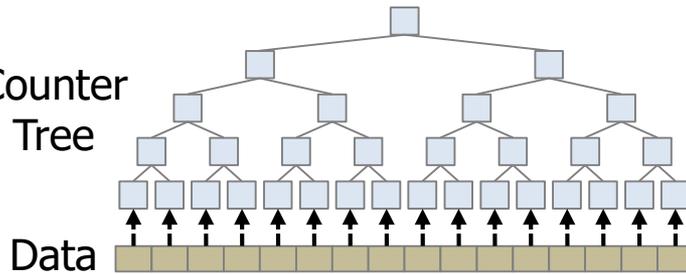
[4] GuardNN: Secure Accelerator Architecture for Privacy-preserving Deep Learning (DAC 2022)

[5] SoftVN: Efficient Memory Protection via Software-provided Version Numbers (ISCA 2022)

# Multi-granular Counter Integrity Tree

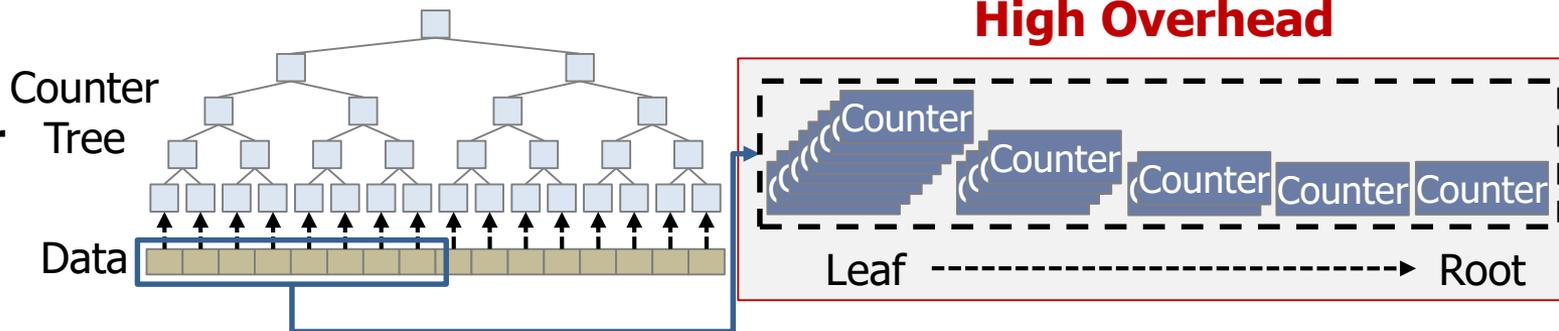
---

**Conventional Counter Tree**  
**Fine-granular Tree**



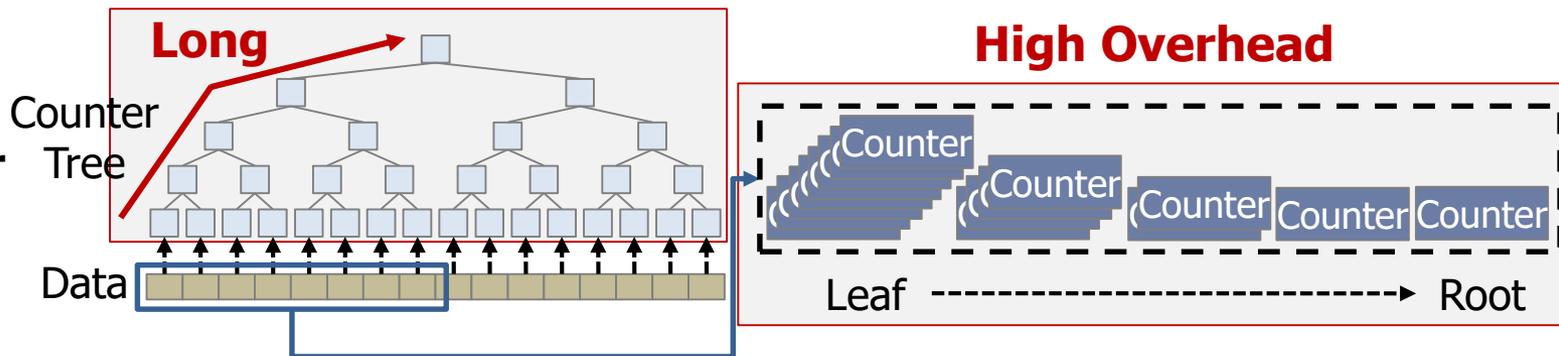
# Multi-granular Counter Integrity Tree

## Conventional Fine-granular Tree



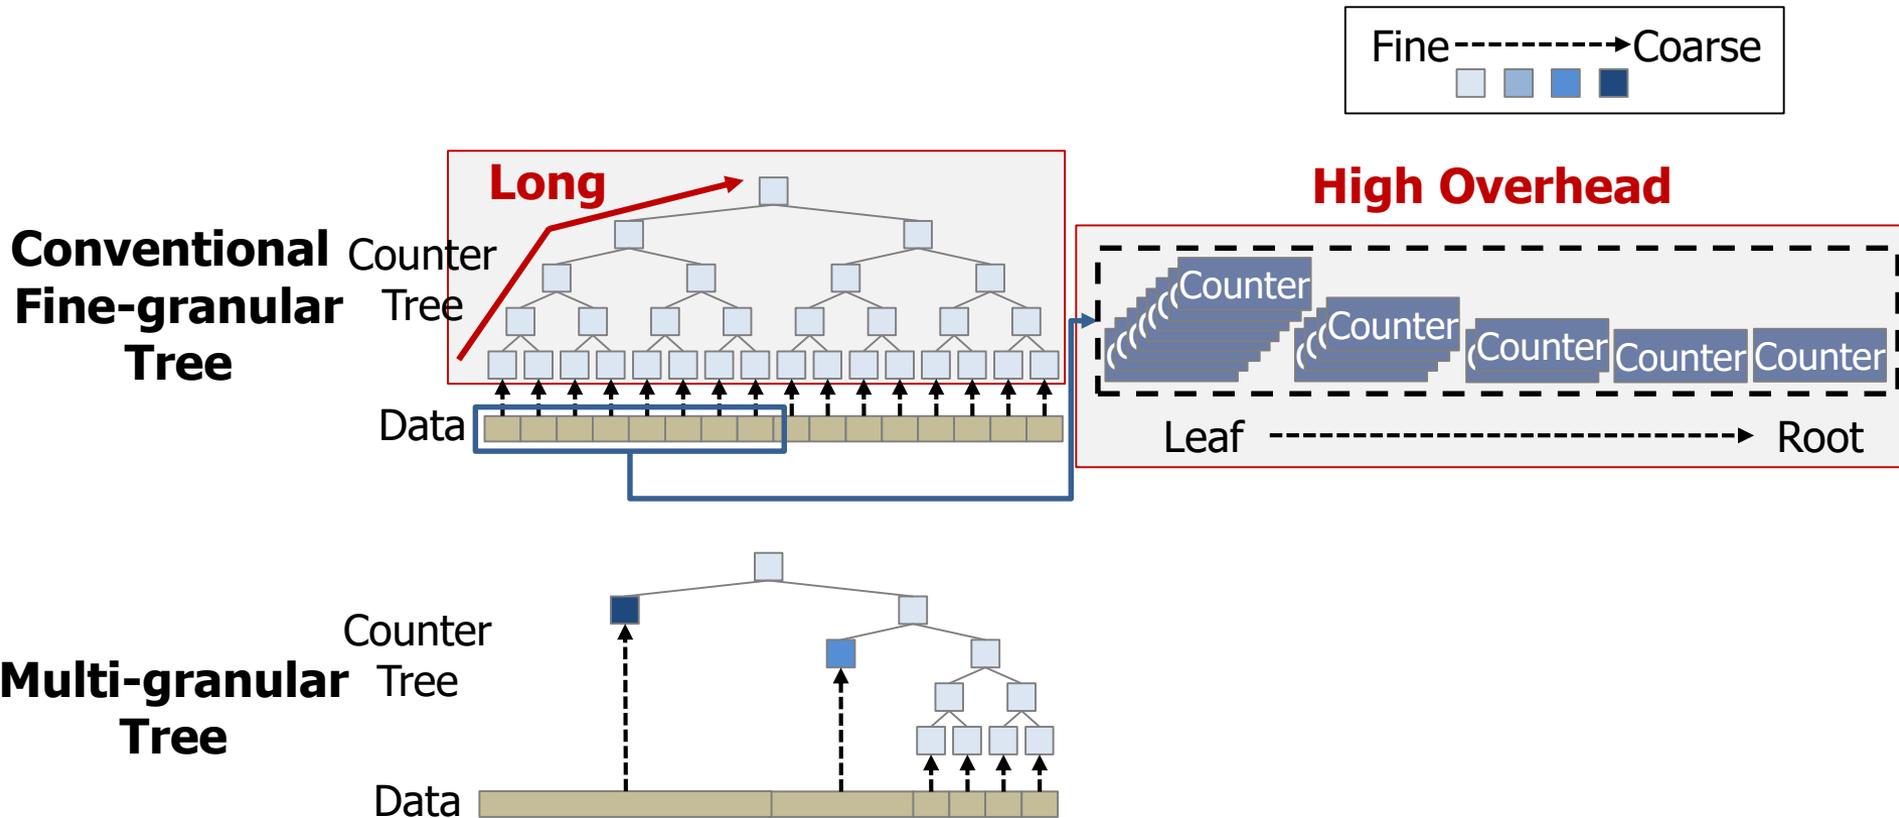
# Multi-granular Counter Integrity Tree

**Conventional  
Fine-granular  
Tree**



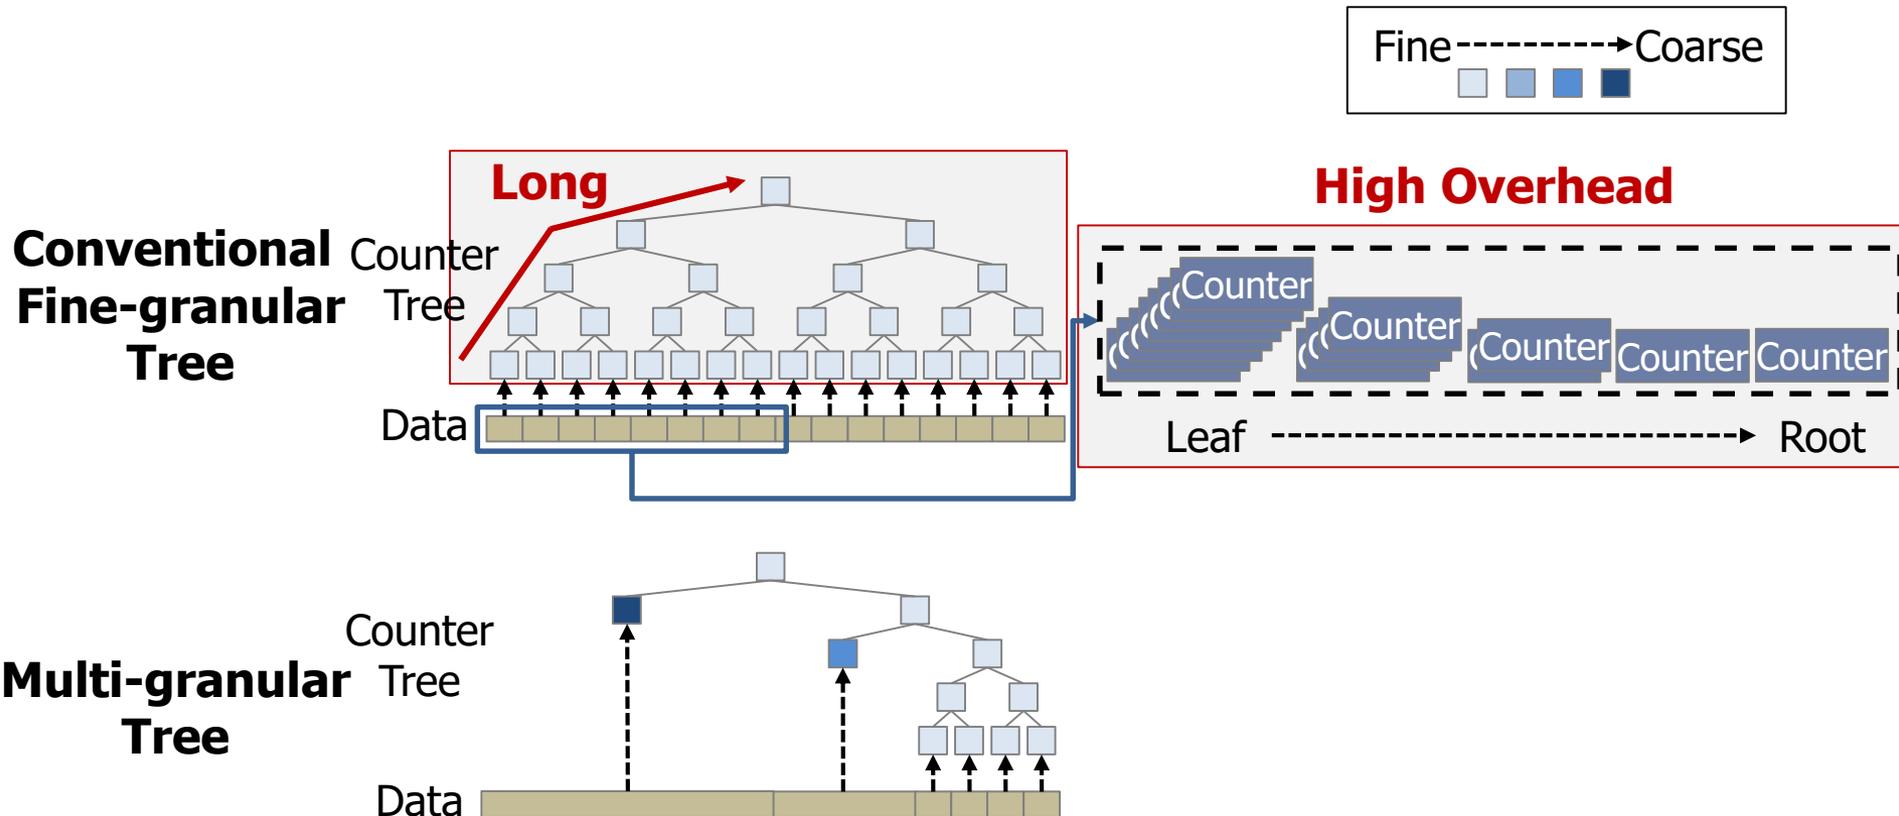
# Multi-granular Counter Integrity Tree

- Multi-granular tree



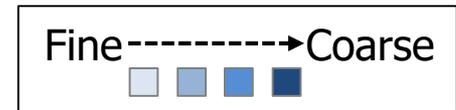
# Multi-granular Counter Integrity Tree

- Multi-granular tree
  - Counters w/ varying granularities are mapped to different levels

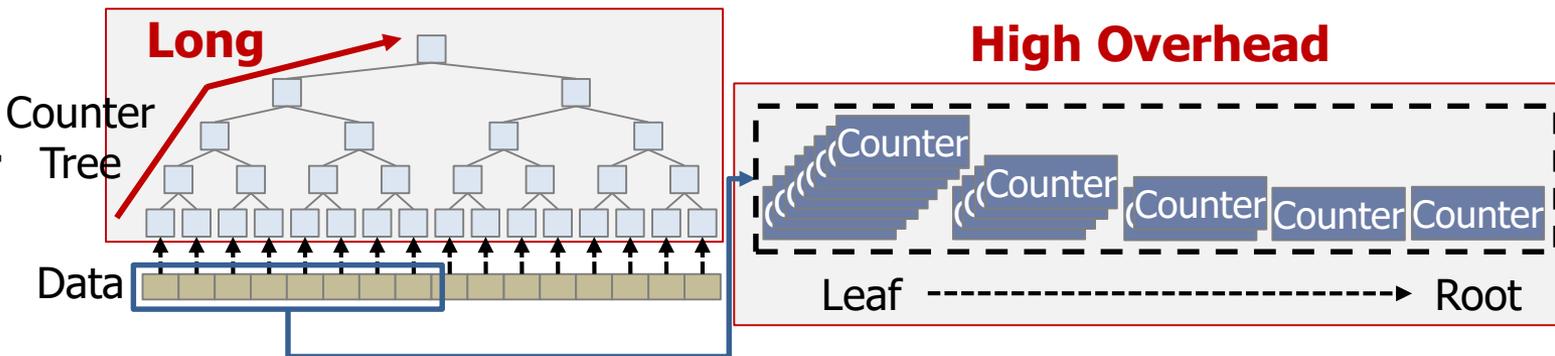


# Multi-granular Counter Integrity Tree

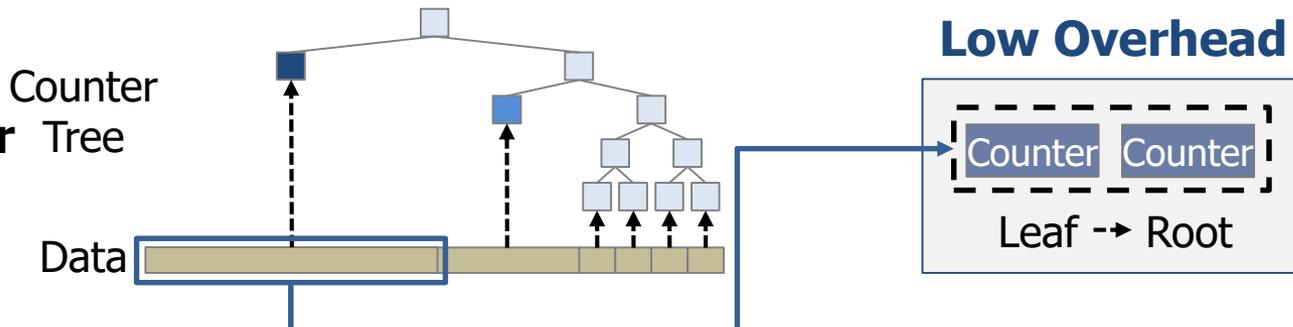
- Multi-granular tree
  - Counters w/ varying granularities are mapped to different levels
  - Fetches **fewer counters**



## Conventional Fine-granular Tree

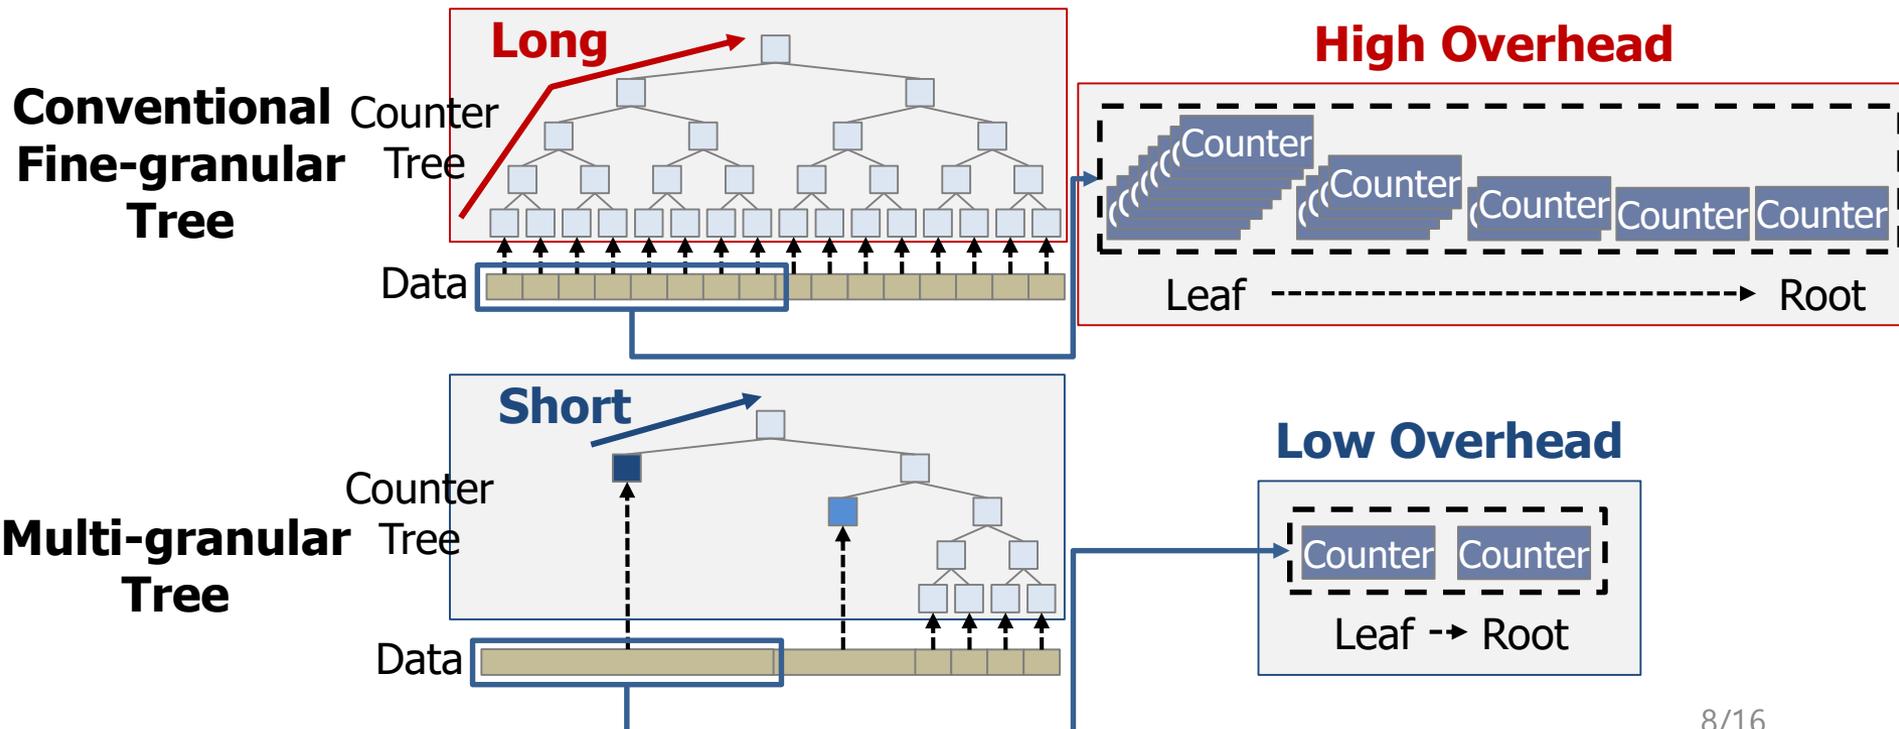
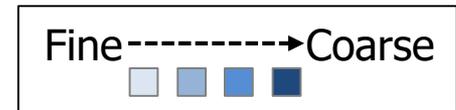


## Multi-granular Tree



# Multi-granular Counter Integrity Tree

- Multi-granular tree
  - Counters w/ varying granularities are mapped to different levels
  - Fetches **fewer counters**
  - **Shortens recursive validation path**



# Multi-granular MAC&Tree

---

- Multi-granular MAC&Tree

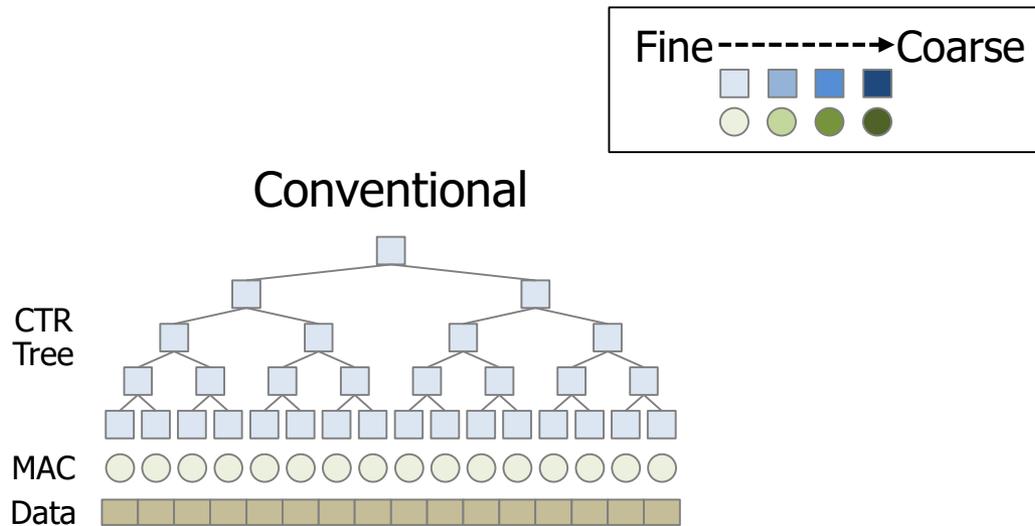
# Multi-granular MAC&Tree

---

- Multi-granular MAC&Tree
  - Dynamically supports multi-granular MACs and a counter tree

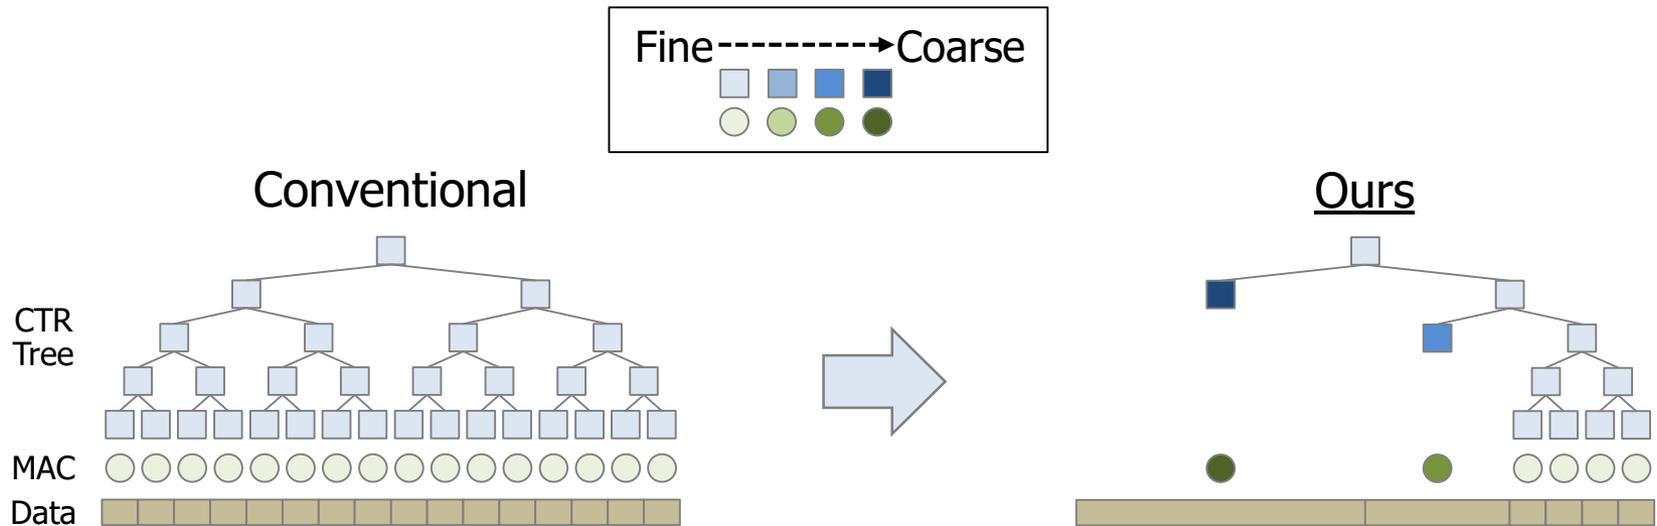
# Multi-granular MAC&Tree

- Multi-granular MAC&Tree
  - Dynamically supports multi-granular MACs and a counter tree



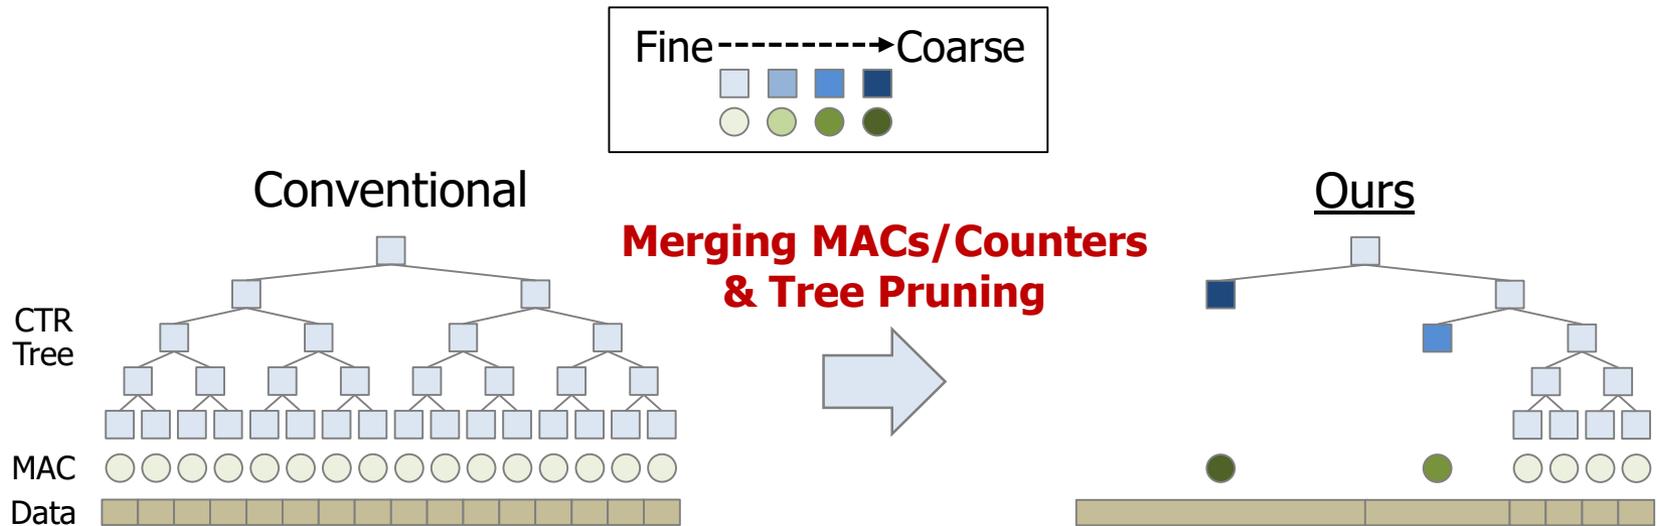
# Multi-granular MAC&Tree

- Multi-granular MAC&Tree
  - Dynamically supports multi-granular MACs and a counter tree



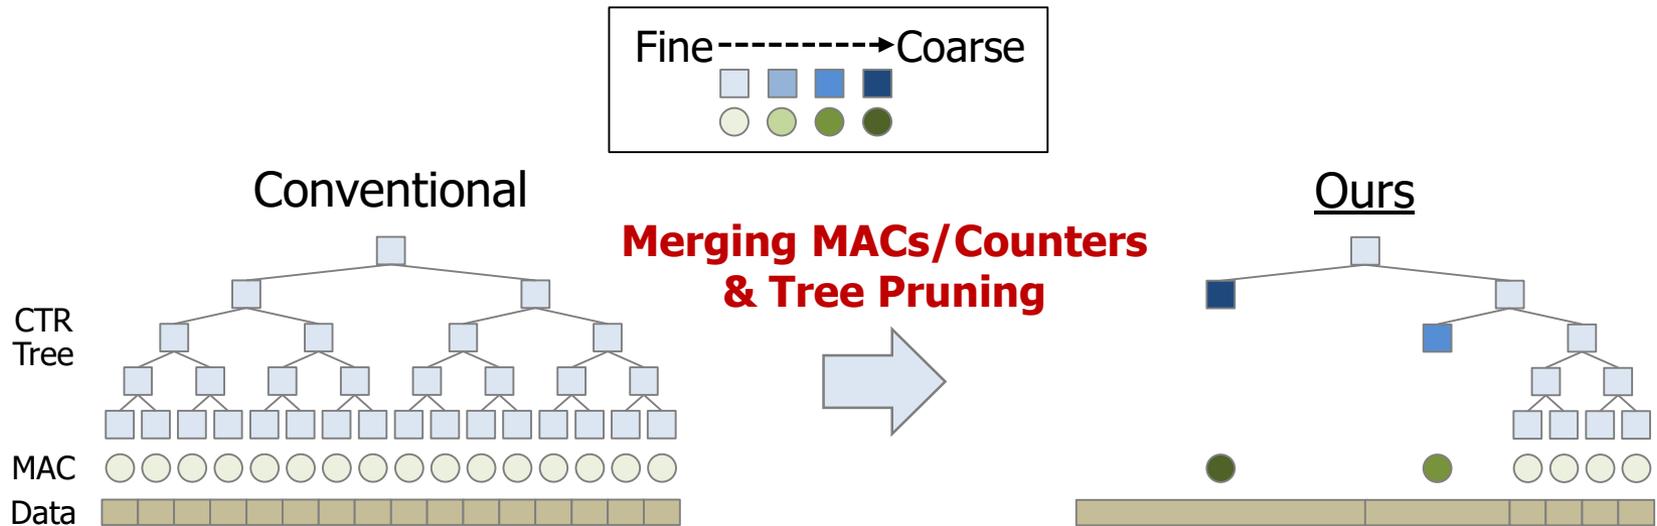
# Multi-granular MAC&Tree

- Multi-granular MAC&Tree
  - Dynamically supports multi-granular MACs and a counter tree
  - Key Idea: Merging MACs/counters & pruning a counter tree



# Multi-granular MAC&Tree

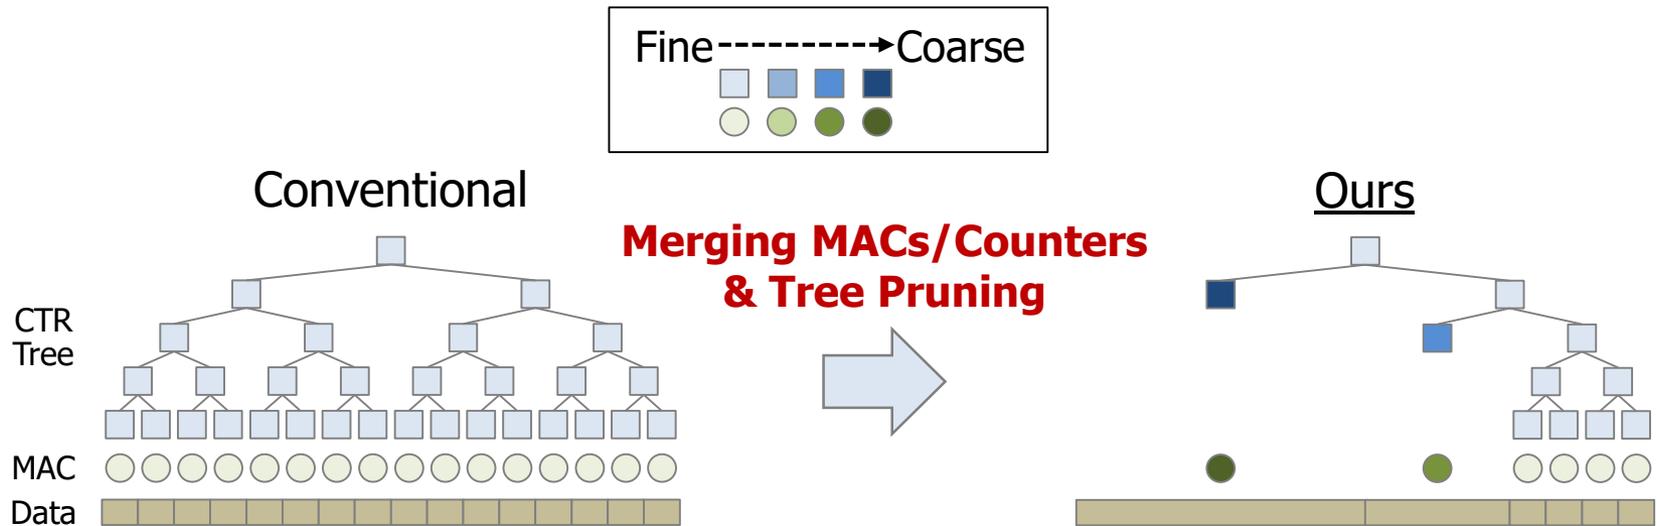
- Multi-granular MAC&Tree
  - Dynamically supports multi-granular MACs and a counter tree
  - Key Idea: Merging MACs/counters & pruning a counter tree



## 1. How to dynamically detect granularity

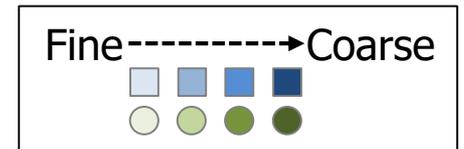
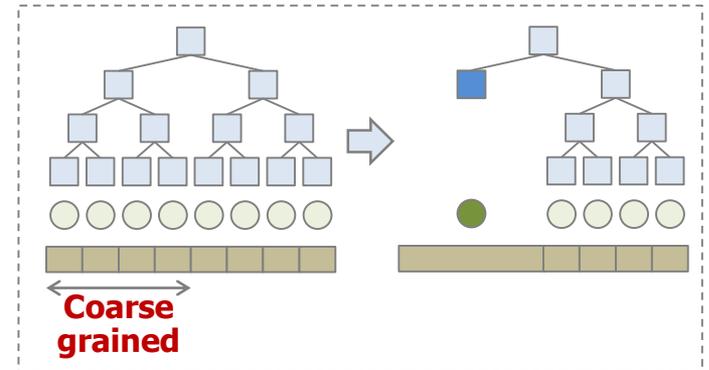
# Multi-granular MAC&Tree

- Multi-granular MAC&Tree
  - Dynamically supports multi-granular MACs and a counter tree
  - Key Idea: Merging MACs/counters & pruning a counter tree



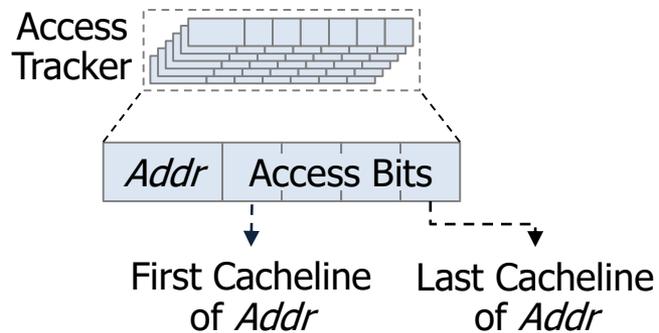
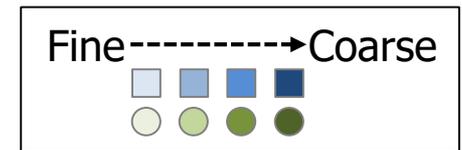
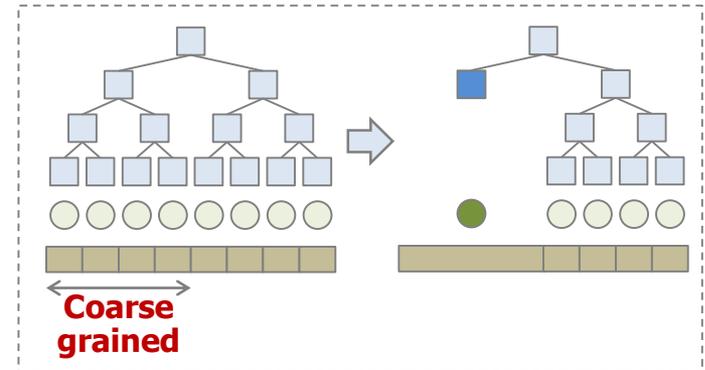
1. How to dynamically detect granularity
2. How to switch granularity

# Granularity Detection (Fine → Coarse)



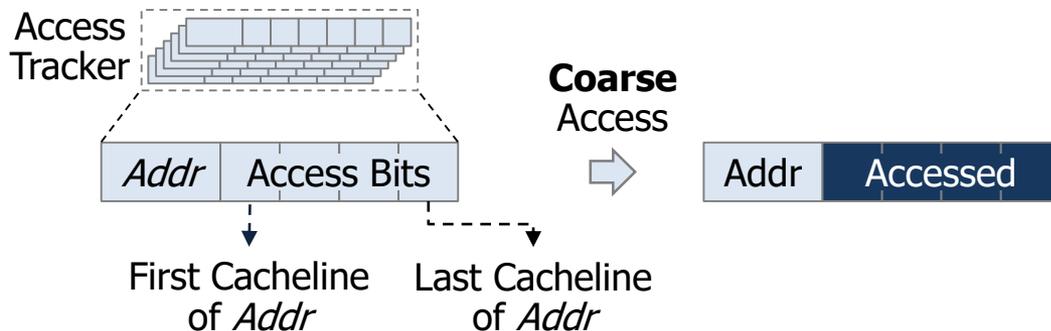
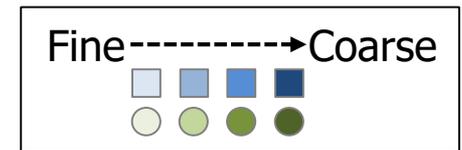
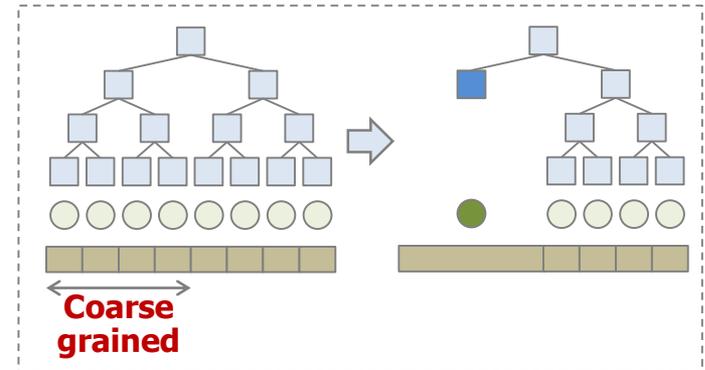
# Granularity Detection (Fine → Coarse)

- Access tracker
  - Records accessed addresses



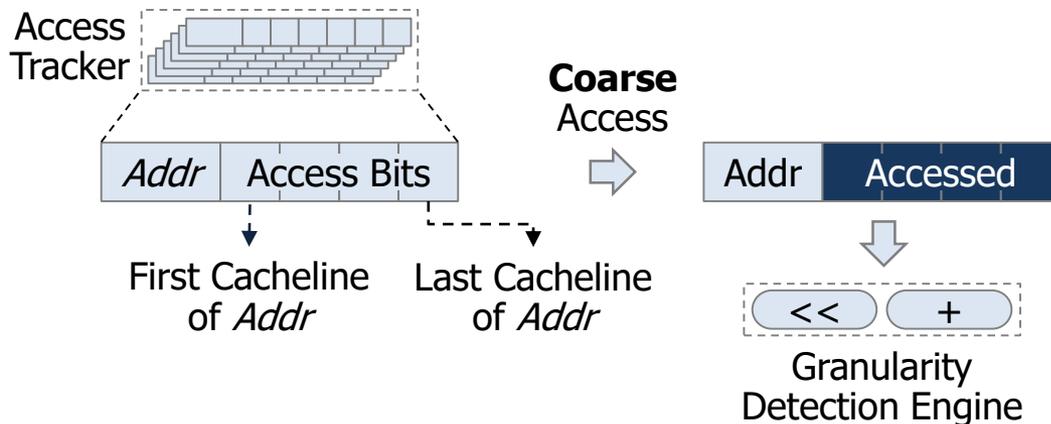
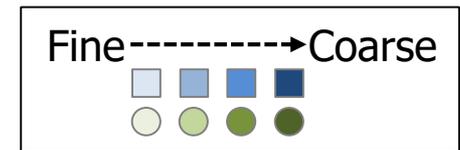
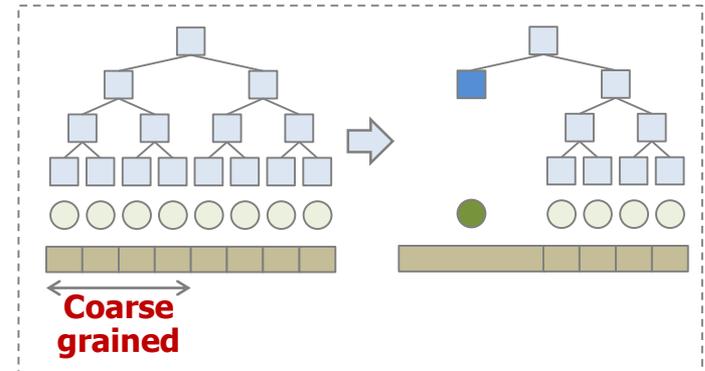
# Granularity Detection (Fine → Coarse)

- Access tracker
  - Records accessed addresses
  - Consecutive access bits are set



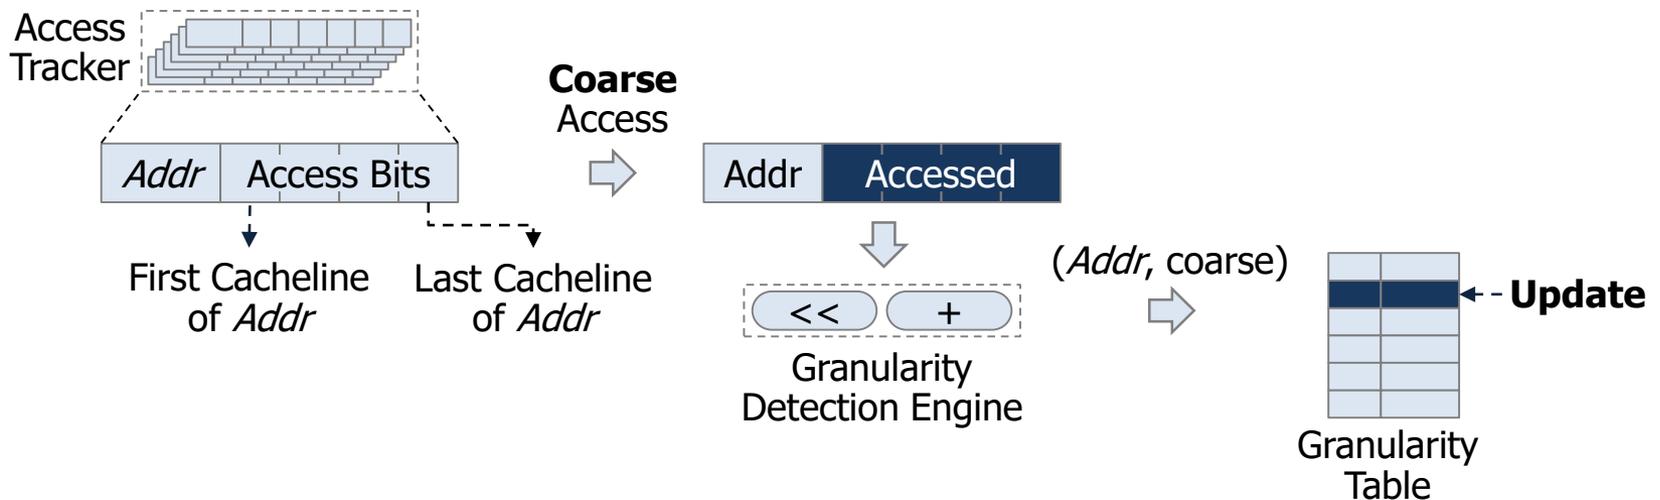
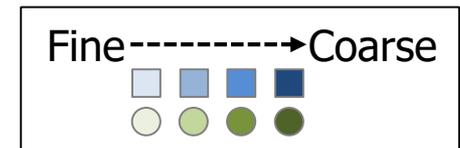
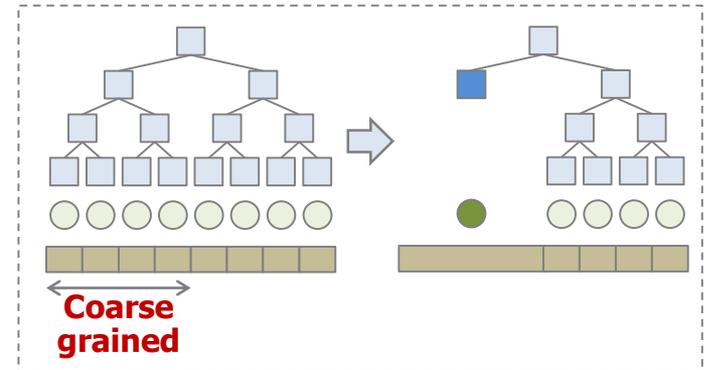
# Granularity Detection (Fine → Coarse)

