

# (Offload) ASAN via Software Managed Virtual Memory

Johannes Doerfert <jdoerfert@llnl.gov>  
Ethan McDonough <ethanluismcdonough@gmail.com>  
Vidush Singhal <singhav@purdue.edu>

# Setting the Stage

Work in Progress → Room for Improvement → Approximate & Limited Results

# Thanks

Aaron Jarmusch <jarmusch@udel.edu> — NVIDIA measurements

Rahulkumar Gayatri <rgayatri@lbl.gov> — KOKKOS testing & measurements

# Prologue

```
>>> nvim src/simulation.cpp
```

```
>>> nvim src/simulation.cpp  
>>> ninja -C build
```

```
>>> nvim src/simulation.cpp  
>>> ninja -C build  
>>> ./build/bin/science -max
```



```
>>> nvim src/simulation.cpp  
>>> ninja -C build  
>>> ./build/bin/science -max
```

Starting to produce **maximal** science...

```
>>> nvim src/simulation.cpp
>>> ninja -C build
>>> ./build/bin/science -max
```

Starting to produce `maximal` science...

**AMDGPU fatal error 1: Memory access fault by GPU 2 (agent 0x557136cb9fd0)  
at virtual address 0x7f3780422000. Reasons: Unknown (0)**

```
>>> nvim src/simulation.cpp
>>> ninja -C build
>>> ./build/bin/science -max
```

Starting to produce **maximal** science...

Display last 2 kernels launched:

Kernel 0: 'omp target in openmc::process\_collision\_events() @ 352

Kernel 1: 'omp target in openmc::process\_surface\_crossing\_events() @ 323

**AMDGPU fatal error 1: Memory access fault by GPU 2 (agent 0x557136cb9fd0)  
at virtual address 0x7f3780422000. Reasons: Unknown (0)**

# Main Cast

- instrumentation + runtime
- shadow memory & redzones
- staple since LLVM 3.1 (~2012)

# ASAN

the code snitch reporting stale  
pointers and buffer overflows

# GPUs

the AI accelerators that can do graphics and general compute

- powerhouse of modern HPC
- we initially focus on AMD GPUs
- dedicated memory is assumed