- Access tracker
  - Records accessed addresses
  - Consecutive access bits are set
- Granularity detection engine
  - Computes a new granularity



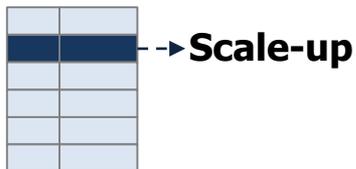
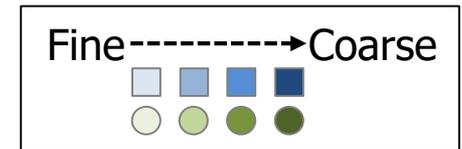
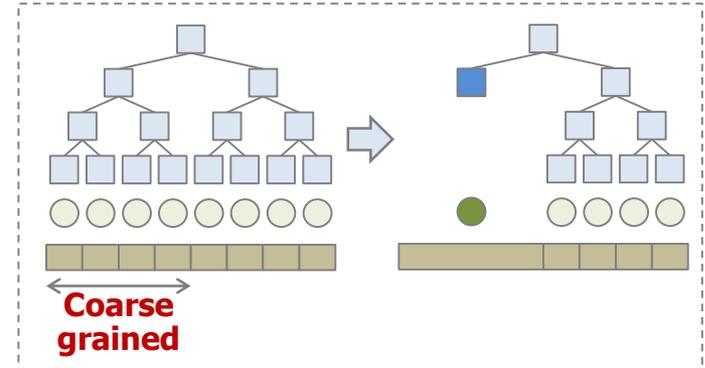
# Granularity Detection (Fine → Coarse)

- Access tracker
  - Records accessed addresses
  - Consecutive access bits are set
- Granularity detection engine
  - Computes a new granularity
  - Updates granularity table



# Granularity Switching (Fine → Coarse)

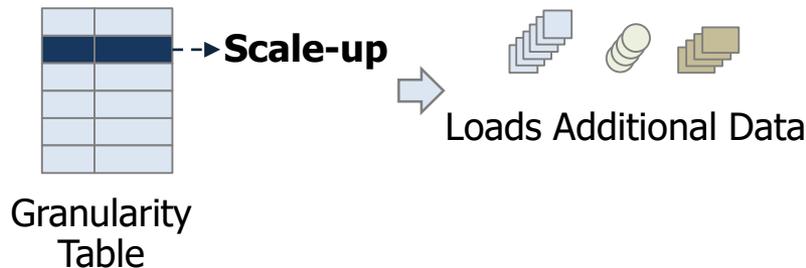
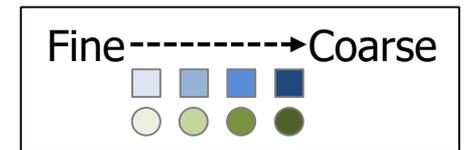
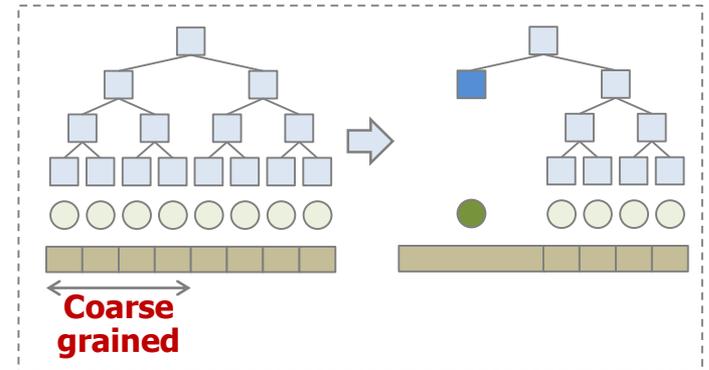
- Granularity switching engine



Granularity Table

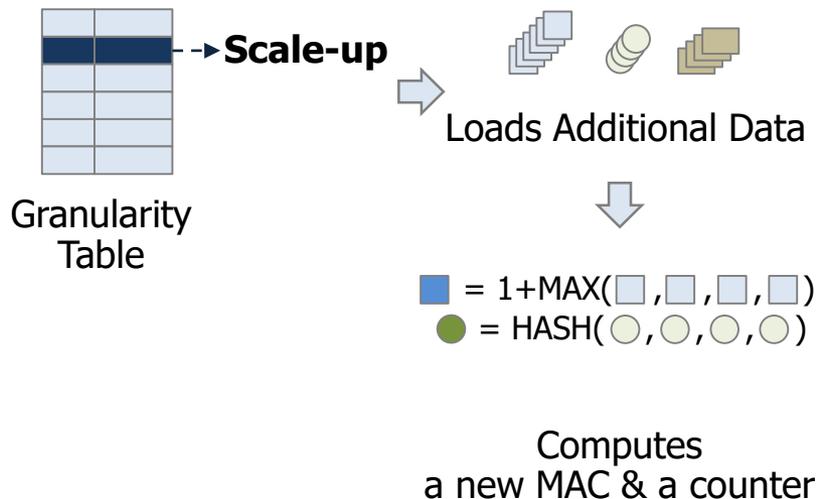
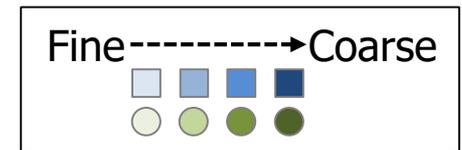
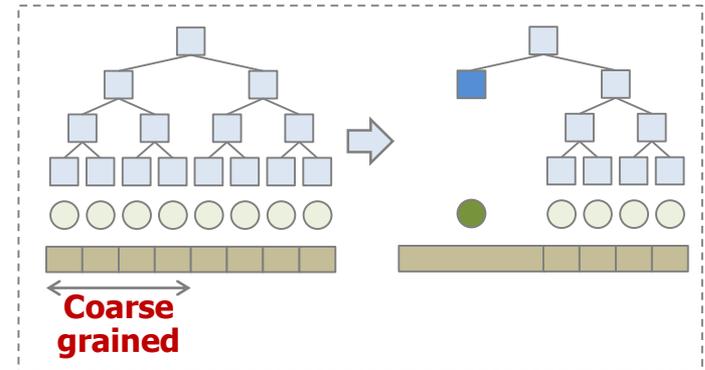
# Granularity Switching (Fine → Coarse)

- Granularity switching engine
  - Loads additional data  
→ Old counters, MACs, data blocks



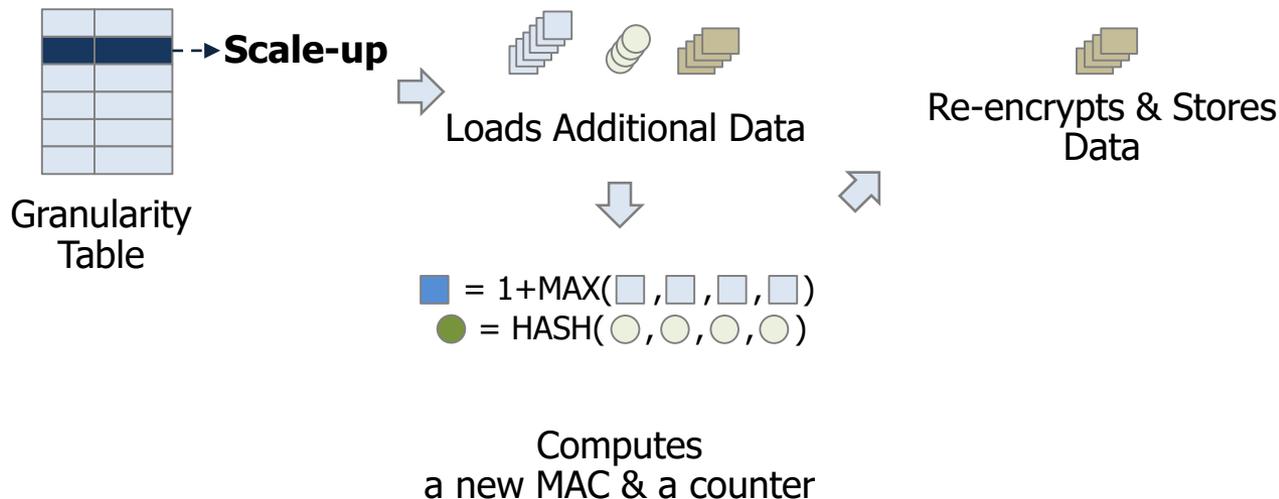
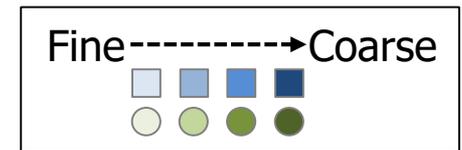
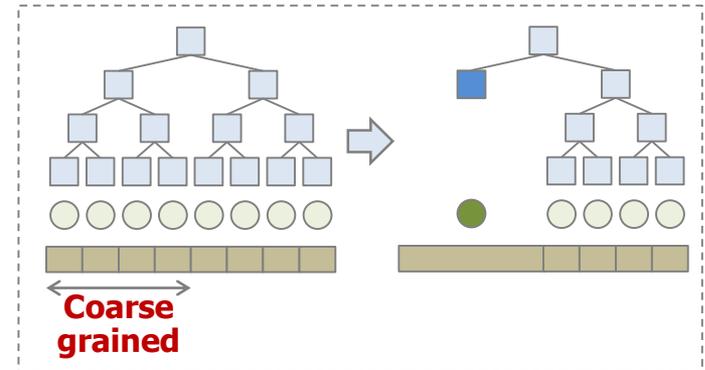
# Granularity Switching (Fine → Coarse)

- Granularity switching engine
  - Loads additional data
    - Old counters, MACs, data blocks
  - Computes counters, MACs



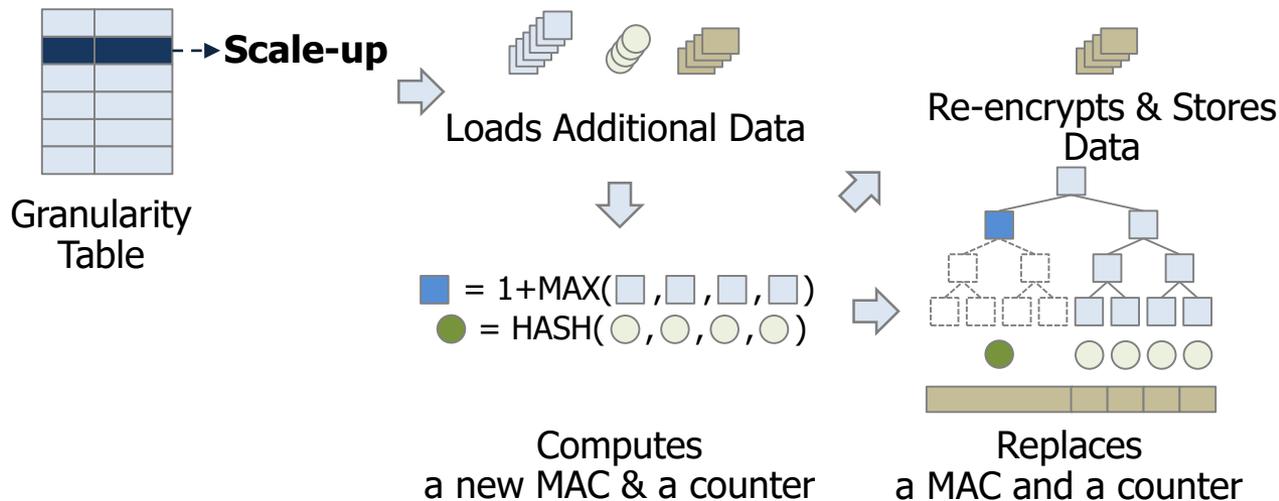
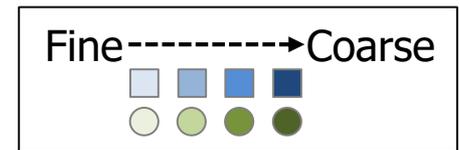
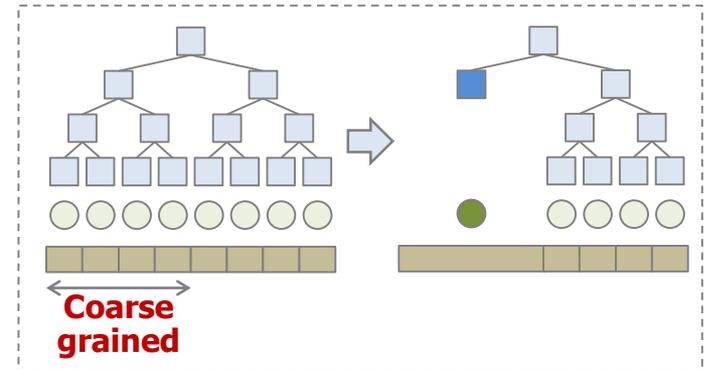
# Granularity Switching (Fine → Coarse)

- Granularity switching engine
  - Loads additional data
    - Old counters, MACs, data blocks
  - Computes counters, MACs
  - Re-encrypts old data



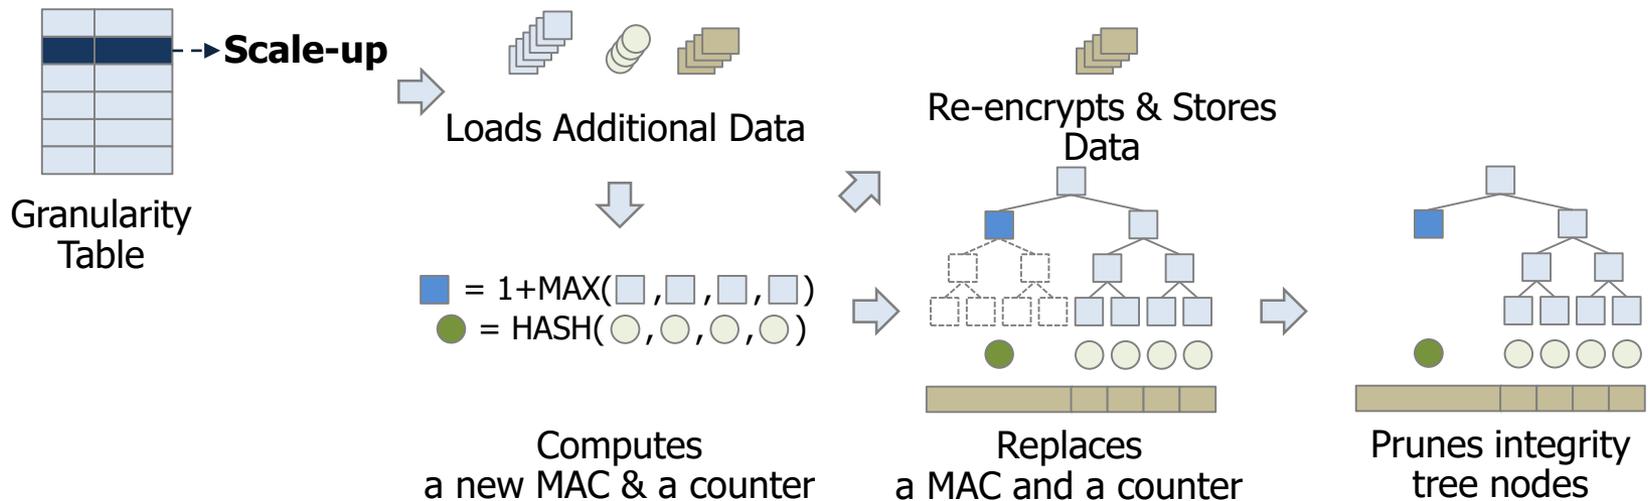
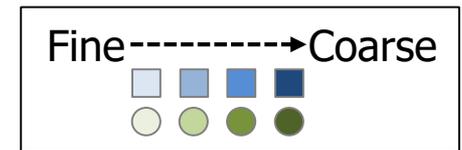
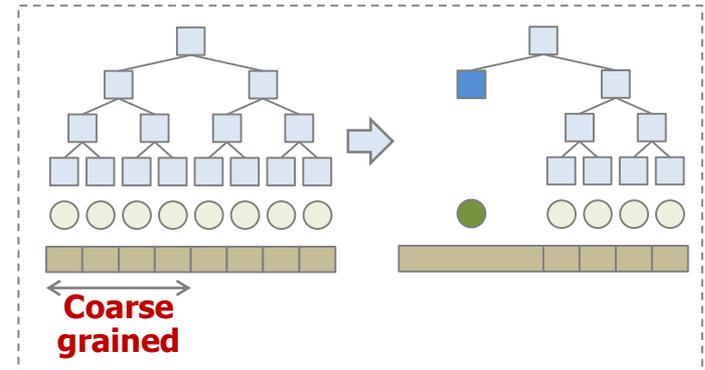
# Granularity Switching (Fine → Coarse)

- Granularity switching engine
  - Loads additional data
    - Old counters, MACs, data blocks
  - Computes counters, MACs
  - Re-encrypts old data
  - Updates & prunes integrity tree



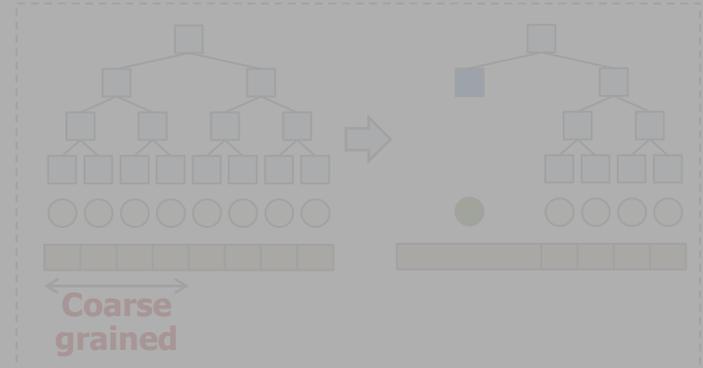
# Granularity Switching (Fine → Coarse)

- Granularity switching engine
  - Loads additional data
    - Old counters, MACs, data blocks
  - Computes counters, MACs
  - Re-encrypts old data
  - Updates & prunes integrity tree

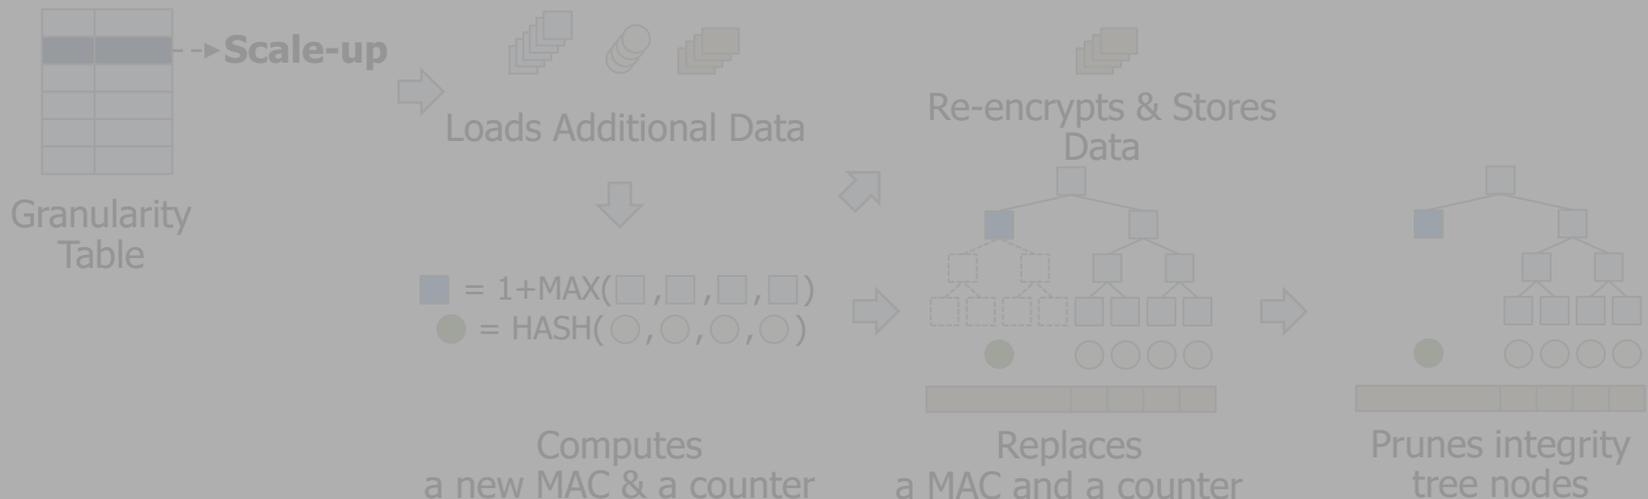


# Granularity Switching (Fine → Coarse)

- Granularity switching engine
  - Loads additional data
    - Old counters, MACs, data blocks
  - Computes counters, MACs



**Granularity switching requires significant overhead → Lazy switching**



# Evaluation Environment

---

# Evaluation Environment

---

- ChampSim (CPU) + MGPUSim (GPU) + mNPUSim (NPU)

# Evaluation Environment

---

- ChampSim (CPU) + MGPUSim (GPU) + mNPUsim (NPU)
- Configuration: Similar to NVIDIA Orin
  - ARM Cortex CPU + Ampere GPU + 2 x NVDLA with LPDDR4

|                 | CPU (Jetson AGX Orin ARM Cortex)     | GPU (Jetson AGX Orin Ampere) | NPU (NVDLA)                        |
|-----------------|--------------------------------------|------------------------------|------------------------------------|
| Compute Engine  | 8-core                               | 14 SMs                       | 45 x 45 Systolic Array             |
| On-chip Storage | Cache (L1: 64KB, L2: 2MB)            | Cache (L1: 192 KB, L2: 4MB)  | Scratchpad Memory (2.2MB in total) |
| Frequency       | 2.2GHz                               | 1GHz                         | 1GHz                               |
| Memory System   | 2.4GHz, 17GB/s, LPDDR4 Memory System |                              |                                    |

---

# Evaluation Environment

---

- Workloads & Scenarios
  - 14 workloads, 250 scenarios (all combinations)

# Evaluation Environment

---

- Workloads & Scenarios
  - 14 workloads, 250 scenarios (all combinations)
  - Access pattern: Fine – ff – f – c – cc – Coarse | Diverse (d)
  - Traffic per cycles: Small (s) – Medium (m) – Large (l)

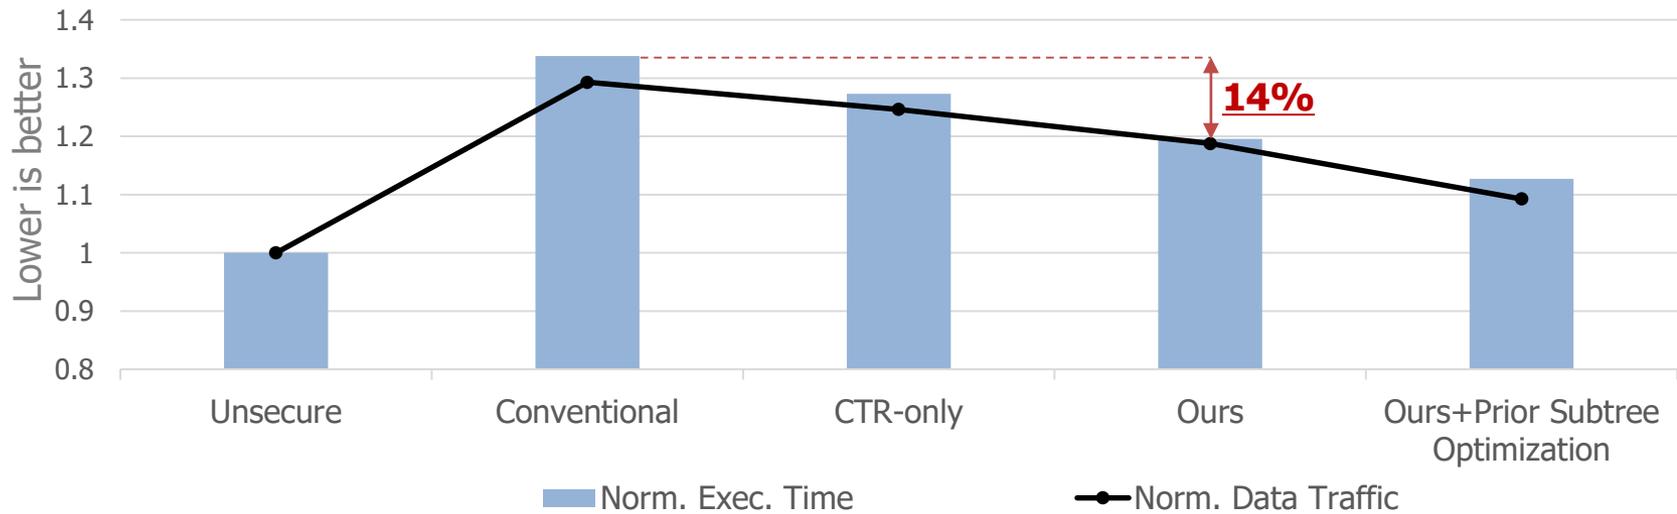
| Workloads (access pattern-traffic per cycles) |   |
|---|---|
| CPU   | bw (ff-s), gcc (ff-s), mcf (ff-m), xal (f-m), ray (ff-s)    |
| GPU   | syr2k (ff-m), pr (f-m), sten (c-l), mm (cc-m), floyd (d-s), |
| NPU   | ncf (c-s), dlrm (c-s), sfrnn (c-l), alex (cc-m)             |

# Evaluation Result (Avg of Processing Units)

---

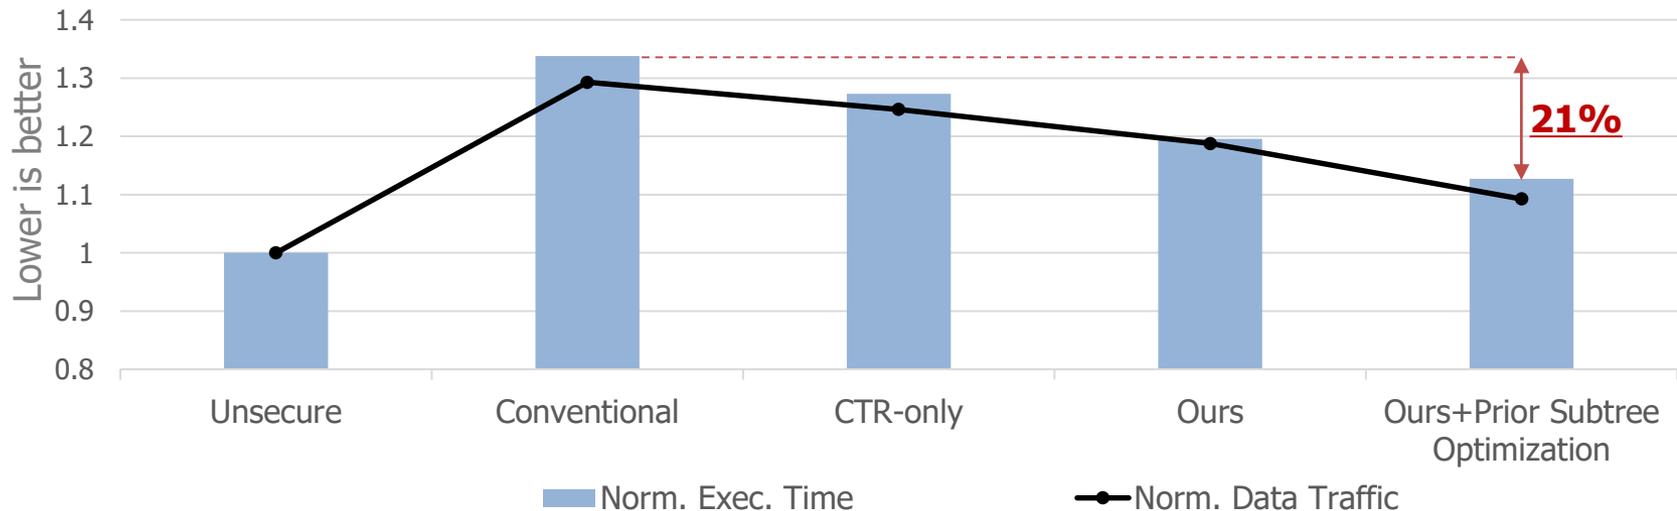
# Evaluation Result (Avg of Processing Units)

- 14% improvement with 11% data reduction



# Evaluation Result (Avg of Processing Units)

- **14%** improvement with **11%** data reduction
- Combining prior subtree optimization <sup>[1-4]</sup>
  - Performance improvement: **14%** → **21%**



[1] Bonsai Merkle Forests: Efficiently Achieving Crash Consistency in Secure Persistent Memory (MICRO 2021)

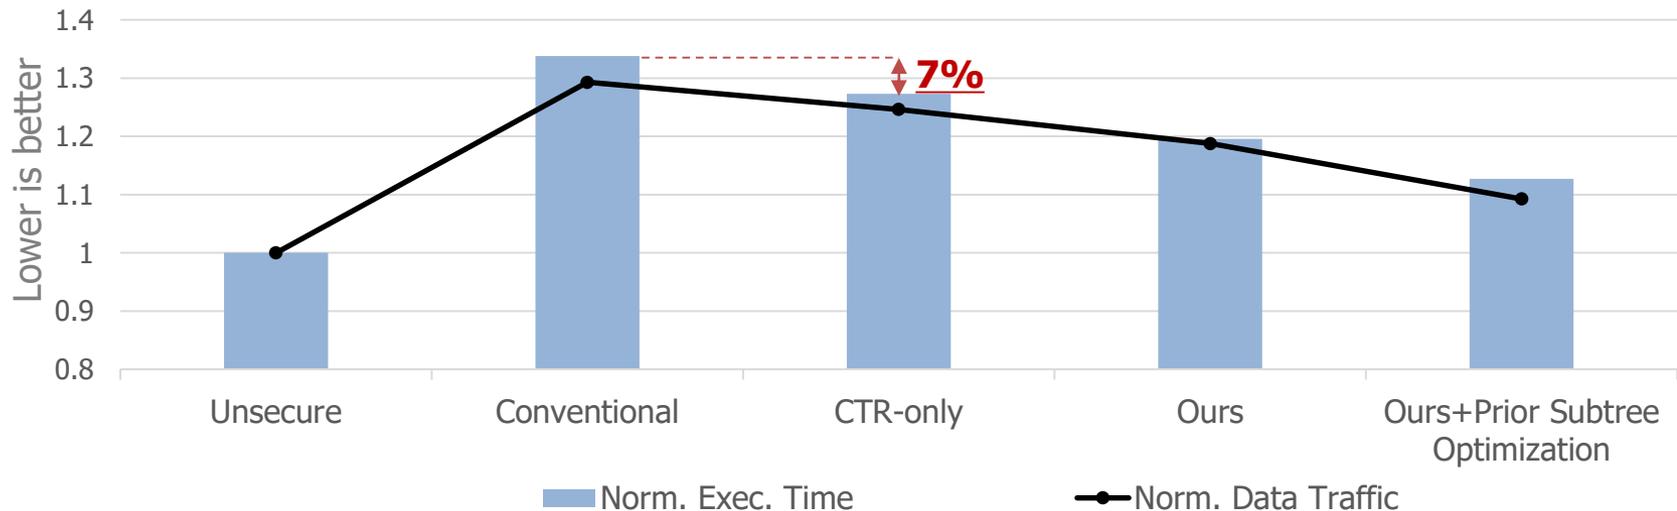
[2] Scalable Memory Protection in the PENGLAI Enclave (OSDI 2021)

[3] Efficient Distributed Secure Memory with Migratable Merkle Tree (HPCA 2023)

[4] Data Enclave: A Data-Centric Trusted Execution Environment (HPCA 2024)

# Evaluation Result (Avg of Processing Units)

- **14%** improvement with **11%** data reduction
- Combining prior subtree optimization <sup>[1-4]</sup>
  - Performance improvement: **14%** → **21%**
  - CTR-only (**7%**), +MAC (**7%**), +Prior tree optimization (**7%**)



[1] Bonsai Merkle Forests: Efficiently Achieving Crash Consistency in Secure Persistent Memory (MICRO 2021)

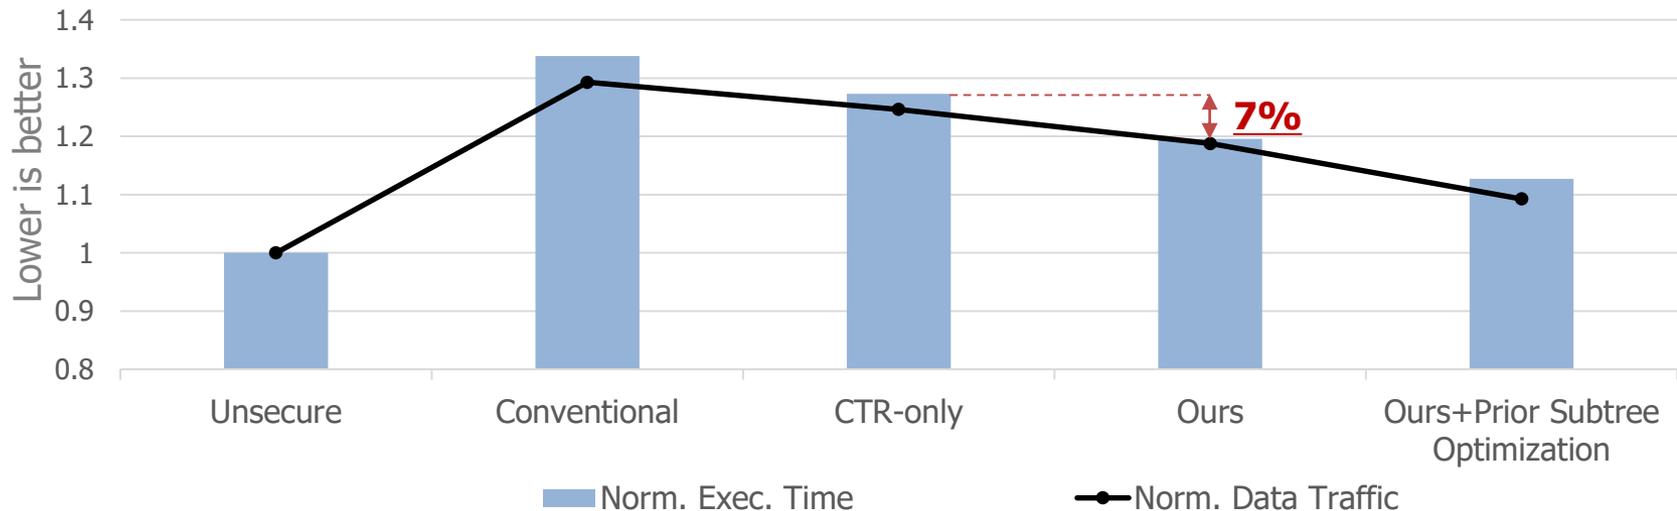
[2] Scalable Memory Protection in the PENGLAI Enclave (OSDI 2021)

[3] Efficient Distributed Secure Memory with Migratable Merkle Tree (HPCA 2023)

[4] Data Enclave: A Data-Centric Trusted Execution Environment (HPCA 2024)

# Evaluation Result (Avg of Processing Units)

- **14%** improvement with **11%** data reduction
- Combining prior subtree optimization <sup>[1-4]</sup>
  - Performance improvement: **14%** → **21%**
  - CTR-only (**7%**), +MAC (**7%**), +Prior tree optimization (**7%**)



[1] Bonsai Merkle Forests: Efficiently Achieving Crash Consistency in Secure Persistent Memory (MICRO 2021)

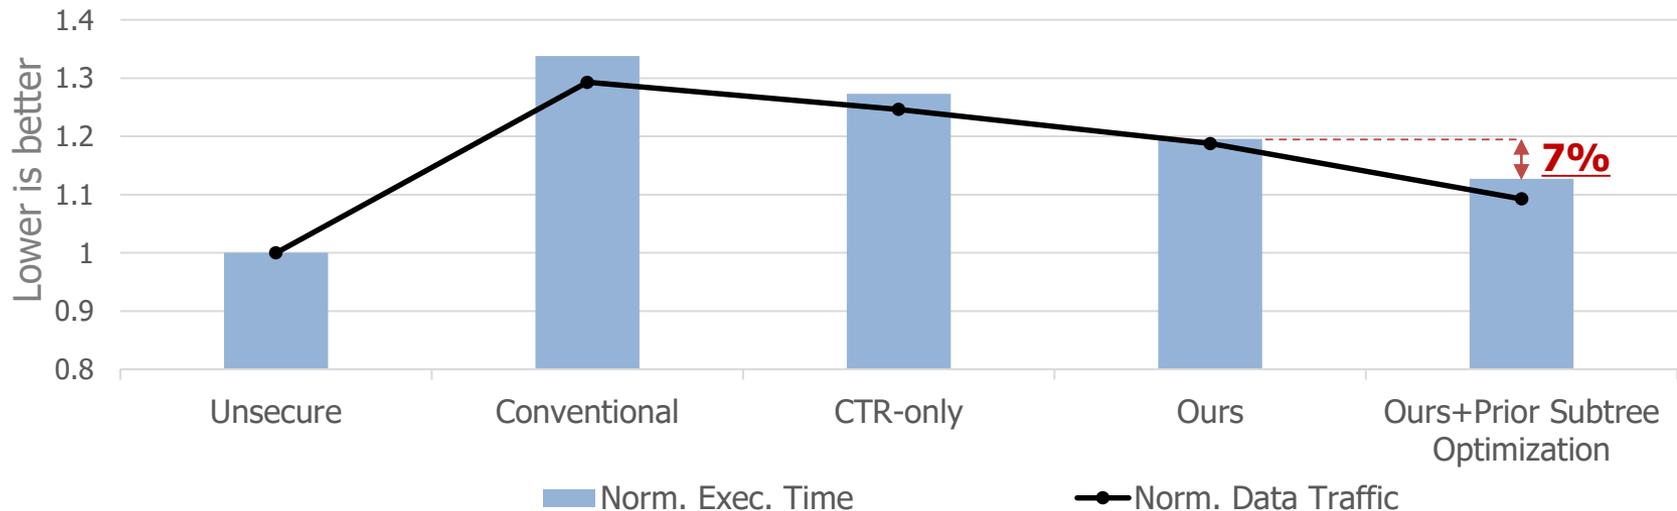
[2] Scalable Memory Protection in the PENGLAI Enclave (OSDI 2021)

[3] Efficient Distributed Secure Memory with Migratable Merkle Tree (HPCA 2023)

[4] Data Enclave: A Data-Centric Trusted Execution Environment (HPCA 2024)

# Evaluation Result (Avg of Processing Units)

- **14%** improvement with **11%** data reduction
- Combining prior subtree optimization <sup>[1-4]</sup>
  - Performance improvement: **14%** → **21%**
  - CTR-only (**7%**), +MAC (**7%**), +Prior tree optimization (**7%**)



[1] Bonsai Merkle Forests: Efficiently Achieving Crash Consistency in Secure Persistent Memory (MICRO 2021)

[2] Scalable Memory Protection in the PENGLAI Enclave (OSDI 2021)

[3] Efficient Distributed Secure Memory with Migratable Merkle Tree (HPCA 2023)

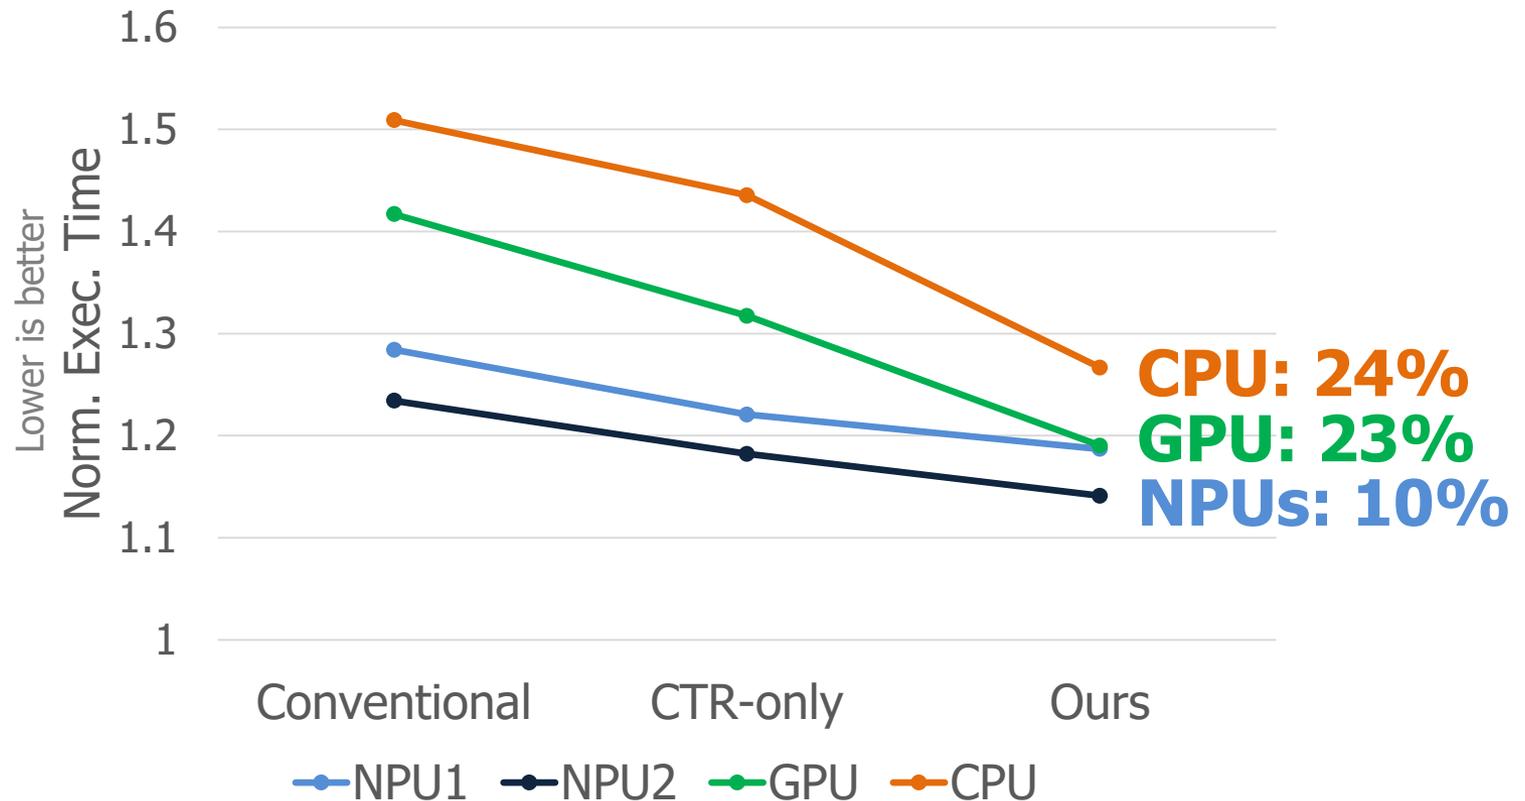
[4] Data Enclave: A Data-Centric Trusted Execution Environment (HPCA 2024)

# Evaluation Result (Per Processing Unit)

---

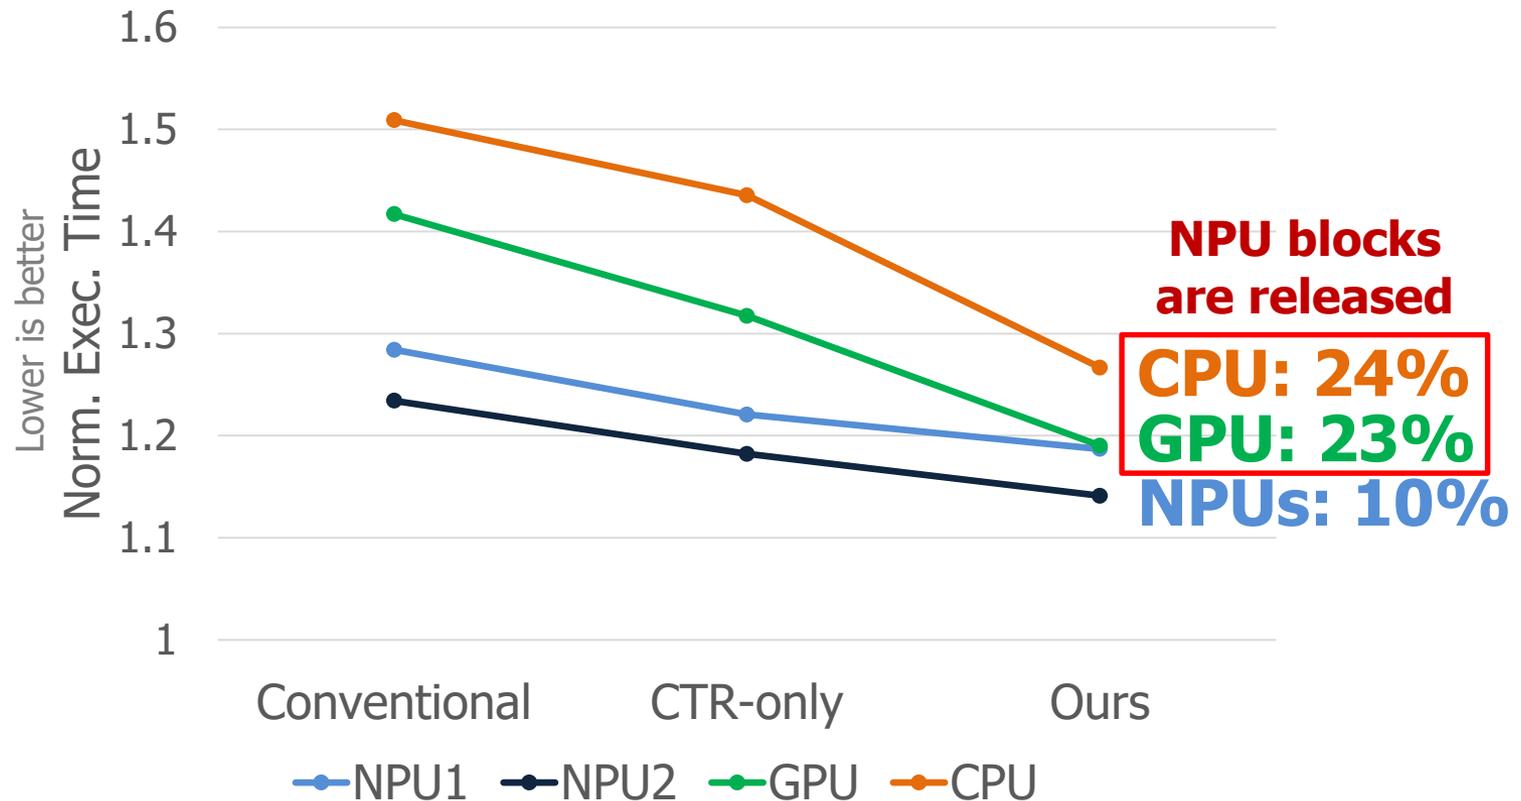
# Evaluation Result (Per Processing Unit)

- Performance improvement of each processing unit
  - CPU (**24%**), GPU (**23%**), NPU (**10%**)



# Evaluation Result (Per Processing Unit)

- Performance improvement of each processing unit
  - CPU (**24%**), GPU (**23%**), NPU (**10%**)



# Conclusion

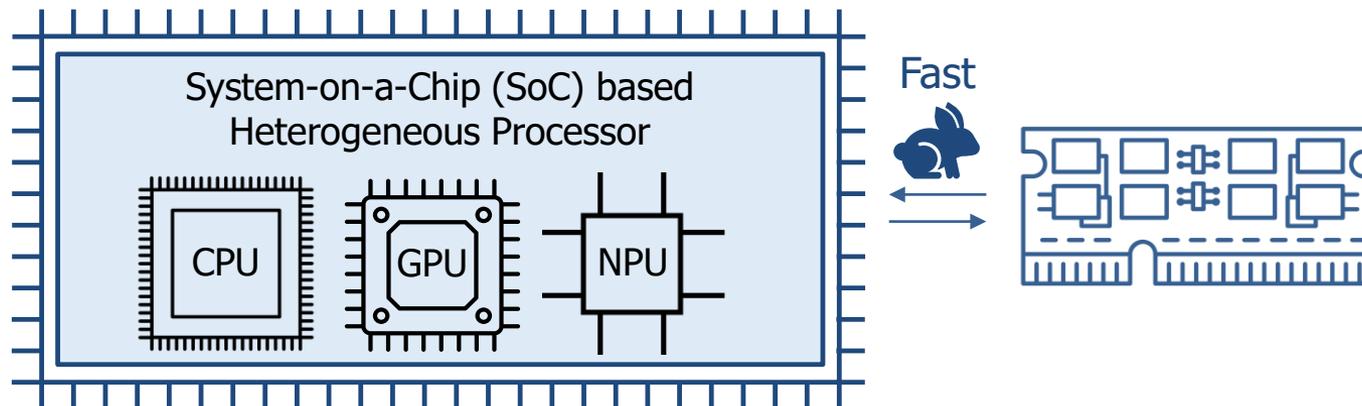
---

# Conclusion

---

- Unified memory protection for heterogeneous processor

## Our Unified Memory Protection Scheme

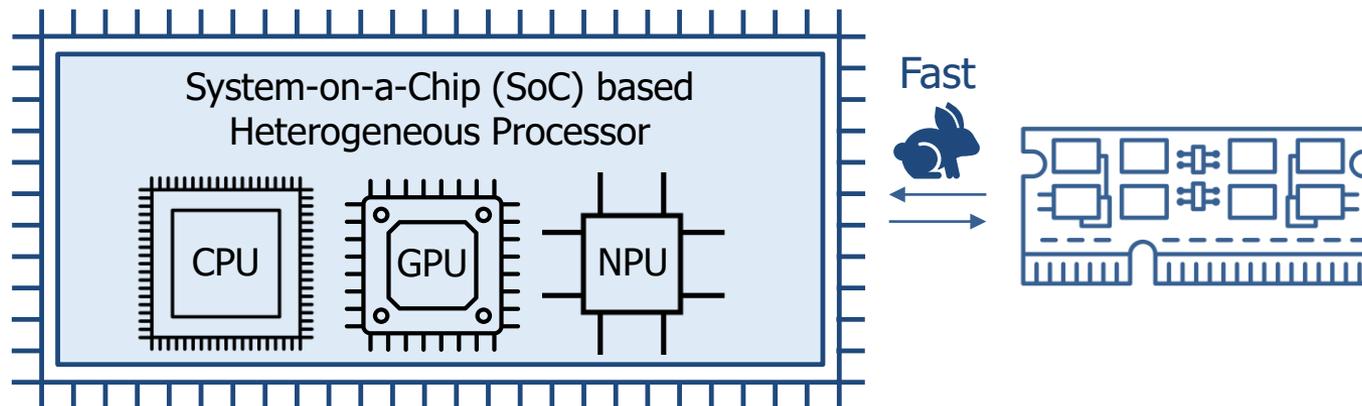


# Conclusion

---

- Unified memory protection for heterogeneous processor
  - **Multi-granular MAC & Integrity Tree**

## Our Unified Memory Protection Scheme

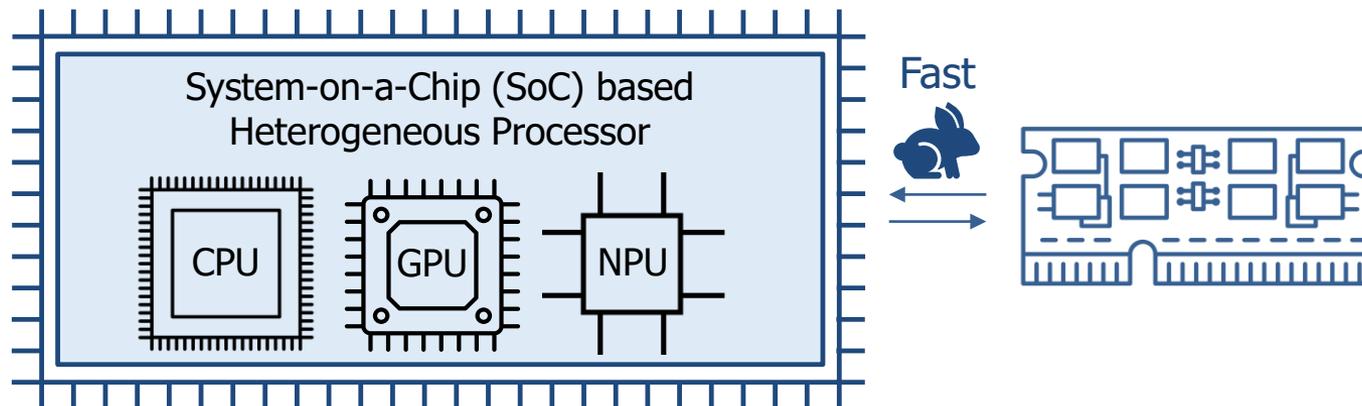


# Conclusion

---

- Unified memory protection for heterogeneous processor
  - **Multi-granular MAC & Integrity Tree**
  - Challenge: Diverse access pattern

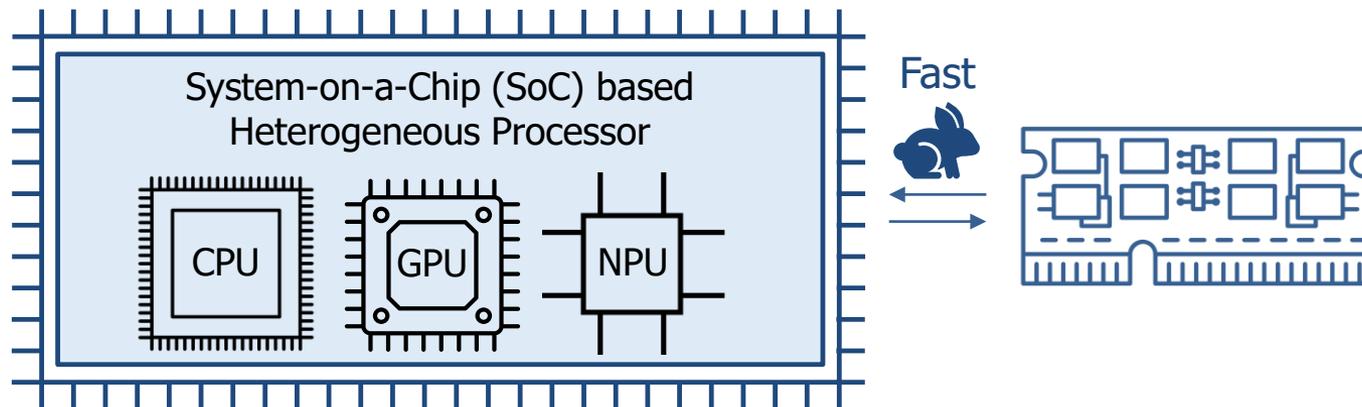
## Our Unified Memory Protection Scheme



# Conclusion

- Unified memory protection for heterogeneous processor
  - **Multi-granular MAC & Integrity Tree**
  - Challenge: Diverse access pattern
- Improvement: **14%** (w/o subtree opt.), **21%** (w/ subtree opt.)

## Our Unified Memory Protection Scheme



**Thank you**

**Backup Slide**

# Detailed Granularity Switching

---

# Detailed Granularity Switching

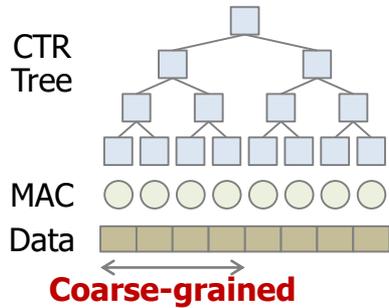
---

\* Scale-up (Fine → Coarse)

# Detailed Granularity Switching

---

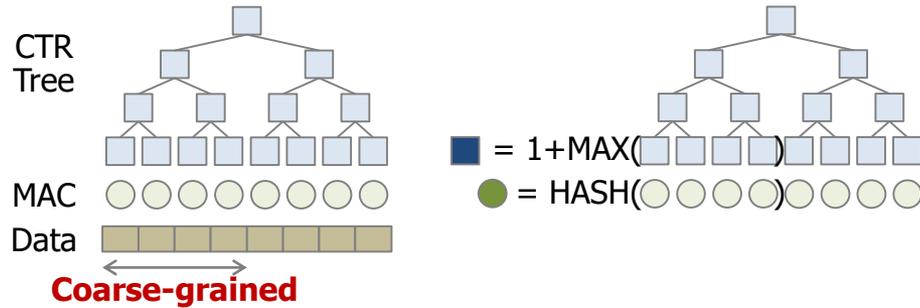
\* Scale-up (Fine  $\rightarrow$  Coarse)



1) Detect scale-up

# Detailed Granularity Switching

\* Scale-up (Fine  $\rightarrow$  Coarse)

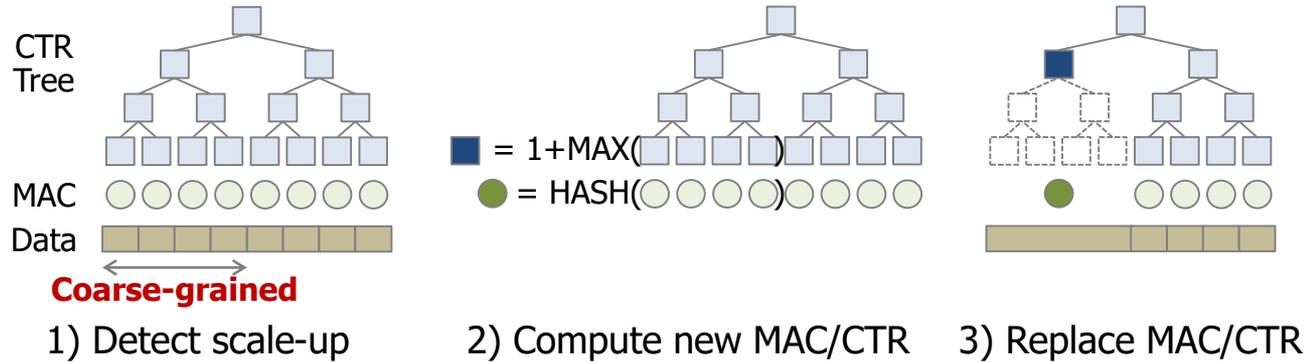


1) Detect scale-up

2) Compute new MAC/CTR

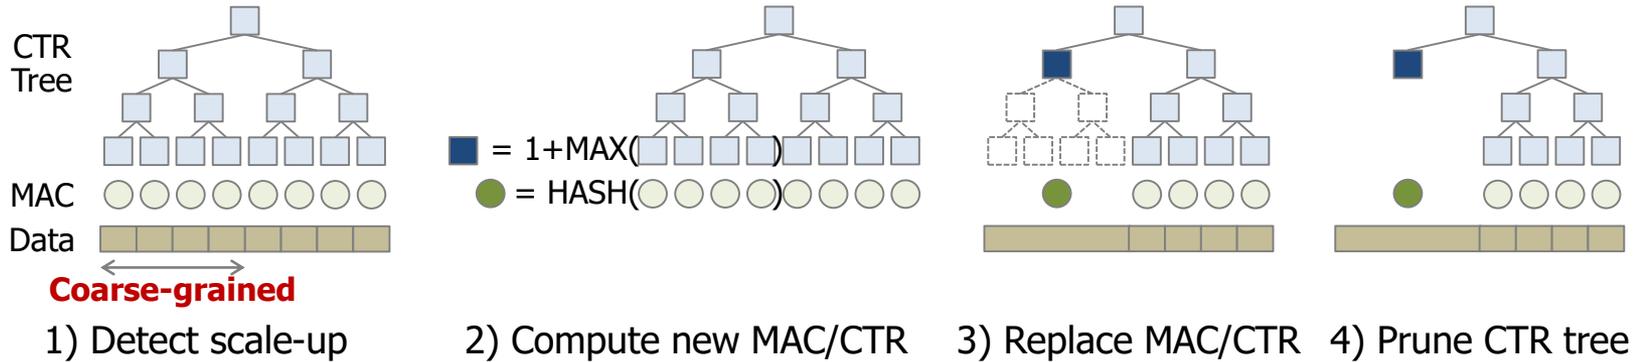
# Detailed Granularity Switching

\* Scale-up (Fine  $\rightarrow$  Coarse)



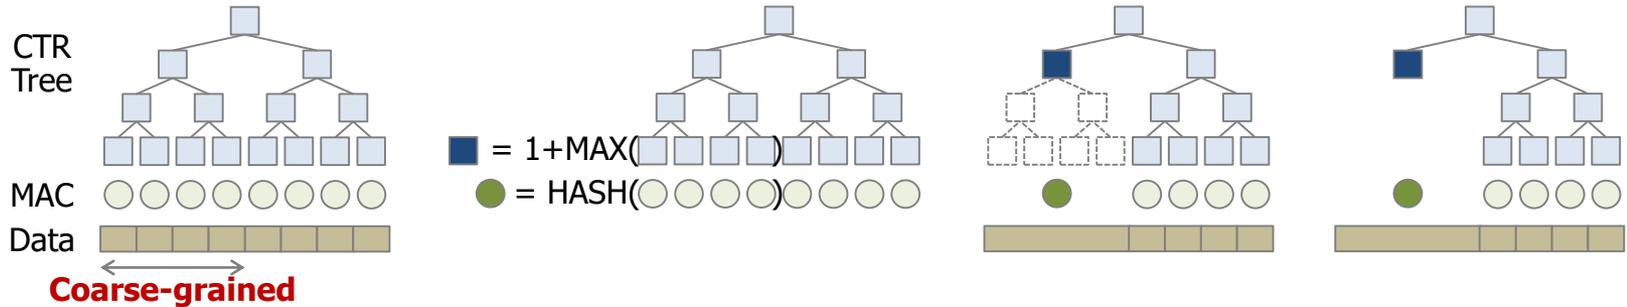
# Detailed Granularity Switching

\* Scale-up (Fine  $\rightarrow$  Coarse)



# Detailed Granularity Switching

## \* Scale-up (Fine → Coarse)



1) Detect scale-up

2) Compute new MAC/CTR

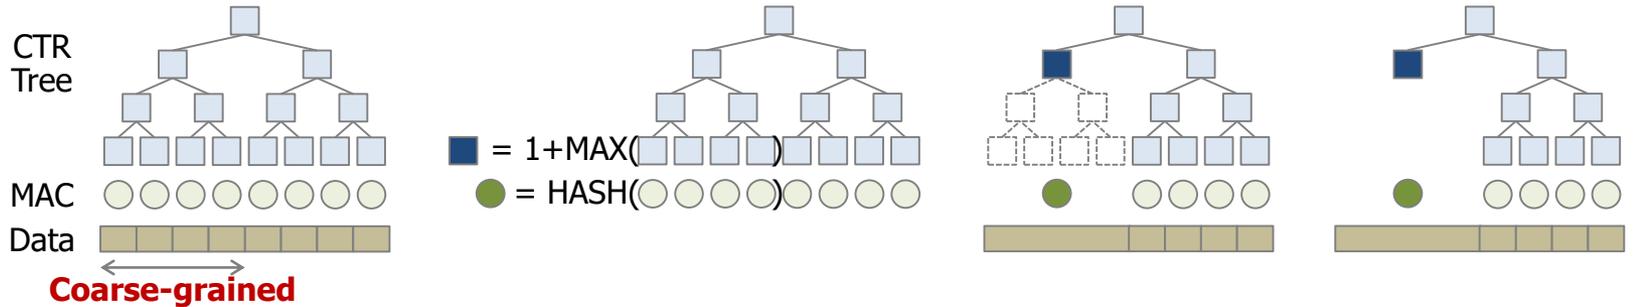
3) Replace MAC/CTR

4) Prune CTR tree

## \* Scale-down (Coarse → Fine)

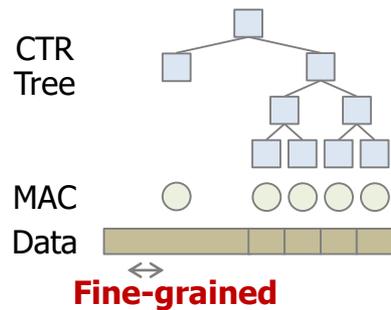
# Detailed Granularity Switching

## \* Scale-up (Fine → Coarse)



- 1) Detect scale-up
- 2) Compute new MAC/CTR
- 3) Replace MAC/CTR
- 4) Prune CTR tree

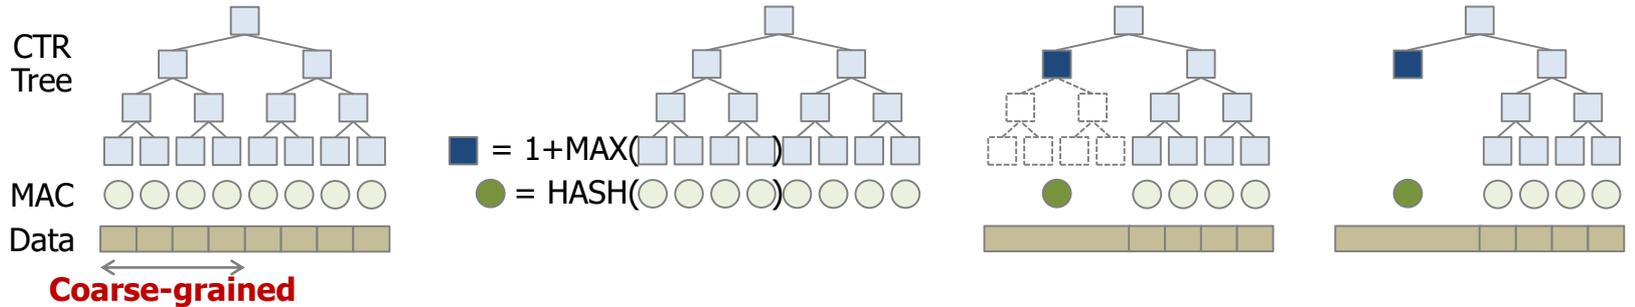
## \* Scale-down (Coarse → Fine)



- 1) Detect scale-down

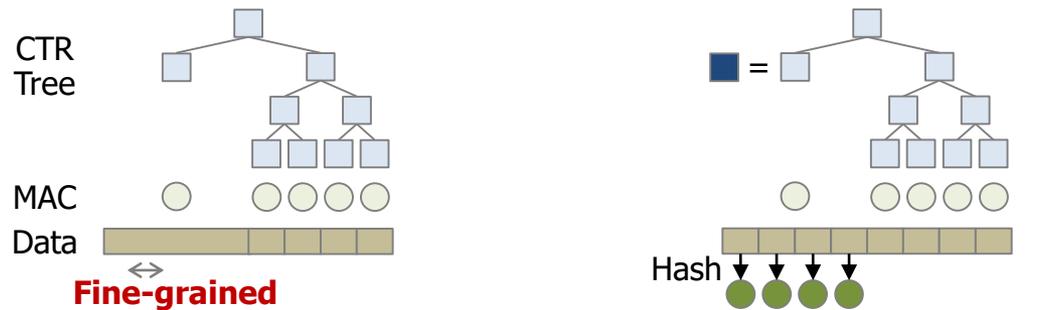
# Detailed Granularity Switching

## \* Scale-up (Fine → Coarse)



- 1) Detect scale-up
- 2) Compute new MAC/CTR
- 3) Replace MAC/CTR
- 4) Prune CTR tree

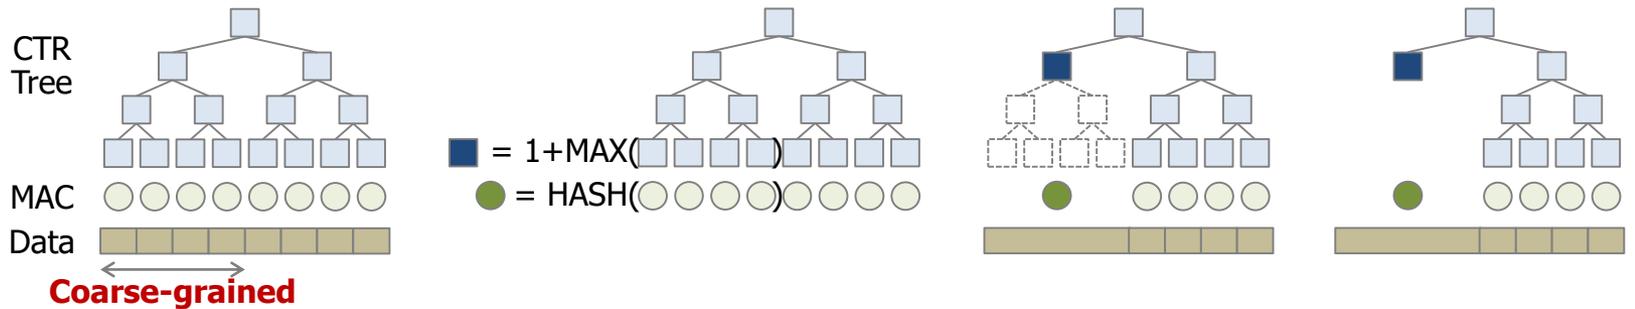
## \* Scale-down (Coarse → Fine)



- 1) Detect scale-down
- 2) Compute new MAC/CTR

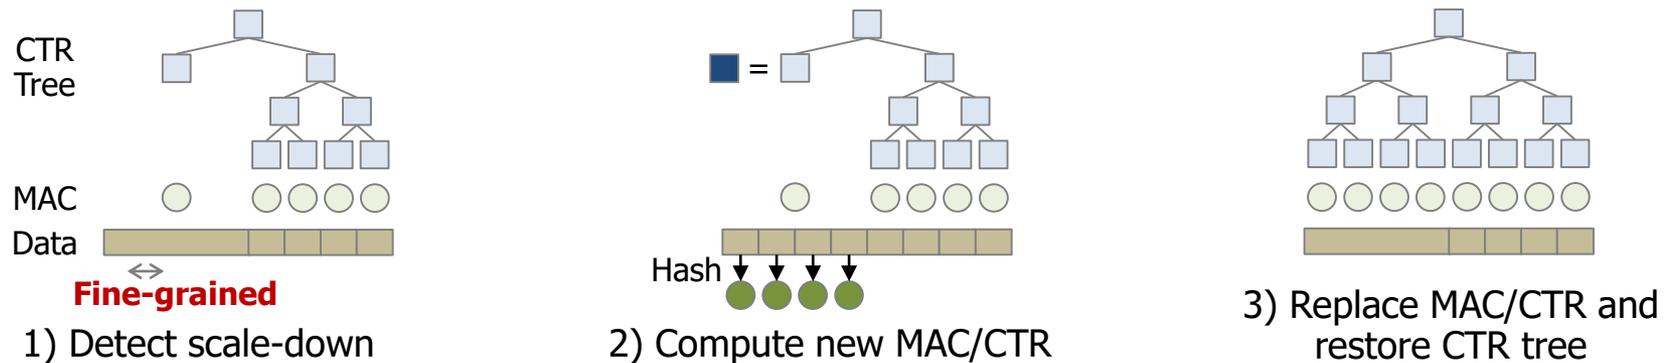
# Detailed Granularity Switching

## \* Scale-up (Fine → Coarse)



- 1) Detect scale-up      2) Compute new MAC/CTR      3) Replace MAC/CTR      4) Prune CTR tree

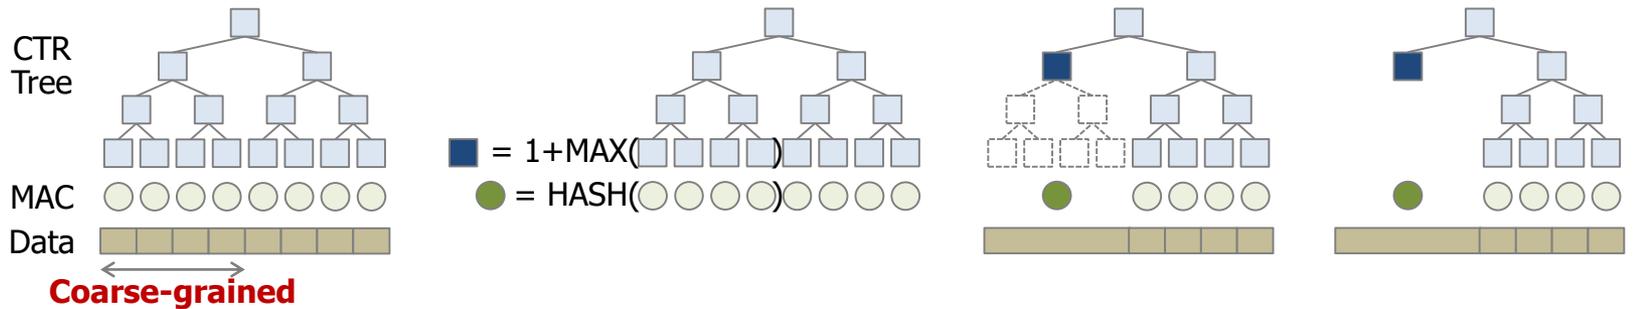
## \* Scale-down (Coarse → Fine)



- 3) Replace MAC/CTR and restore CTR tree

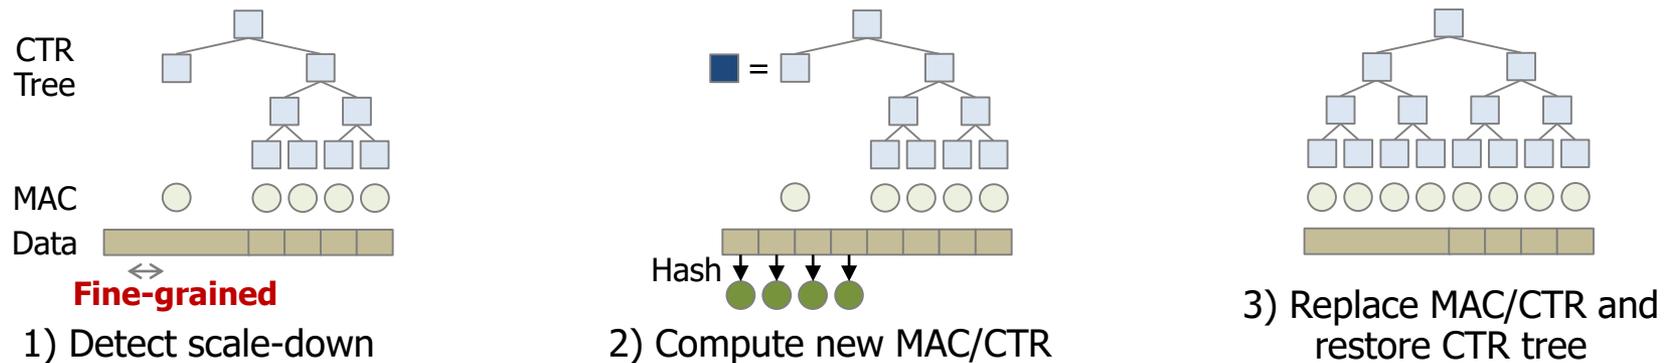
# Detailed Granularity Switching

## \* Scale-up (Fine → Coarse)



- 1) Detect scale-up
- 2) Compute new MAC/CTR
- 3) Replace MAC/CTR
- 4) Prune CTR tree

## \* Scale-down (Coarse → Fine)



- 3) Replace MAC/CTR and restore CTR tree

**Granularity switching requires significant overhead! → Lazy switching**

# Lazy Switching Overhead by MAC

---

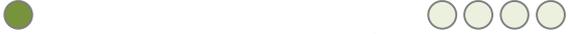
# Lazy Switching Overhead by MAC

---

 *Coarse* MAC = HASH ( *Fine* MACs)

# Lazy Switching Overhead by MAC

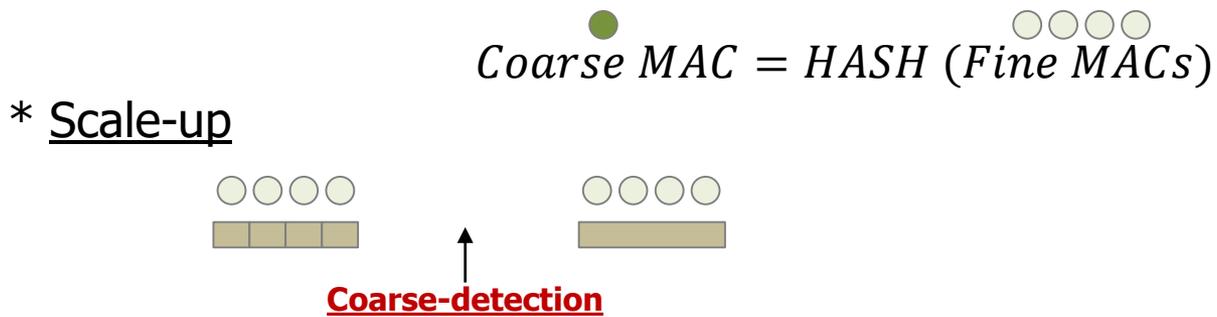
---

  
*Coarse MAC = HASH (Fine MACs)*

\* Scale-up

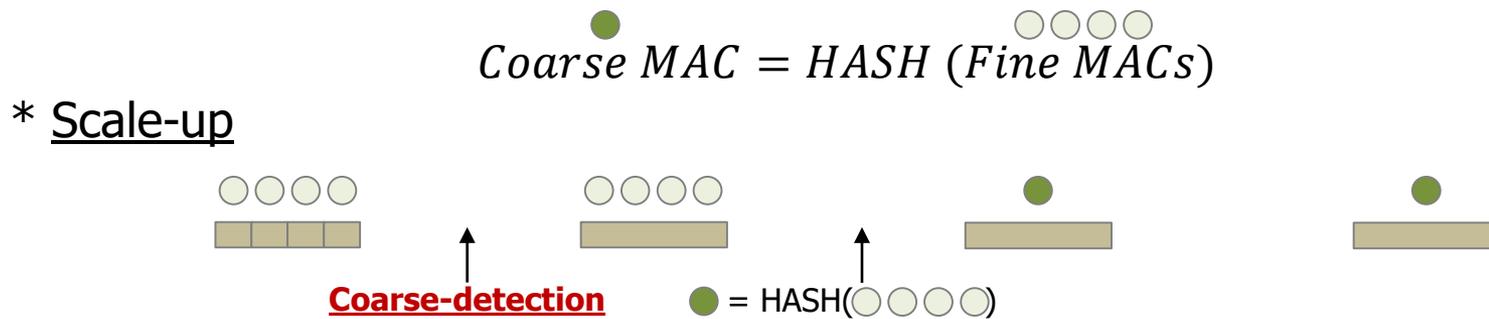
# Lazy Switching Overhead by MAC

---



# Lazy Switching Overhead by MAC

---

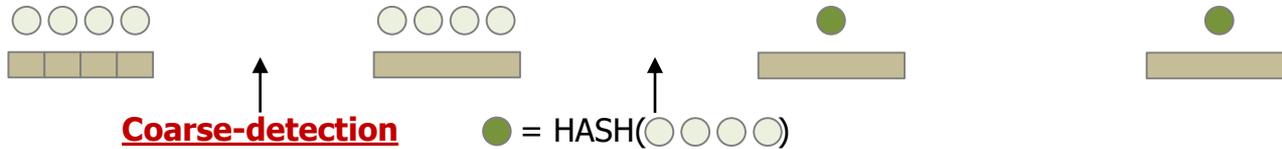


# Lazy Switching Overhead by MAC

---

$$\text{Coarse MAC} = \text{HASH}(\text{Fine MACs})$$

\* Scale-up



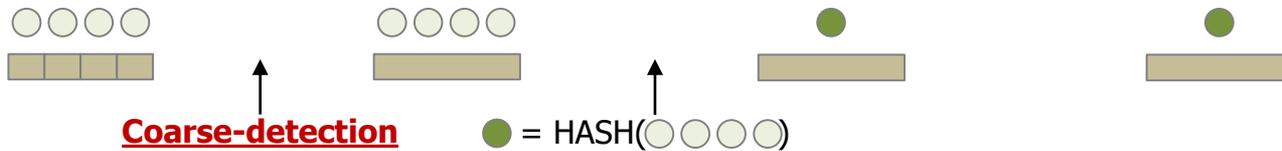
\* Scale-down

# Lazy Switching Overhead by MAC

---

$$\text{Coarse MAC} = \text{HASH}(\text{Fine MACs})$$

\* Scale-up



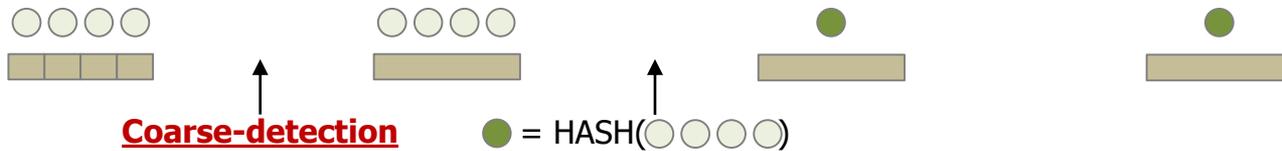
\* Scale-down



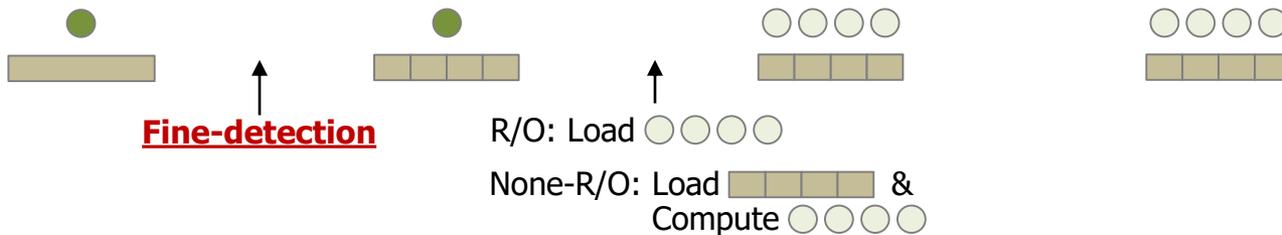
# Lazy Switching Overhead by MAC

$$\text{Coarse MAC} = \text{HASH}(\text{Fine MACs})$$

## \* Scale-up



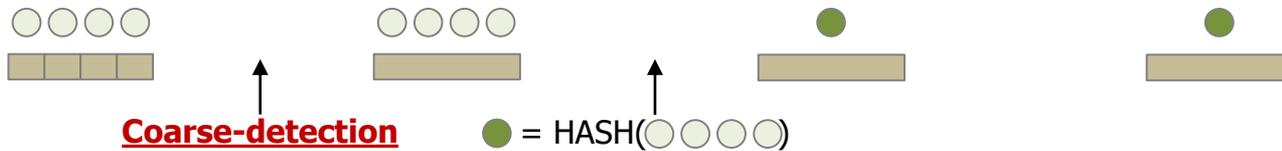
## \* Scale-down



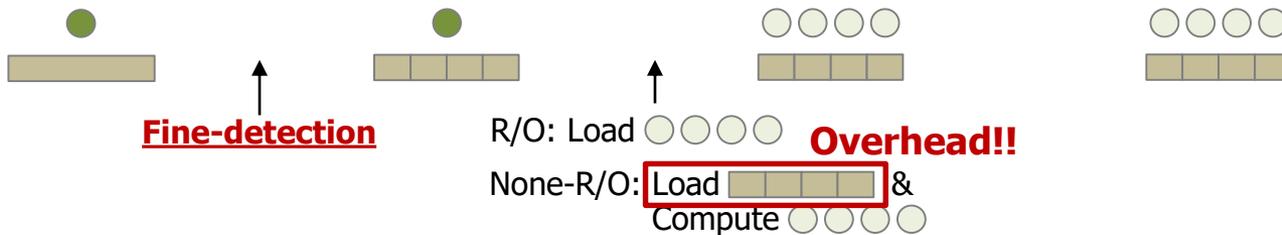
# Lazy Switching Overhead by MAC

$$\text{Coarse MAC} = \text{HASH}(\text{Fine MACs})$$

\* Scale-up



\* Scale-down

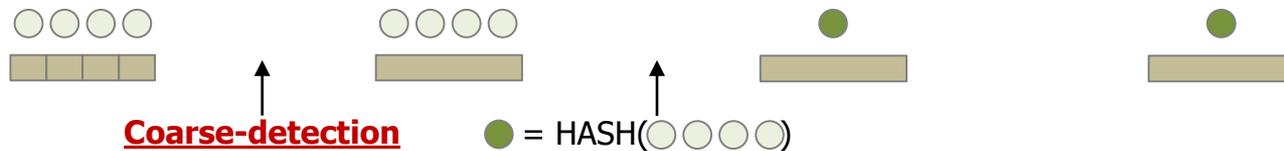


# Lazy Switching Overhead by MAC

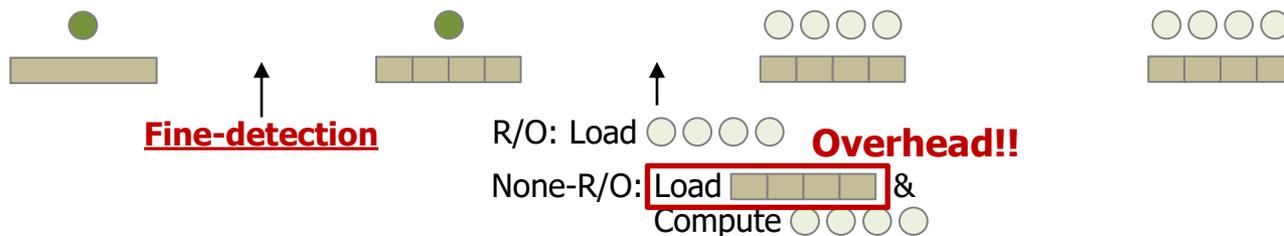
- 97.2% of reqs → Hidden by lazy switching & R/O
  - Only 2.8% of reqs makes moderate overhead (ld data chunks)

Coarse MAC = HASH (Fine MACs)

## \* Scale-up



## \* Scale-down

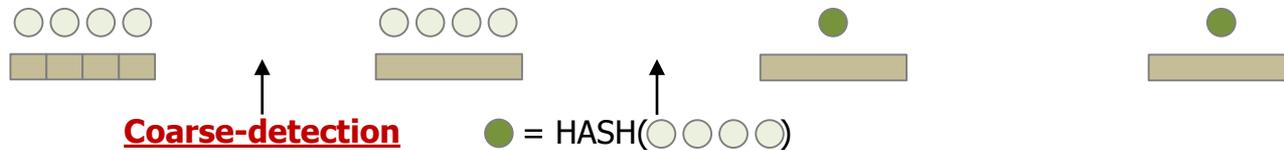


# Lazy Switching Overhead by MAC

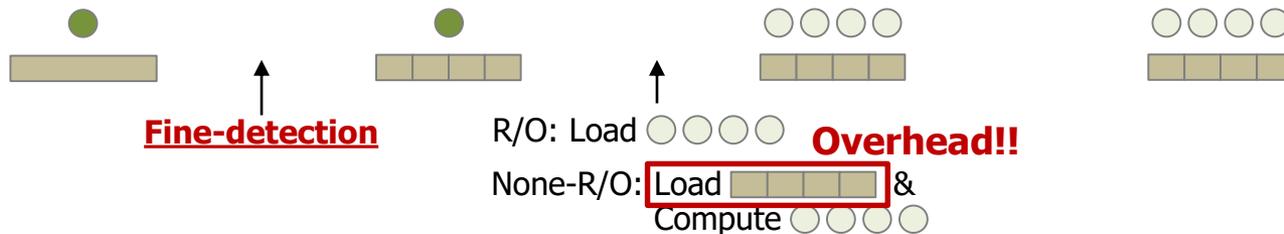
- 97.2% of reqs → Hidden by lazy switching & R/O
  - Only 2.8% of reqs makes moderate overhead (ld data chunks)

Coarse MAC = HASH (Fine MACs)

\* Scale-up



\* Scale-down



**Lazy switching considerably reduces switching overhead!!**

# Lazy Switching Overhead by CTR Tree

---

# Lazy Switching Overhead by CTR Tree

---

Granul.  
detection

# Lazy Switching Overhead by CTR Tree

---

Granul. → Store  
detection → next granul.

# Lazy Switching Overhead by CTR Tree

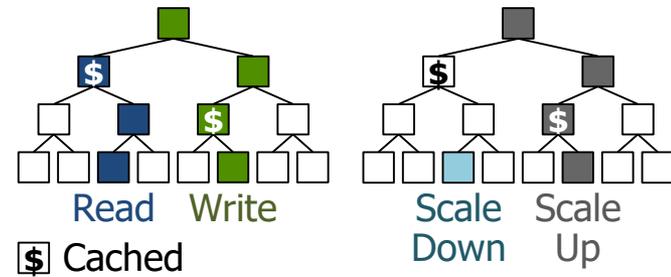
---

Granul. detection → Store next granul. → Granul. switch after next access

# Lazy Switching Overhead by CTR Tree

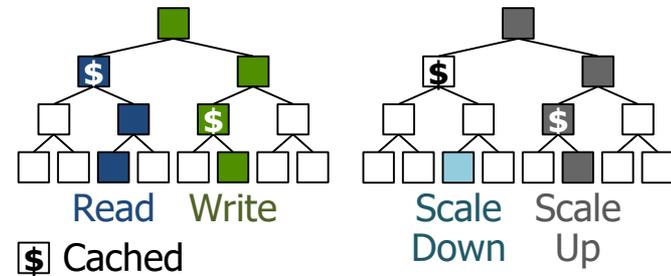
---

Granul. detection → Store next granul. → Granul. switch after next access

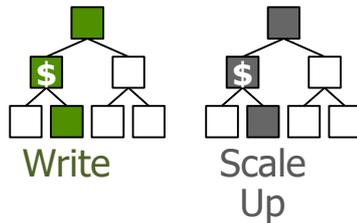


# Lazy Switching Overhead by CTR Tree

Granul. detection → Store next granul. → Granul. switch after next access

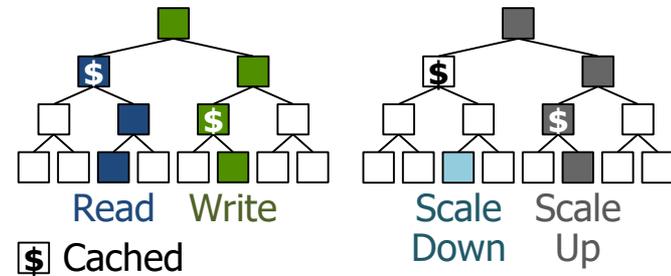


\* Scale-up with WAR/WAW

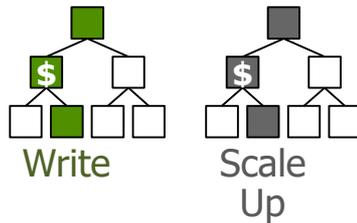


# Lazy Switching Overhead by CTR Tree

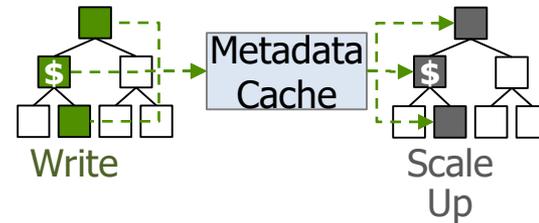
Granul. detection → Store next granul. → Granul. switch after next access



\* Scale-up with WAR/WAW



\* Scale-up with RAW

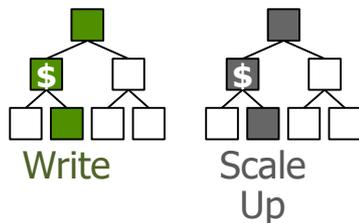


# Lazy Switching Overhead by CTR Tree

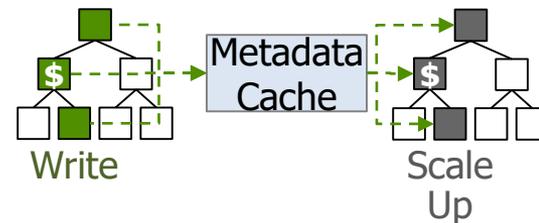
Granul. detection → Store next granul. → Granul. switch after next access



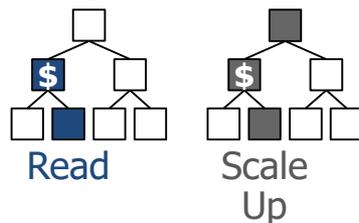
\* Scale-up with WAR/WAW



\* Scale-up with RAW

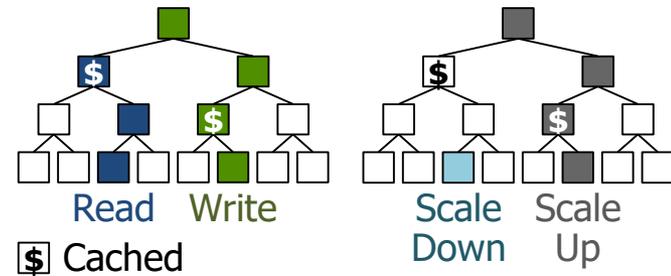


\* Scale-up with RAR

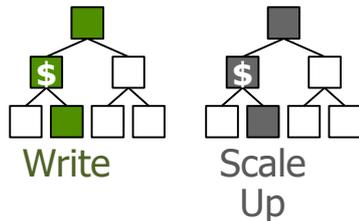


# Lazy Switching Overhead by CTR Tree

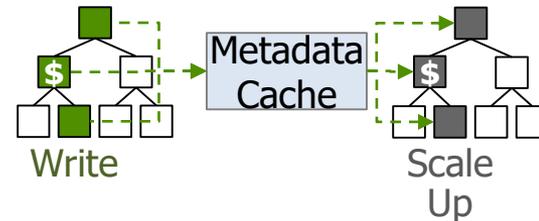
Granul. detection → Store next granul. → Granul. switch after next access



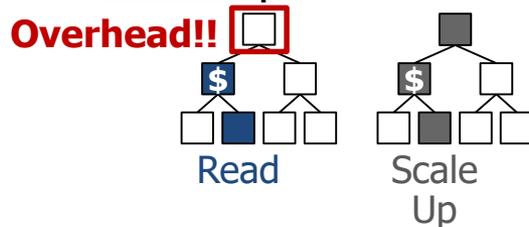
\* Scale-up with WAR/WAW



\* Scale-up with RAW



\* Scale-up with RAR

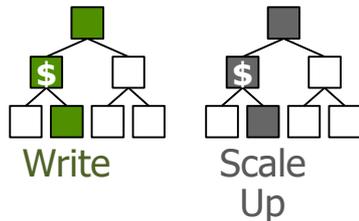


# Lazy Switching Overhead by CTR Tree

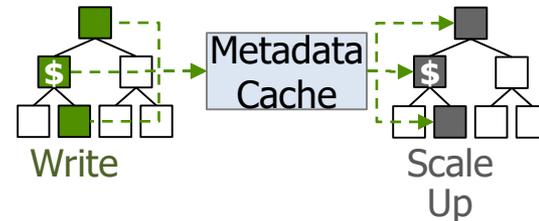
Granul. detection → Store next granul. → Granul. switch after next access



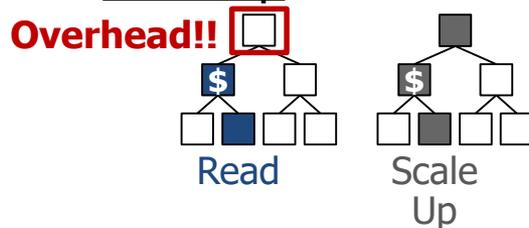
\* Scale-up with WAR/WAW



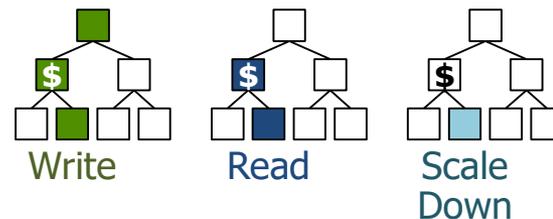
\* Scale-up with RAW



\* Scale-up with RAR



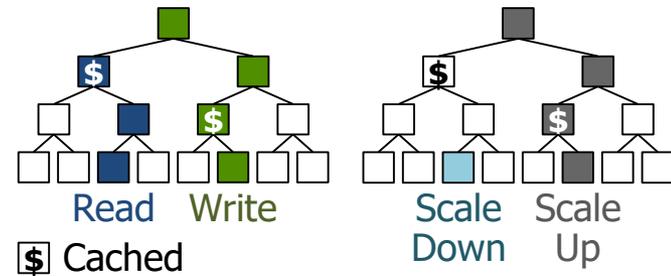
\* Scale-down



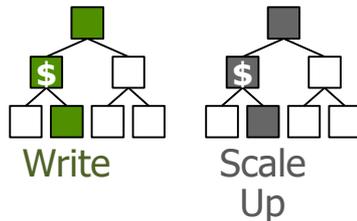
# Lazy Switching Overhead by CTR Tree

- 91.2% of reqs → Hidden by lazy switching
  - Only 8.2% of reqs makes low overhead (read req → write req)

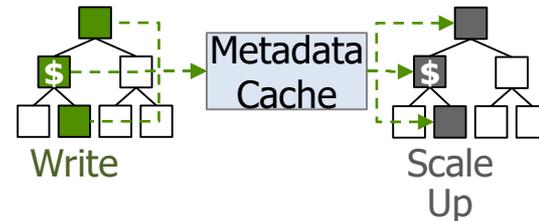
Granul. detection → Store next granul. → Granul. switch after next access