Plot

# Architecture

Code

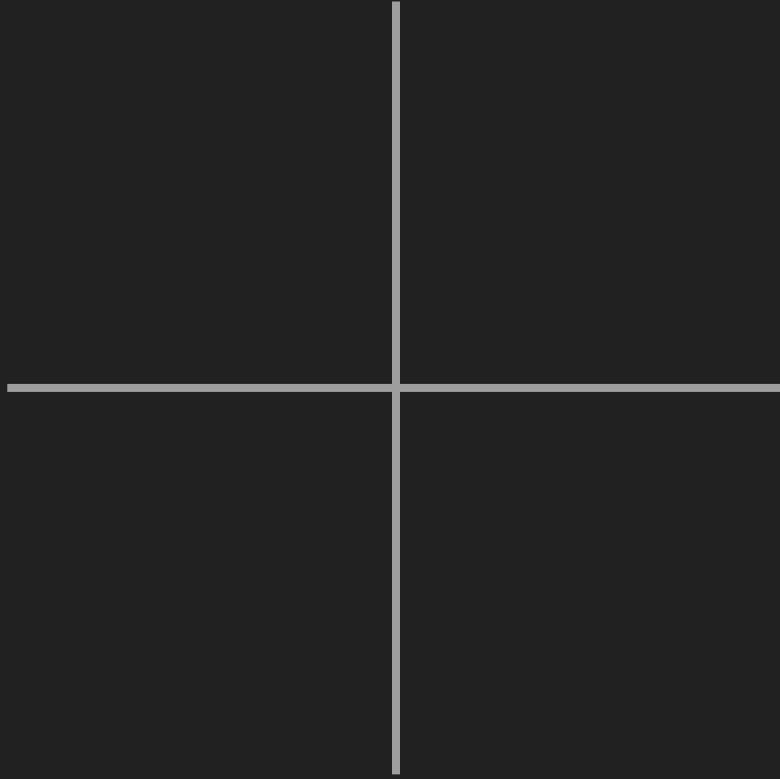


# Architecture

CPU

Code

Correct



# Architecture

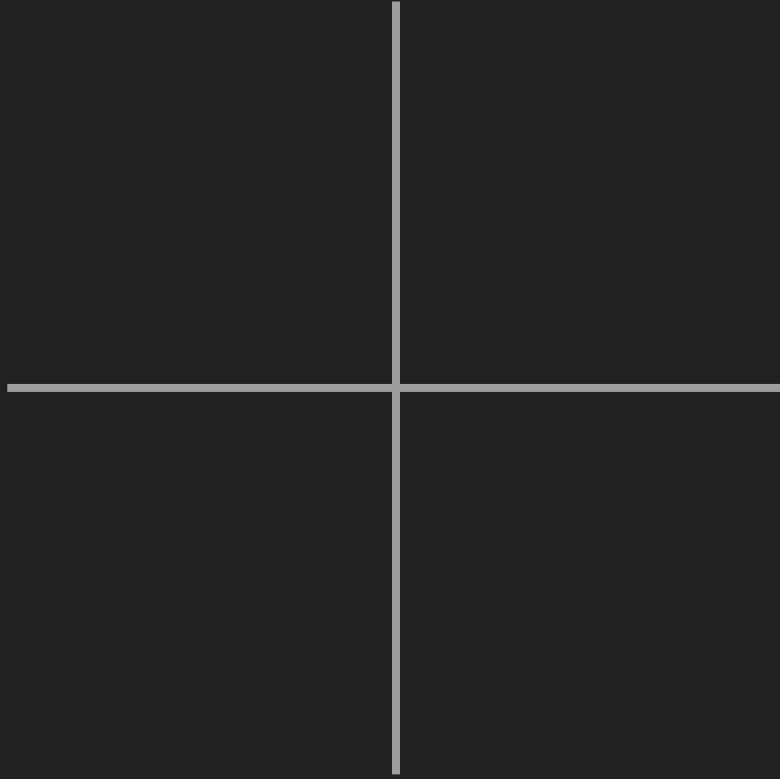
CPU

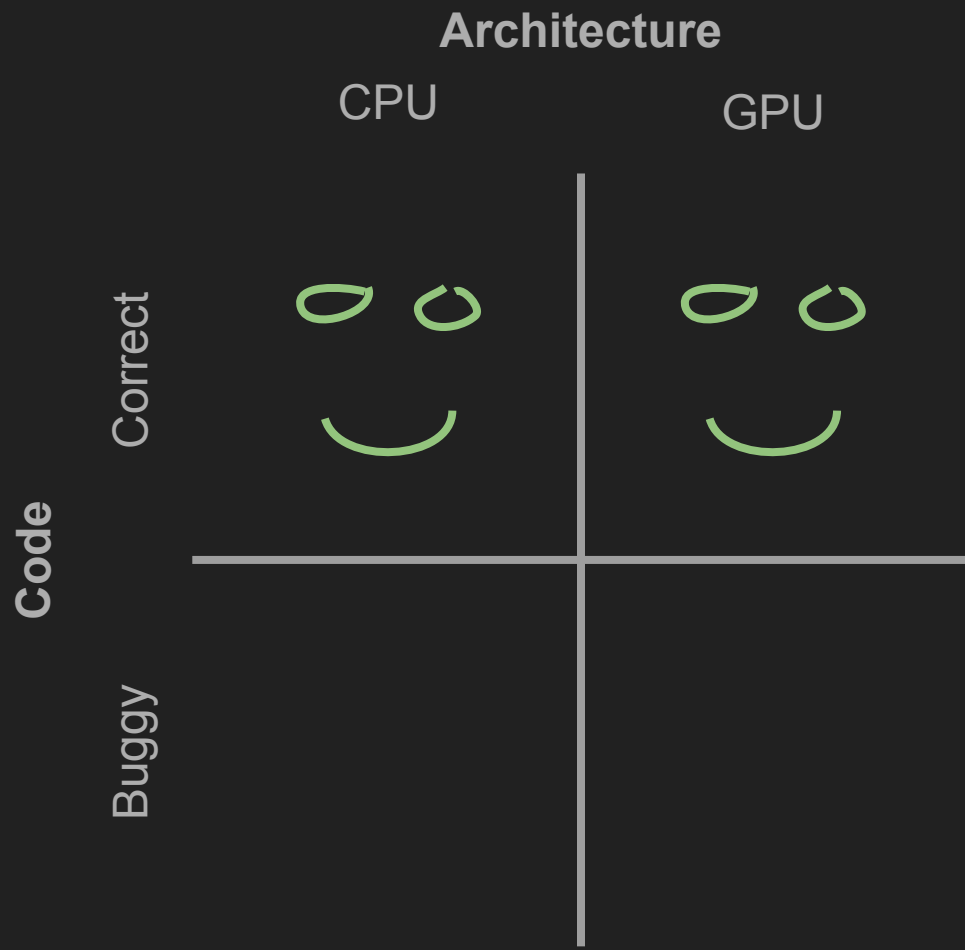
GPU




Code

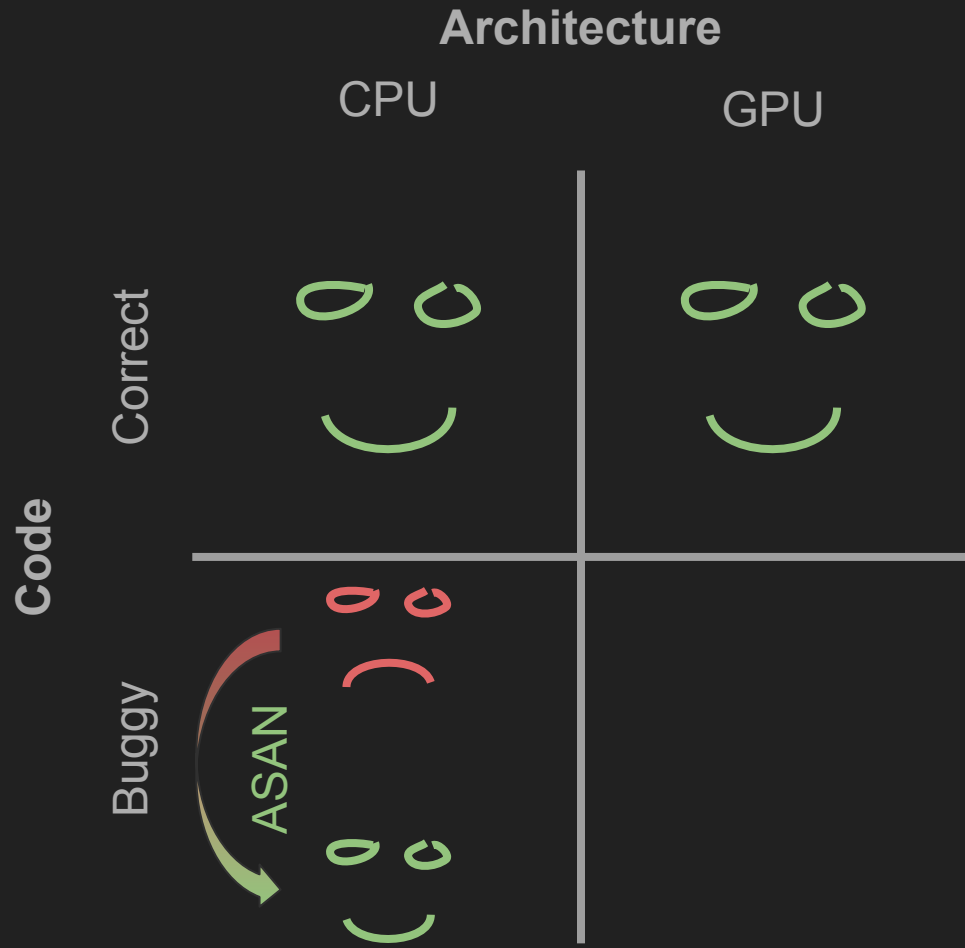
Correct

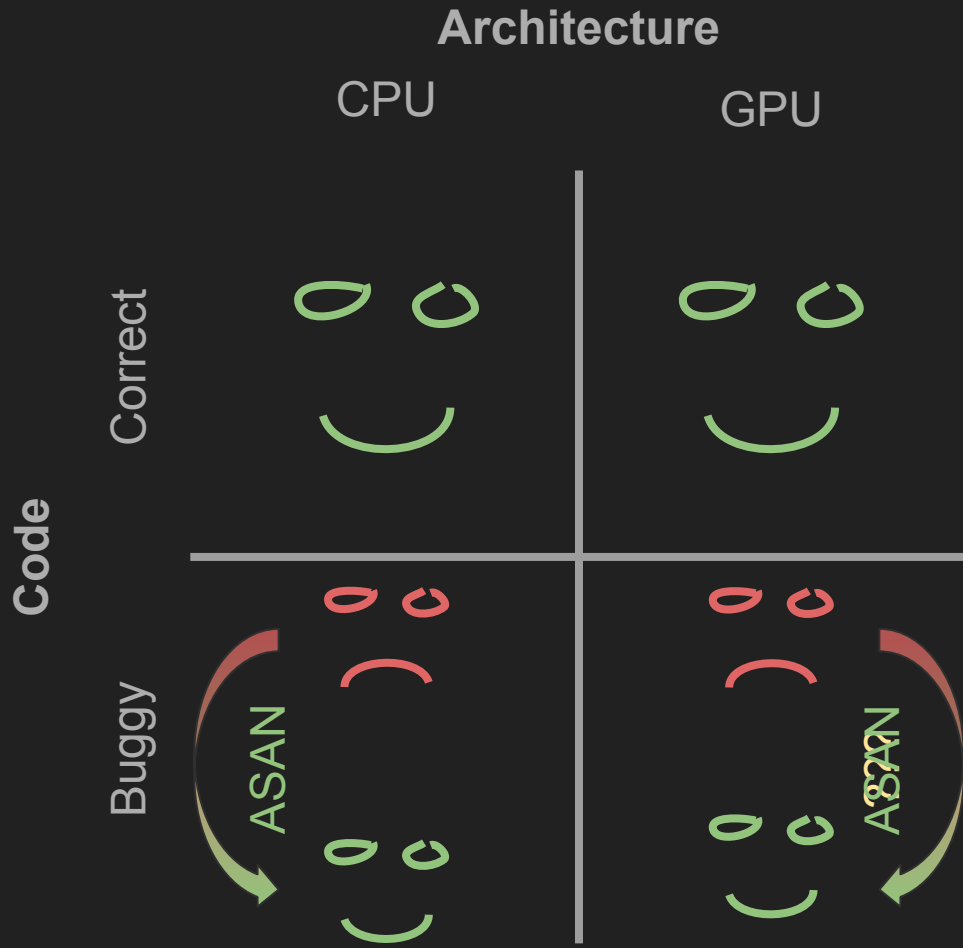
Buggy



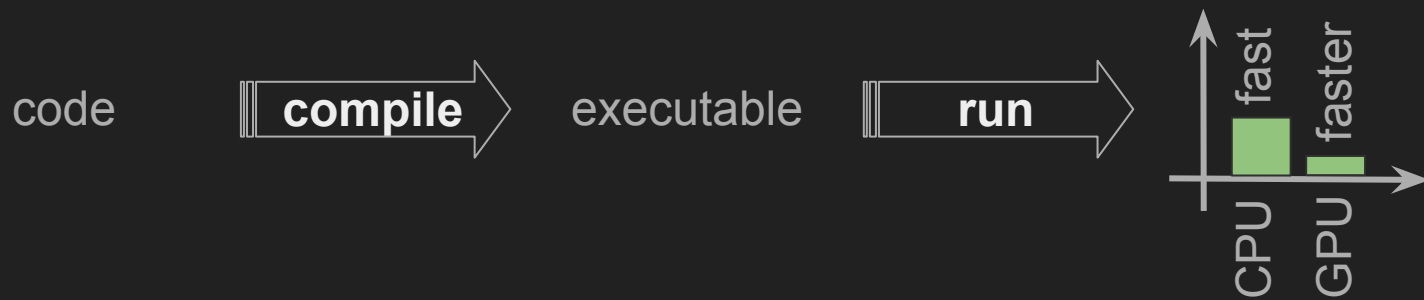


		Architecture	
		CPU	GPU
Code	Correct		
	Buggy		

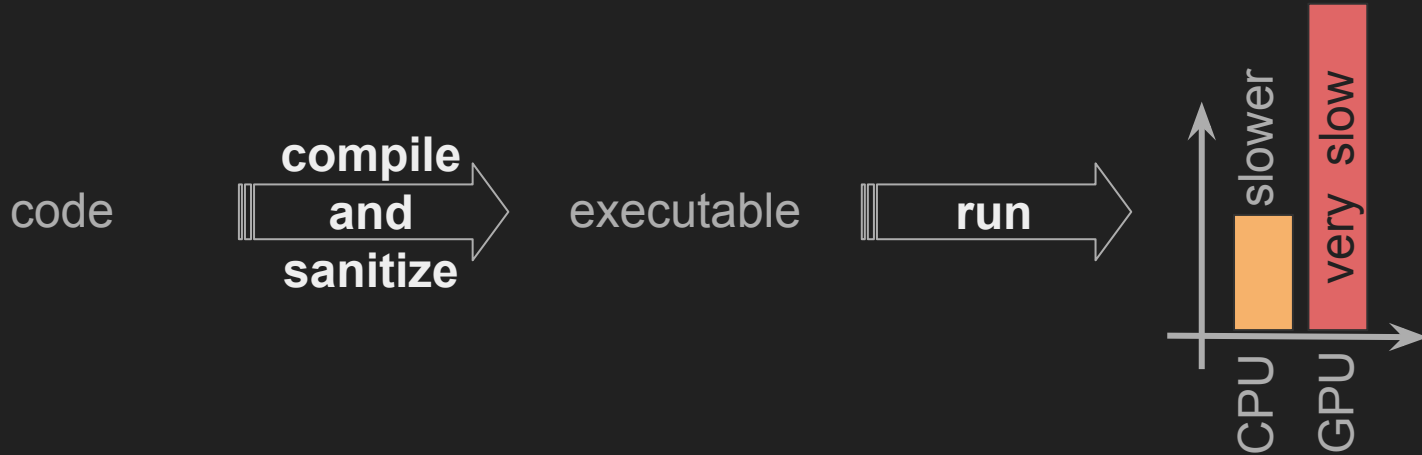
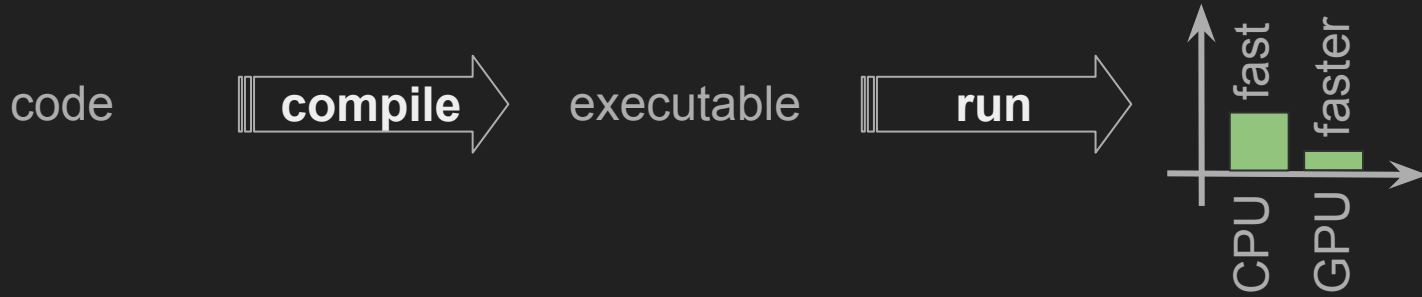




Complication





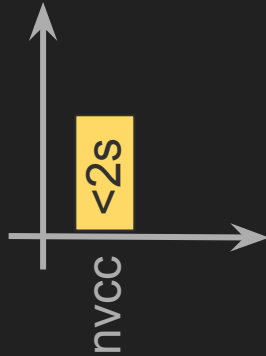


Situation

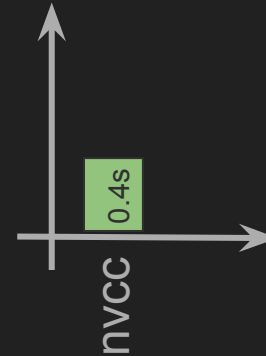
## **XSbench**

- continuous energy macroscopic neutron cross section lookup
- event-based lookup, unionized grid, large H-M size (~5.5GB)