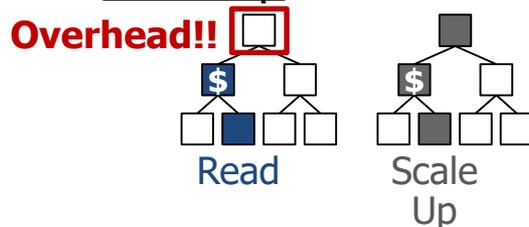
## \* Scale-up with WAR/WAW



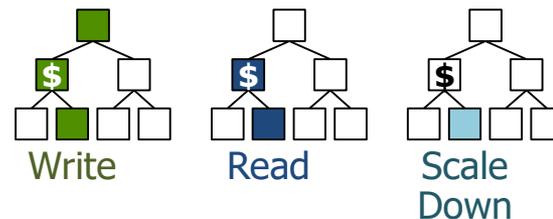
## \* Scale-up with RAW



## \* Scale-up with RAR



## \* Scale-down



# Pros of Matching Granularity

---

# Pros of Matching Granularity

---

- Proper granularity → Reduce security metadata

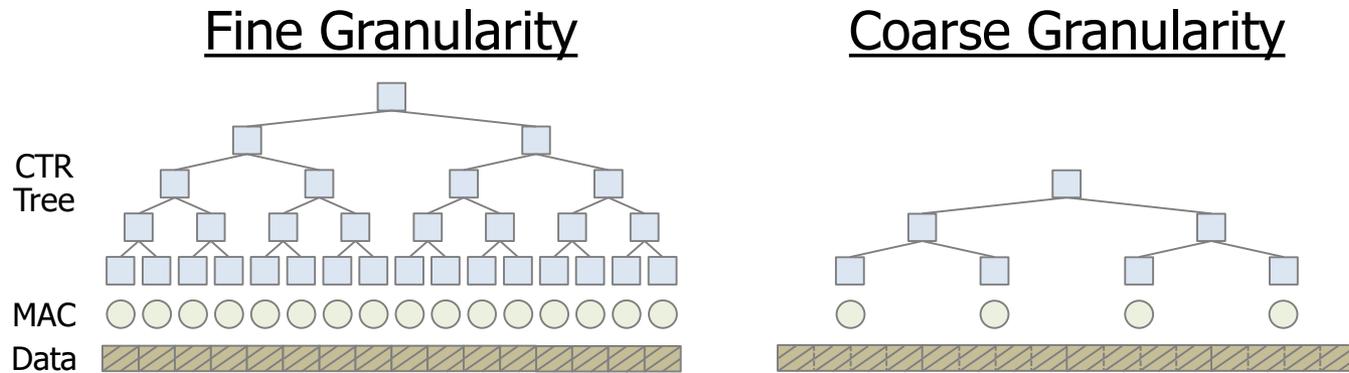
# Pros of Matching Granularity

---

- Proper granularity → Reduce security metadata
- Wrong granularity → Data load penalty

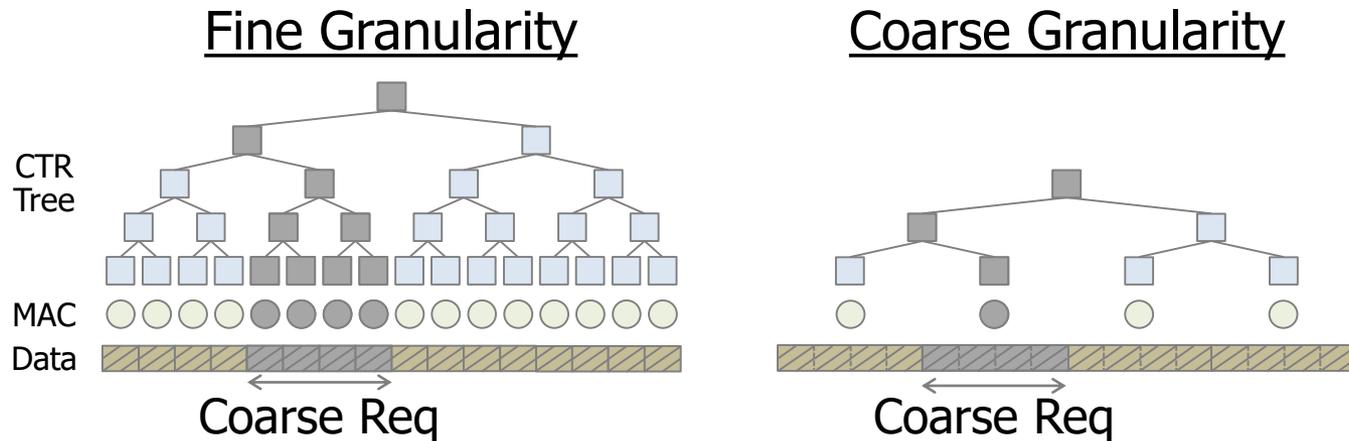
# Pros of Matching Granularity

- Proper granularity → Reduce security metadata
- Wrong granularity → Data load penalty



# Pros of Matching Granularity

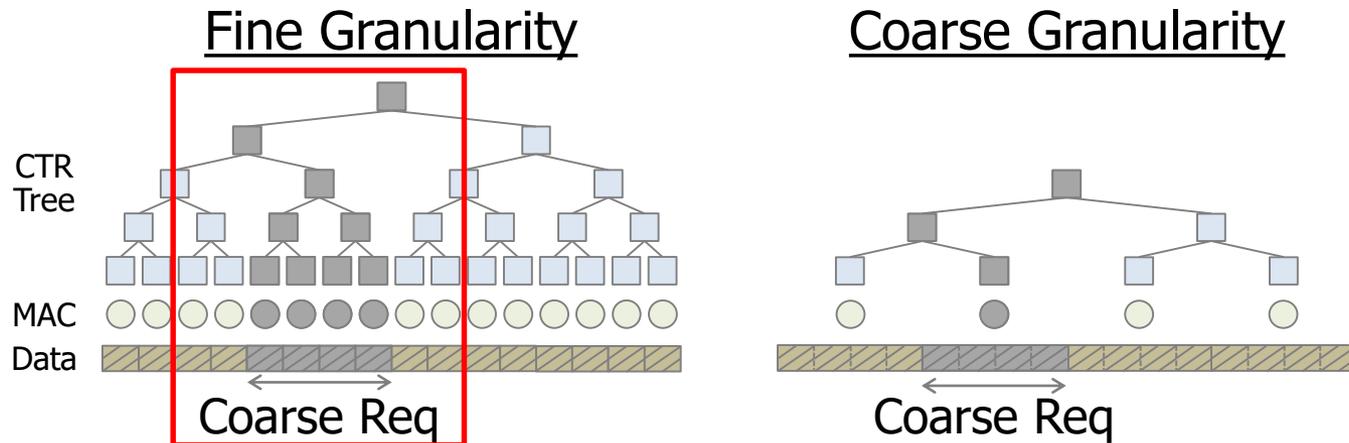
- Proper granularity → Reduce security metadata
- Wrong granularity → Data load penalty



Coarse Reqs

# Pros of Matching Granularity

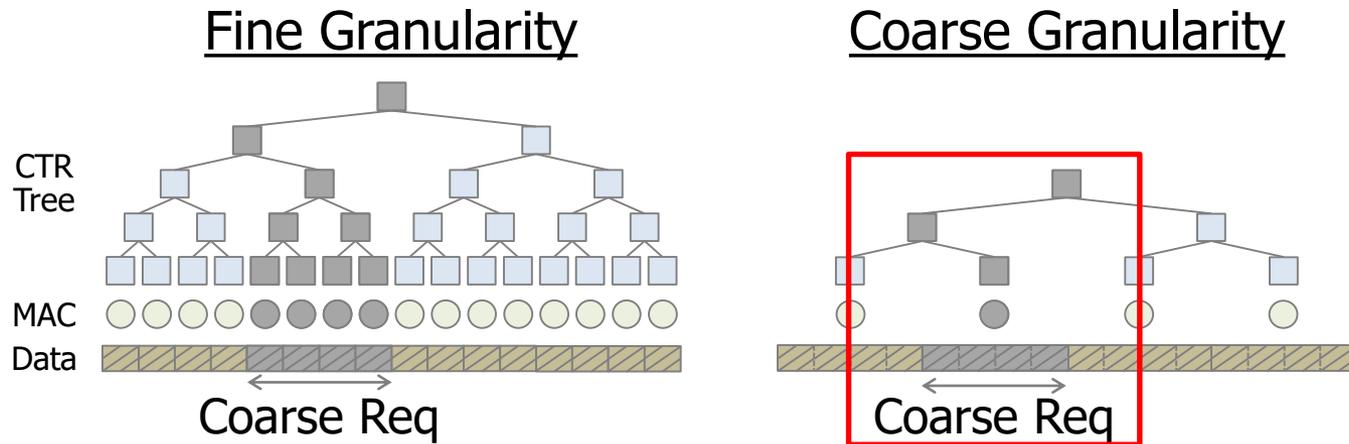
- Proper granularity → Reduce security metadata
- Wrong granularity → Data load penalty



Coarse Reqs    ▨ X 4    ● X 4    ■ X 9

# Pros of Matching Granularity

- Proper granularity → Reduce security metadata
- Wrong granularity → Data load penalty

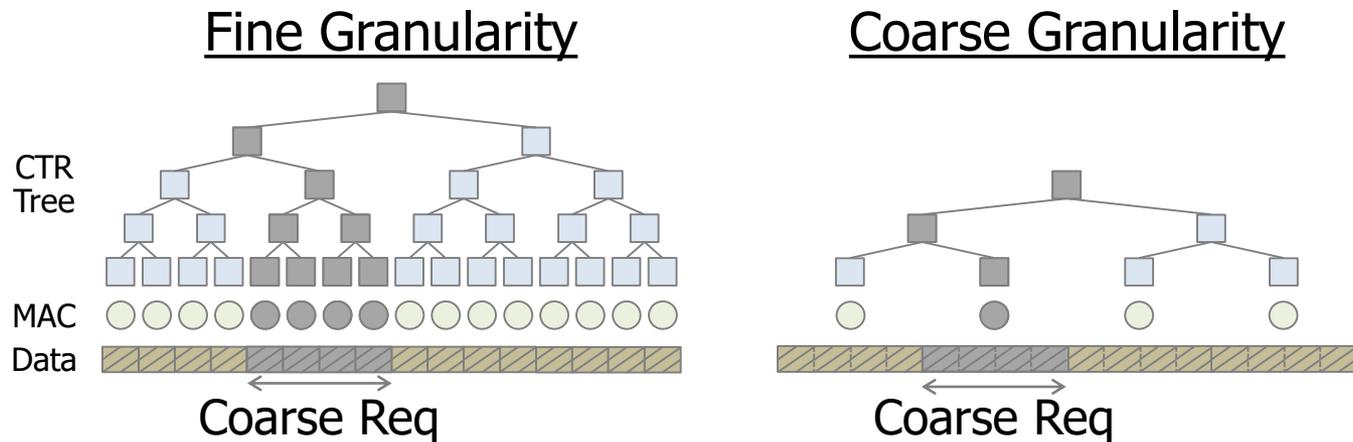


Coarse Reqs    ▨ X 4    ● X 4    ■ X 9

▨ X 4    ● X 1    ■ X 3

# Pros of Matching Granularity

- Proper granularity → Reduce security metadata
- Wrong granularity → Data load penalty

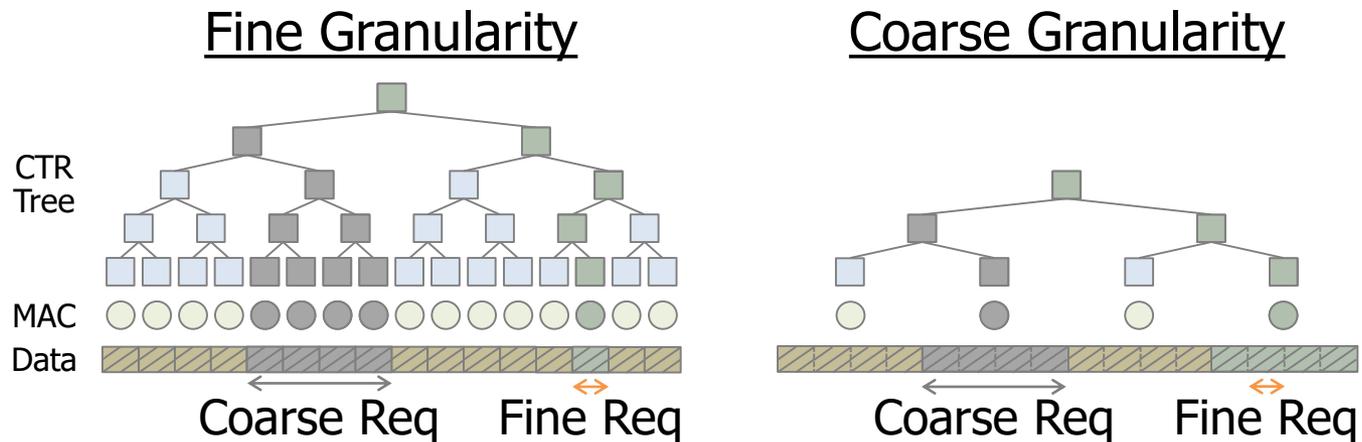


Coarse Reqs    ▨ X 4    ● X 4    ■ X 9

▨ X 4    ● X 1    ■ X 3 **Good**

# Pros of Matching Granularity

- Proper granularity → Reduce security metadata
- Wrong granularity → Data load penalty



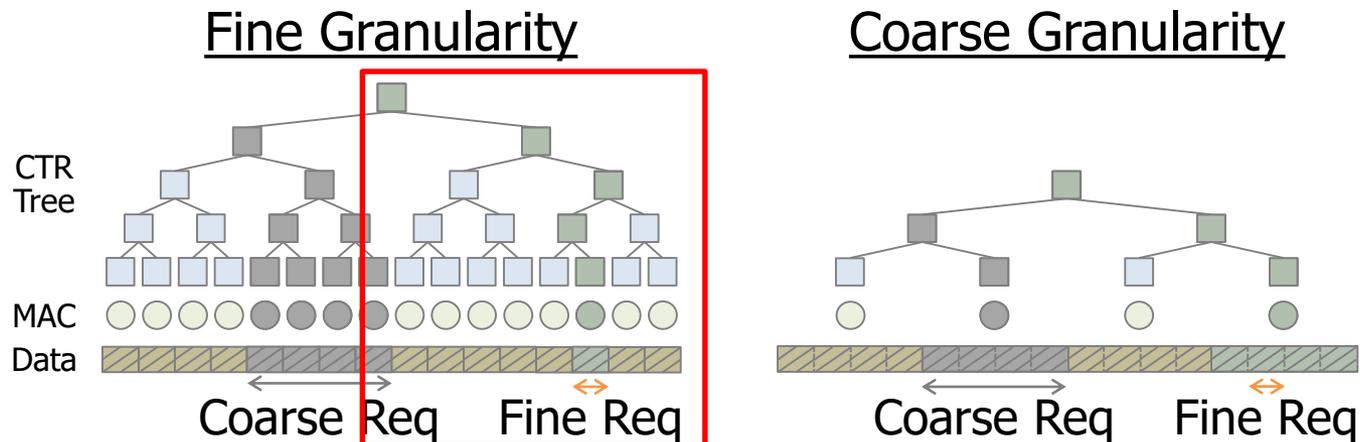
Coarse Reqs    ▨ X 4    ● X 4    ■ X 9

▨ X 4    ● X 1    ■ X 3 **Good**

Fine Reqs

# Pros of Matching Granularity

- Proper granularity → Reduce security metadata
- Wrong granularity → Data load penalty



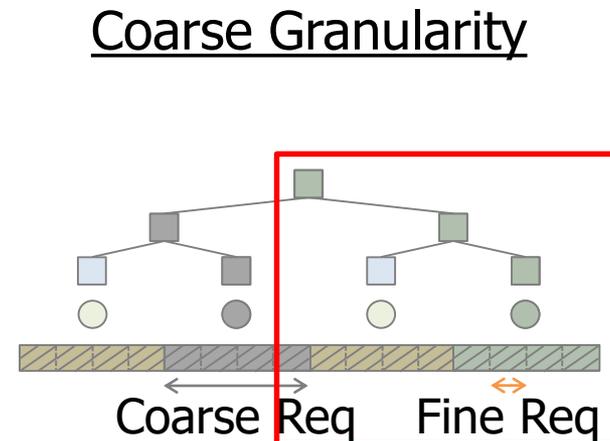
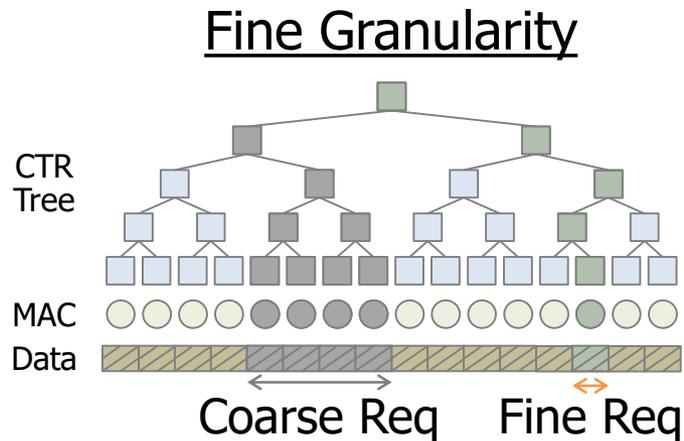
Coarse Reqs    ▨ X 4    ● X 4    ■ X 9

Fine Reqs     ▩ X 1    ● X 1    ■ X 5

▨ X 4    ● X 1    ■ X 3 **Good**

# Pros of Matching Granularity

- Proper granularity → Reduce security metadata
- Wrong granularity → Data load penalty



Coarse Reqs    ▨ X 4    ● X **4**    ■ X **9**

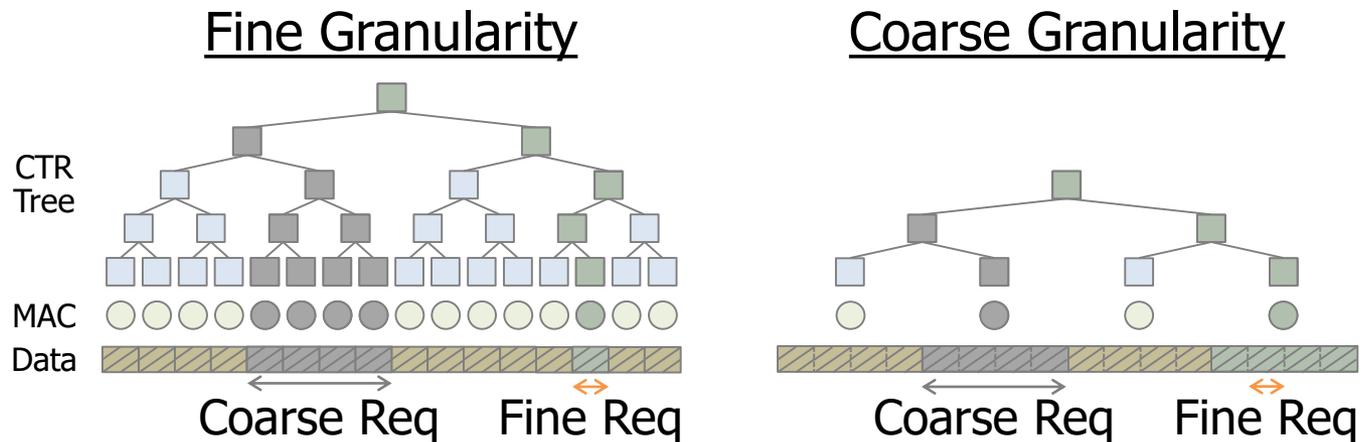
Fine Reqs      ▩ X 1    ● X 1    ■ X **5**

▨ X 4    ● X 1    ■ X 3    **Good**

▩ X **4**    ● X 1    ■ X 3

# Pros of Matching Granularity

- Proper granularity → Reduce security metadata
- Wrong granularity → Data load penalty



Coarse Reqs

▨ X 4   ● X 4   ■ X 9

▨ X 4   ● X 1   ■ X 3 **Good**

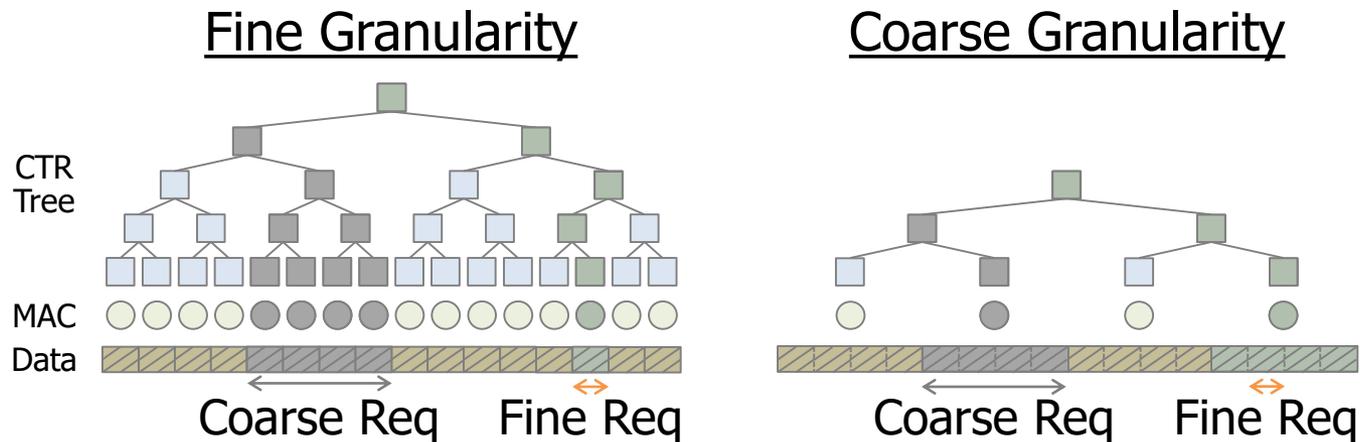
Fine Reqs

▨ X 1   ● X 1   ■ X 5 **Good**

▨ X 4   ● X 1   ■ X 3

# Pros of Matching Granularity

- Proper granularity → Reduce security metadata
- Wrong granularity → Data load penalty



Coarse Reqs

▨ X 4   ● X 4   ■ X 9

▨ X 4   ● X 1   ■ X 3 **Good**

Fine Reqs

▨ X 1   ● X 1   ■ X 5 **Good**

▨ X 4   ● X 1   ■ X 3

**Granularity-managed MAC&tree makes efficient memory protection**

# Prior Domain-specific Memory Protections

---

- No prior study using **integrity tree pruning** or **multi-granular MAC&counter**

# Prior Domain-specific Memory Protections

---

- No prior study using **integrity tree pruning** or **multi-granular MAC&counter**

## 1. Dual-granular & GPU-optimized Counter [1]



[1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)

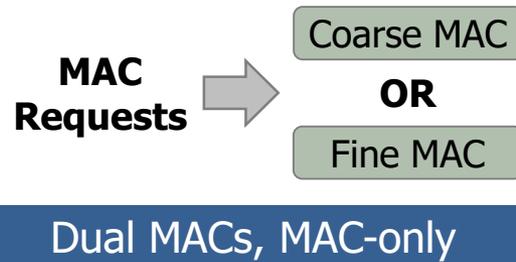
# Prior Domain-specific Memory Protections

- No prior study using **integrity tree pruning** or **multi-granular MAC&counter**

## 1. Dual-granular & GPU-optimized Counter [1]



## 2. Dual-granular MAC [2]



[1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)

[2] Adaptive Security Support for Heterogeneous Memory on GPUs (HPCA 2022)

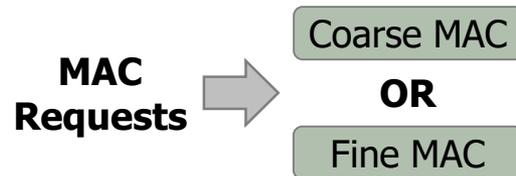
# Prior Domain-specific Memory Protections

- No prior study using **integrity tree pruning** or **multi-granular MAC&counter**

## 1. Dual-granular & GPU-optimized Counter [1]



## 2. Dual-granular MAC [2]



Dual MACs, MAC-only

## 3. S/W Counter [3-6]



Domain-specific

[1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)

[2] Adaptive Security Support for Heterogeneous Memory on GPUs (HPCA 2022)

[3] TNPU: Supporting Trusted Execution with Tree-less Integrity Protection for Neural Processing Unit (HPCA 2022)

[4] MGX: Near-zero Overhead Memory Protection for Data-intensive Accelerators (ISCA 2022)

[5] GuardNN: Secure Accelerator Architecture for Privacy-preserving Deep Learning (DAC 2022)

[6] SoftVN: Efficient Memory Protection via Software-provided Version Numbers (ISCA 2022)

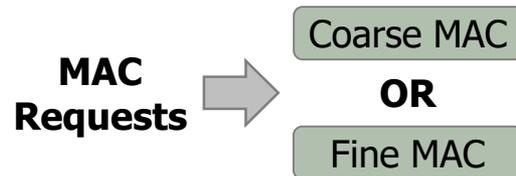
# Prior Domain-specific Memory Protections

- No prior study using **integrity tree pruning** or **multi-granular MAC&counter**

## 1. Dual-granular & GPU-optimized Counter [1]



## 2. Dual-granular MAC [2]



Dual MACs, MAC-only

## 3. S/W Counter [3-6]



Domain-specific

## 4. S/W MAC [4]



Domain-specific

[1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)

[2] Adaptive Security Support for Heterogeneous Memory on GPUs (HPCA 2022)

[3] TNPU: Supporting Trusted Execution with Tree-less Integrity Protection for Neural Processing Unit (HPCA 2022)

[4] MGX: Near-zero Overhead Memory Protection for Data-intensive Accelerators (ISCA 2022)

[5] GuardNN: Secure Accelerator Architecture for Privacy-preserving Deep Learning (DAC 2022)

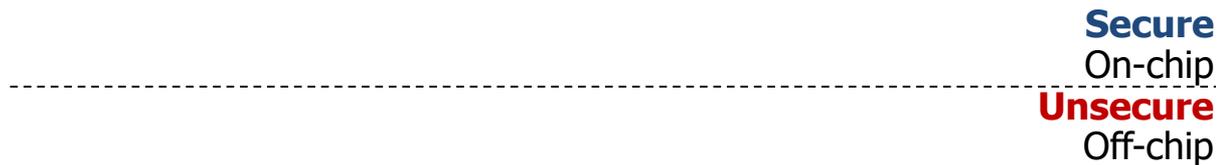
[6] SoftVN: Efficient Memory Protection via Software-provided Version Numbers (ISCA 2022)

# Detailed of Counter-mode Protection

---

# Detailed of Counter-mode Protection

---



**Ciphertext**

# Detailed of Counter-mode Protection

---

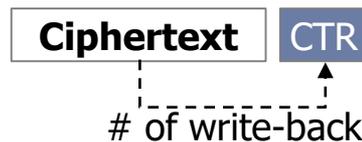
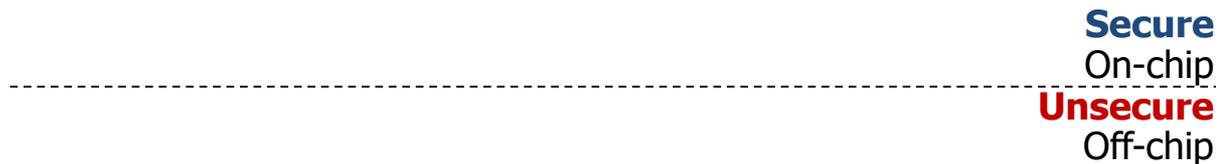
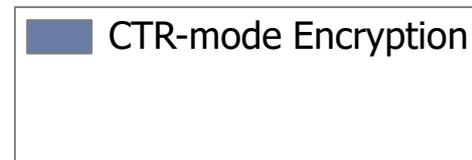
- CTR-mode encryption: confidentiality



# Detailed of Counter-mode Protection

---

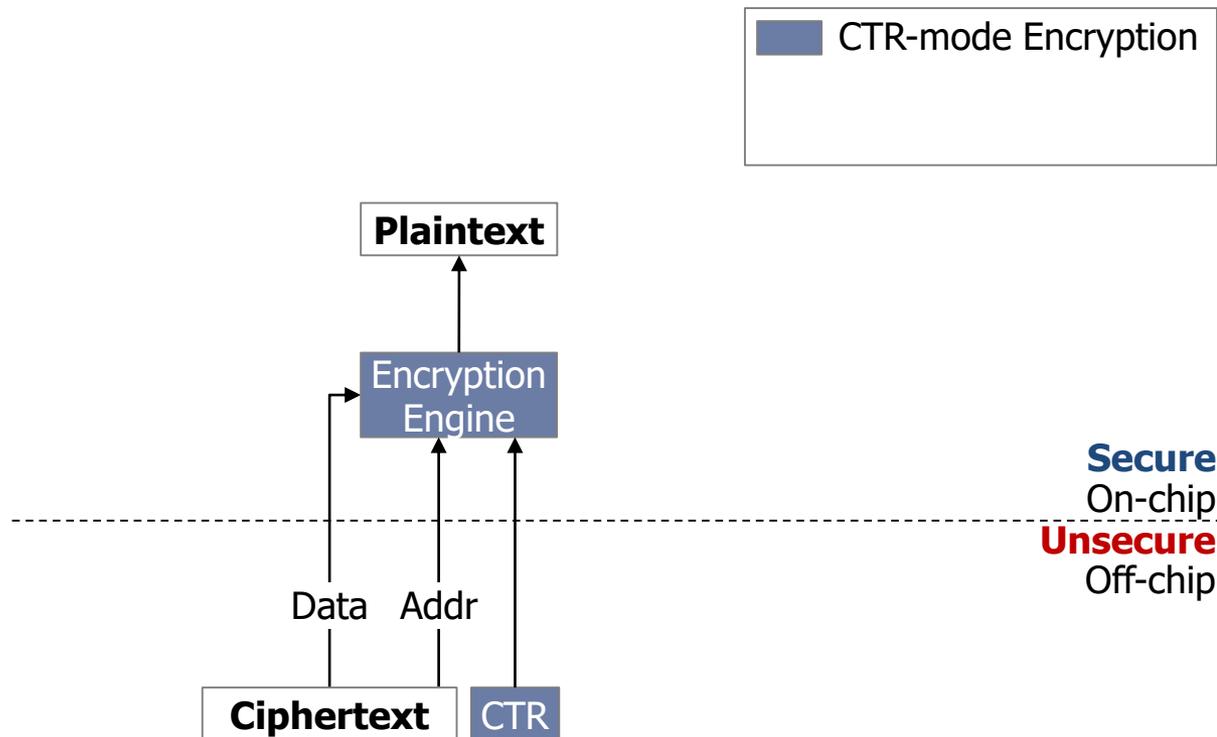
- CTR-mode encryption: confidentiality



# Detailed of Counter-mode Protection

---

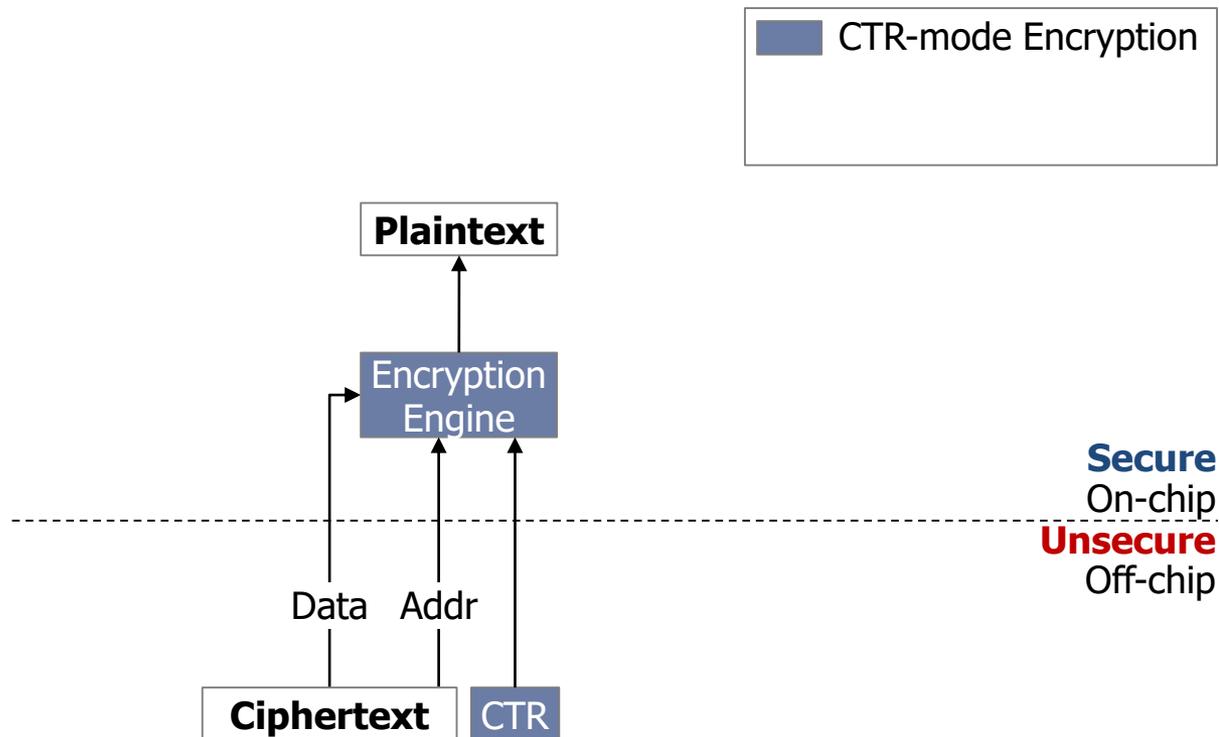
- CTR-mode encryption: confidentiality



# Detailed of Counter-mode Protection

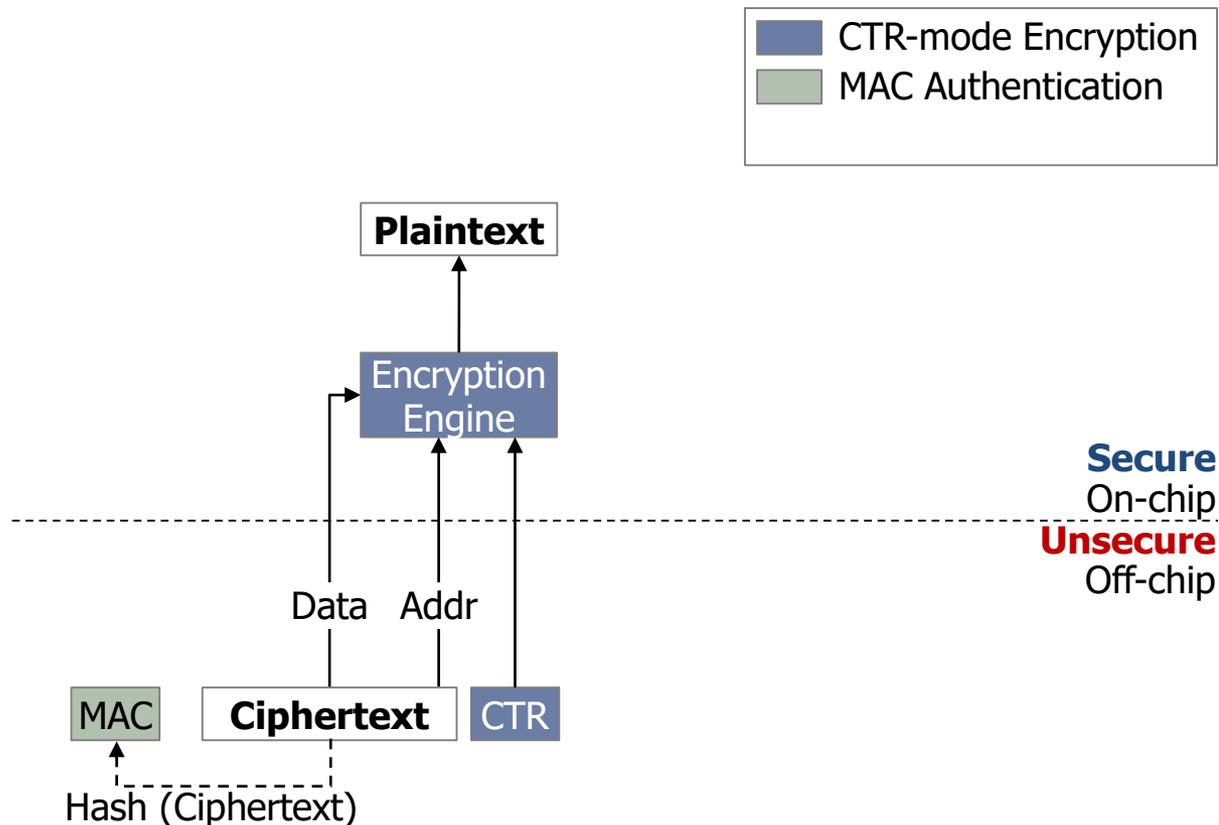
---

- CTR-mode encryption: confidentiality
- MAC authentication: value-integrity



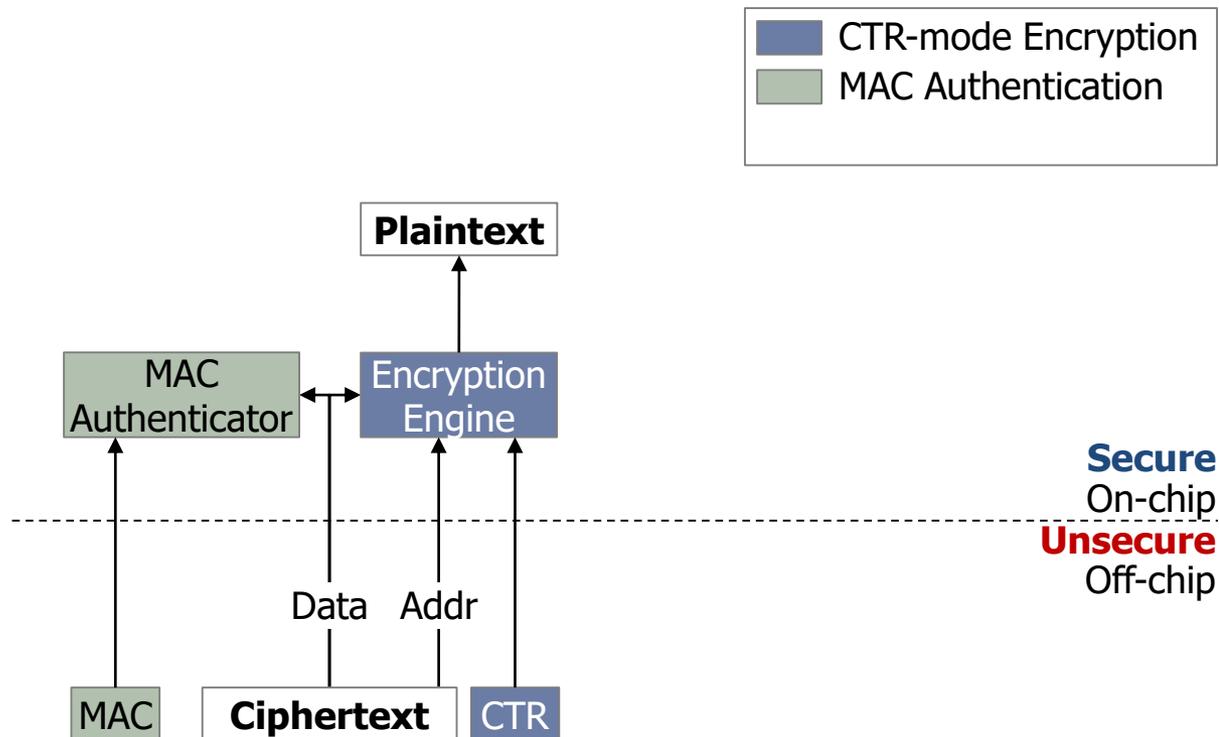
# Detailed of Counter-mode Protection

- CTR-mode encryption: confidentiality
- MAC authentication: value-integrity



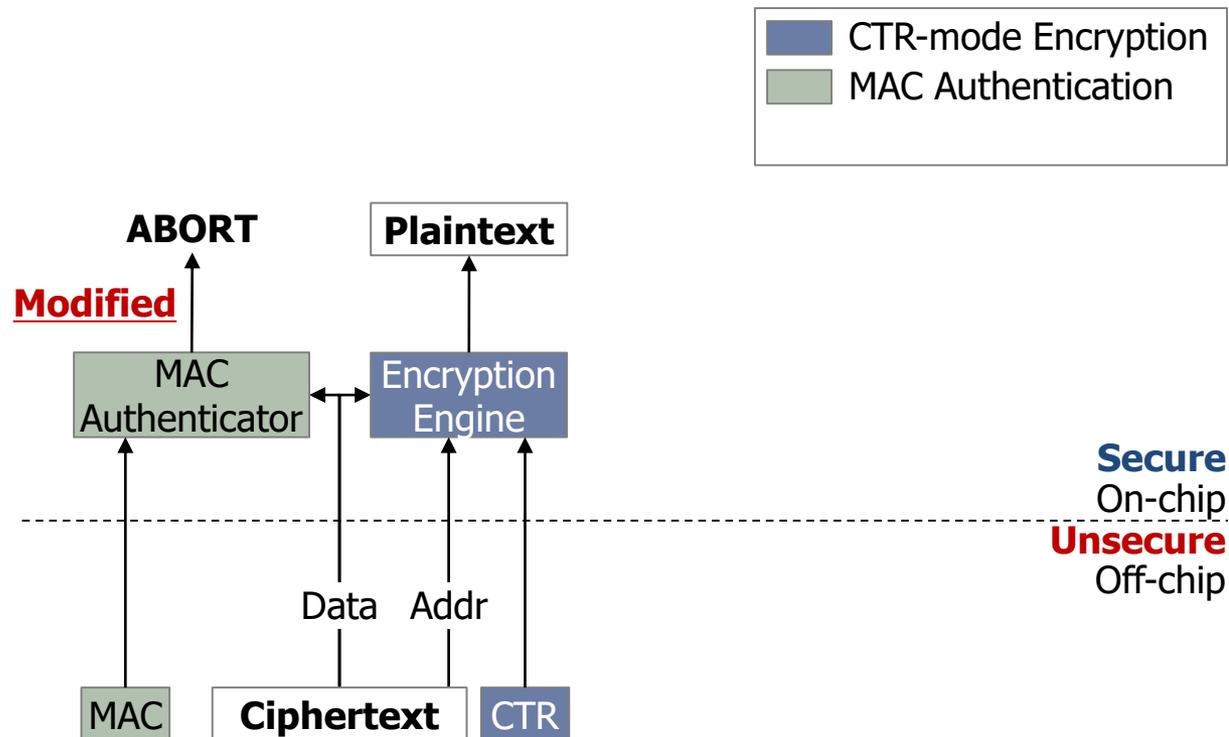
# Detailed of Counter-mode Protection

- CTR-mode encryption: confidentiality
- MAC authentication: value-integrity



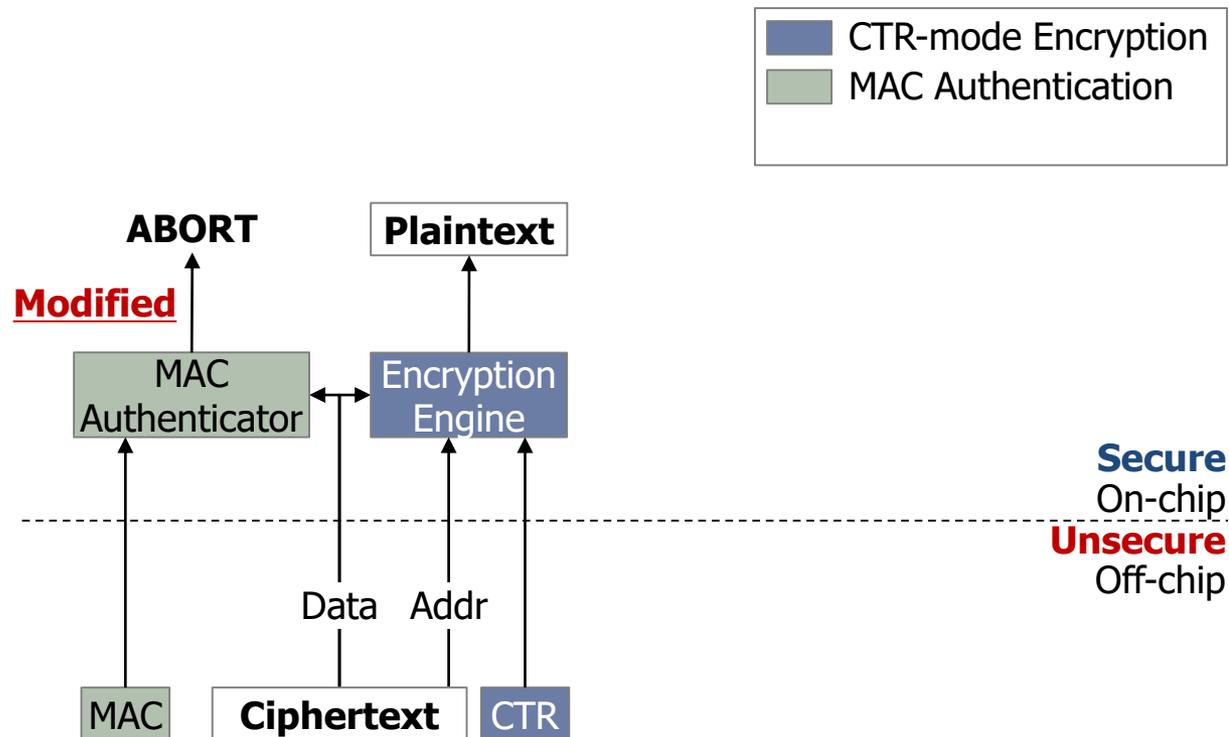
# Detailed of Counter-mode Protection

- CTR-mode encryption: confidentiality
- MAC authentication: value-integrity



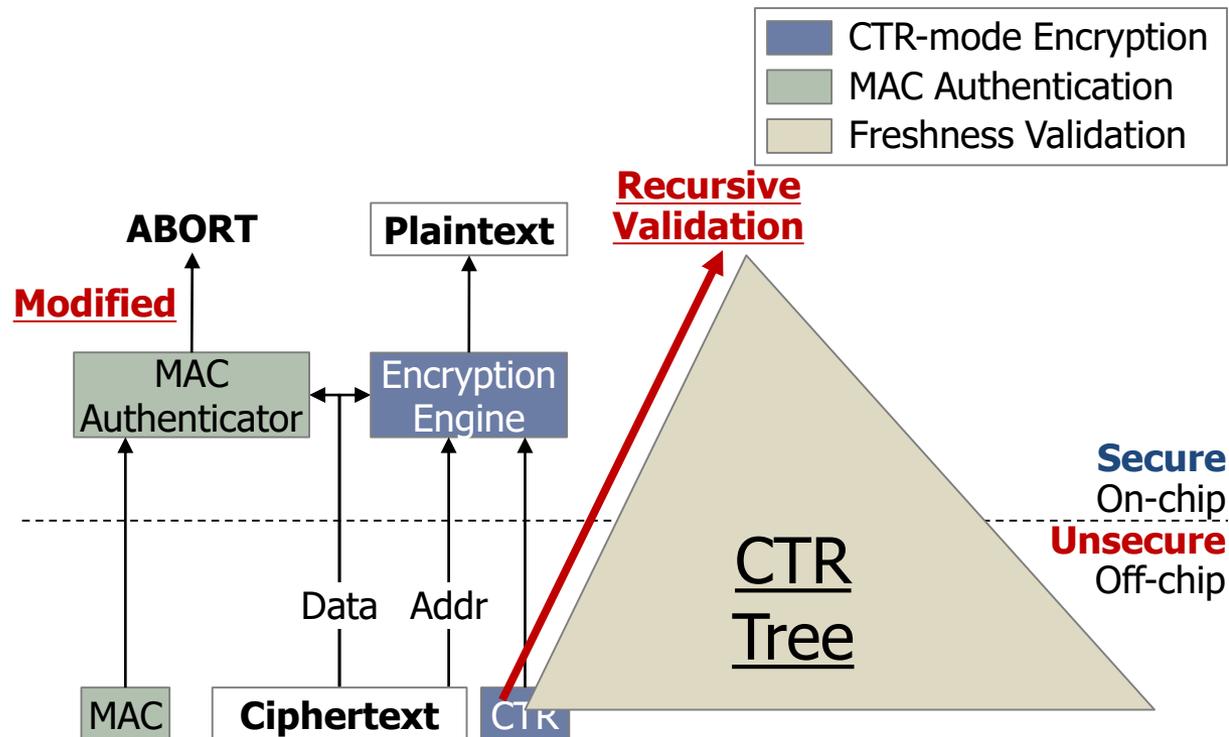
# Detailed of Counter-mode Protection

- CTR-mode encryption: confidentiality
- MAC authentication: value-integrity
- Freshness validation: replay-attack