- proxy application for OpenMC
- CUDA/HIP/OpenMP offload versions

V100 - OpenMP



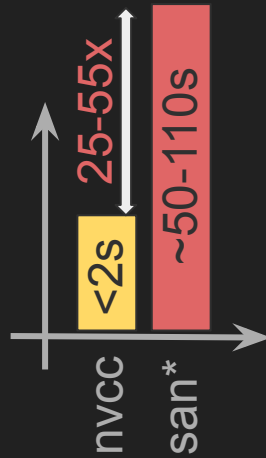
V100 - CUDA



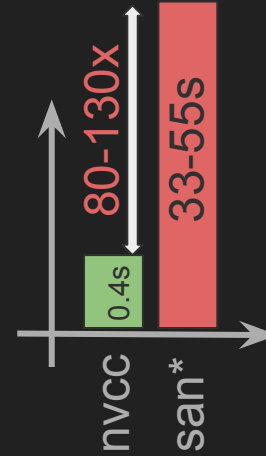
## XSbench

- continuous energy macroscopic neutron cross section lookup
- event-based lookup, unionized grid, large H-M size (~5.5GB)
- proxy application for OpenMC
- CUDA/HIP/OpenMP offload versions

## V100 - OpenMP



## V100 - CUDA

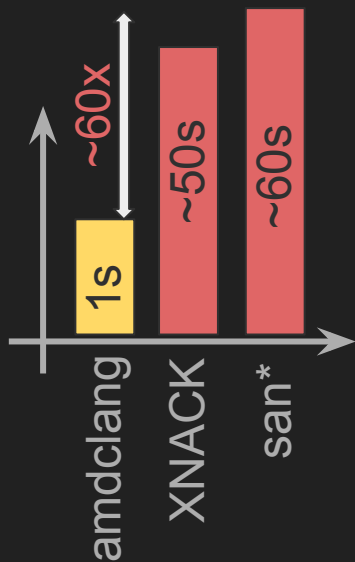


\* NVIDIA Compute Sanitizer

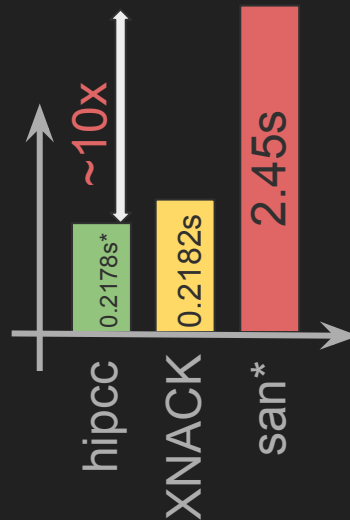
## XSbench

- continuous energy macroscopic neutron cross section lookup
- event-based lookup, unionized grid, large H-M size (~5.5GB)
- proxy application for OpenMC
- CUDA/HIP/OpenMP offload versions