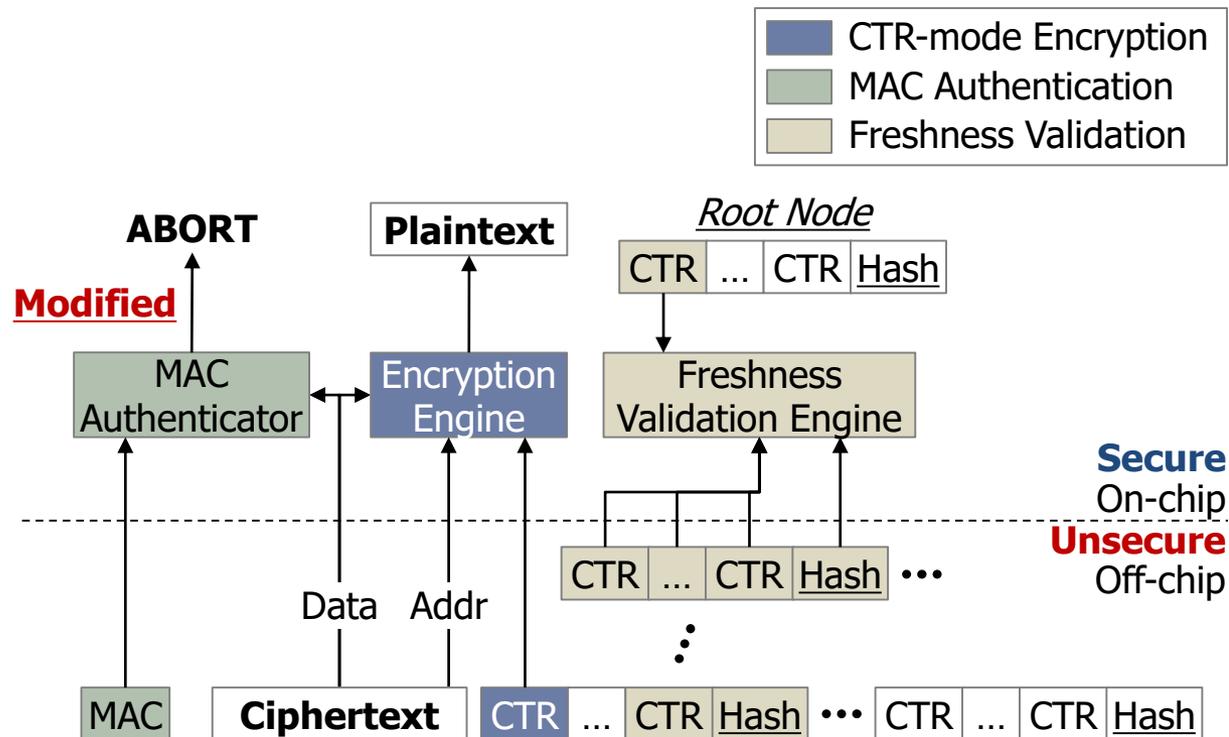
# Detailed of Counter-mode Protection

- CTR-mode encryption: confidentiality
- MAC authentication: value-integrity
- Freshness validation: replay-attack



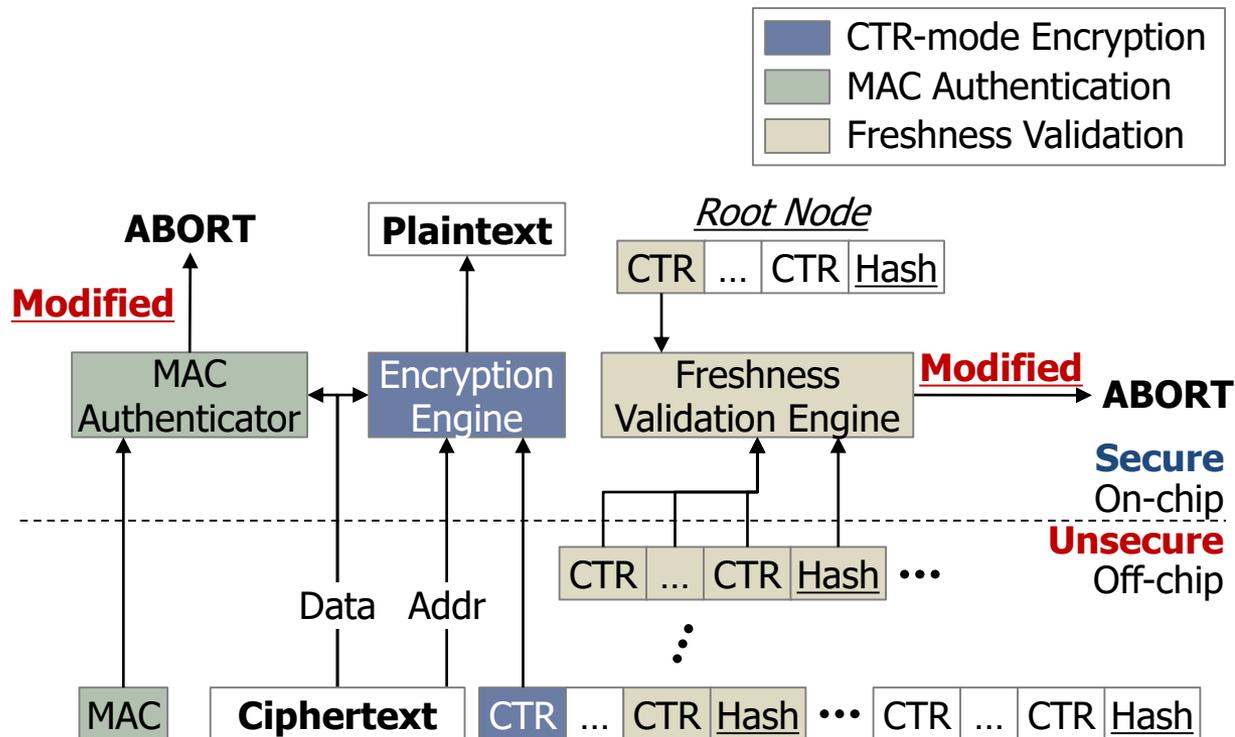
# Detailed of Counter-mode Protection

- CTR-mode encryption: confidentiality
- MAC authentication: value-integrity
- Freshness validation: replay-attack



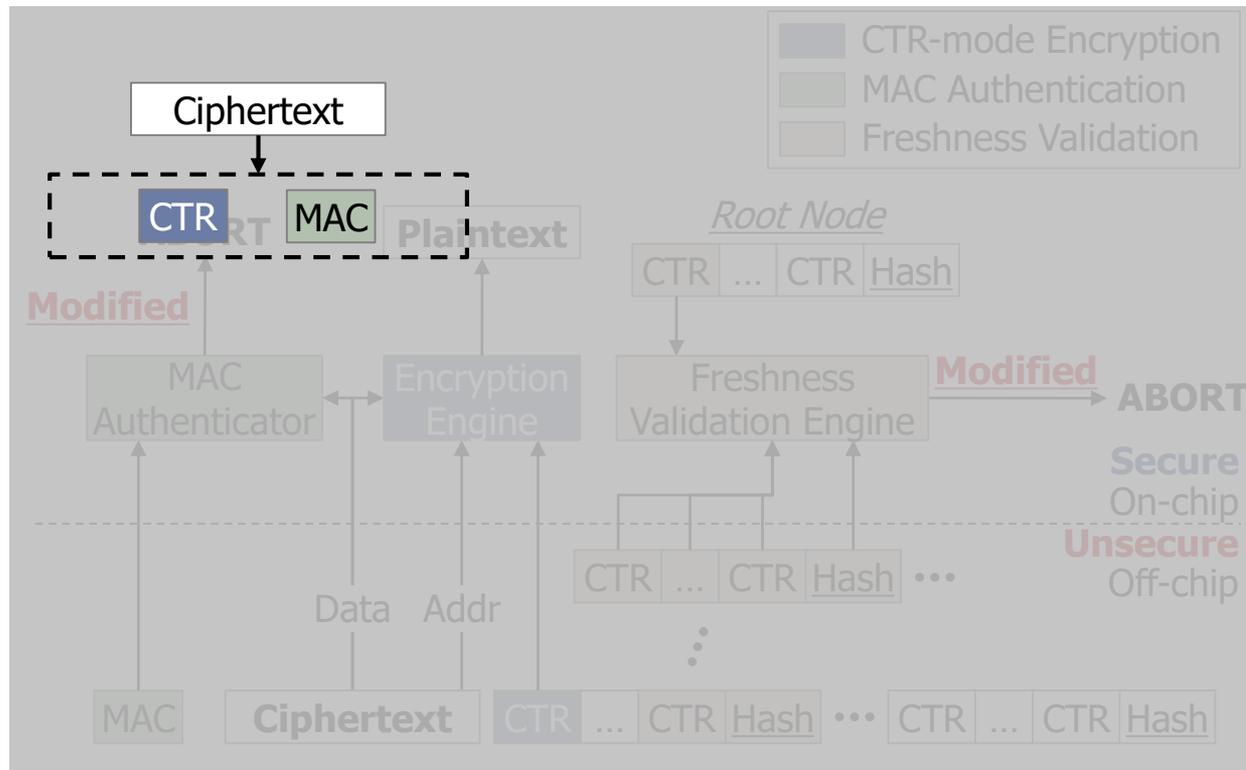
# Detailed of Counter-mode Protection

- CTR-mode encryption: confidentiality
- MAC authentication: value-integrity
- Freshness validation: replay-attack



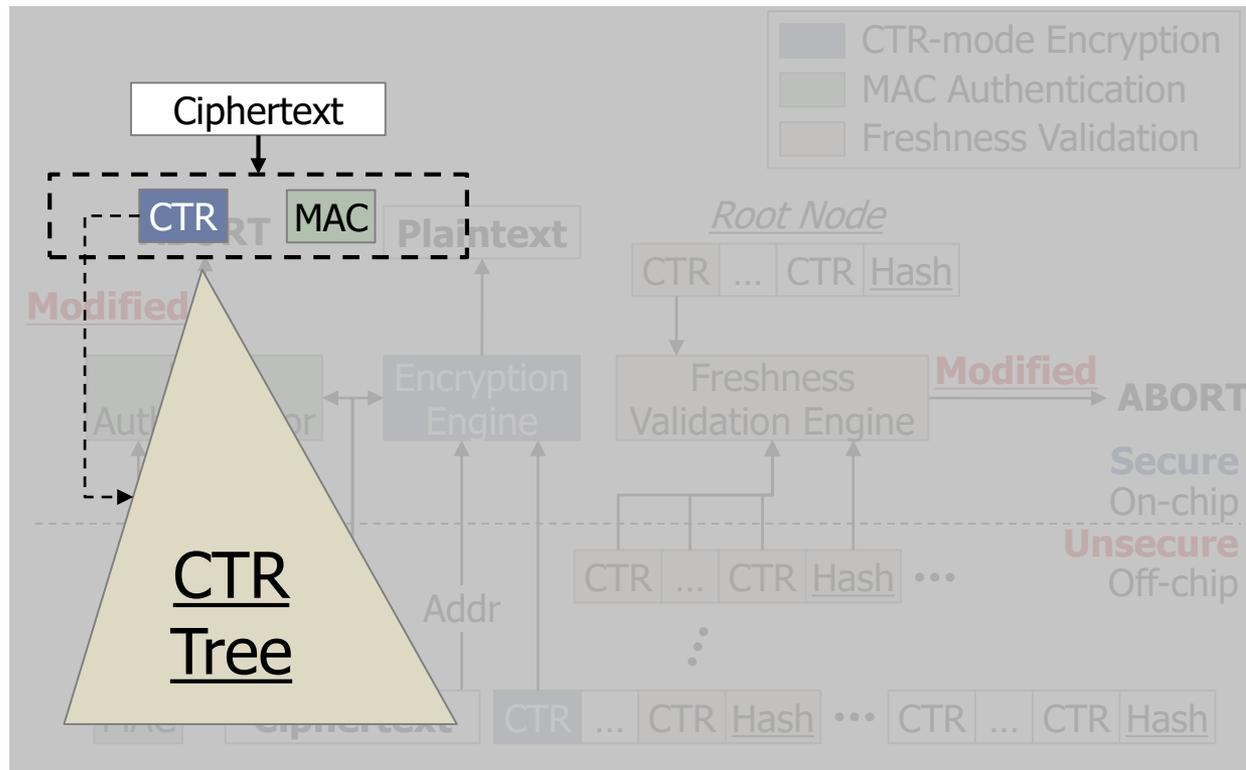
# Detailed of Counter-mode Protection

- CTR-mode encryption: confidentiality
- MAC authentication: value-integrity
- Freshness validation: replay-attack



# Detailed of Counter-mode Protection

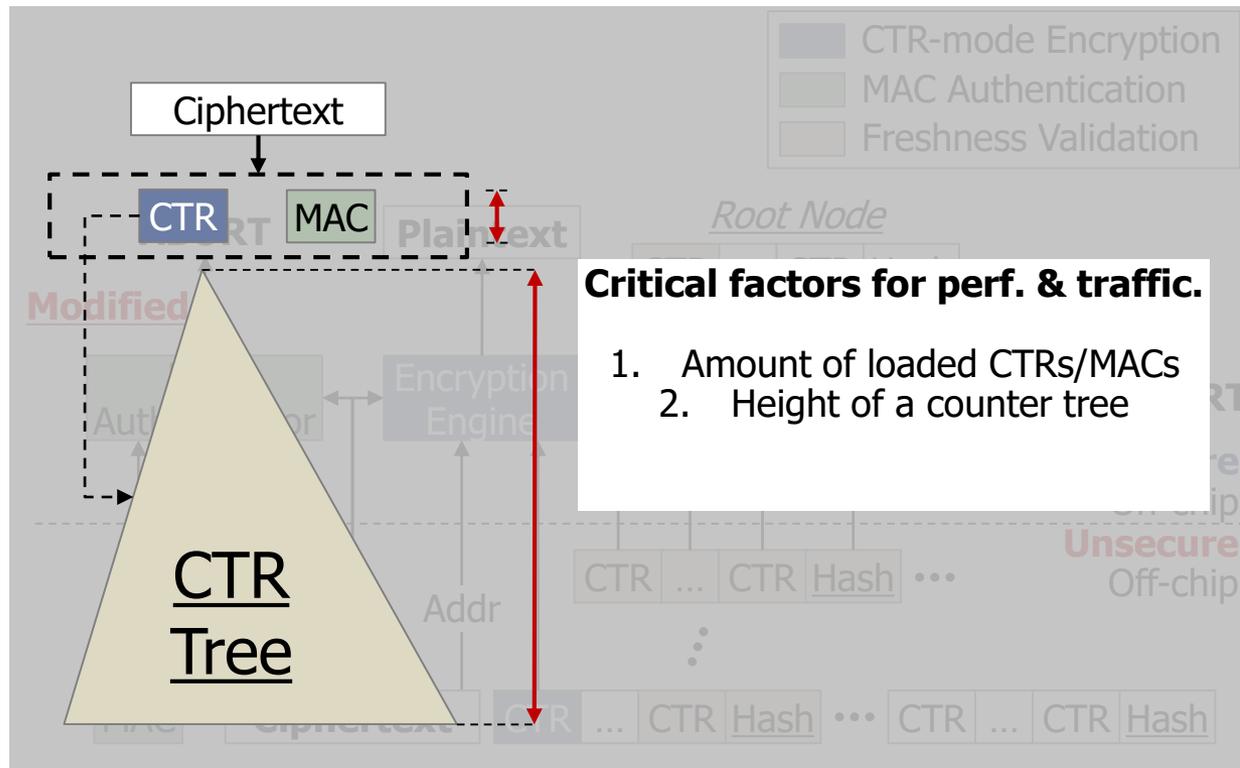
- CTR-mode encryption: confidentiality
- MAC authentication: value-integrity
- Freshness validation: replay-attack





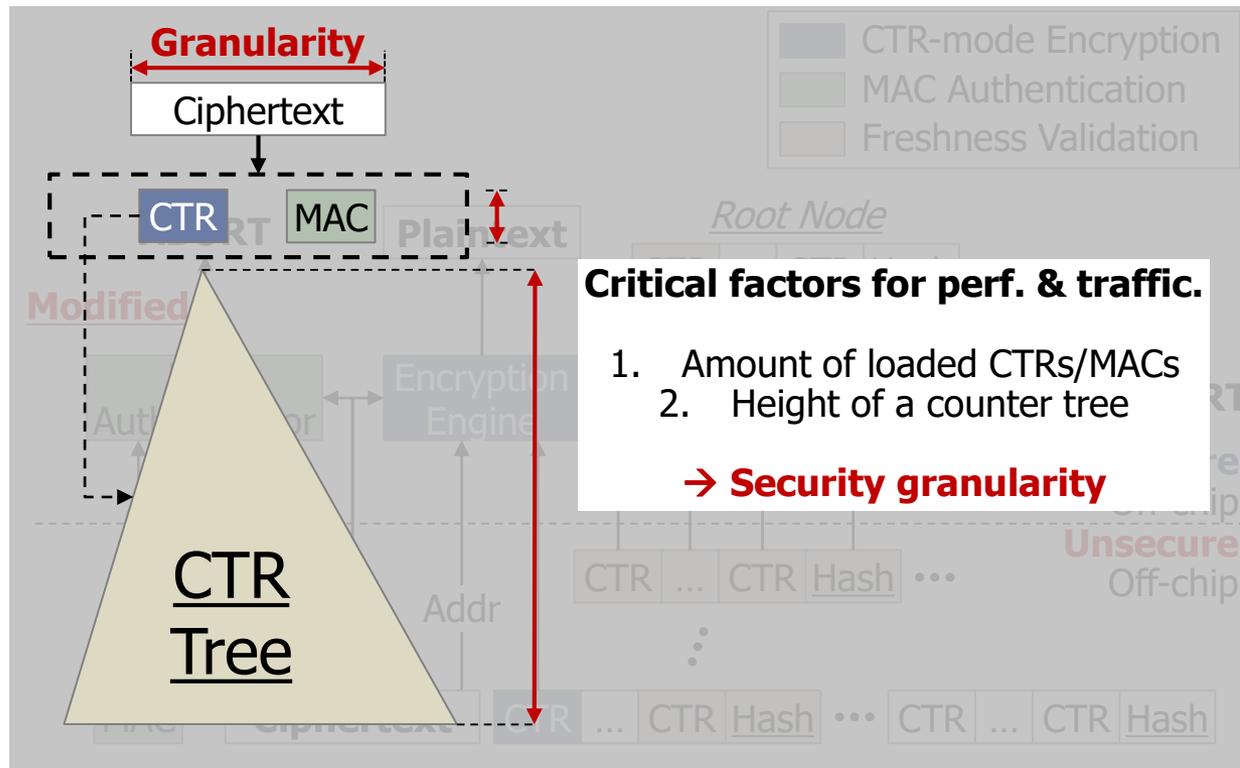
# Detailed of Counter-mode Protection

- CTR-mode encryption: confidentiality
- MAC authentication: value-integrity
- Freshness validation: replay-attack



# Detailed of Counter-mode Protection

- CTR-mode encryption: confidentiality
- MAC authentication: value-integrity
- Freshness validation: replay-attack



# Combined with Prior CTR Tree Optimization

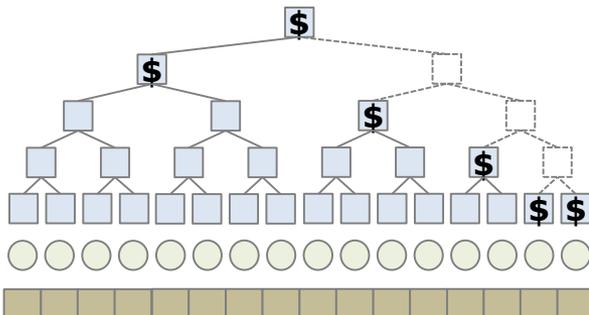
---



# Combined with Prior CTR Tree Optimization

- Prior hotness-based integrity tree optimization scheme (Subtree optimization)<sup>[1-4]</sup>
  - Caching highly used roots of subtrees

## Prior Subtree Optimization



[1] Bonsai Merkle Forests: Efficiently Achieving Crash Consistency in Secure Persistent Memory (MICRO 2021)

[2] Scalable Memory Protection in the PENGLAI Enclave (OSDI 2021)

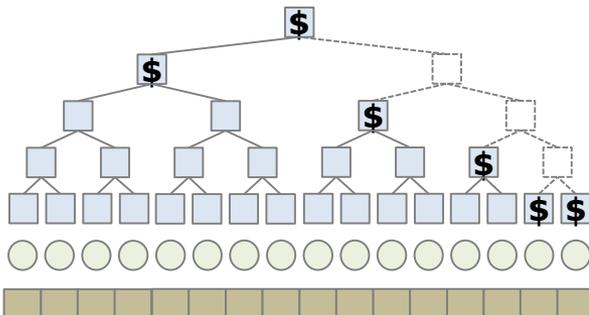
[3] Efficient Distributed Secure Memory with Migratable Merkle Tree (HPCA 2023)

[4] Data Enclave: A Data-Centric Trusted Execution Environment (HPCA 2024)

# Combined with Prior CTR Tree Optimization

- Prior hotness-based integrity tree optimization scheme (Subtree optimization)<sup>[1-4]</sup>
  - Caching highly used roots of subtrees
  - Pruned unused nodes

## Prior Subtree Optimization



[1] Bonsai Merkle Forests: Efficiently Achieving Crash Consistency in Secure Persistent Memory (MICRO 2021)

[2] Scalable Memory Protection in the PENGLAI Enclave (OSDI 2021)

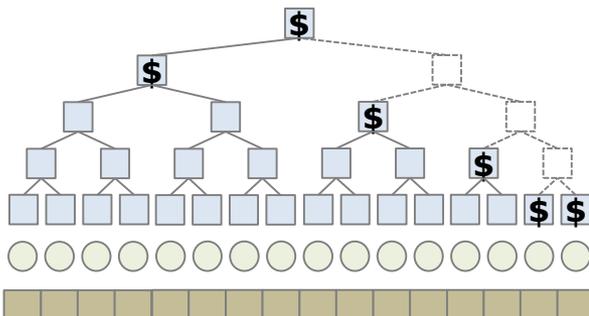
[3] Efficient Distributed Secure Memory with Migratable Merkle Tree (HPCA 2023)

[4] Data Enclave: A Data-Centric Trusted Execution Environment (HPCA 2024)

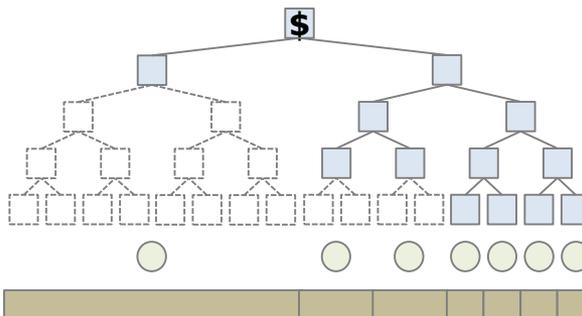
# Combined with Prior CTR Tree Optimization

- Prior hotness-based integrity tree optimization scheme (Subtree optimization)<sup>[1-4]</sup>
  - Caching highly used roots of subtrees
  - Pruned unused nodes

Prior Subtree Optimization



Ours



[1] Bonsai Merkle Forests: Efficiently Achieving Crash Consistency in Secure Persistent Memory (MICRO 2021)

[2] Scalable Memory Protection in the PENGLAI Enclave (OSDI 2021)

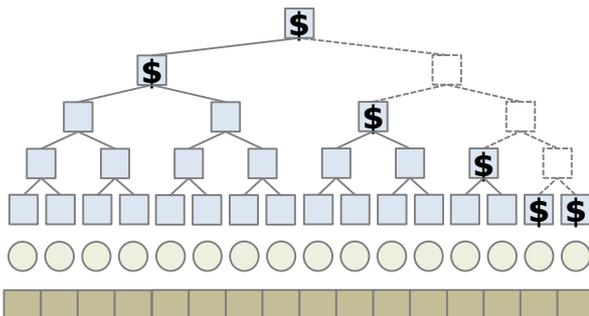
[3] Efficient Distributed Secure Memory with Migratable Merkle Tree (HPCA 2023)

[4] Data Enclave: A Data-Centric Trusted Execution Environment (HPCA 2024)

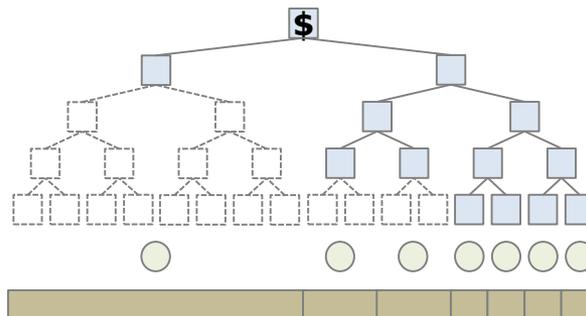
# Combined with Prior CTR Tree Optimization

- Prior hotness-based integrity tree optimization scheme (Subtree optimization)<sup>[1-4]</sup>
  - Caching highly used roots of subtrees
  - Pruned unused nodes

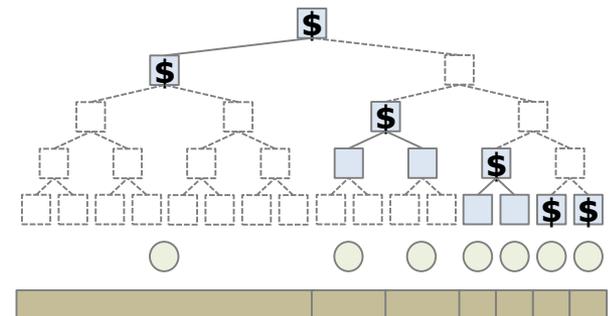
Prior Subtree Optimization



Ours



Ours + Prior



[1] Bonsai Merkle Forests: Efficiently Achieving Crash Consistency in Secure Persistent Memory (MICRO 2021)

[2] Scalable Memory Protection in the PENGLAI Enclave (OSDI 2021)

[3] Efficient Distributed Secure Memory with Migratable Merkle Tree (HPCA 2023)

[4] Data Enclave: A Data-Centric Trusted Execution Environment (HPCA 2024)



# Design of Multi-granular MAC&Tree

---

# Design of Multi-granular MAC&Tree

---

## 1. Granularity Switching

# Design of Multi-granular MAC&Tree

---

1. Granularity Switching
2. Granularity Detection

# Design of Multi-granular MAC&Tree

---

1. Granularity Switching

2. Granularity Detection

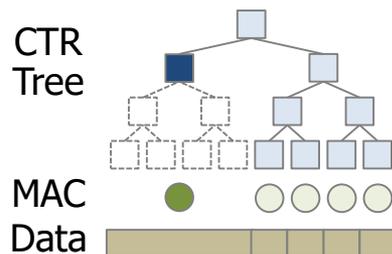
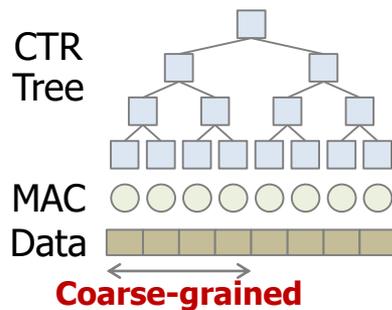
3. Multi-granularity  
Memory Protection

# Design of Multi-granular MAC&Tree

1. Granularity Switching

2. Granularity Detection

3. Multi-granularity  
Memory Protection



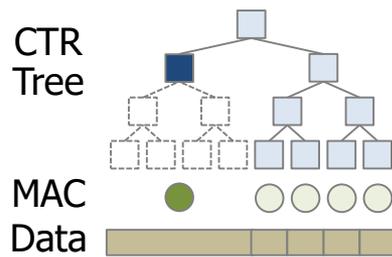
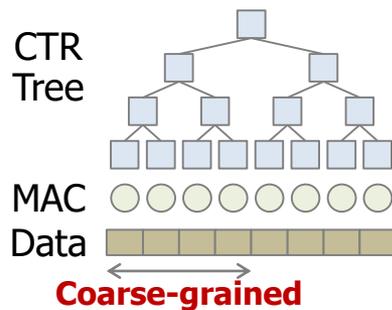
MAC/CTR Merging  
& CTR Tree Pruning

# Design of Multi-granular MAC&Tree

1. Granularity Switching

2. Granularity Detection

3. Multi-granularity  
Memory Protection

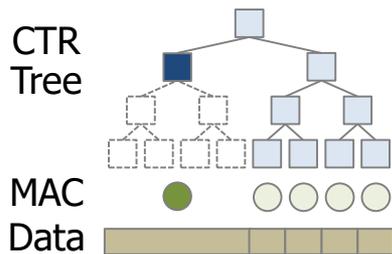
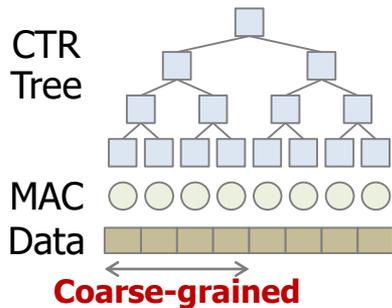


MAC/CTR Merging  
& CTR Tree Pruning

Dynamic Access Tracking

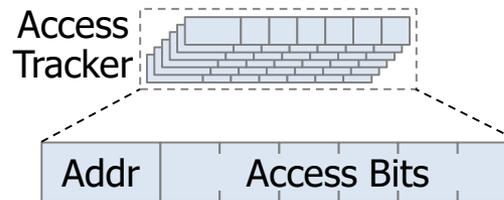
# Design of Multi-granular MAC&Tree

## 1. Granularity Switching



MAC/CTR Merging  
& CTR Tree Pruning

## 2. Granularity Detection

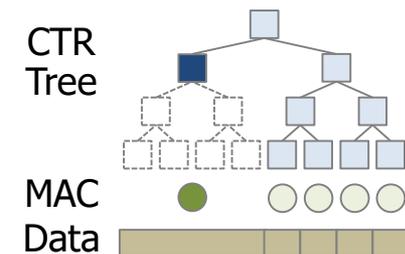
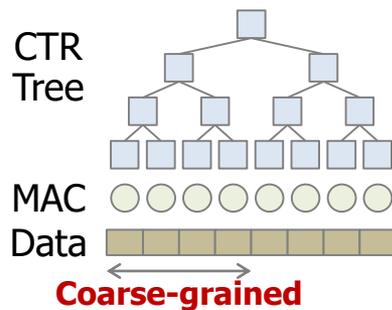


Dynamic Access Tracking

## 3. Multi-granularity Memory Protection

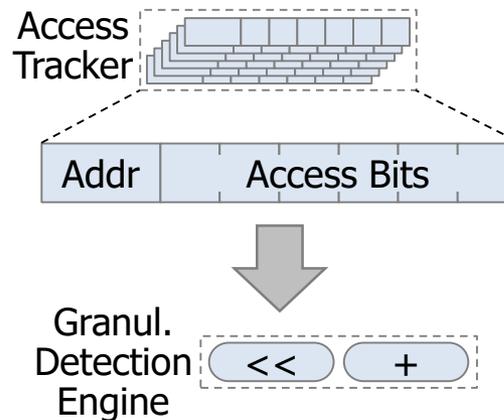
# Design of Multi-granular MAC&Tree

## 1. Granularity Switching



MAC/CTR Merging  
& CTR Tree Pruning

## 2. Granularity Detection

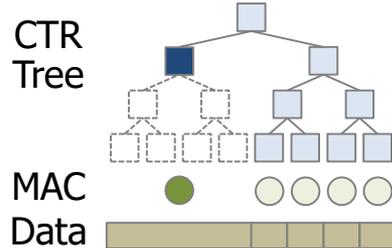
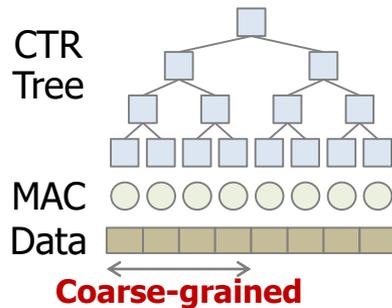


Dynamic Access Tracking

## 3. Multi-granularity Memory Protection

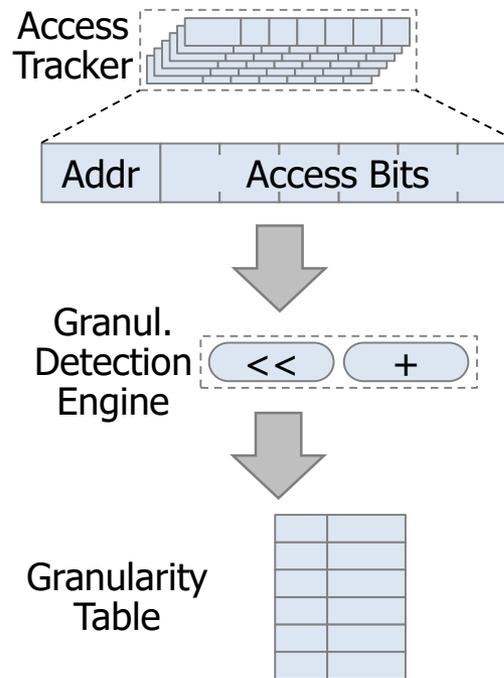
# Design of Multi-granular MAC&Tree

## 1. Granularity Switching



MAC/CTR Merging  
& CTR Tree Pruning

## 2. Granularity Detection

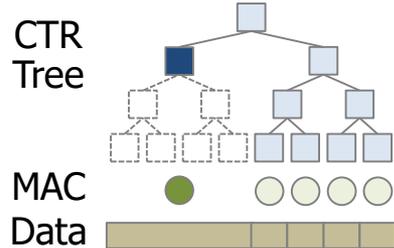
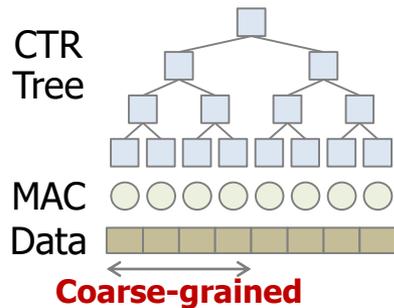


Dynamic Access Tracking

## 3. Multi-granularity Memory Protection

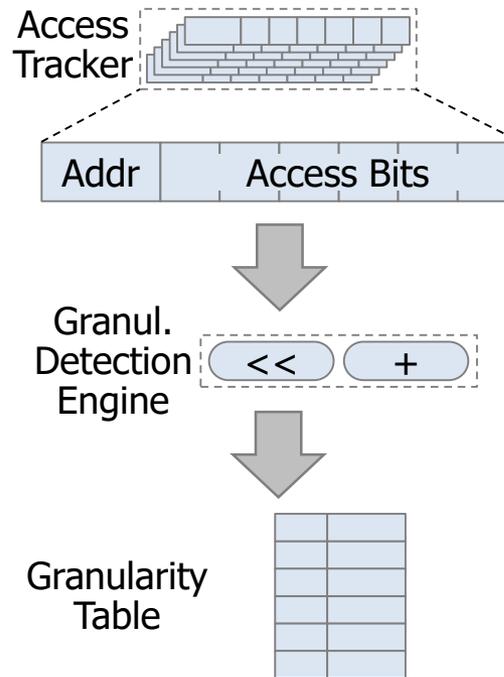
# Design of Multi-granular MAC&Tree

## 1. Granularity Switching



MAC/CTR Merging  
& CTR Tree Pruning

## 2. Granularity Detection



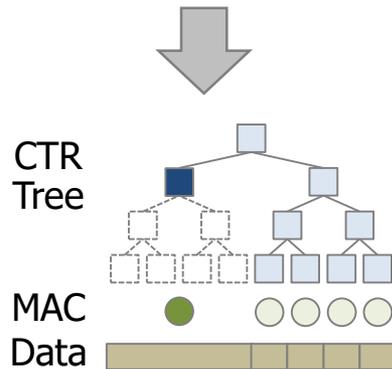
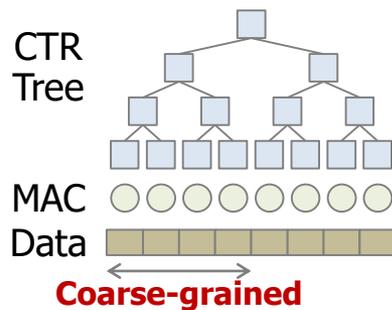
Dynamic Access Tracking

## 3. Multi-granularity Memory Protection

Granularity-aware Protection

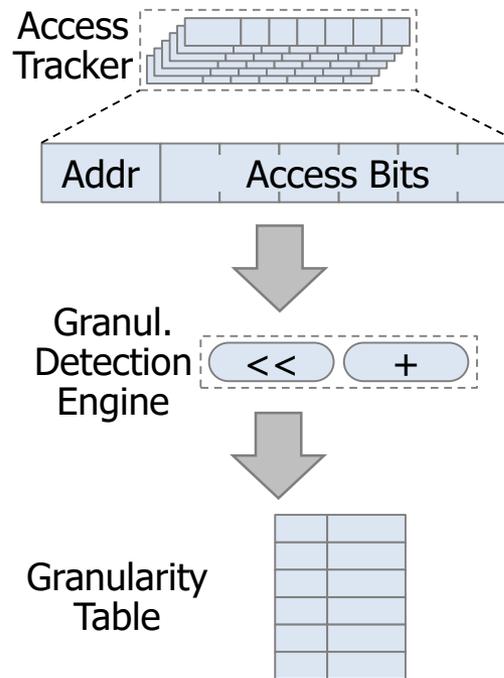
# Design of Multi-granular MAC&Tree

## 1. Granularity Switching



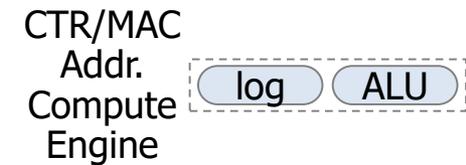
MAC/CTR Merging  
& CTR Tree Pruning

## 2. Granularity Detection



Dynamic Access Tracking

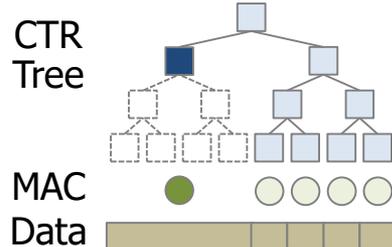
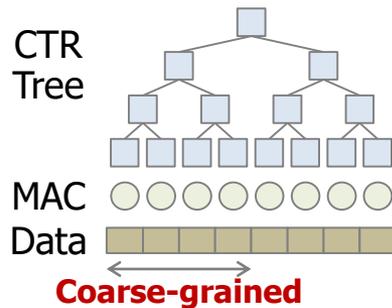
## 3. Multi-granularity Memory Protection



Granularity-aware Protection

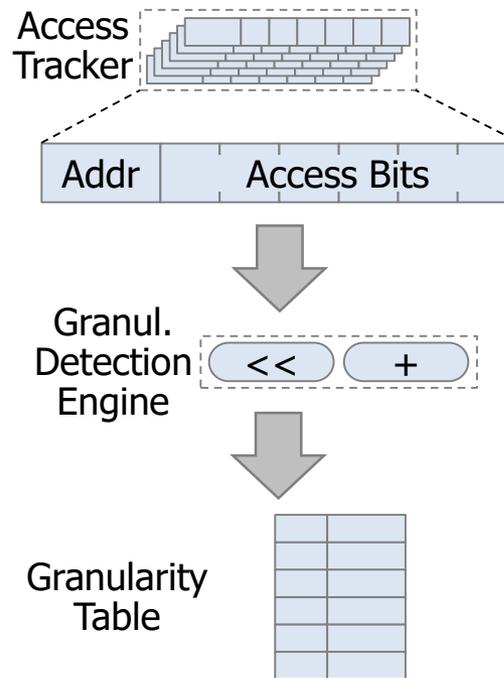
# Design of Multi-granular MAC&Tree

## 1. Granularity Switching



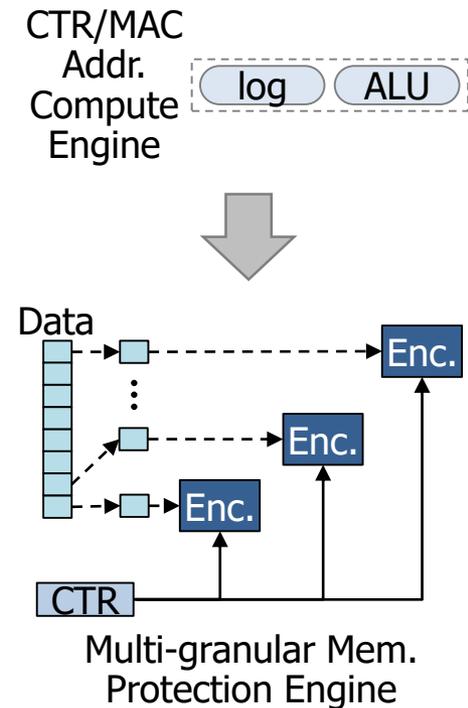
MAC/CTR Merging  
& CTR Tree Pruning

## 2. Granularity Detection



Dynamic Access Tracking

## 3. Multi-granularity Memory Protection



Granularity-aware Protection

# Design of Multi-granular MAC&Tree

---

# Design of Multi-granular MAC&Tree

---

Granularity  
Switching



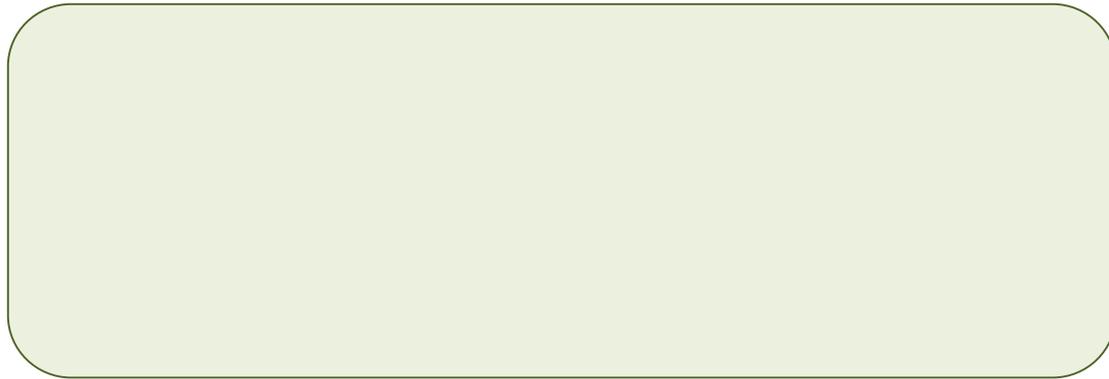
# Design of Multi-granular MAC&Tree

---

Granularity  
Switching



Granularity Detection



# Design of Multi-granular MAC&Tree

---

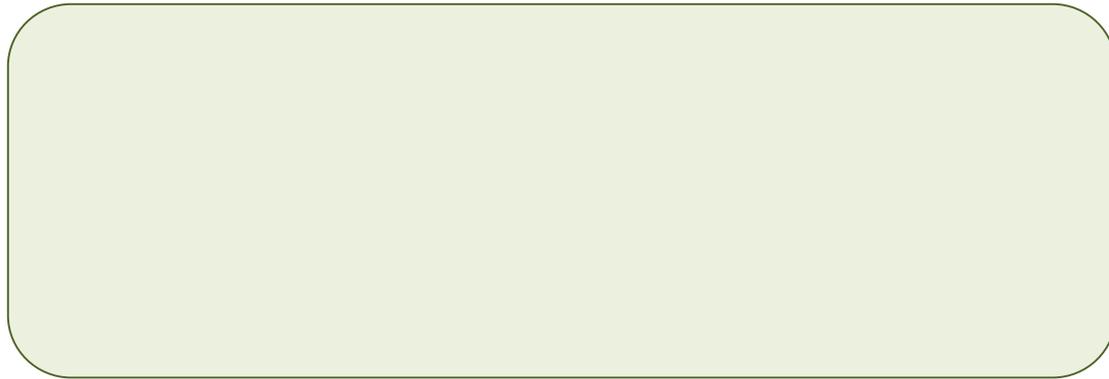
Granularity  
Switching



Multi-granularity based Memory Protection



Granularity Detection



# Design of Multi-granular MAC&Tree

---

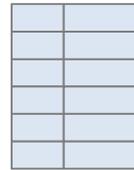
Granularity  
Switching



Multi-granularity based Memory Protection



Granul.  
Table



Granularity Detection



# Design of Multi-granular MAC&Tree

---

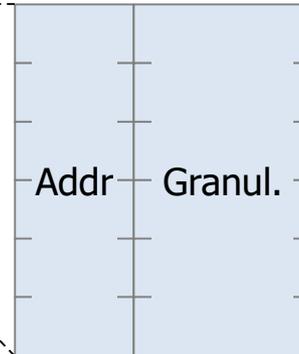
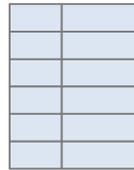
Granularity Switching