## MI-250X - OpenMP



## MI-250X - HIP

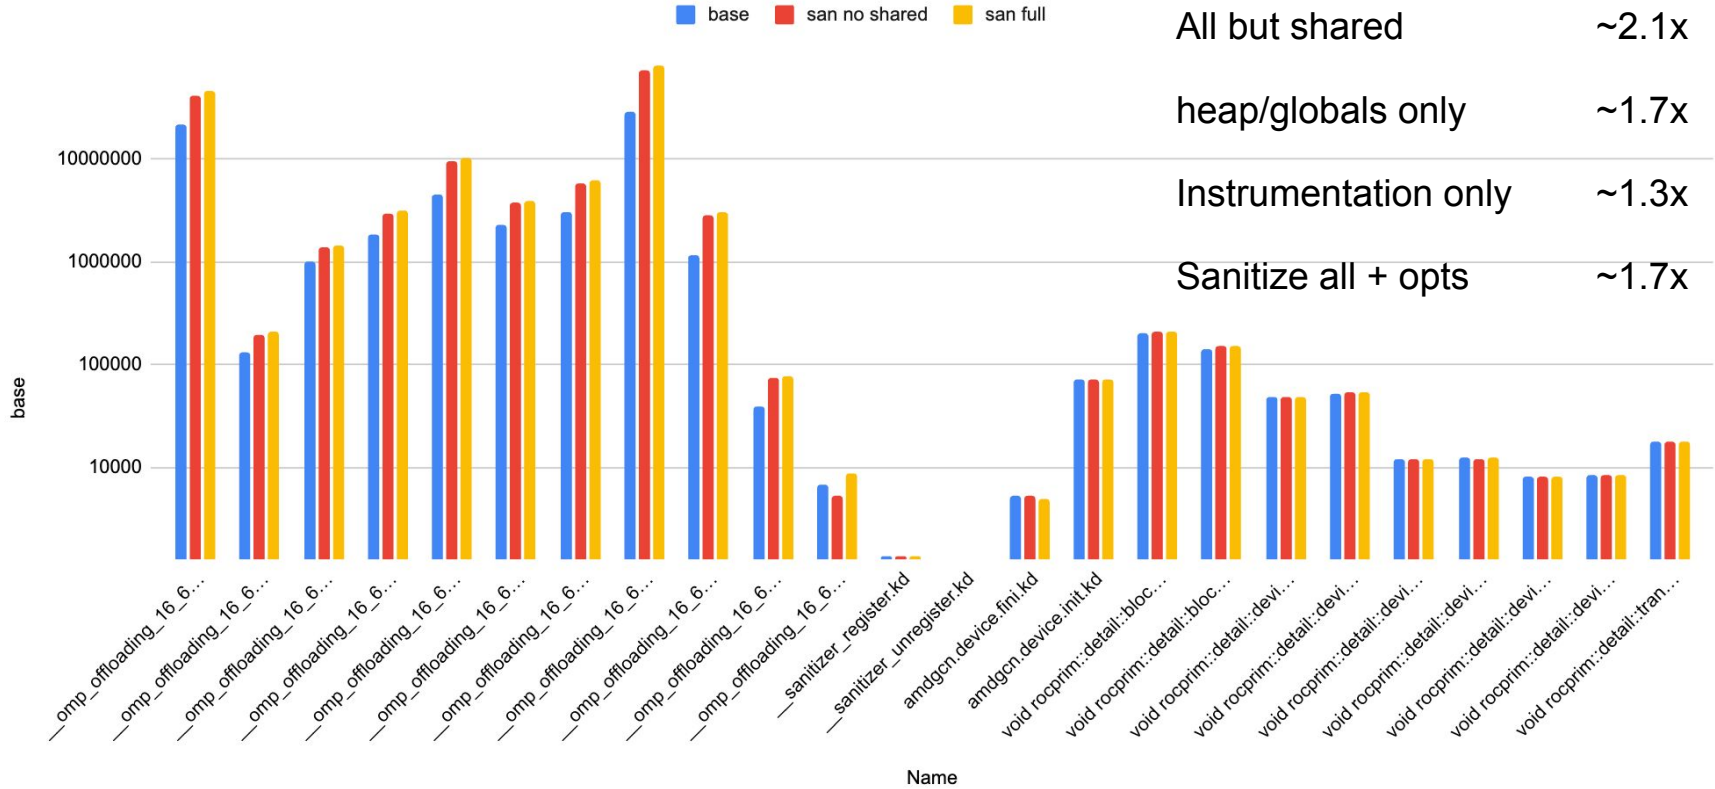


\* AMD ASAN for GPU incl. XNACK

### XSbench

- continuous energy macroscopic neutron cross section lookup
- event-based lookup, unionized grid, large H-M size (~5.5GB)
- proxy application for OpenMC
- CUDA/HIP/OpenMP offload versions

# OpenMC kernel times



# Backstory



fixed-sized  
red zones

...



double  
A[1000]



...



long x



...

fixed-sized  
red zones

...



double  
A[1000]

...



long x

...

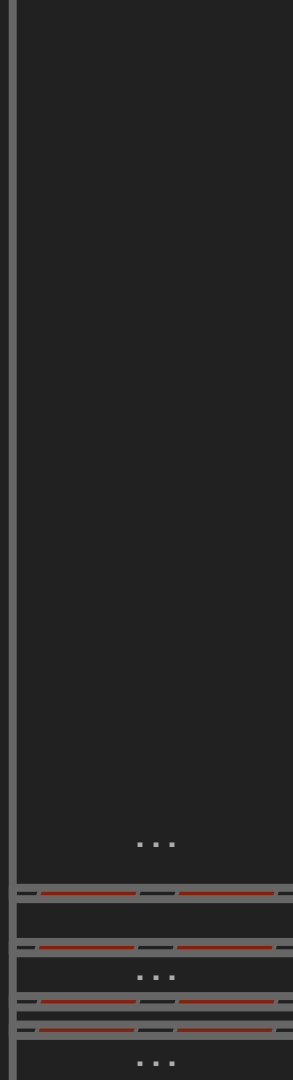
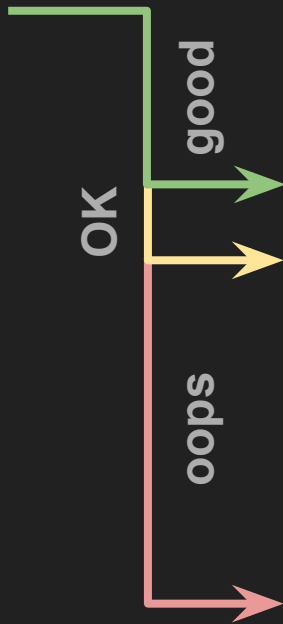
1/8 mapping

...

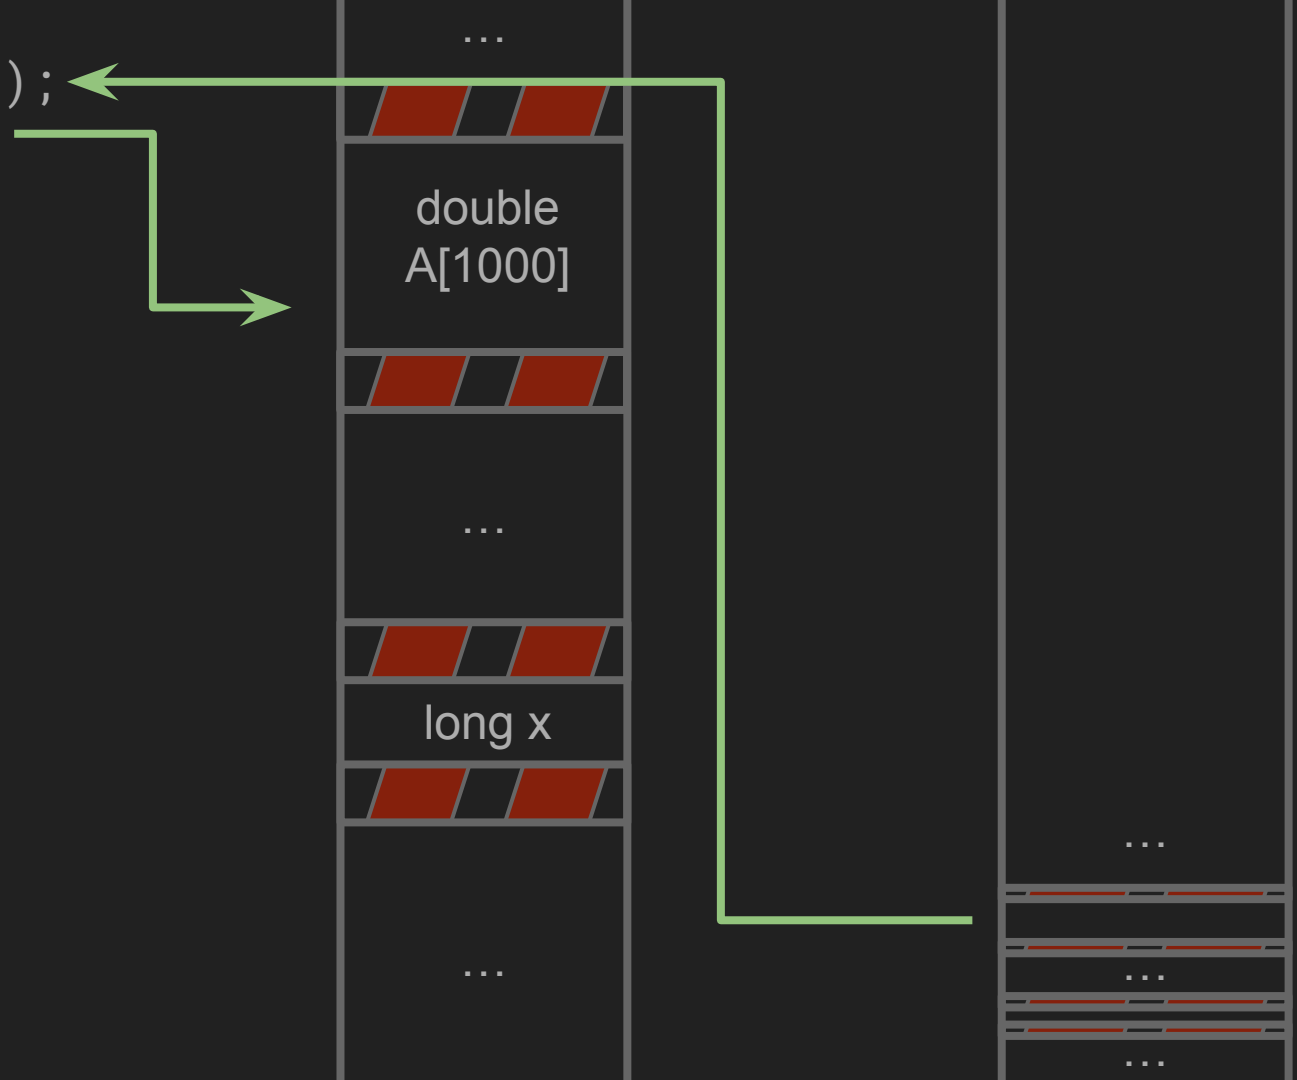
...

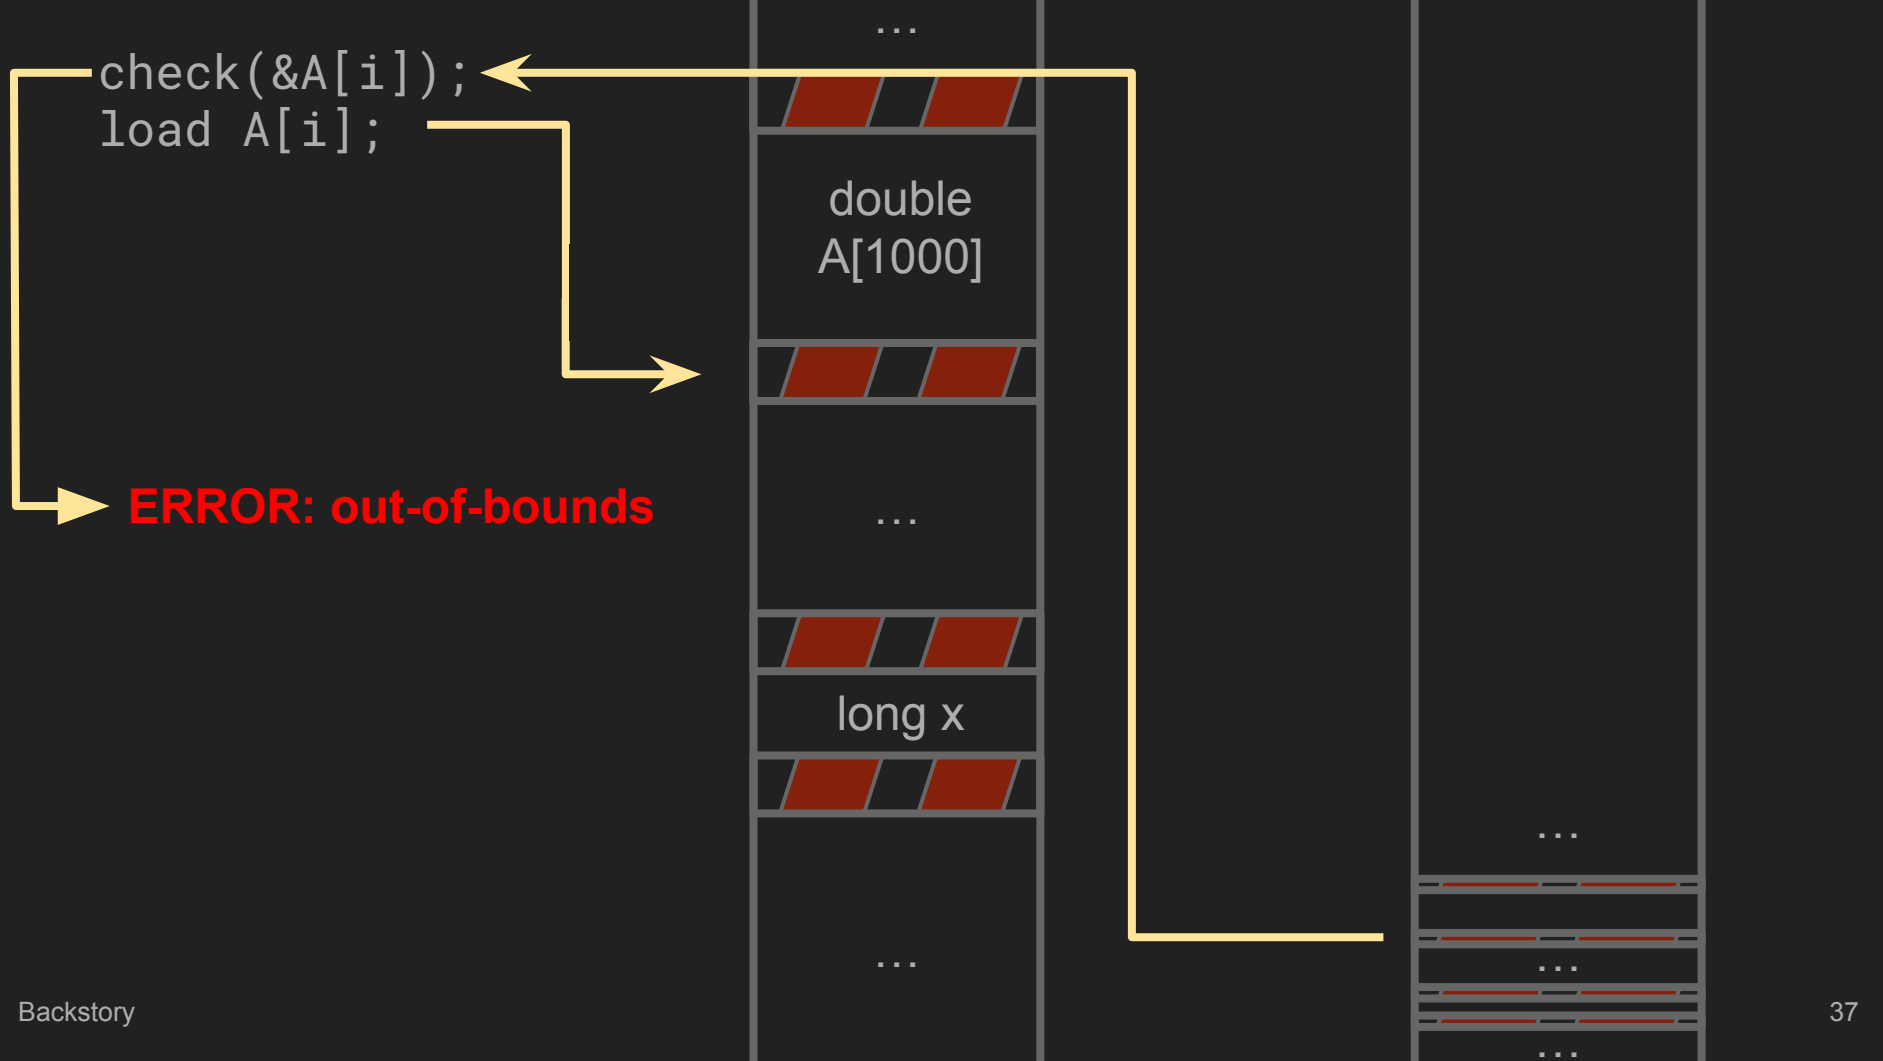
...