Multi-granularity based Memory Protection



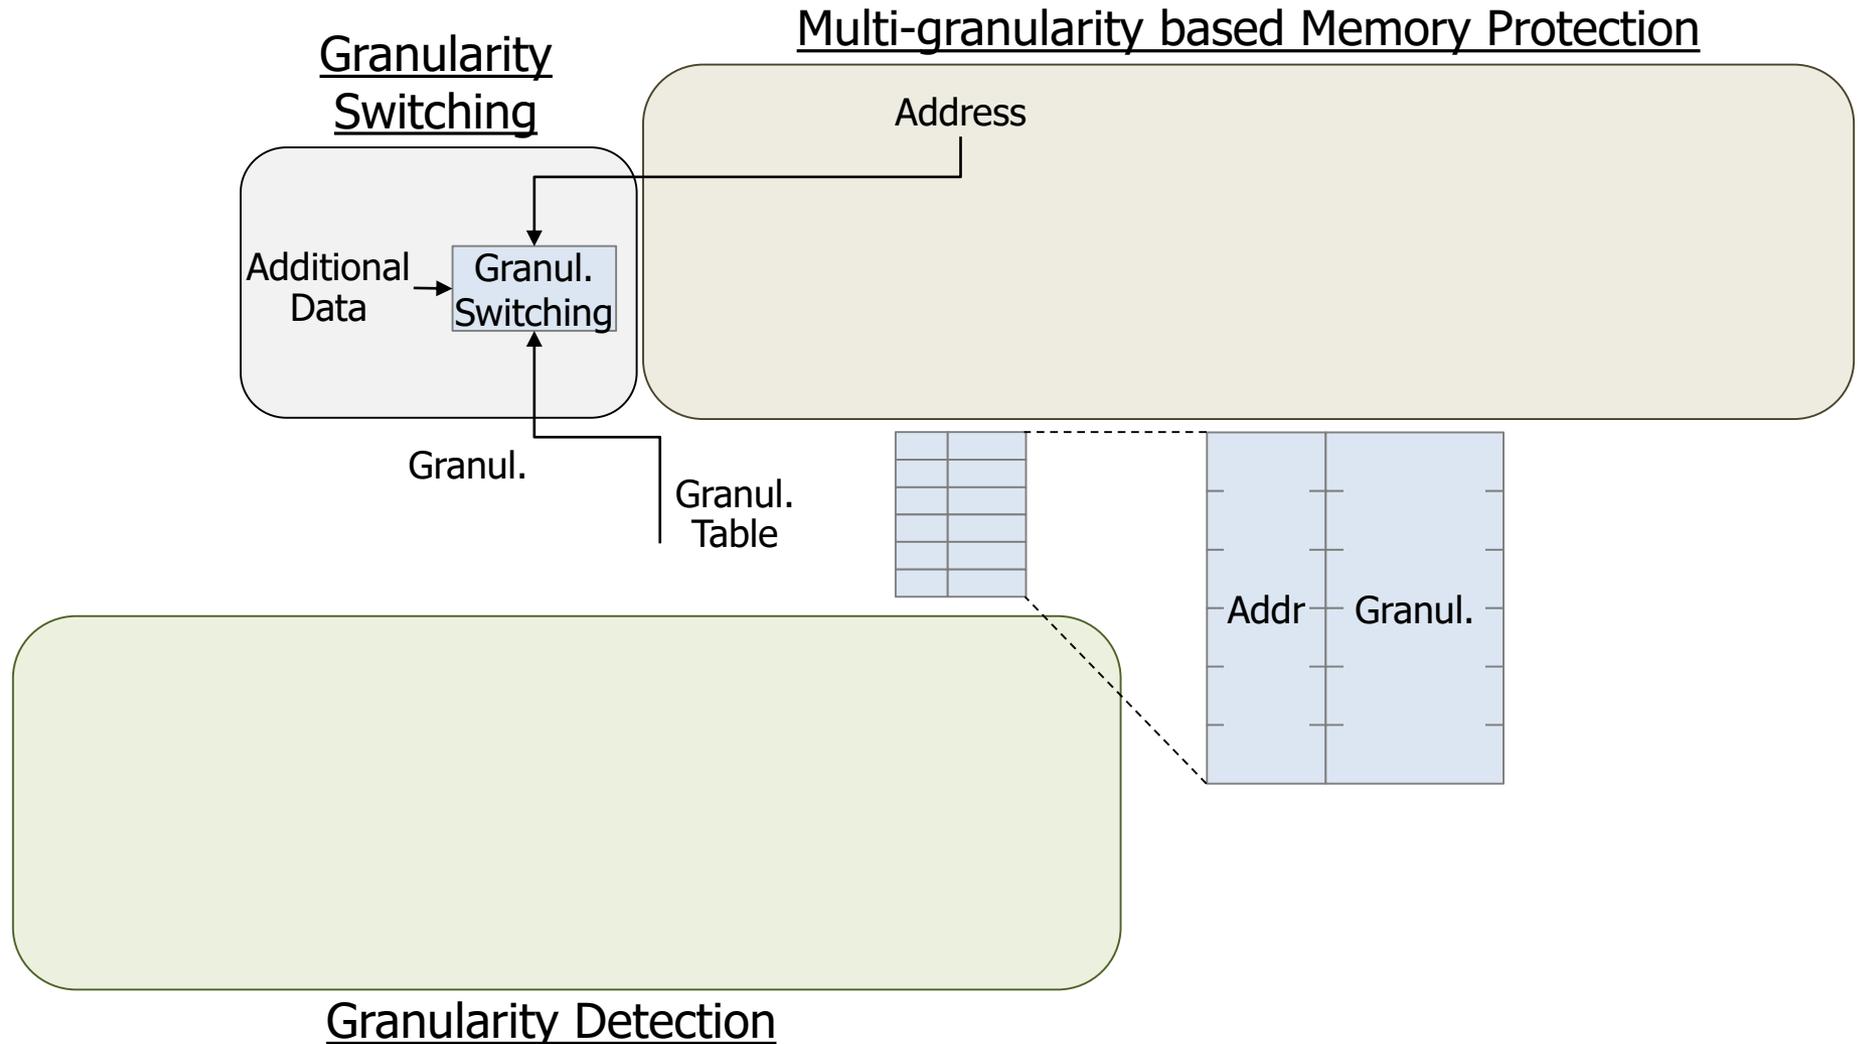
Granul.  
Table



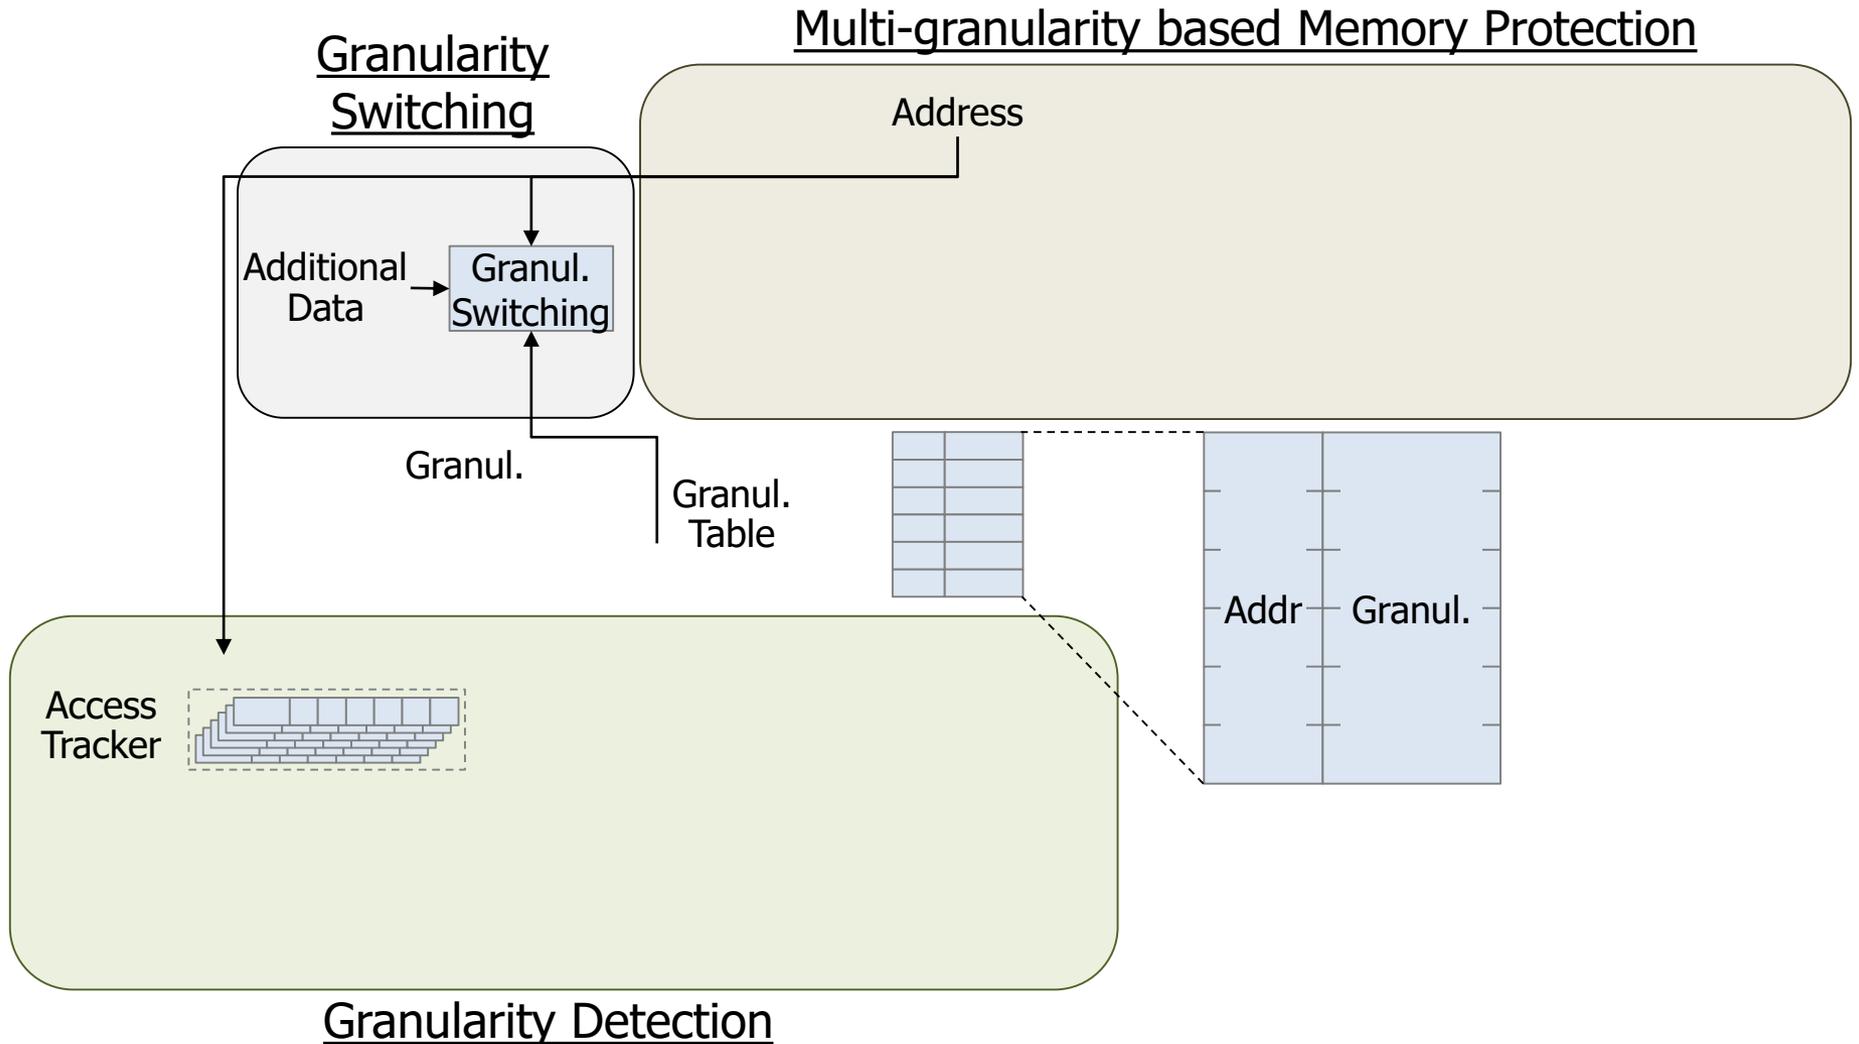
Granularity Detection



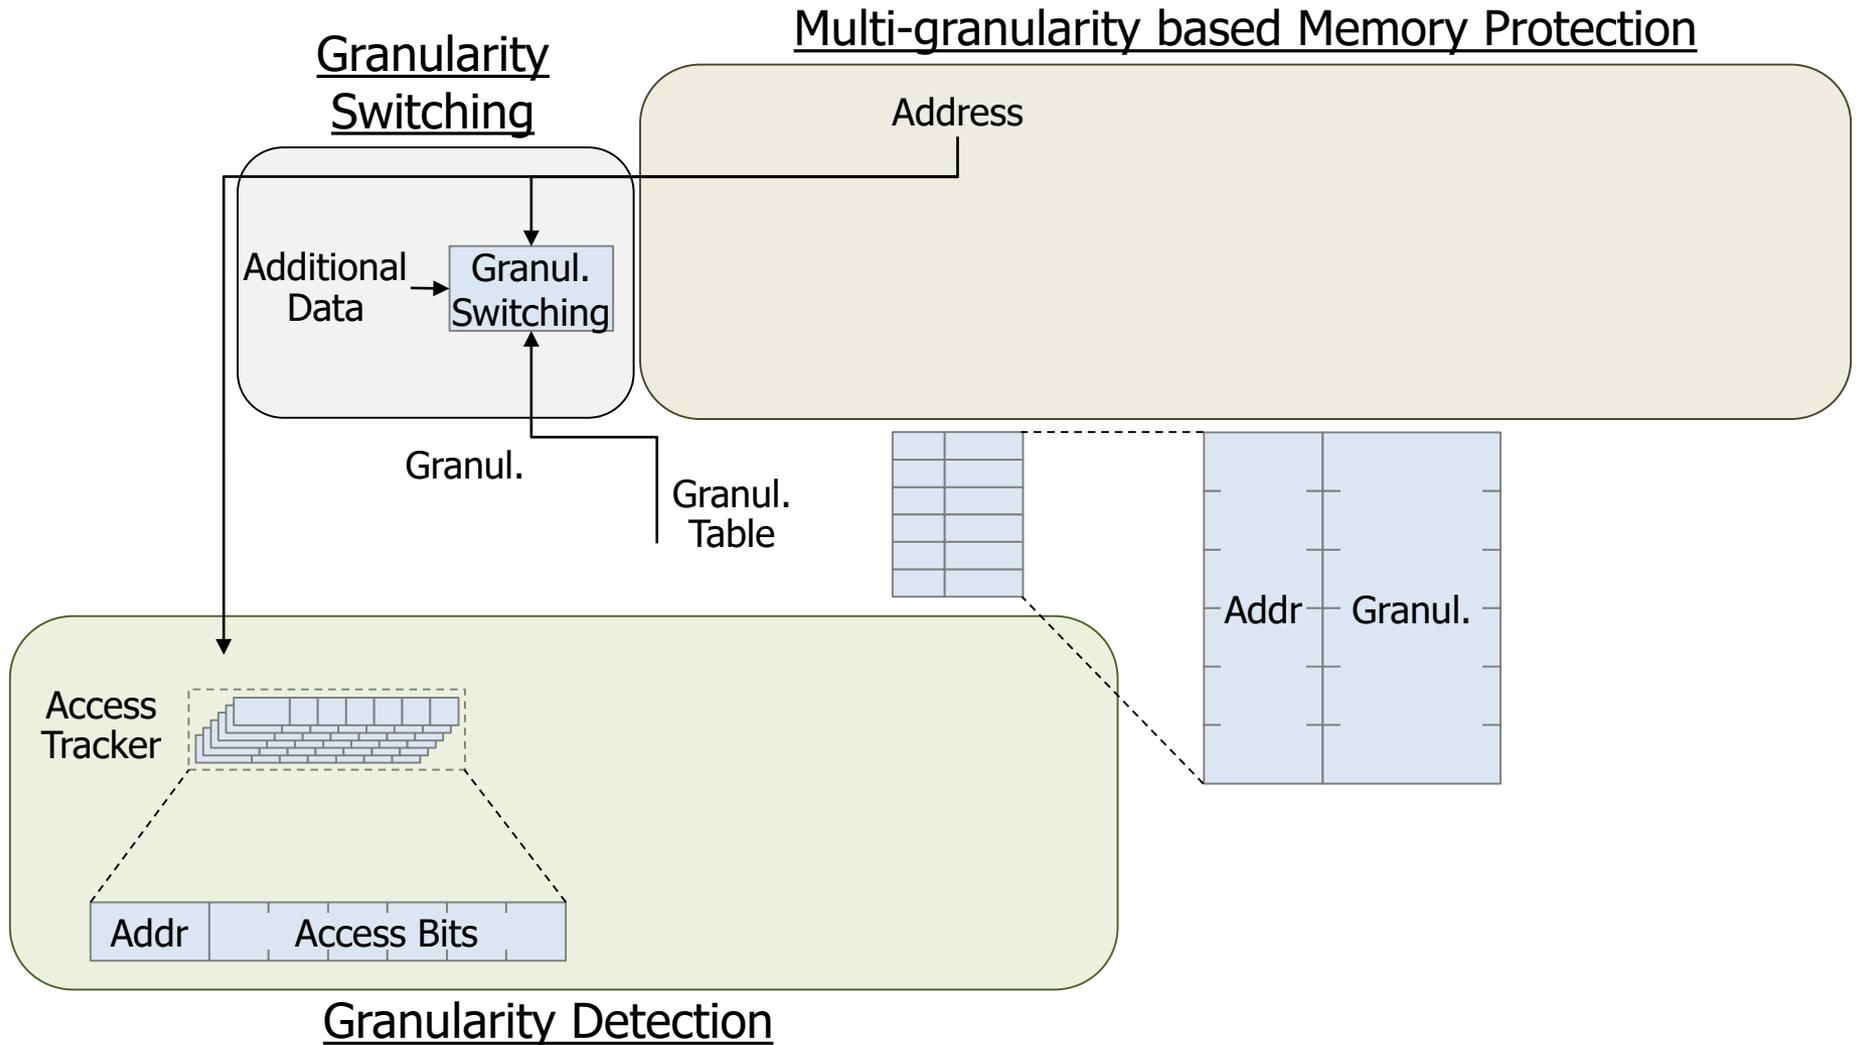
# Design of Multi-granular MAC&Tree



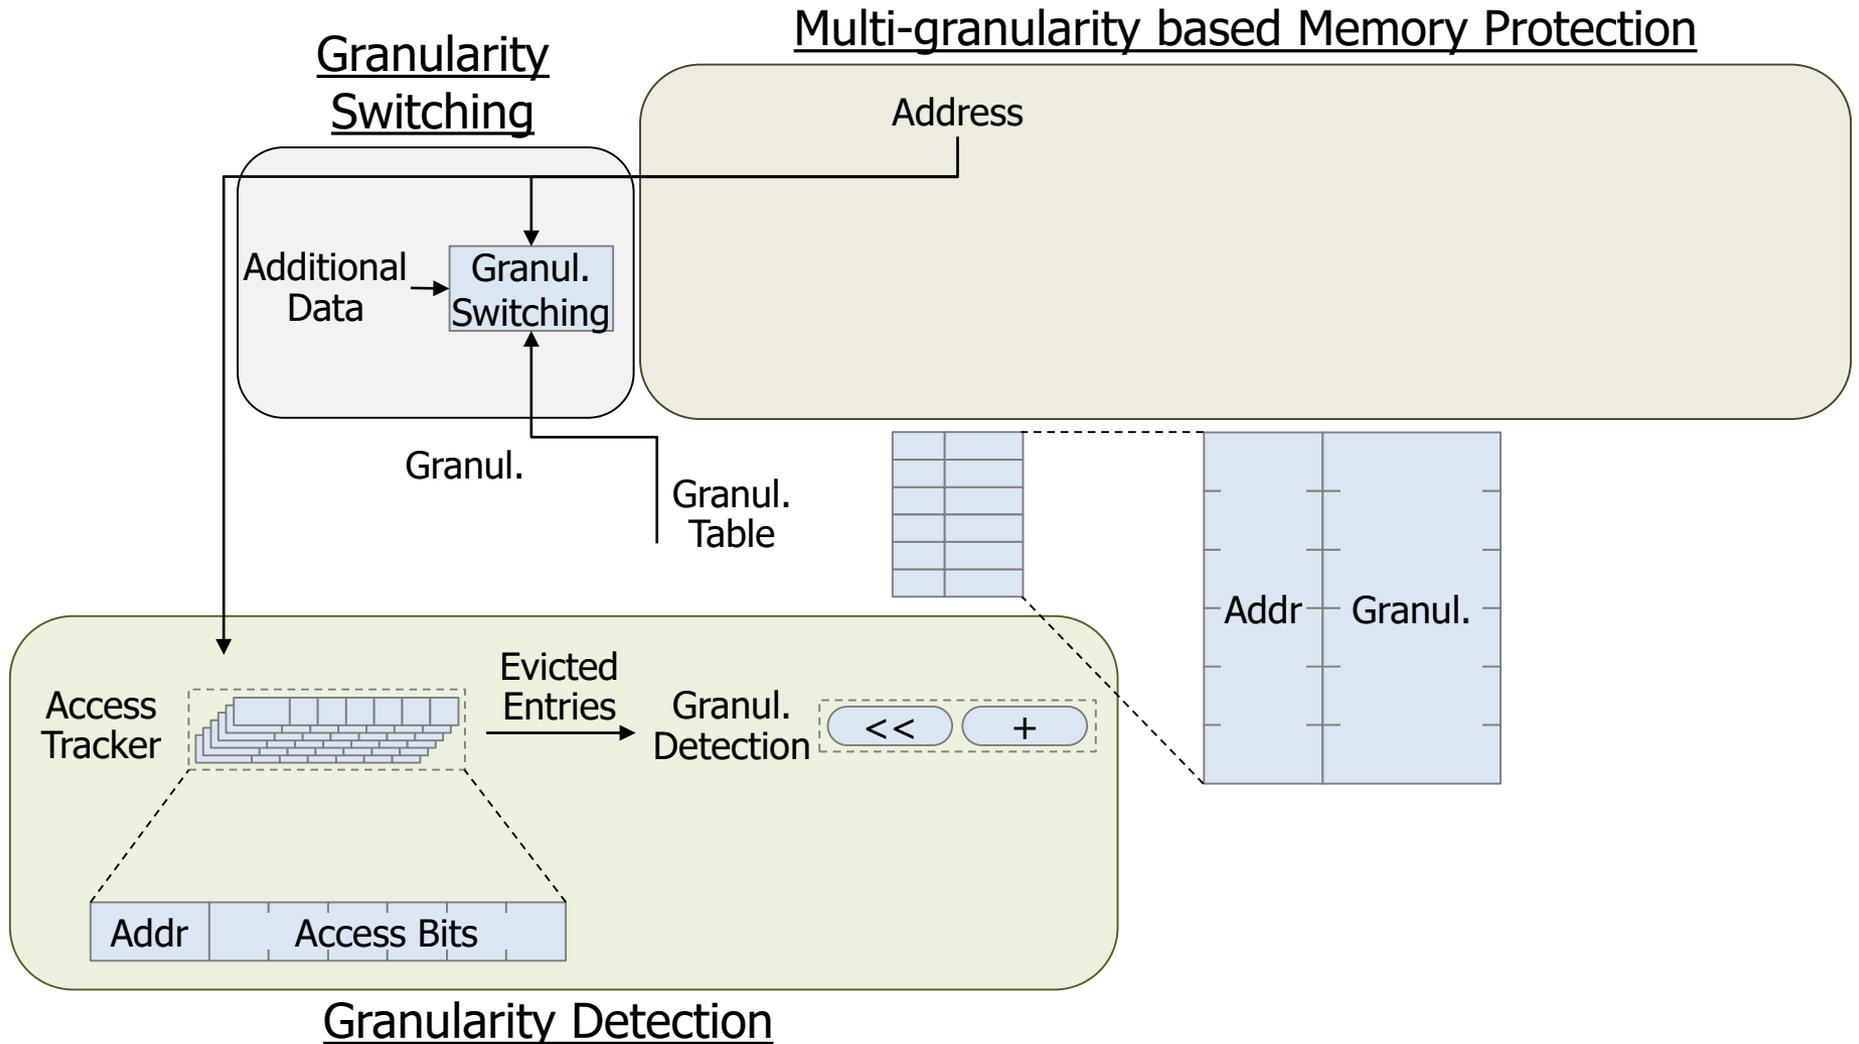
# Design of Multi-granular MAC&Tree



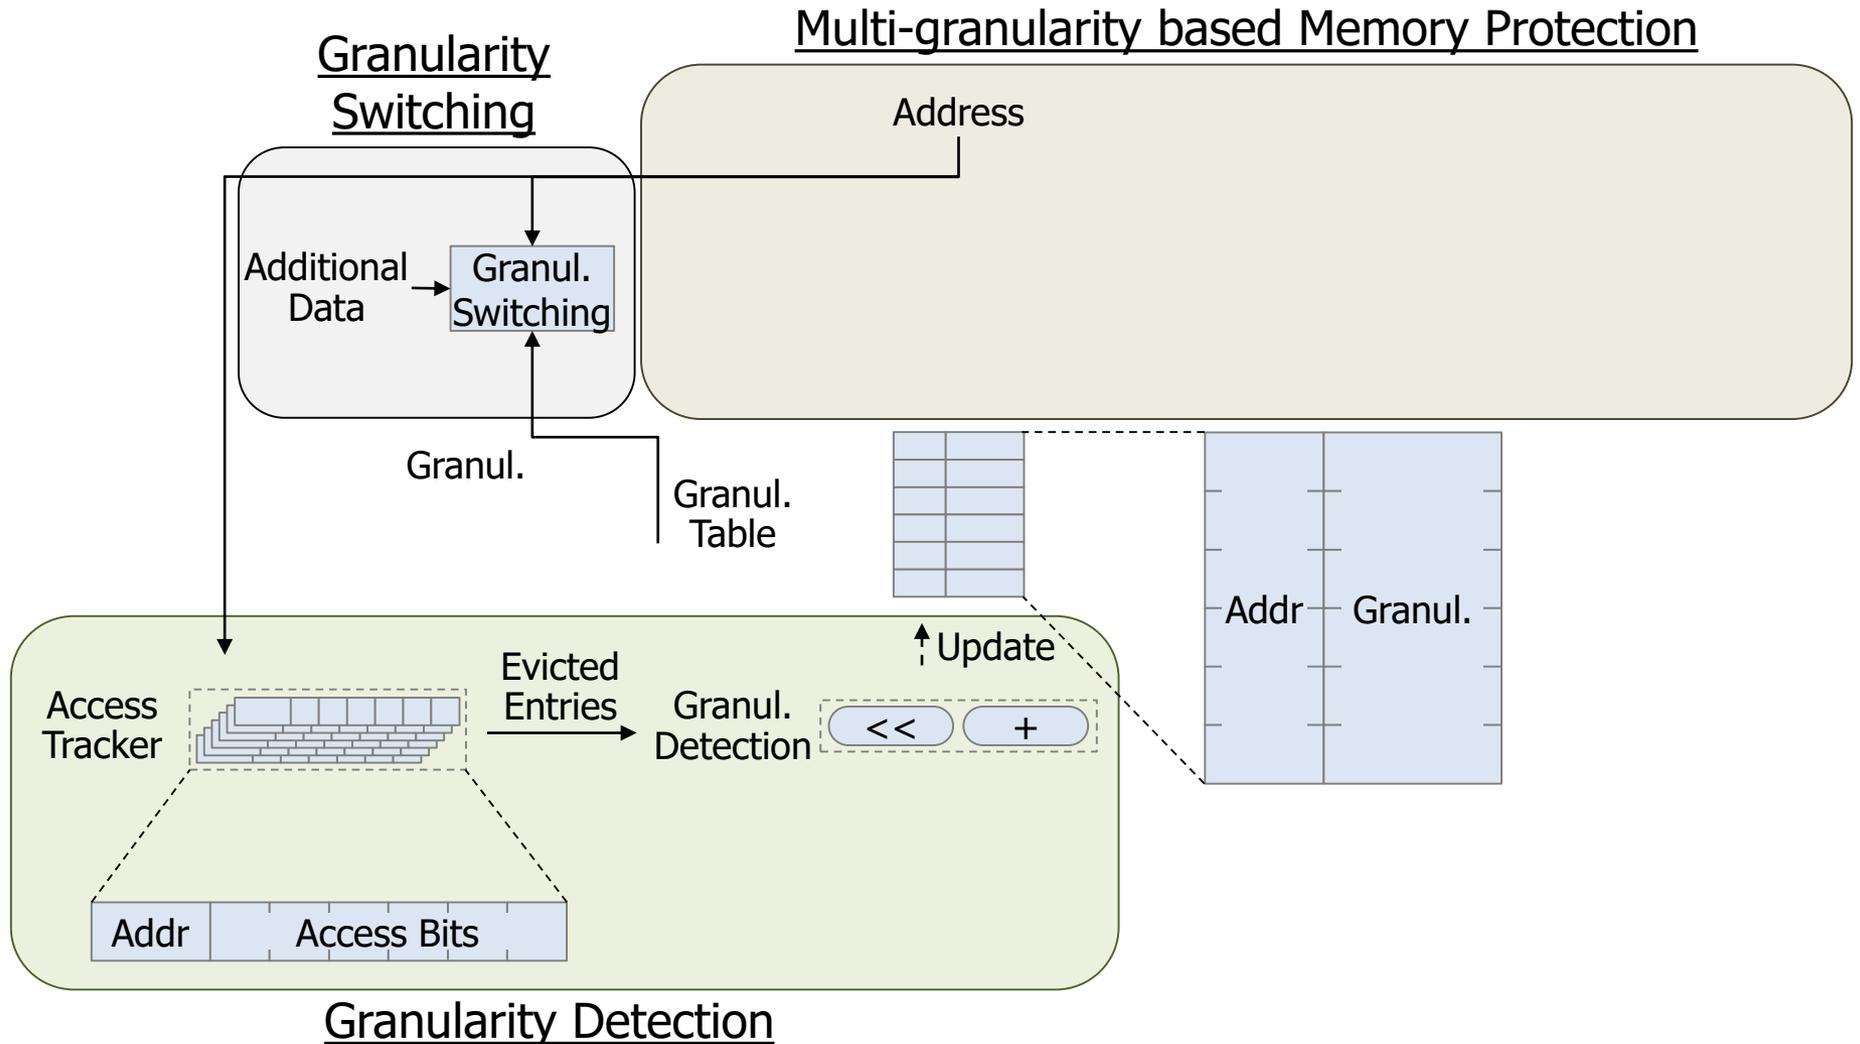
# Design of Multi-granular MAC&Tree



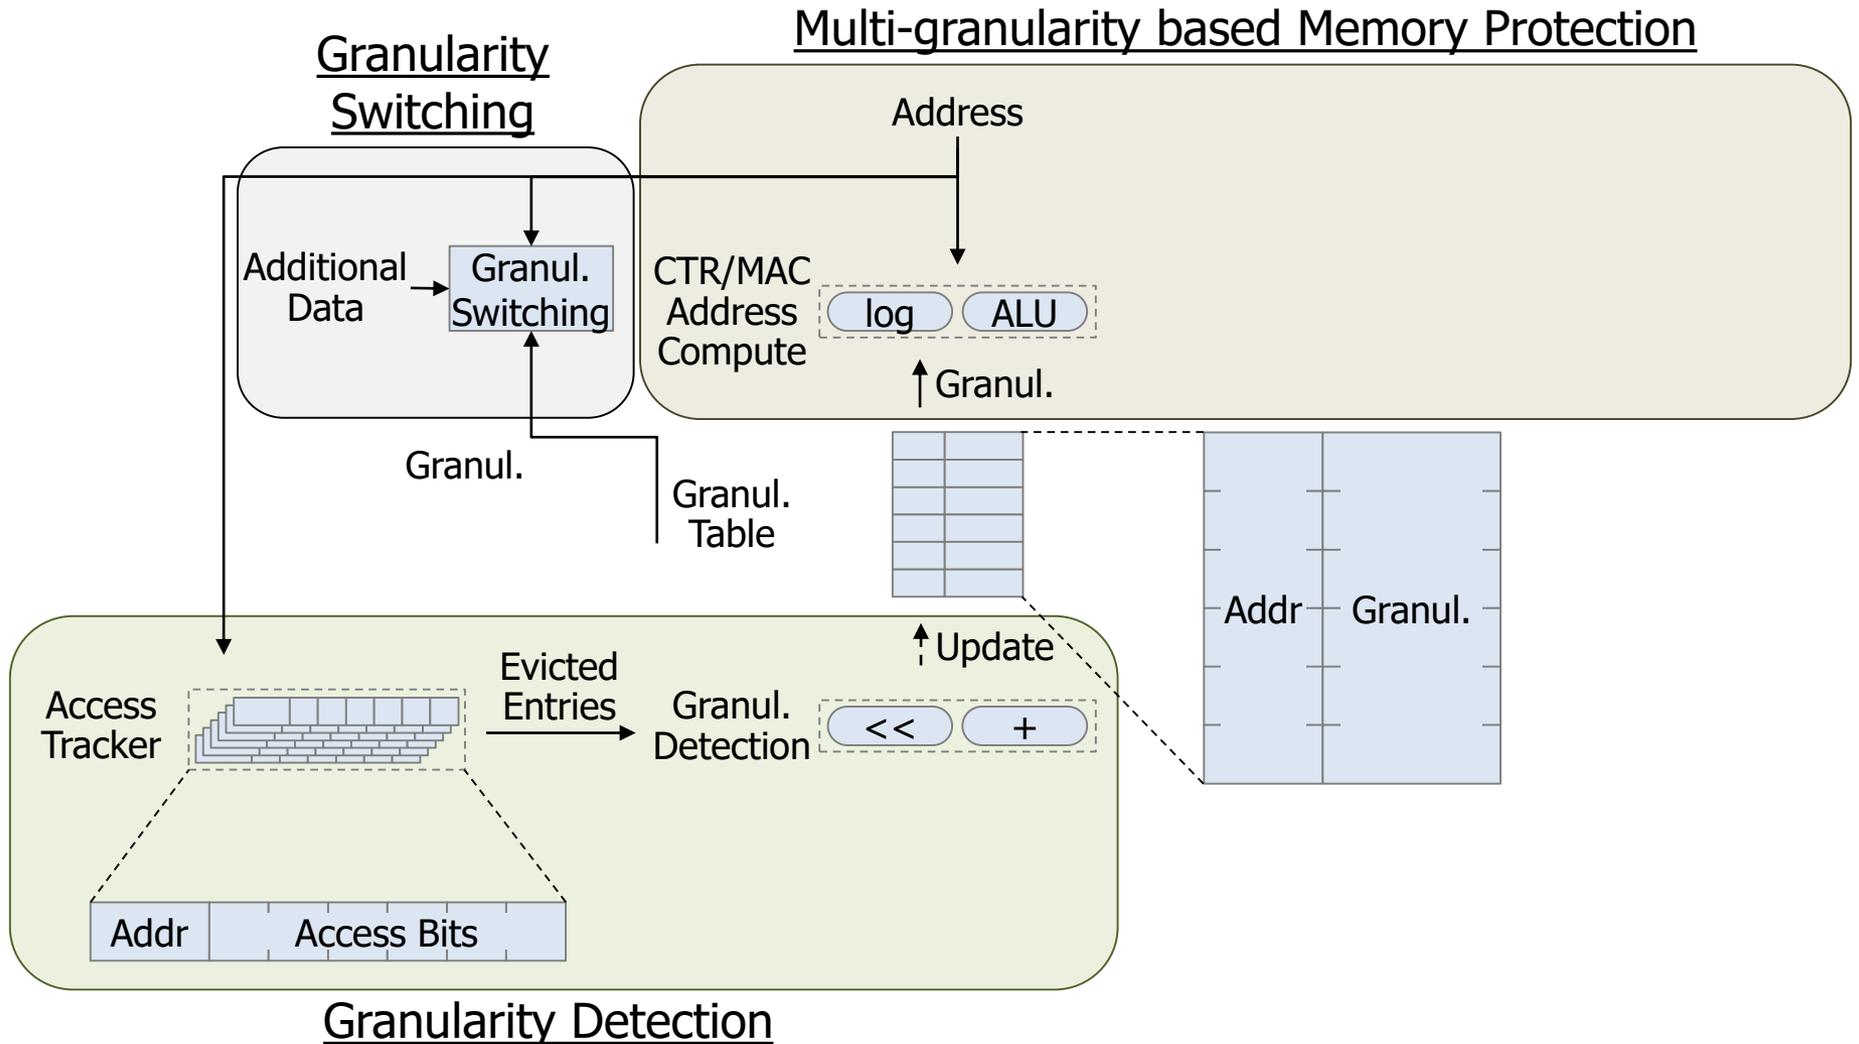
# Design of Multi-granular MAC&Tree



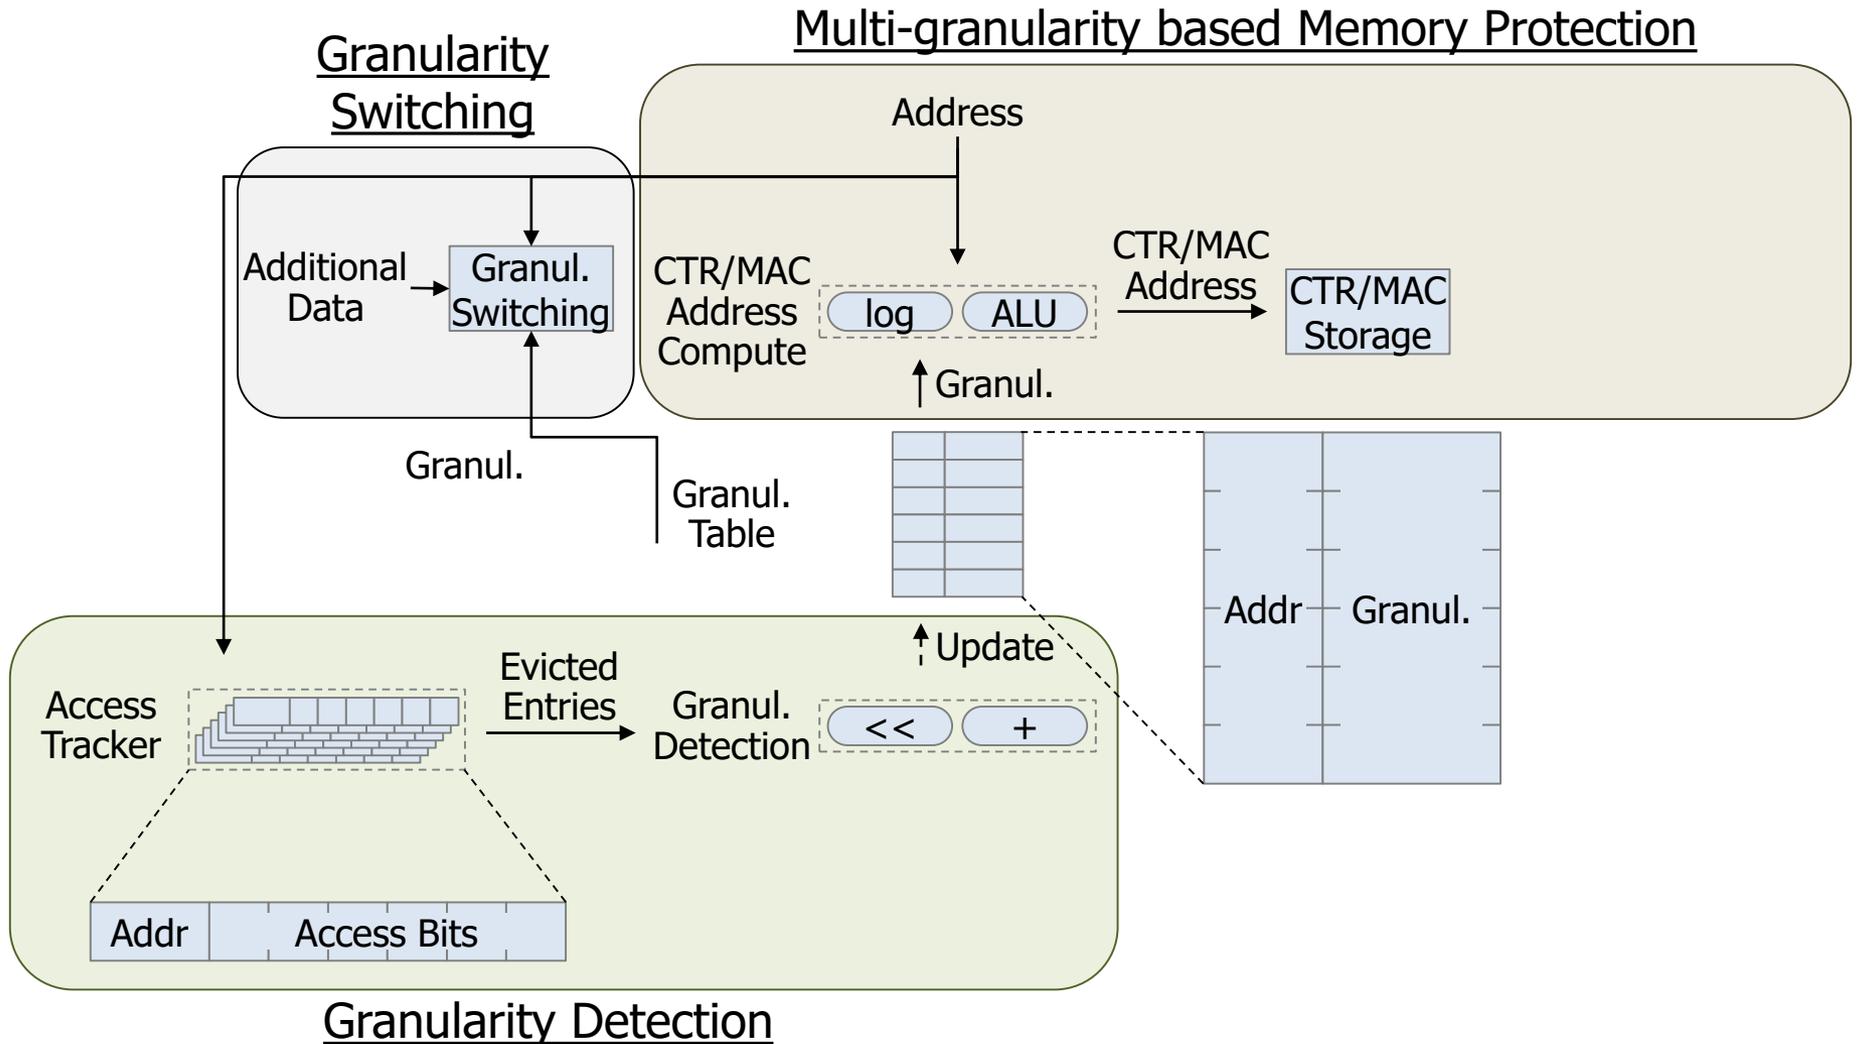
# Design of Multi-granular MAC&Tree



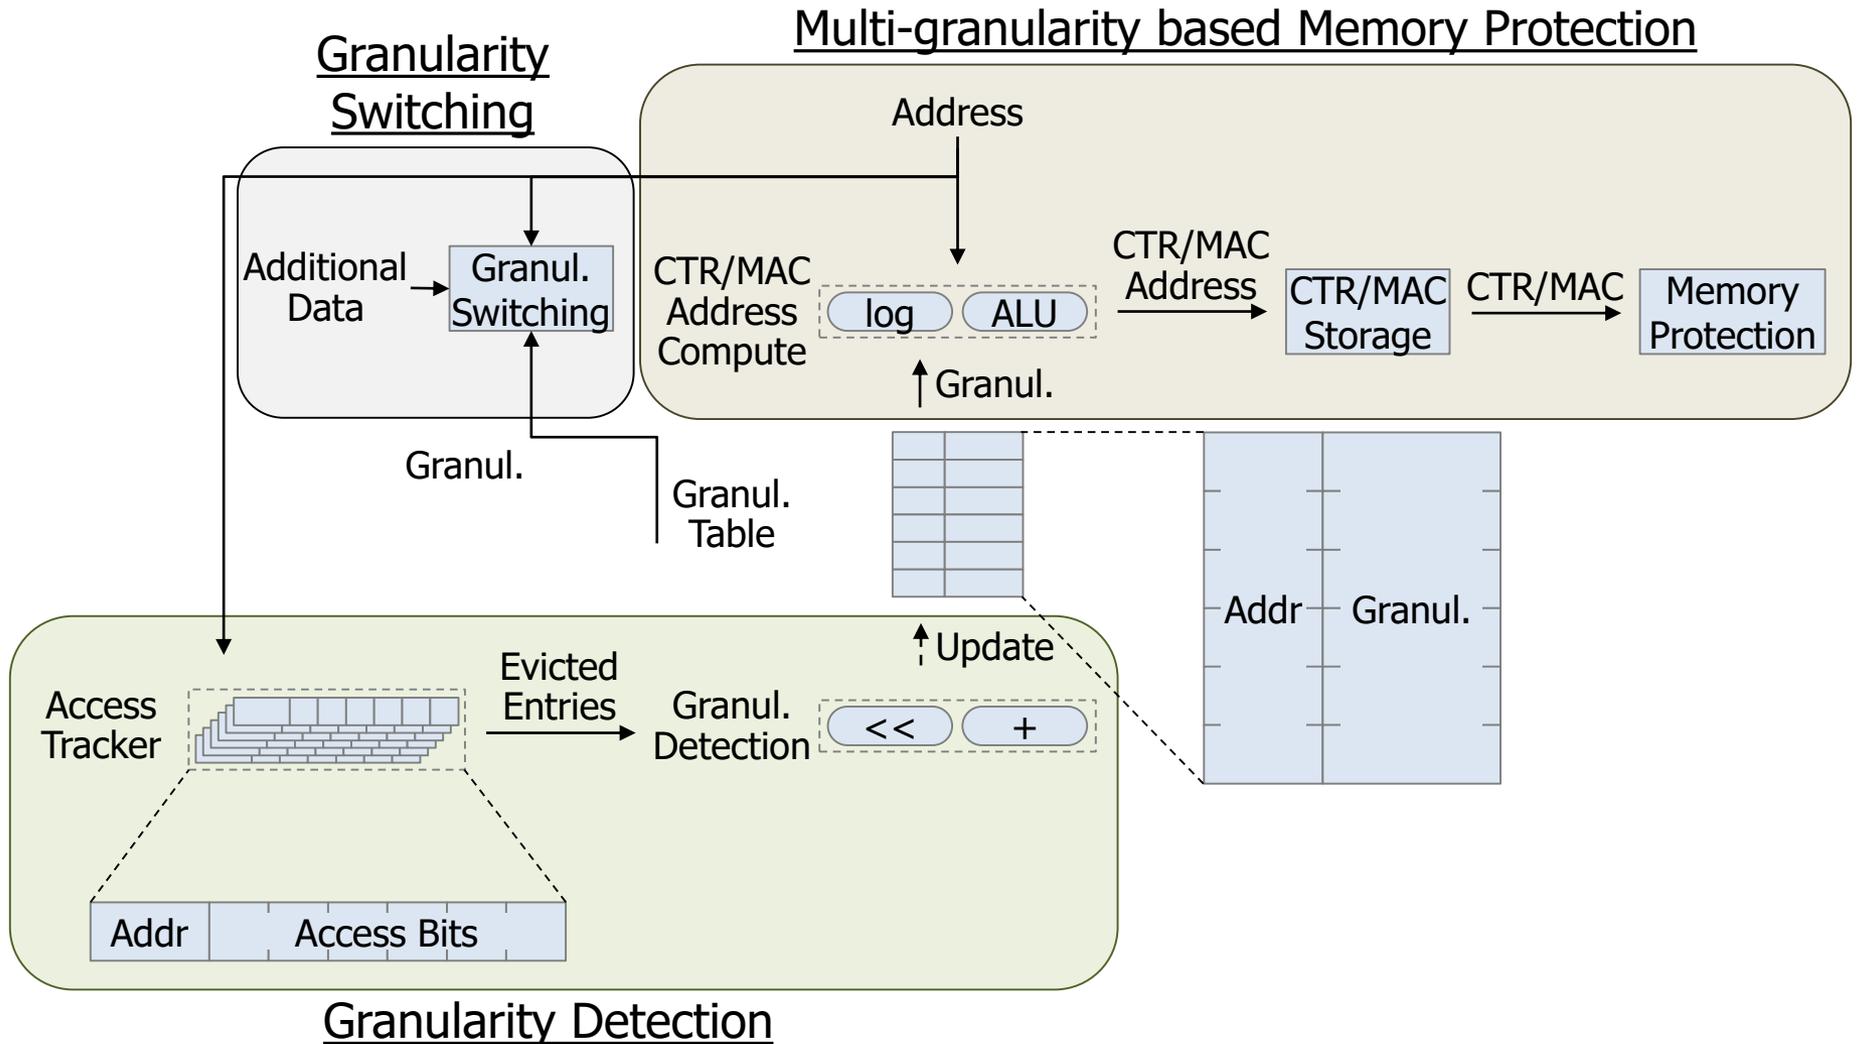
# Design of Multi-granular MAC&Tree



# Design of Multi-granular MAC&Tree



# Design of Multi-granular MAC&Tree



# Recent Memory Protection Studies

---

# Recent Memory Protection Studies

| Study                  | Target             | Multi CTR | Int. Tree Opt. | Multi MAC | Dynamic Update | Target App.    |
|------------------------|--------------------|-----------|----------------|-----------|----------------|----------------|
| <i>SoftVN</i>          | CPU                | O         | X              | X         | X              | ML-specific    |
| <i>Common Counters</i> | GPU                | Dual      | X              | X         | X              | General        |
| <i>Adaptive</i>        | GPU                | X         | X              | Dual      | O              | General        |
| <i>TNPU</i>            | NPU                | O         | X              | X         | X              | ML-specific    |
| <i>Tunable Tree</i>    | NPU                | O         | Sub Optimal    | X         | X              | General        |
| <i>MGX</i>             | NPU                | O         | X              | O         | X              | ML-specific    |
| <i>GuardNN</i>         | NPU                | O         | X              | X         | X              | ML-specific    |
| <i>TensorTEE</i>       | CPU+NPU            | O         | X              | O         | O              | ML-specific    |
| <b>Ours</b>            | <b>CPU+GPU+NPU</b> | <b>O</b>  | <b>Optimal</b> | <b>O</b>  | <b>O</b>       | <b>General</b> |

# Prior Integrity Tree Optimization

---

# Prior Integrity Tree Optimization

| Study                         | Target      | Multi CTR | Int. Tree Opt. | Multi MAC | Dynamic Update | Target App. |
|-------------------------------|-------------|-----------|----------------|-----------|----------------|-------------|
| <i>Bonsai Merkle Forests</i>  | CPU         | X         | Sub Optimal    | X         | X              | General     |
| <i>PENGLAI</i>                | GPU         | X         | Sub Optimal    | X         | X              | General     |
| <i>Migratable Merkle Tree</i> | GPU         | X         | Sub Optimal    | X         | X              | General     |
| <i>Data Enclave</i>           | NPU         | X         | Sub Optimal    | X         | X              | General     |
| <b>Ours</b>                   | CPU+GPU+NPU | O         | Optimal        | O         | O              | General     |