load A[i];



```
check(&A[i]);  
load A[i];
```

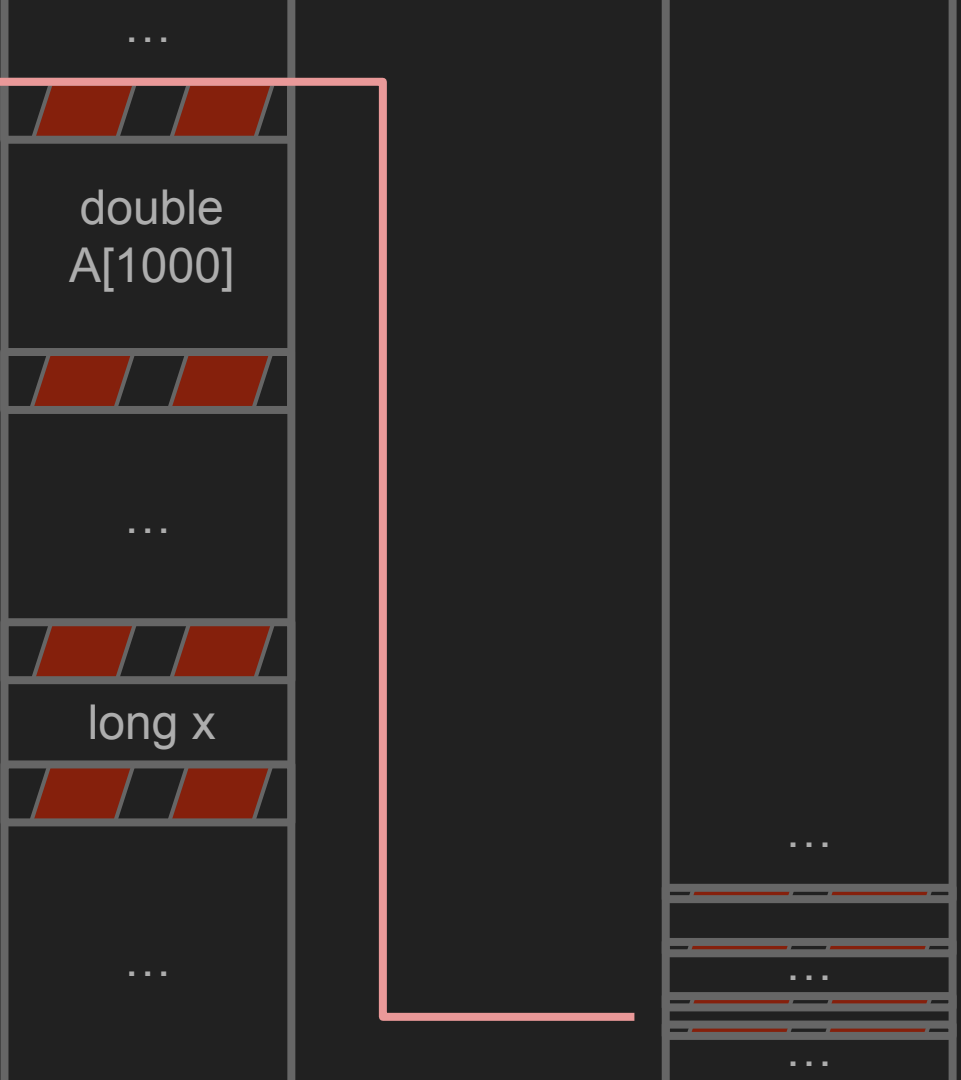




```
check(&A[i]);  
load A[i];
```

**ERROR: out-of-bounds**

```
check(&A[i]);  
load A[i];
```



Recap

CPU and GPU  
are different

- open vs closed environment
- high vs low total memory
- 1 memory vs 3+ memories

ASAN uses extra  
memory, accesses

- >1.25x memory per allocation
- 2x number of accesses
- false negatives possible

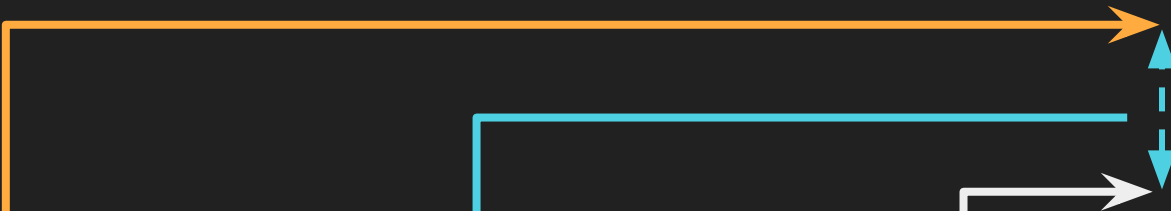
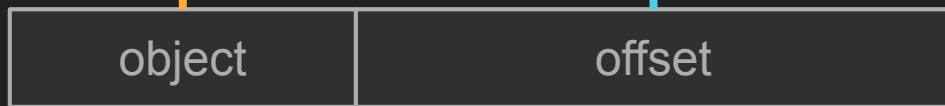
overheads of 1-2  
orders of magnitude

- ASAN increases compile time
- vendor products only,  
non-portable, varying results



Protagonist

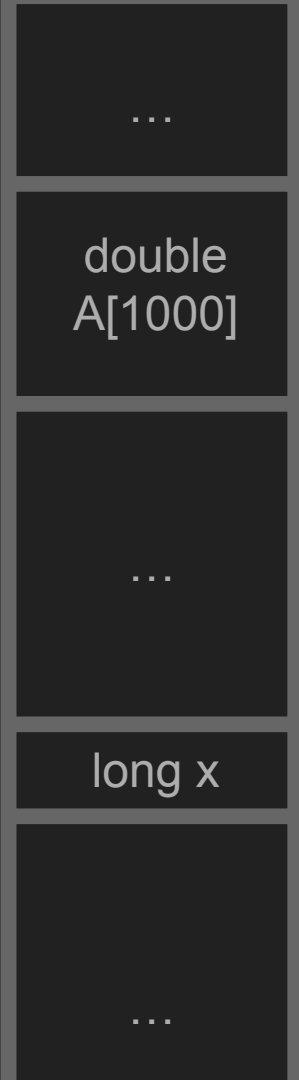
PTR



allocation record (16 bytes each)



PTR



```
void increment(double *A, int N) {  
    for (int I = 0; I < N; ++I) {  
        A[I]++;  
    }  
}
```

```
void increment(double *A, int N) {  
    auto [Base, Size, Offset] = __lookup(A);  
    for (int I = 0; I < N; ++I) {  
  
        A[I]++;  
    }  
}
```

```
void increment(double *A, int N) {  
    auto [Base, Size, Offset] = __lookup(A);  
    for (int I = 0; I < N; ++I) {  
        __check(Offset + I, Size);  
        A[I]++;  
    }  
}
```

```
void increment(double *A, int N) {  
    auto [Base, Size, Offset] = __lookup(A);  
    for (int I = 0; I < N; ++I) {  
        __check(Offset + I, Size);  
        (Base + Offset + I)[I]++;  
    }  
}
```

allocation record (16 bytes each)

...	&A   8000	&x   8	...
-----	-----------	--------	-----

PTR

object	offset
--------	--------



allocation record (16 bytes each)

...	&A   8000	&x   8	...
-----	-----------	--------	-----

- on-device registration for each allocation
- 16 bytes load at base pointer definition (`__lookup`)
- bit operations and comparison for each access (`__check`)
- **bad** for 32-bit shared and stack pointers

PTR

object	magic	offset
--------	-------	--------

allocation record (16 bytes each)

...	&A   8000	&x   8	...
-----	-----------	--------	-----

- on-device registration for each allocation
- 16 bytes load at base pointer definition (`__lookup`)
- bit operations and comparison for each access (`__check`)
- **bad** for 32-bit shared and stack pointers

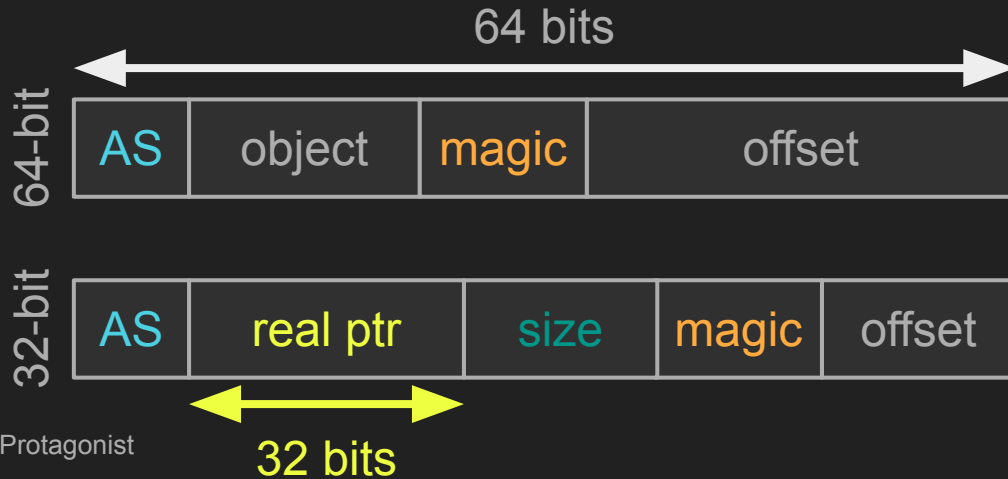
PTR

AS	object	magic	offset
----	--------	-------	--------

allocation record (16 bytes each)



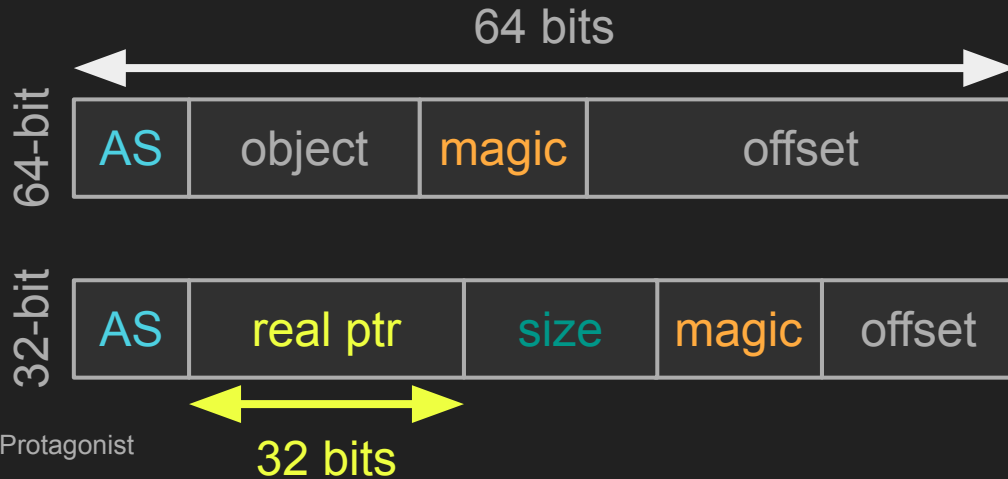
- on-device registration for each allocation
- 16 bytes load at base pointer definition (`__lookup`)
- bit operations and comparison for each access (`__check`)
- **bad** for 32-bit shared and stack pointers