# Encryption & Integrity Validation using Coarse-grained MAC & Counter

---

# Encryption & Integrity Validation using Coarse-grained MAC & Counter

---

- Coarse-MAC & counter encryption, integrity validation

# Encryption & Integrity Validation using Coarse-grained MAC & Counter

---

- Coarse-MAC & counter encryption, integrity validation
  - $CTR_{course(1 \dots k)} = \text{MAX}(CTR_{fine 1}, CTR_{fine 2}, \dots, CTR_{fine k})$

# Encryption & Integrity Validation using Coarse-grained MAC & Counter

---

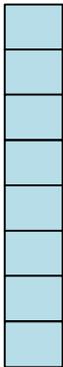
- Coarse-MAC & counter encryption, integrity validation
  - $CTR_{course(1 \dots k)} = \text{MAX}(CTR_{fine 1}, CTR_{fine 2}, \dots, CTR_{fine k})$
  - $MAC_{course(1 \dots k)} =$   
 $\text{HASH}(\dots \text{HASH}(\text{HASH}(MAC_{fine 1}), MAC_{fine 2}), \dots, MAC_{fine k})$

# Encryption & Integrity Validation using Coarse-grained MAC & Counter

---

- Coarse-MAC & counter encryption, integrity validation
  - $CTR_{course(1 \dots k)} = \text{MAX}(CTR_{fine 1}, CTR_{fine 2}, \dots, CTR_{fine k})$
  - $MAC_{course(1 \dots k)} = \text{HASH}(\dots \text{HASH}(\text{HASH}(MAC_{fine 1}), MAC_{fine 2}), \dots, MAC_{fine k})$

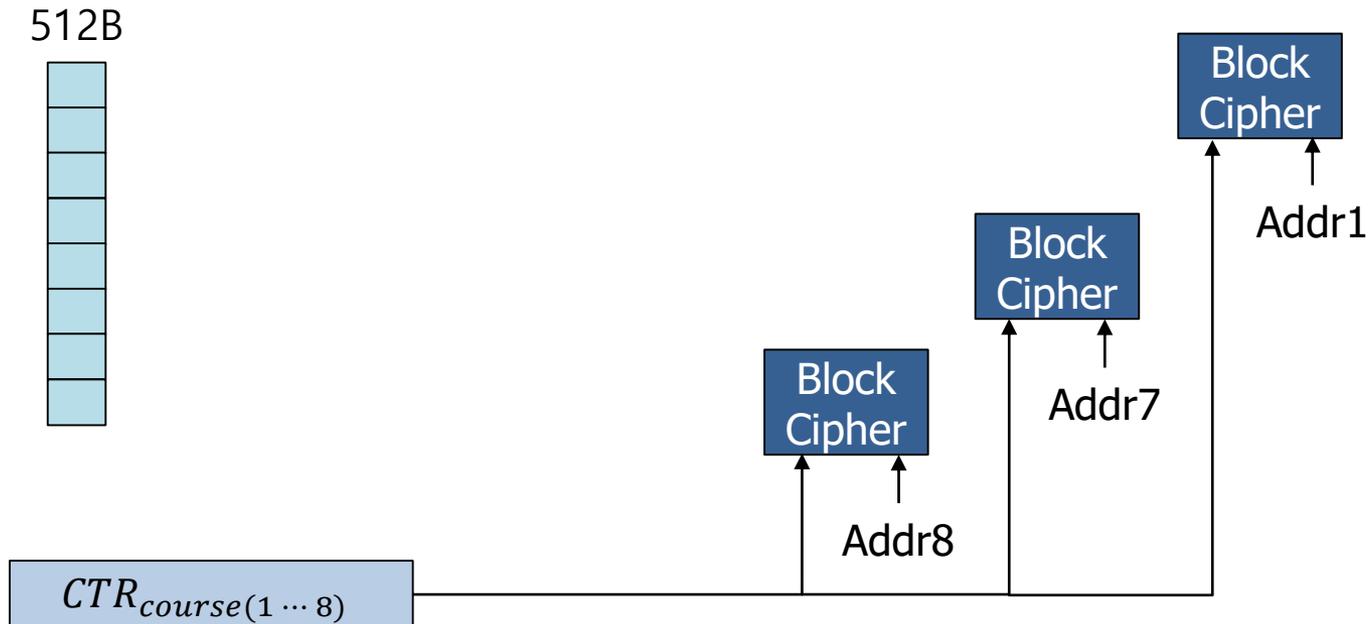
512B



$CTR_{course(1 \dots 8)}$

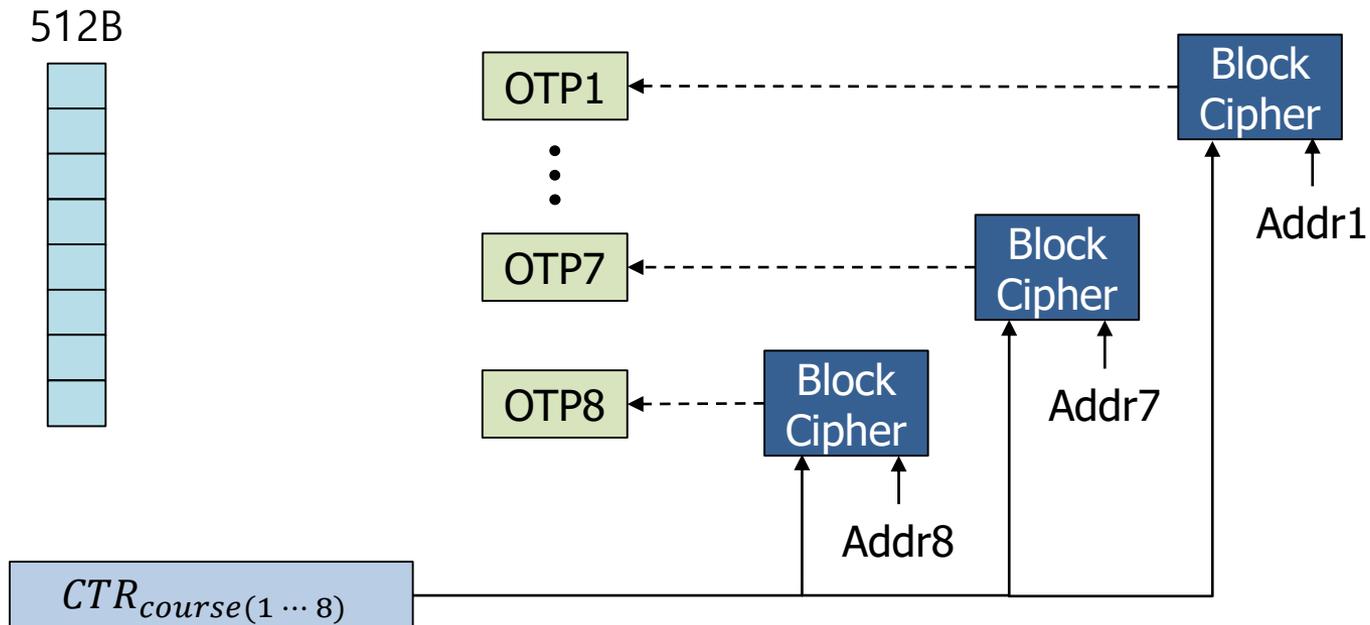
# Encryption & Integrity Validation using Coarse-grained MAC & Counter

- Coarse-MAC & counter encryption, integrity validation
  - $CTR_{course(1 \dots k)} = \text{MAX}(CTR_{fine 1}, CTR_{fine 2}, \dots, CTR_{fine k})$
  - $MAC_{course(1 \dots k)} = \text{HASH}(\dots \text{HASH}(\text{HASH}(MAC_{fine 1}), MAC_{fine 2}), \dots, MAC_{fine k})$



# Encryption & Integrity Validation using Coarse-grained MAC & Counter

- Coarse-MAC & counter encryption, integrity validation
  - $CTR_{course(1 \dots k)} = \text{MAX}(CTR_{fine 1}, CTR_{fine 2}, \dots, CTR_{fine k})$
  - $MAC_{course(1 \dots k)} = \text{HASH}(\dots \text{HASH}(\text{HASH}(MAC_{fine 1}), MAC_{fine 2}), \dots, MAC_{fine k})$



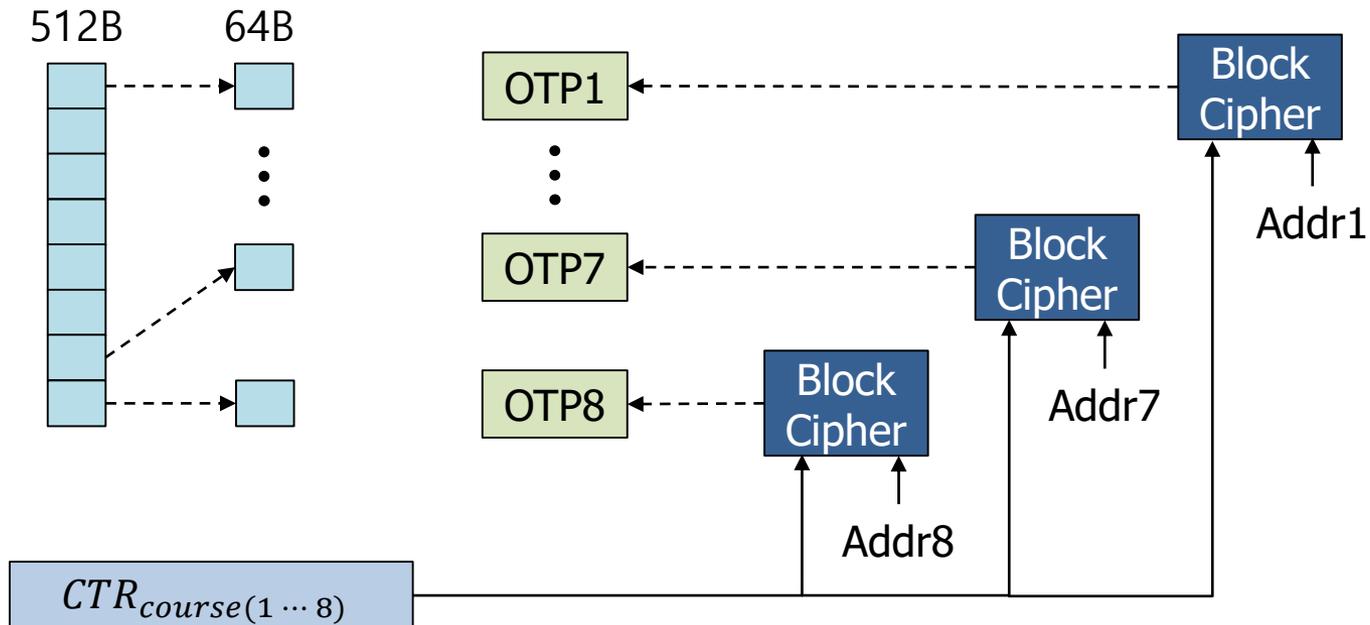
# Encryption & Integrity Validation using Coarse-grained MAC & Counter

- Coarse-MAC & counter encryption, integrity validation

- $CTR_{course(1 \dots k)} = \text{MAX}(CTR_{fine 1}, CTR_{fine 2}, \dots, CTR_{fine k})$

- $MAC_{course(1 \dots k)} =$

$\text{HASH}(\dots \text{HASH}(\text{HASH}(MAC_{fine 1}), MAC_{fine 2}), \dots, MAC_{fine k})$



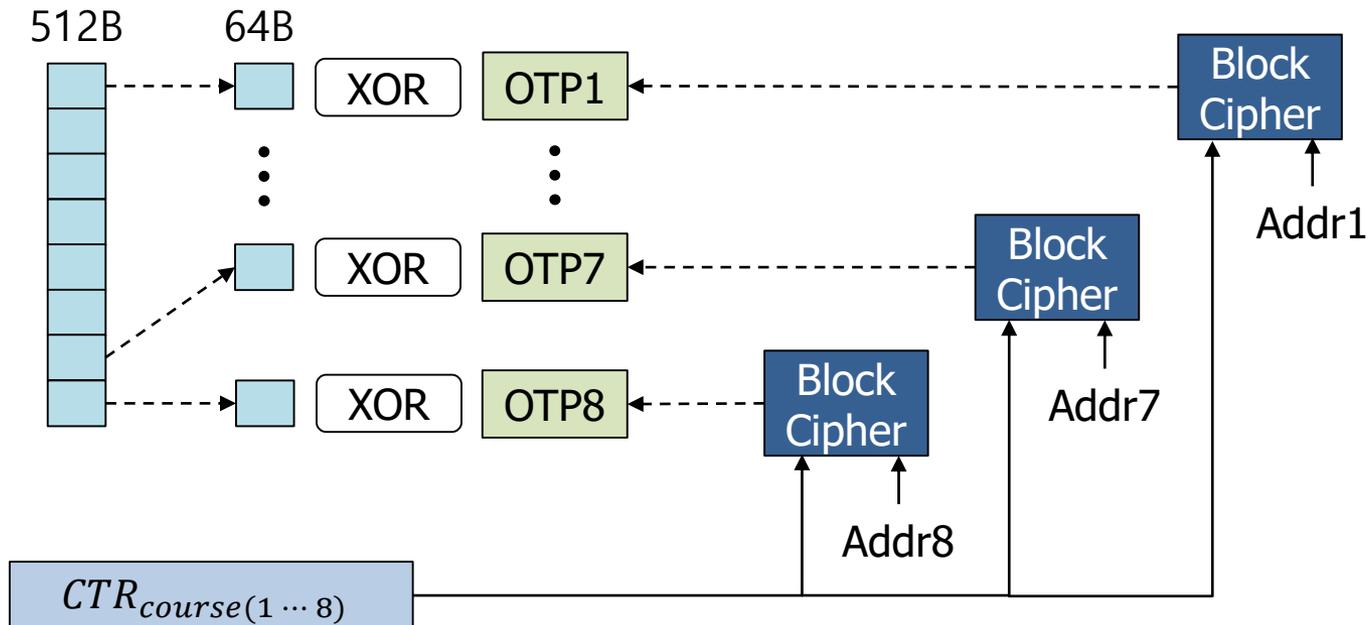
# Encryption & Integrity Validation using Coarse-grained MAC & Counter

- Coarse-MAC & counter encryption, integrity validation

- $CTR_{course(1 \dots k)} = \text{MAX}(CTR_{fine 1}, CTR_{fine 2}, \dots, CTR_{fine k})$

- $MAC_{course(1 \dots k)} =$

$\text{HASH}(\dots \text{HASH}(\text{HASH}(MAC_{fine 1}), MAC_{fine 2}), \dots, MAC_{fine k})$



# Address of MAC & Counter

---

# Address of MAC & Counter

---

- Chunk-level index computation

# Address of MAC & Counter

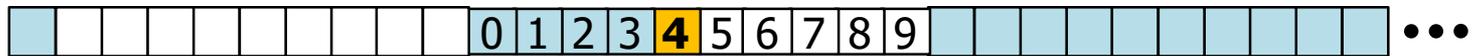
---

- Chunk-level index computation
- Recursive parent call from leaf counters

# Address of MAC & Counter

---

- Chunk-level index computation
- Recursive parent call from leaf counters



# Address of MAC & Counter

---

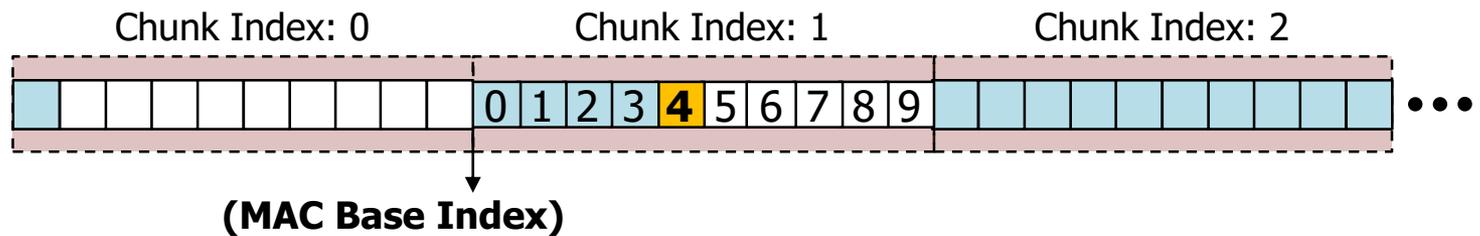
- Chunk-level index computation
- Recursive parent call from leaf counters



# Address of MAC & Counter

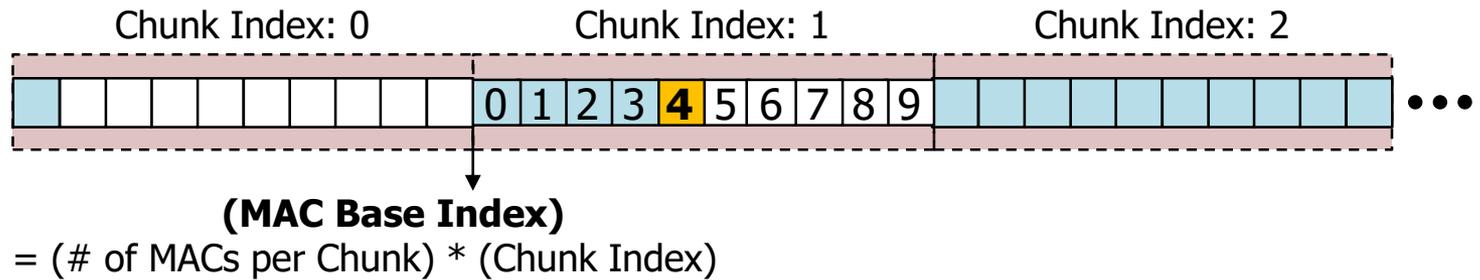
---

- Chunk-level index computation
- Recursive parent call from leaf counters



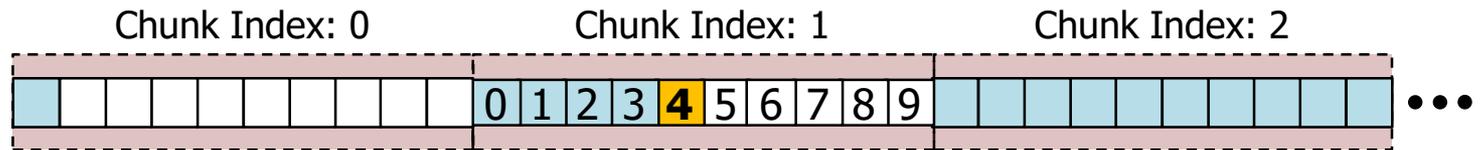
# Address of MAC & Counter

- Chunk-level index computation
- Recursive parent call from leaf counters



# Address of MAC & Counter

- Chunk-level index computation
- Recursive parent call from leaf counters

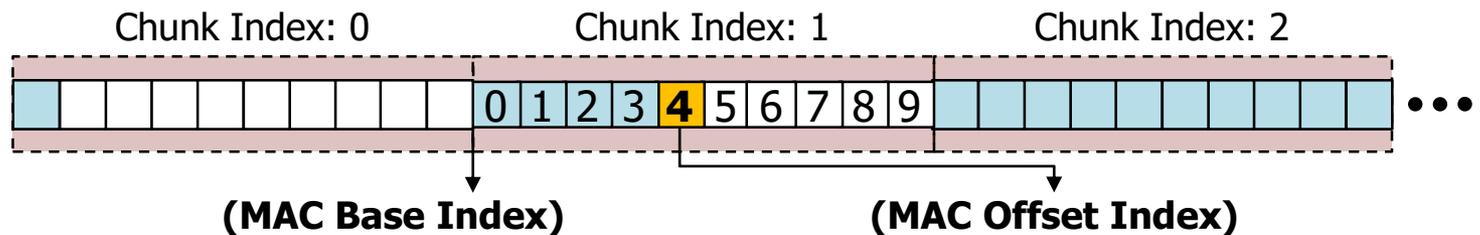


**(MAC Base Index)**

$$\begin{aligned} &= (\# \text{ of MACs per Chunk}) * (\text{Chunk Index}) \\ &= 10 * 1 = 10 \end{aligned}$$

# Address of MAC & Counter

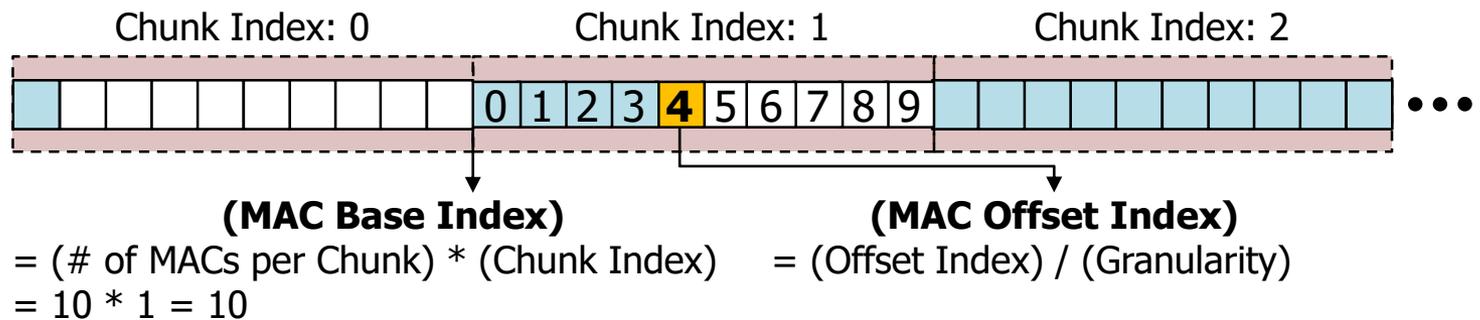
- Chunk-level index computation
- Recursive parent call from leaf counters



$$\begin{aligned} &= (\# \text{ of MACs per Chunk}) * (\text{Chunk Index}) \\ &= 10 * 1 = 10 \end{aligned}$$

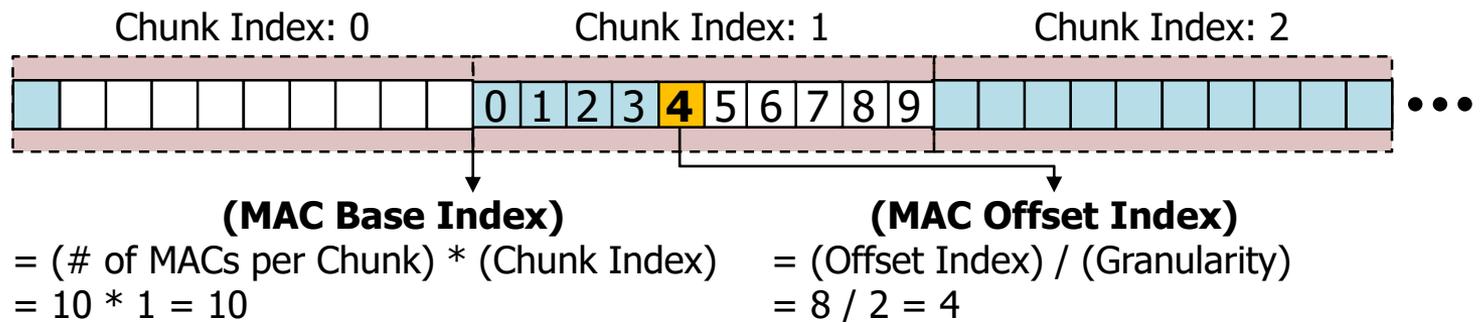
# Address of MAC & Counter

- Chunk-level index computation
- Recursive parent call from leaf counters



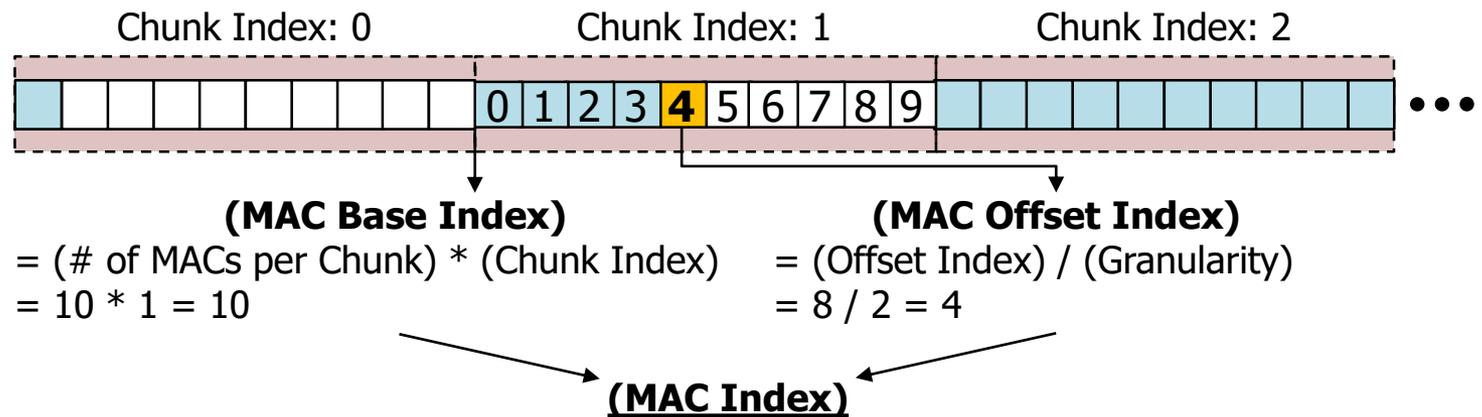
# Address of MAC & Counter

- Chunk-level index computation
- Recursive parent call from leaf counters



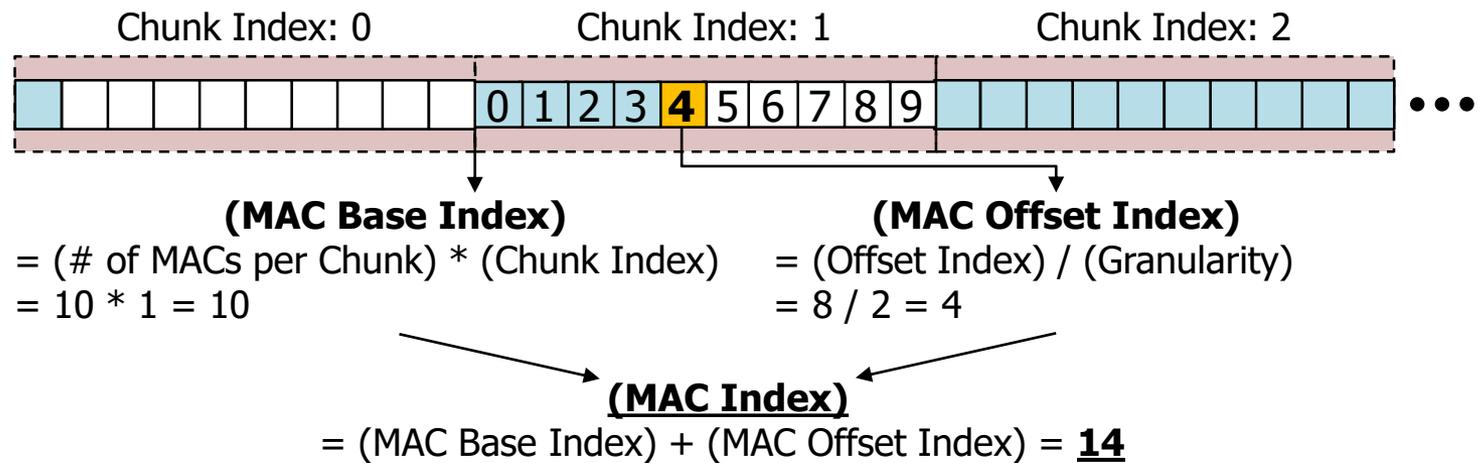
# Address of MAC & Counter

- Chunk-level index computation
- Recursive parent call from leaf counters



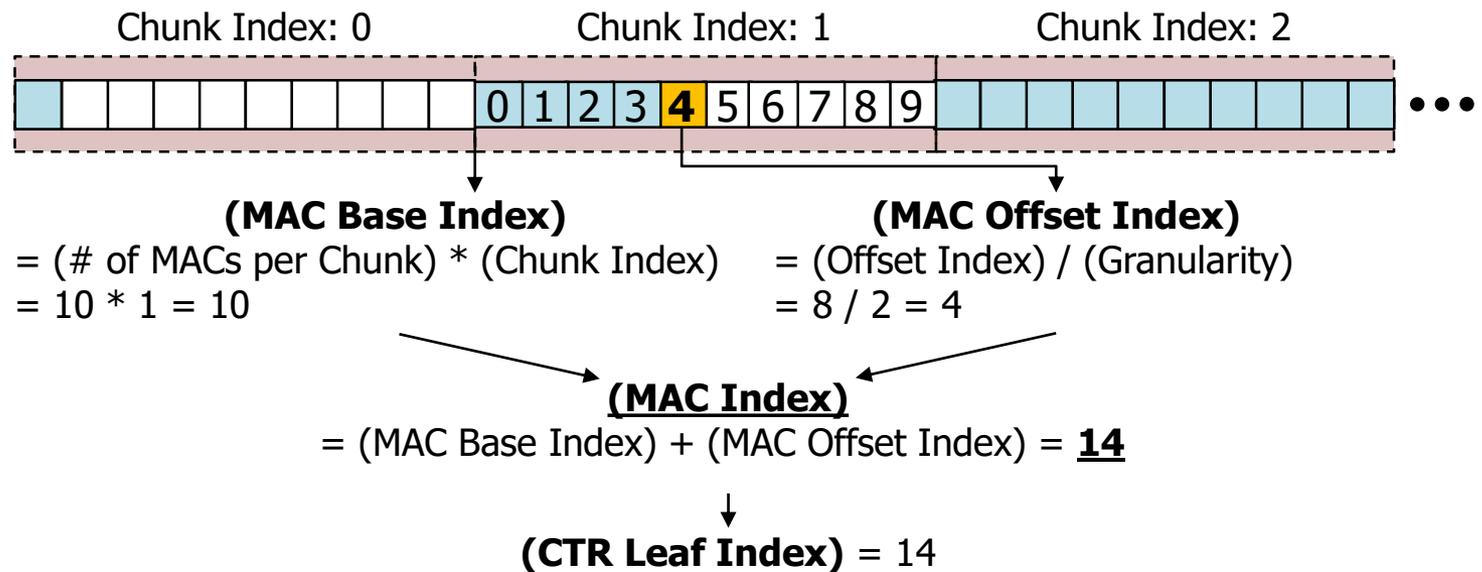
# Address of MAC & Counter

- Chunk-level index computation
- Recursive parent call from leaf counters



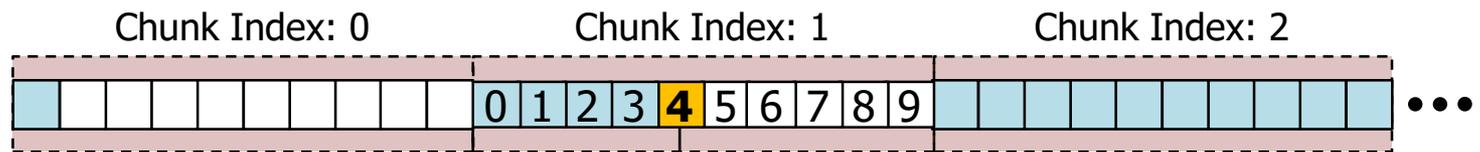
# Address of MAC & Counter

- Chunk-level index computation
- Recursive parent call from leaf counters



# Address of MAC & Counter

- Chunk-level index computation
- Recursive parent call from leaf counters



**(MAC Base Index)**  
 $= (\# \text{ of MACs per Chunk}) * (\text{Chunk Index})$   
 $= 10 * 1 = 10$

**(MAC Offset Index)**  
 $= (\text{Offset Index}) / (\text{Granularity})$   
 $= 8 / 2 = 4$

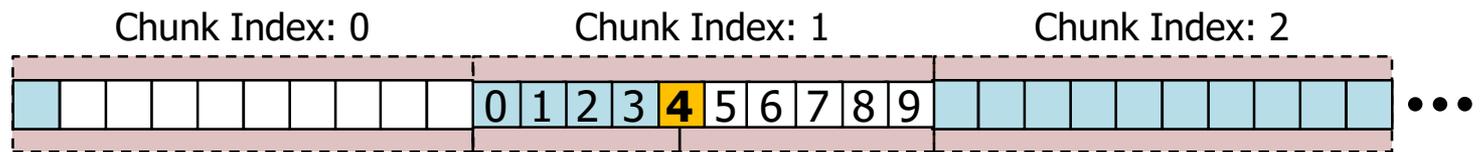
**(MAC Index)**  
 $= (\text{MAC Base Index}) + (\text{MAC Offset Index}) = \mathbf{14}$

**(CTR Leaf Index) = 14**

**(# of Parents)**

# Address of MAC & Counter

- Chunk-level index computation
- Recursive parent call from leaf counters



**(MAC Base Index)**  
= (# of MACs per Chunk) \* (Chunk Index)  
= 10 \* 1 = 10

**(MAC Offset Index)**  
= (Offset Index) / (Granularity)  
= 8 / 2 = 4

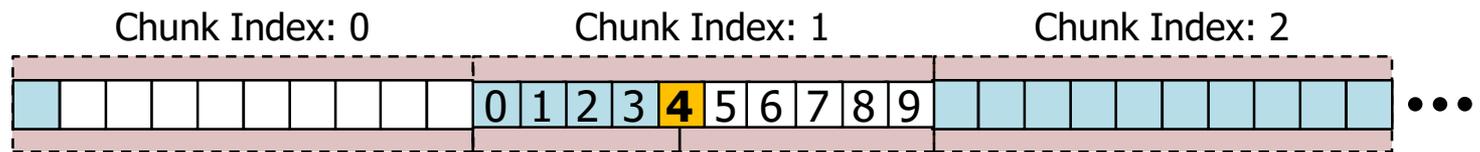
**(MAC Index)**  
= (MAC Base Index) + (MAC Offset Index) = **14**

**(CTR Leaf Index) = 14**

**(# of Parents)**  
=  $\sqrt{\text{Arity}}(\text{Granularity})$

# Address of MAC & Counter

- Chunk-level index computation
- Recursive parent call from leaf counters



**(MAC Base Index)**

$$\begin{aligned} &= (\# \text{ of MACs per Chunk}) * (\text{Chunk Index}) \\ &= 10 * 1 = 10 \end{aligned}$$

**(MAC Offset Index)**

$$\begin{aligned} &= (\text{Offset Index}) / (\text{Granularity}) \\ &= 8 / 2 = 4 \end{aligned}$$

**(MAC Index)**

$$= (\text{MAC Base Index}) + (\text{MAC Offset Index}) = \mathbf{14}$$

**(CTR Leaf Index) = 14**

**(CTR Index)**

**(# of Parents)**

$$= \text{sqrt}\{\text{Arity}\}(\text{Granularity}) \rightarrow$$

# Address of MAC & Counter

- Chunk-level index computation
- Recursive parent call from leaf counters



**(MAC Base Index)**

$$= (\# \text{ of MACs per Chunk}) * (\text{Chunk Index})$$

$$= 10 * 1 = 10$$

**(MAC Offset Index)**

$$= (\text{Offset Index}) / (\text{Granularity})$$

$$= 8 / 2 = 4$$

**(MAC Index)**

$$= (\text{MAC Base Index}) + (\text{MAC Offset Index}) = \mathbf{14}$$

**(CTR Leaf Index) = 14**

**(CTR Index)**

$$= \sqrt{\text{Arity}}(\text{Granularity}) \rightarrow = \text{Parent}(\text{Parent}(\dots(\text{Parent}(\text{CTR Leaf Index}))))$$

# Workload Analysis & Selected Scenarios

---

# Workload Analysis & Selected Scenarios

---

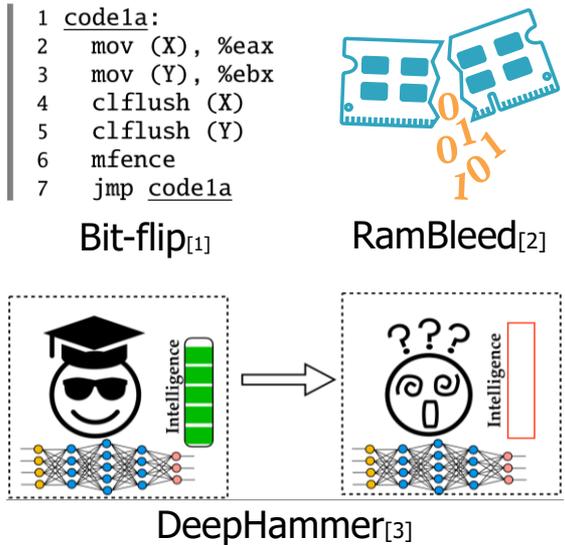
- Workloads & Scenarios

# Workload Analysis & Selected Scenarios

- Workloads & Scenarios
  - 14 workloads, 250 scenarios (all combinations)
  - Access pattern: Fine – ff – f – c – cc – Coarse | Diverse (d)
  - Traffic per cycles: Small (s) – Medium (m) – Large (l)

| Workloads (access pattern-traffic per cycles) |   |
|---|---|
| CPU   | bw (ff-s), gcc (ff-s), mcf (ff-m), xal (f-m), ray (ff-s)                    |
| GPU   | syr2k (ff-m), pr (f-m), sten (c-l), mm (cc-m), floyd (d-s),                 |
| NPU   | ncf (c-s), dlrm (c-s), sfrnn (c-l), alex (cc-m)                             |
| ID  | (CPU, GPU, NPU1, NPU2)  |
| <b>ff</b>                                     | (bw, syr2k, ncf, dlrm), (mcf, syr2k, sfrnn, dlrm), (gcc, floyd, sfrnn, ncf) |
| <b>f</b>                                      | (xal, pr, sfrnn, ncf), (xal, pr, ncf, ncf)                                  |
| <b>c</b>                                      | (gcc, sten, alex, dlrm), (bw, sten, ncf, ncf), (mcf, sten, sfrnn, sfrnn)    |
| <b>cc</b>                                     | (xal, mm, alex, dlrm), (ray, mm, alex, alex), (ray, Floyd, alex, alex)      |

# Rowhammer Attacks



[1] Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors (ISCA 2014)

[2] RAMBleed: Reading Bits in Memory Without Accessing Them (S&P 2020)

[3] DeepHammer: Depleting the Intelligence of Deep Neural Networks through Targeted Chain of Bit Flips (USENIX Security 2020)

# More Design Descriptions in Our Paper

---

- Lazy-switching analysis
- Cacheline fragmentation issue
- CTR/MAC addressing for multi-granularity
- Coarse-grained memory protection engine using parallel counter sharing and nested MAC hasing
- Misprediction handler
- Efficient granularity representation
- Hardware overhead
- Comparison to prior subtree optimization schemes

# More Results in Our Paper

---

- The ratio of stream chunks
- Performance analysis of selected scenarios
- End-to-end performance
- Drawbacks of the per-device (static) granularity
- Performance comparison with dual-granularity
- Switching overhead measurement
- Security cache hit ratio improvement
- Hardware overhead

# **Temp Slide**

# Research Objective

---

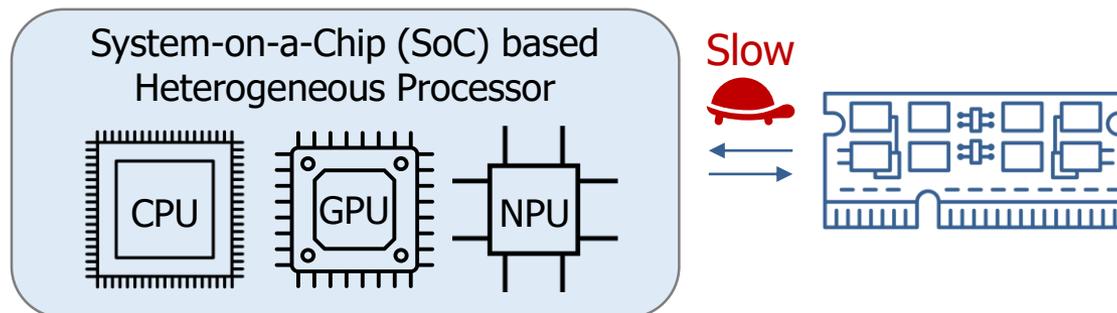
Constructs a general and efficient memory protection scheme for heterogeneous processors

# Research Objective

Constructs a general and efficient memory protection scheme for heterogeneous processors

- Challenge 1: Heterogeneous processors have **diverse access pattern**

## Conventional Memory Protection: High Overhead

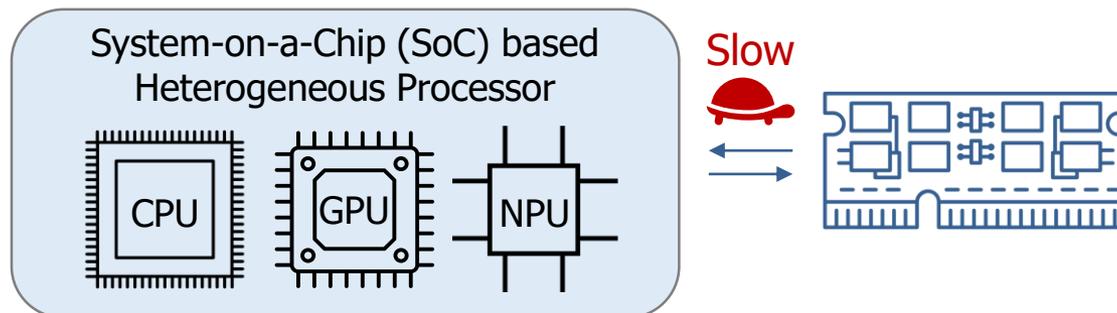


# Research Objective

Constructs a **general** and efficient memory protection scheme for heterogeneous processors

- Challenge 1: Heterogeneous processors have **diverse access pattern**
- Challenge 2: Each prior protection **only for a specific access pattern**

## Conventional Memory Protection: High Overhead

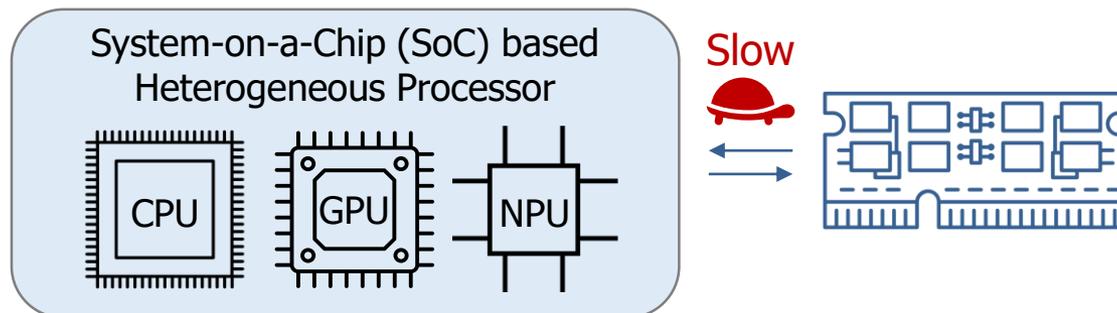


# Research Objective

Constructs a **general** and efficient memory protection scheme for heterogeneous processors

- Challenge 1: Heterogeneous processors have **diverse access pattern**
- Challenge 2: Each prior protection **only for a specific access pattern**
  - For example, GPU coarse-grained pattern, NPU software-detected pattern

## Conventional Memory Protection: High Overhead



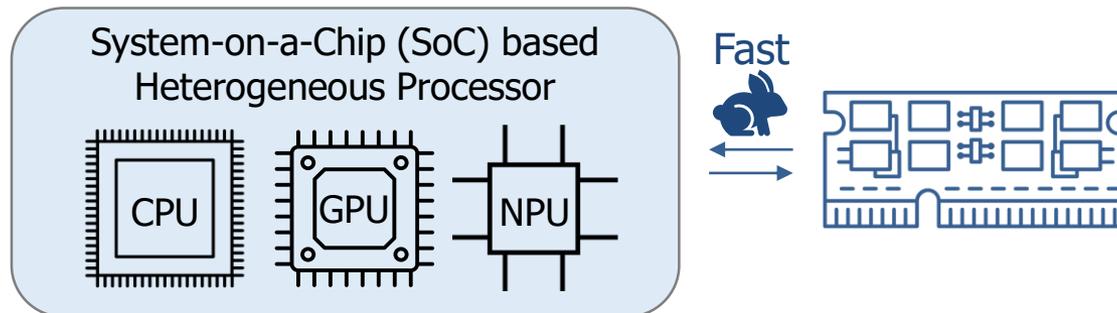
# Research Objective

Constructs a **general** and efficient memory protection scheme for heterogeneous processors

- Challenge 1: Heterogeneous processors have **diverse access pattern**
- Challenge 2: Each prior protection **only for a specific access pattern**
  - For example, GPU coarse-grained pattern, NPU software-detected pattern

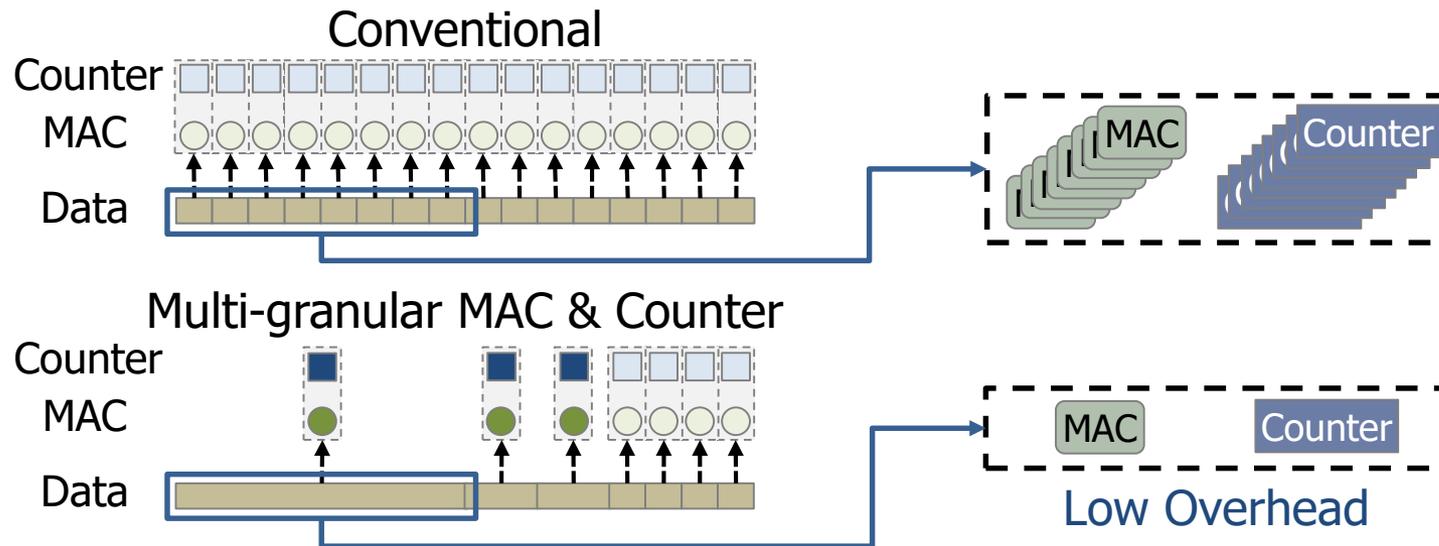
→ We **unified** prior studies with our novel multi-granular tree

## Our Unified Memory Protection Scheme



# Multi-granular MAC and Counter

- Multi-granular MAC and counter
  - Multi-granular MAC and counter fetches **small # of MACs and counters** for coarse-grained access



**Multi-granularity can reduce memory protection overhead  
However, how maintain a counter integrity tree?**

# Prior Memory Protection Schemes

---

# Prior Memory Protection Schemes

---

## Counter-mode Protection

# Prior Memory Protection Schemes

---

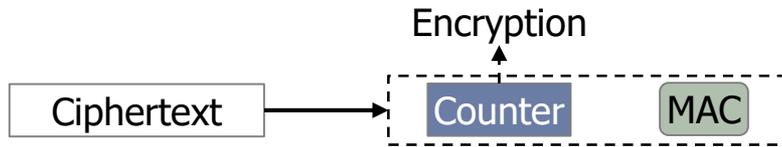
## Counter-mode Protection



# Prior Memory Protection Schemes

---

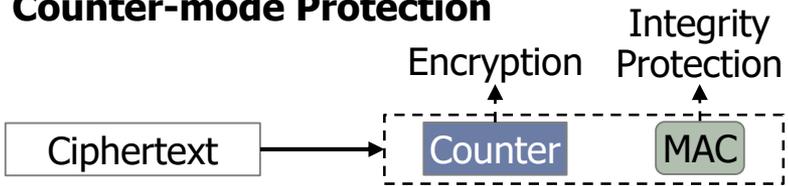
## Counter-mode Protection



# Prior Memory Protection Schemes

---

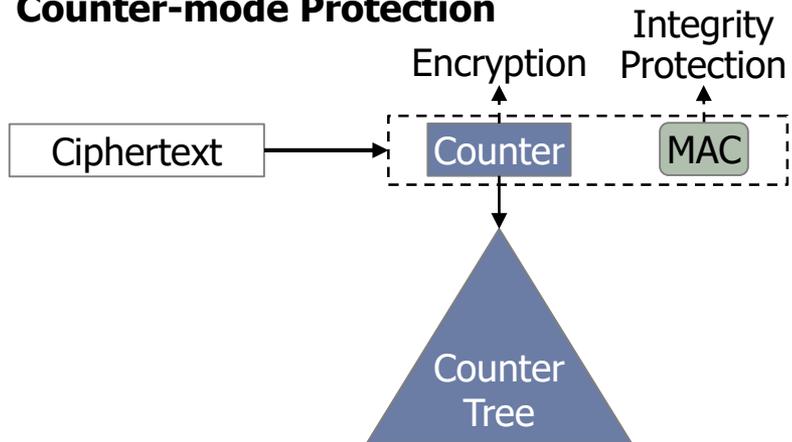
## Counter-mode Protection



# Prior Memory Protection Schemes

---

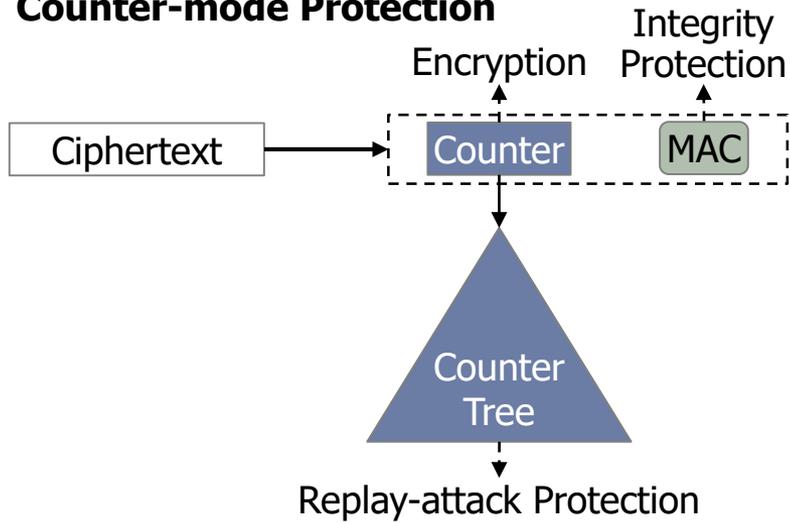
## Counter-mode Protection



# Prior Memory Protection Schemes

---

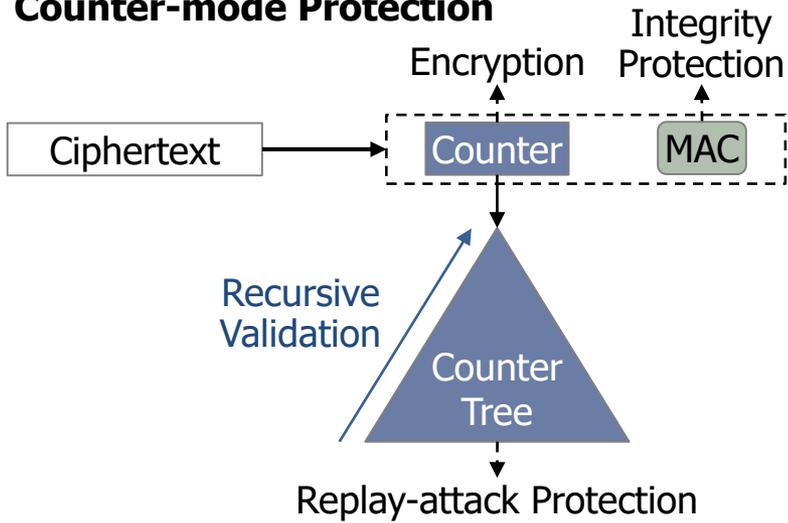
## Counter-mode Protection



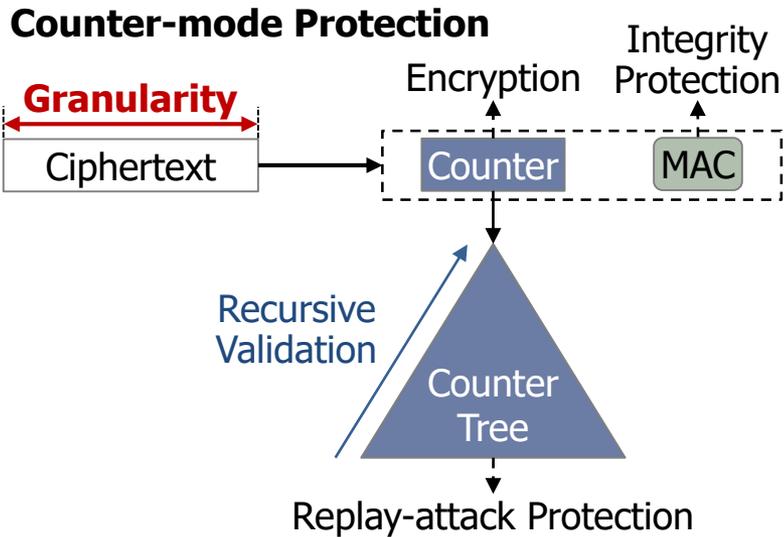
# Prior Memory Protection Schemes

---

## Counter-mode Protection

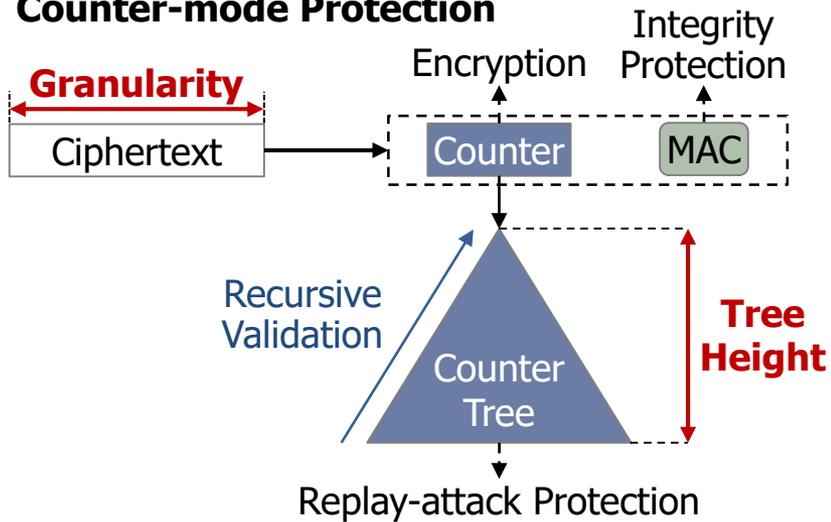


# Prior Memory Protection Schemes



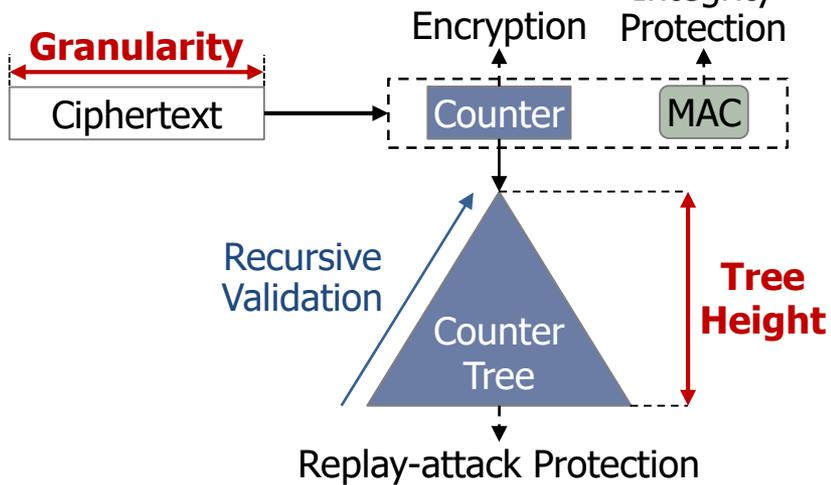
# Prior Memory Protection Schemes

## Counter-mode Protection



# Prior Memory Protection Schemes

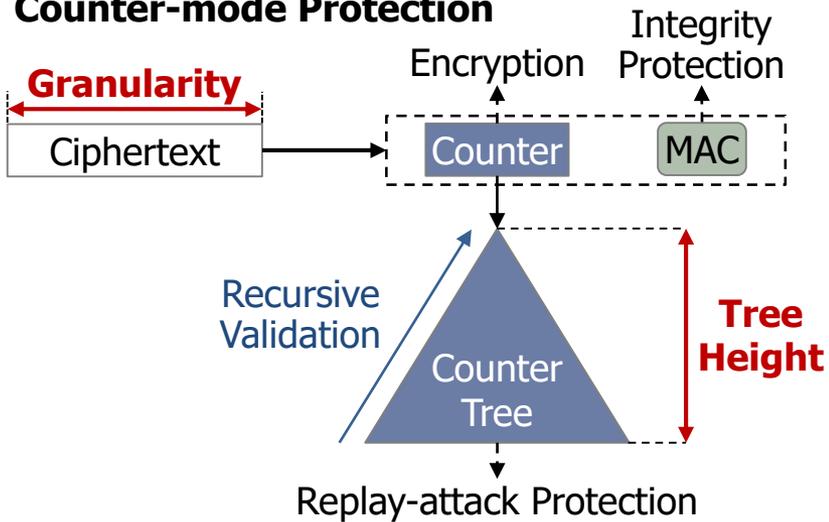
## Counter-mode Protection



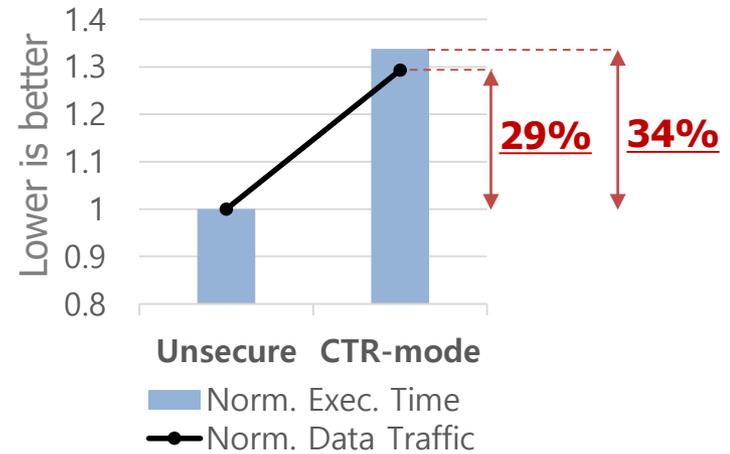
## Conventional 64B-granular Protection

# Prior Memory Protection Schemes

## Counter-mode Protection

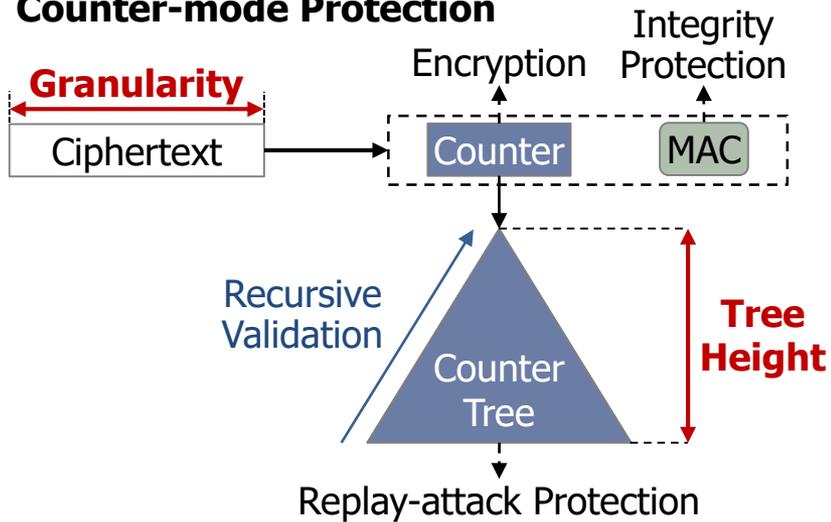


## Conventional 64B-granular Protection

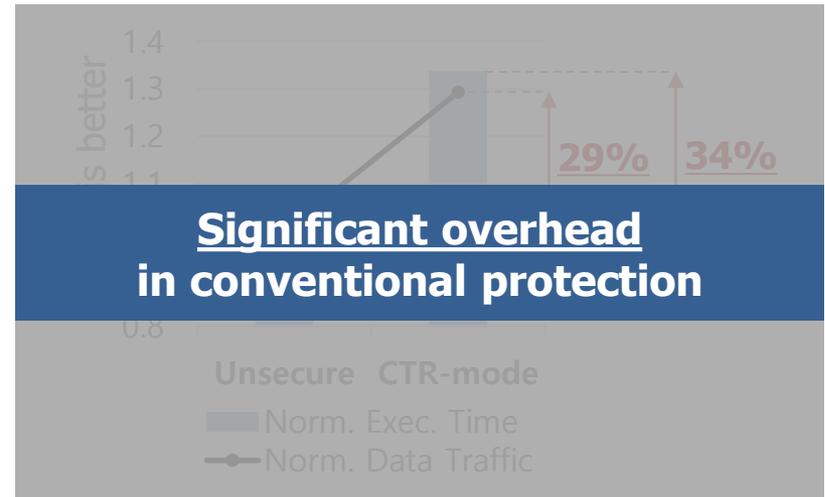


# Prior Memory Protection Schemes

## Counter-mode Protection

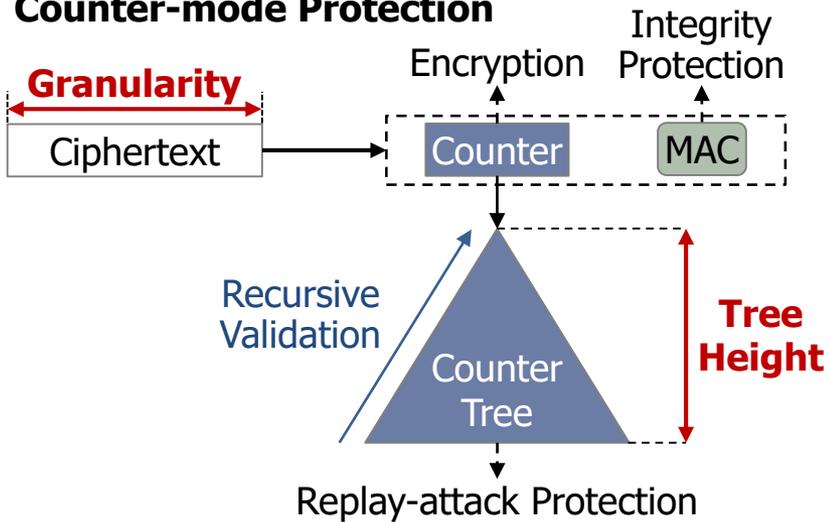


## Conventional 64B-granular Protection



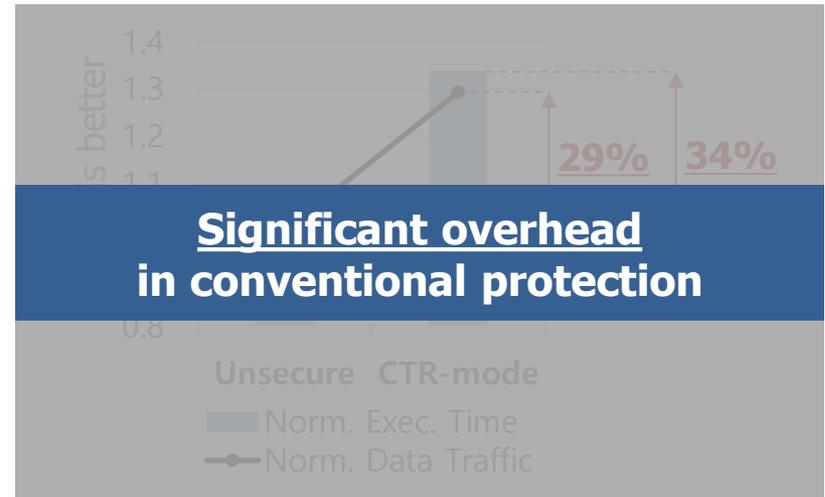
# Prior Memory Protection Schemes

## Counter-mode Protection



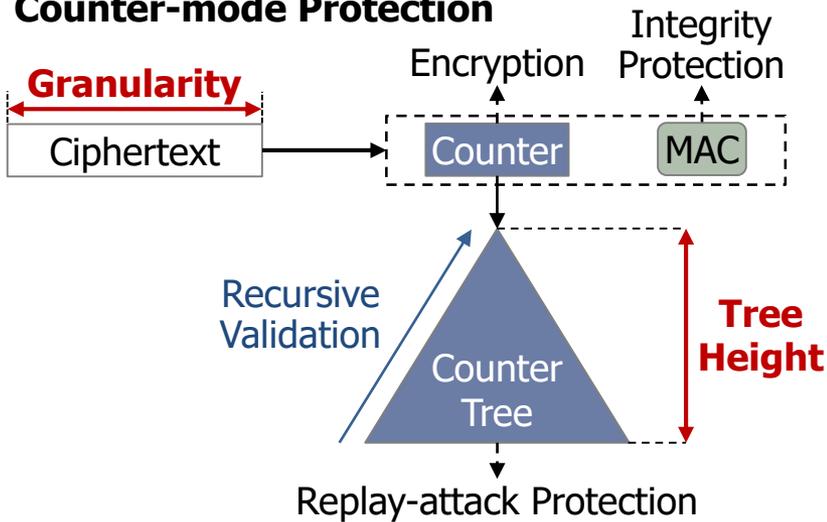
## Prior Domain-specific Memory Protection

## Conventional 64B-granular Protection

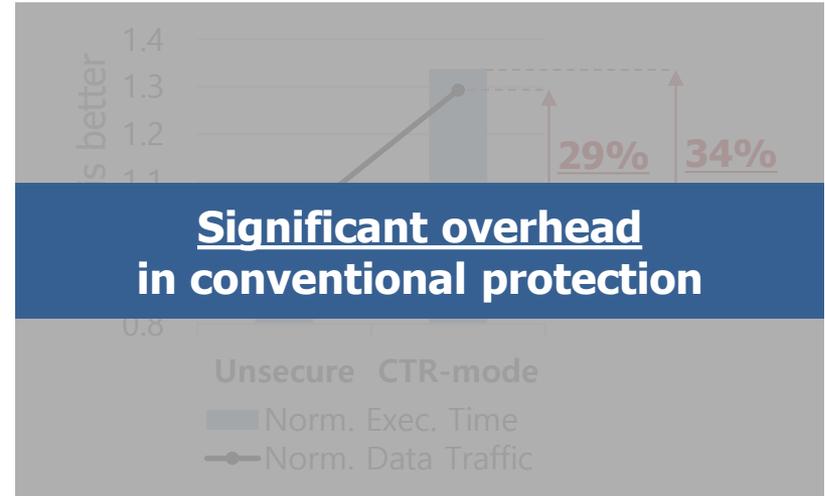


# Prior Memory Protection Schemes

## Counter-mode Protection



## Conventional 64B-granular Protection



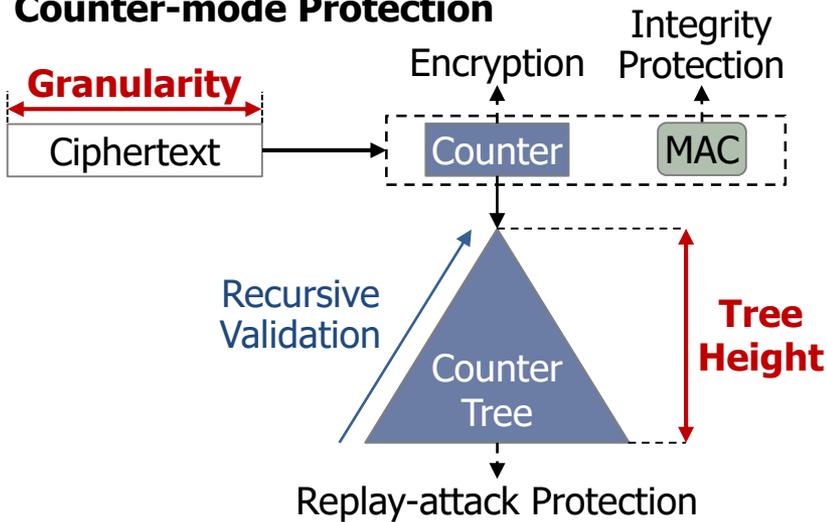
## Prior Domain-specific Memory Protection

1. Common Counters<sup>[1]</sup>

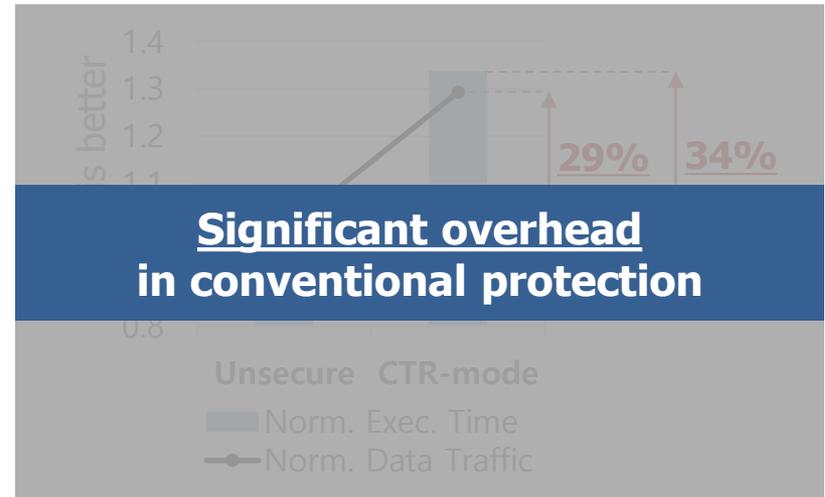
[1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)

# Prior Memory Protection Schemes

## Counter-mode Protection



## Conventional 64B-granular Protection



## Prior Domain-specific Memory Protection

1. Common Counters<sup>[1]</sup>

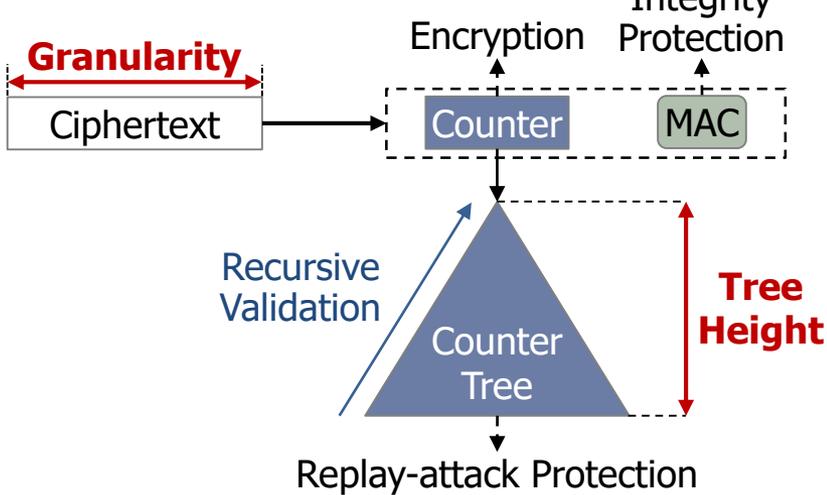
2. Dual-MAC<sup>[2]</sup>

[1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)

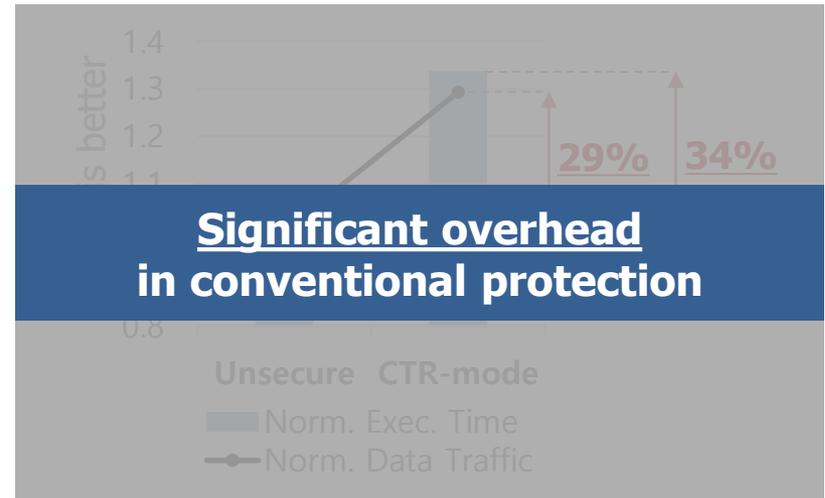
[2] Adaptive Security Support for Heterogeneous Memory on GPUs (HPCA 2022)

# Prior Memory Protection Schemes

## Counter-mode Protection



## Conventional 64B-granular Protection



## Prior Domain-specific Memory Protection

1. Common Counters<sup>[1]</sup>

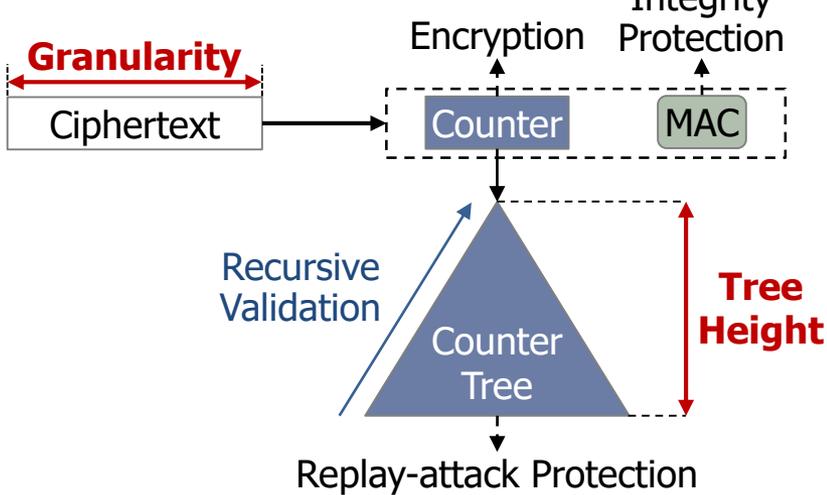
2. Dual-MAC<sup>[2]</sup>

3. Software-managed Granularity<sup>[3-4]</sup>

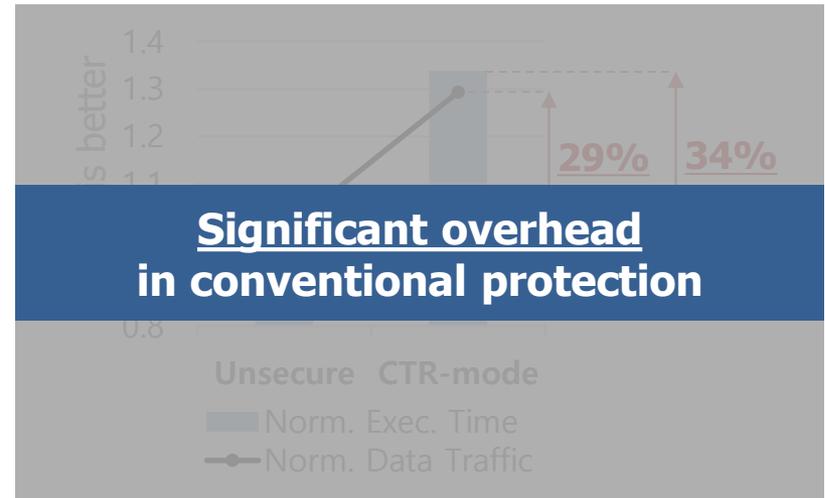
- [1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)
- [2] Adaptive Security Support for Heterogeneous Memory on GPUs (HPCA 2022)
- [3] MGX: Near-zero Overhead Memory Protection for Data-intensive Accelerators (ISCA 2022)
- [4] GuardNN: Secure Accelerator Architecture for Privacy-preserving Deep Learning (DAC 2022)

# Prior Memory Protection Schemes

## Counter-mode Protection



## Conventional 64B-granular Protection



## Prior Domain-specific Memory Protection

1. Common Counters [1]

2. Dual-MAC [2]

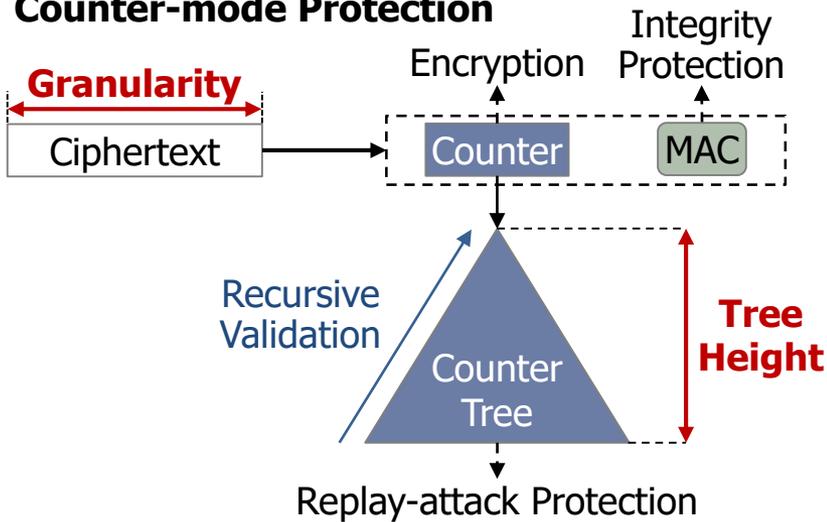
3. Software-managed Granularity [3-4]



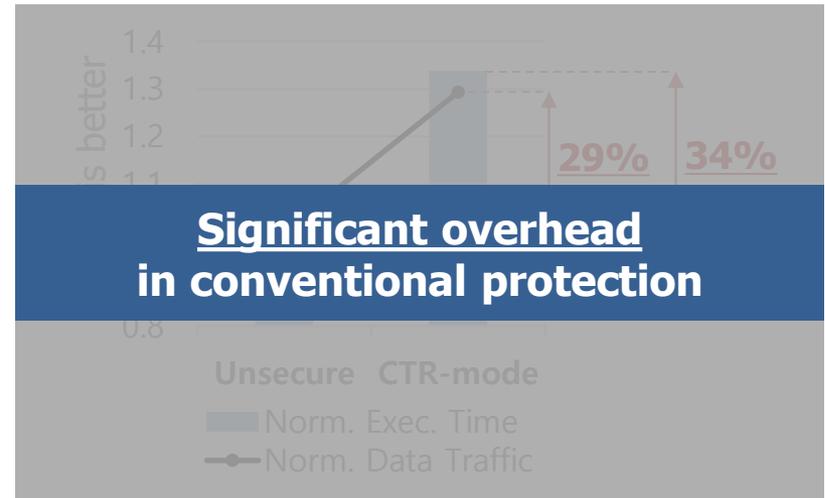
[1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)  
 [2] Adaptive Security Support for Heterogeneous Memory on GPUs (HPCA 2022)  
 [3] MGX: Near-zero Overhead Memory Protection for Data-intensive Accelerators (ISCA 2022)  
 [4] GuardNN: Secure Accelerator Architecture for Privacy-preserving Deep Learning (DAC 2022)

# Prior Memory Protection Schemes

## Counter-mode Protection



## Conventional 64B-granular Protection

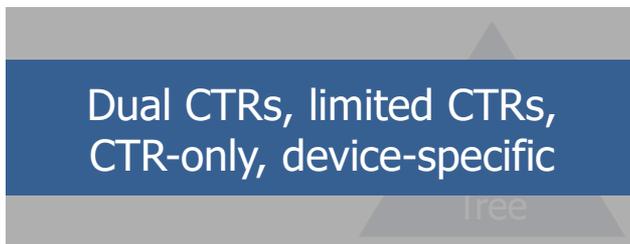


## Prior Domain-specific Memory Protection

1. Common Counters [1]

2. Dual-MAC [2]

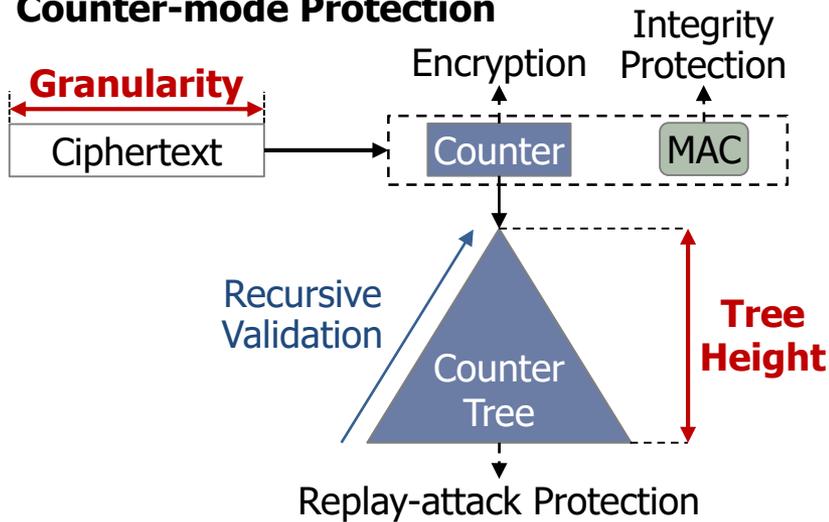
3. Software-managed Granularity [3-4]



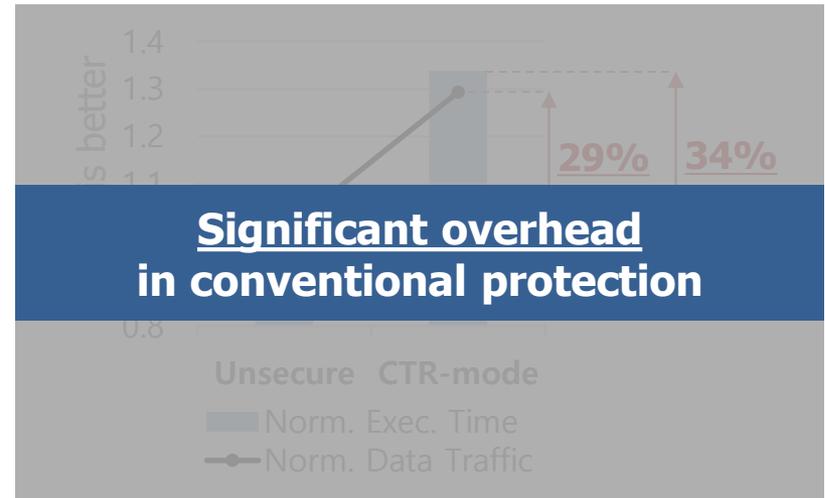
- [1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)
- [2] Adaptive Security Support for Heterogeneous Memory on GPUs (HPCA 2022)
- [3] MGX: Near-zero Overhead Memory Protection for Data-intensive Accelerators (ISCA 2022)
- [4] GuardNN: Secure Accelerator Architecture for Privacy-preserving Deep Learning (DAC 2022)

# Prior Memory Protection Schemes

## Counter-mode Protection



## Conventional 64B-granular Protection

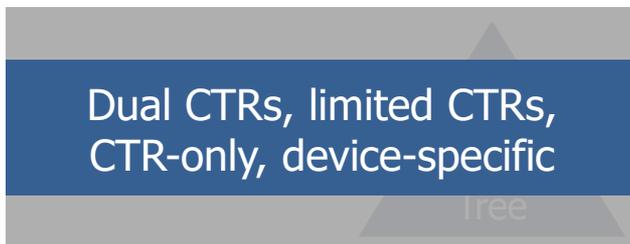


## Prior Domain-specific Memory Protection

1. Common Counters [1]

2. Dual-MAC [2]

3. Software-managed Granularity [3-4]



Coarse MAC

OR

Fine MAC

[1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)

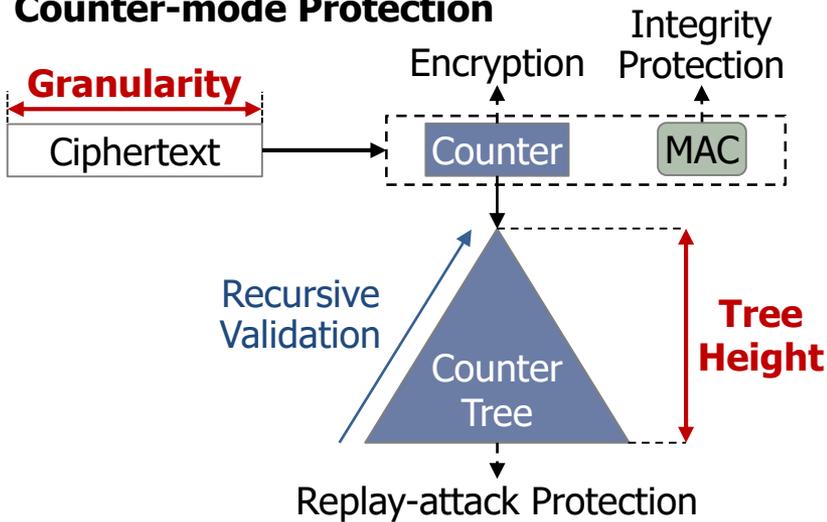
[2] Adaptive Security Support for Heterogeneous Memory on GPUs (HPCA 2022)

[3] MGX: Near-zero Overhead Memory Protection for Data-intensive Accelerators (ISCA 2022)

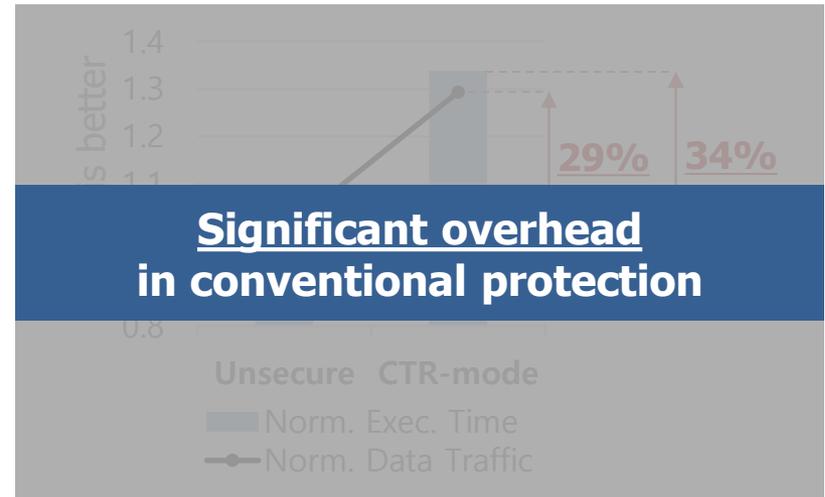
[4] GuardNN: Secure Accelerator Architecture for Privacy-preserving Deep Learning (DAC 2022)

# Prior Memory Protection Schemes

## Counter-mode Protection



## Conventional 64B-granular Protection

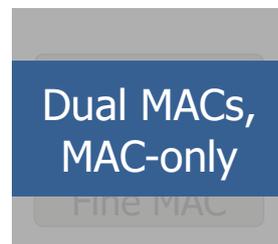
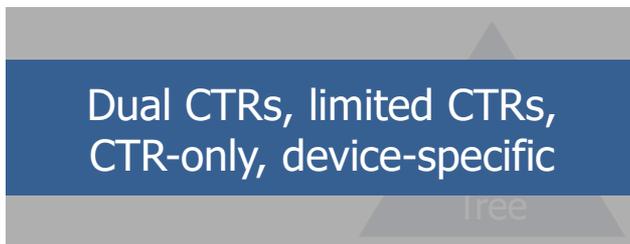


## Prior Domain-specific Memory Protection

1. Common Counters [1]

2. Dual-MAC [2]

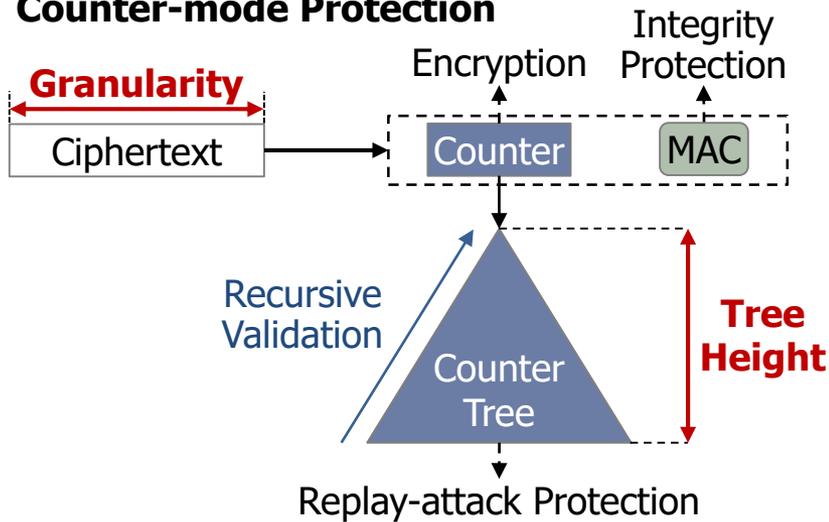
3. Software-managed Granularity [3-4]



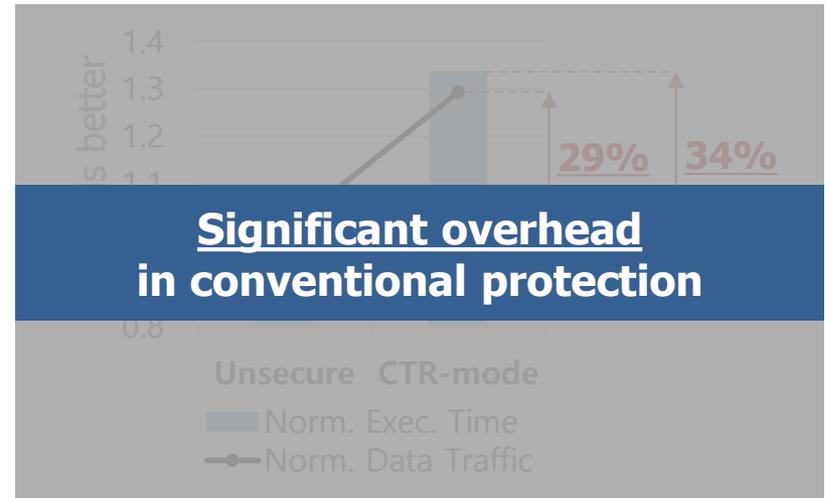
- [1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)
- [2] Adaptive Security Support for Heterogeneous Memory on GPUs (HPCA 2022)
- [3] MGX: Near-zero Overhead Memory Protection for Data-intensive Accelerators (ISCA 2022)
- [4] GuardNN: Secure Accelerator Architecture for Privacy-preserving Deep Learning (DAC 2022)

# Prior Memory Protection Schemes

## Counter-mode Protection



## Conventional 64B-granular Protection

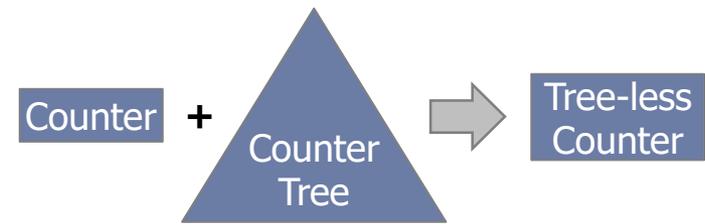
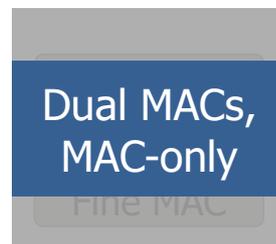


## Prior Domain-specific Memory Protection

### 1. Common Counters [1]

### 2. Dual-MAC [2]

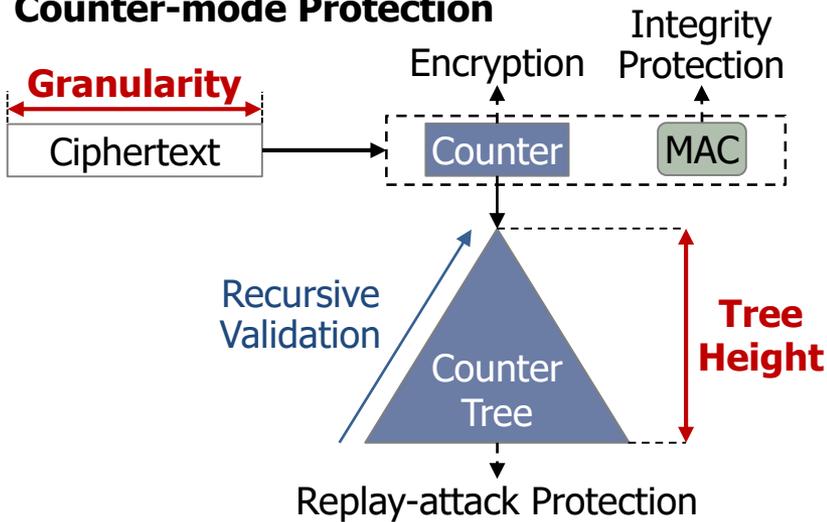
### 3. Software-managed Granularity [3-4]



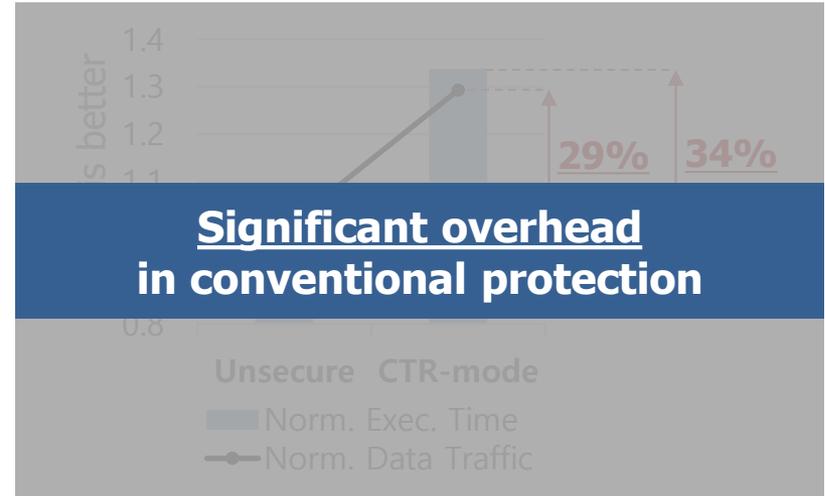
- [1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)
- [2] Adaptive Security Support for Heterogeneous Memory on GPUs (HPCA 2022)
- [3] MGX: Near-zero Overhead Memory Protection for Data-intensive Accelerators (ISCA 2022)
- [4] GuardNN: Secure Accelerator Architecture for Privacy-preserving Deep Learning (DAC 2022)

# Prior Memory Protection Schemes

## Counter-mode Protection

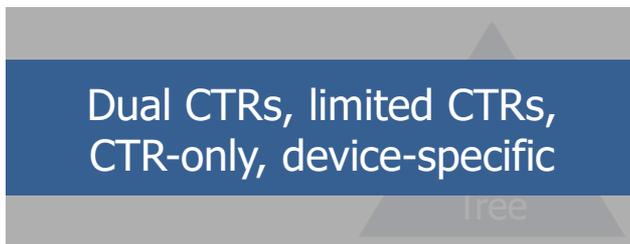


## Conventional 64B-granular Protection

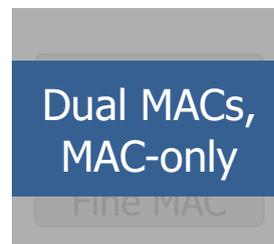


## Prior Domain-specific Memory Protection

### 1. Common Counters [1]



### 2. Dual-MAC [2]



### 3. Software-managed Granularity [3-4]



- [1] Common Counters: Compressed Encryption Counters for Secure GPU Memory (HPCA 2021)
- [2] Adaptive Security Support for Heterogeneous Memory on GPUs (HPCA 2022)
- [3] MGX: Near-zero Overhead Memory Protection for Data-intensive Accelerators (ISCA 2022)
- [4] GuardNN: Secure Accelerator Architecture for Privacy-preserving Deep Learning (DAC 2022)