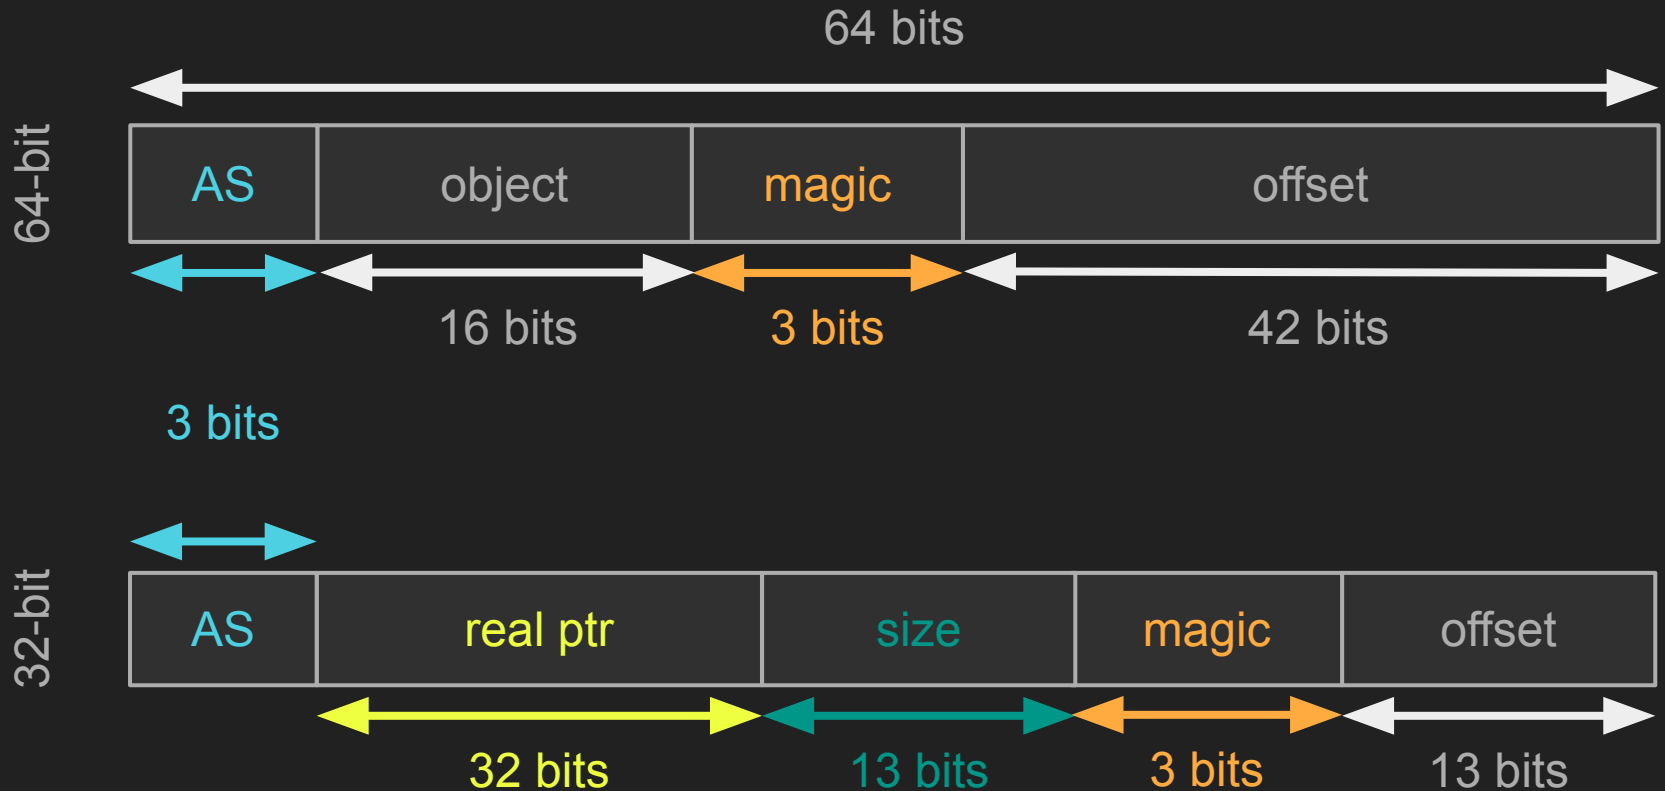


allocation record (16 bytes each)



- on-device registration for each **heap** allocation
- **replace 32-bit pointers with 64-bit pointers** (remove AS)
- **AS check** + 16 bytes load at **heap** base pointer definition (`__lookup`)
- bit operations and comparison for each access (`__check`)
- **no memory overhead** for 32-bit shared and stack pointers





```
void increment(double *A, int N) {
    auto [Base, Size, Offset] = __lookup(A);
    for (int I = 0; I < N; ++I) {
        __check(Offset + I, Size);
        (Base + Offset + I)[I]++;
    }
}
```

```
void increment(double *A, int N) {
    auto [AS, Base, Size, Offset] = __lookup(A);
    for (int I = 0; I < N; ++I) {
        __check(AS, Offset + I, Size);
        (Base + Offset + I)[I]++;
    }
}
```

```
void increment(double AS(3) *A, int N) {
    auto [Base, Size, Offset] = __lookup_as3(A);
    for (int I = 0; I < N; ++I) {
        __check_as3(Offset + I, Size);
        (Base + Offset + I)[I]++;
    }
}
```



```
int* __[cuda]malloc(int N) {  
    auto DevPtr = [cuda]malloc(N);  
    return __register_heap(DevPtr, N);  
}
```

```
double [ __shared__ ] A[1000];

void increment(int N) {
    auto [Base, Size, Offset] = __lookup_asX(A);
    for (int I = 0; I < N; ++I) {
        __check_asX(Offset + I, Size);
        (Base + Offset + I)[I]++;
    }
}
```

```
double [__shared__] A[1000];
double [__shared__] *A.fake;
void increment(int N) {
    auto [Base, Size, Offset] = __lookup_asX(*A.fake);
    for (int I = 0; I < N; ++I) {
        __check_asX(Offset + I, Size);
        (Base + Offset + I)[I]++;
    }
}
```

```
double [__shared__] A[1000];  
double [__shared__] *A.fake;
```

```
void __init_program() {           // called during startup  
    A.fake = __register_global(A, 8000);  
}
```

```
void __init_kernel() {           // called from kernel  
    A.fake = __register_shared(A, 8000);  
}
```

# Supporting Characters

```
void increment(double *A, int N) {
    for (int I = 0; I < N; ++I)
        A[I]++;
}
```

avoid duplicate checks

```
void increment(double *A, int N) {
    auto [Base, Size, Offset] = __lookup(A);
    for (int I = 0; I < N; ++I) {
        __check(Offset + I, Size);
        (Base + Offset + I)[I]++;
    }
}
```

```
void shift(double *A, int N) {  
    for (int I = 0; I < N - 2; ++I)  
        A[I] = A[I+1] + A[I+2];  
}
```

avoid “middle” checks

```
void shift(double *A, int N) {  
    auto [Base, Size, Offset] = __lookup(A);  
    for (int I = 0; I < N - 2; ++I) {  
        __check(Offset + 0, Size);  
        __check(Offset + 2, Size);  
        (... + I)[I] = (... + I)[I+1] + (... + I)[I+2];  
    }  
}
```

```
double [__shared__] A[1000];
```

```
double middle() {  
    return A[500];  
}
```

avoid checks for known good accesses

```
double [__shared__] A[1000];
```

```
double middle() {  
    return A[500];  
}
```



```
OFFLOAD ERROR: execution encountered an out-of-bounds access  
ACCESS of size 8 at 0x3 by thread <0,0,0> block <0,0,0>
```

**OFFLOAD ERROR: execution encountered an out-of-bounds access**

**ACCESS of size 8 at 0x3 by thread <0,0,0> block <0,0,0>**

# 0 omp target in main @ 21 ( \_\_omp\_offloading\_16\_8d8c61c\_main\_l21\_debug\_\_ ) offload/tes

# 1 omp target in main @ 21 ( \_\_omp\_offloading\_16\_8d8c61c\_main\_l21 ) offload/test/saniti





# Synopsis

## Reported Errors

allocation too large	✓	4TB for heap or 8KB for stack and shared
bad pointer	✓	magic mismatch, e.g., due to large out-of-bounds
out-of-bounds	✓	offset negative or larger than size
use-after-free	✓	size was negative due to deallocation (heap only)
wrong address space	✓	address space bits did not match compile time value

no false positives/negatives for out-of-bounds

low memory

avoid memory  
overhead

memory bound

avoid memory  
accesses, especially  
in loops

different memories

fast/cache memory  
⇒ fast verification

stack memory (per thread)  
⇒ no memory overhead

“slow”/large heap  
⇒ fixed overhead per allocation

no false positives/negatives for out-of-bounds