



A Deterministic Memory Allocator for Dynamic Symbolic Execution

Daniel Schemmel, Julian Büning, Frank Busse, Martin Nowack, Cristian Cadar

Dynamic Symbolic Execution

```
int x = input();  
if (x == 0) {  
    abort();  
} else {  
    return x;  
}
```


Dynamic Symbolic Execution

```
int x = input();  
if (x == 0) {  
    abort();  
} else {  
    return x;  
}
```

state 1
({x := λ}, {})

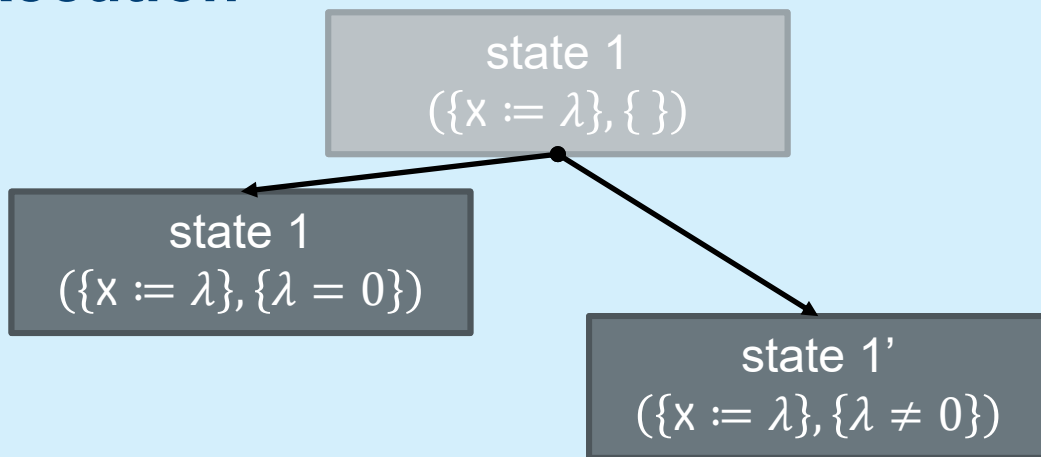
Dynamic Symbolic Execution

```
int x = input();  
if (x == 0) {  
    abort();  
} else {  
    return x;  
}
```



Dynamic Symbolic Execution

```
int x = input();  
if (x == 0) {  
    abort();  
} else {  
    return x;  
}
```



Address Translation in KLEE

- What is the address of x ?

Address Translation in KLEE

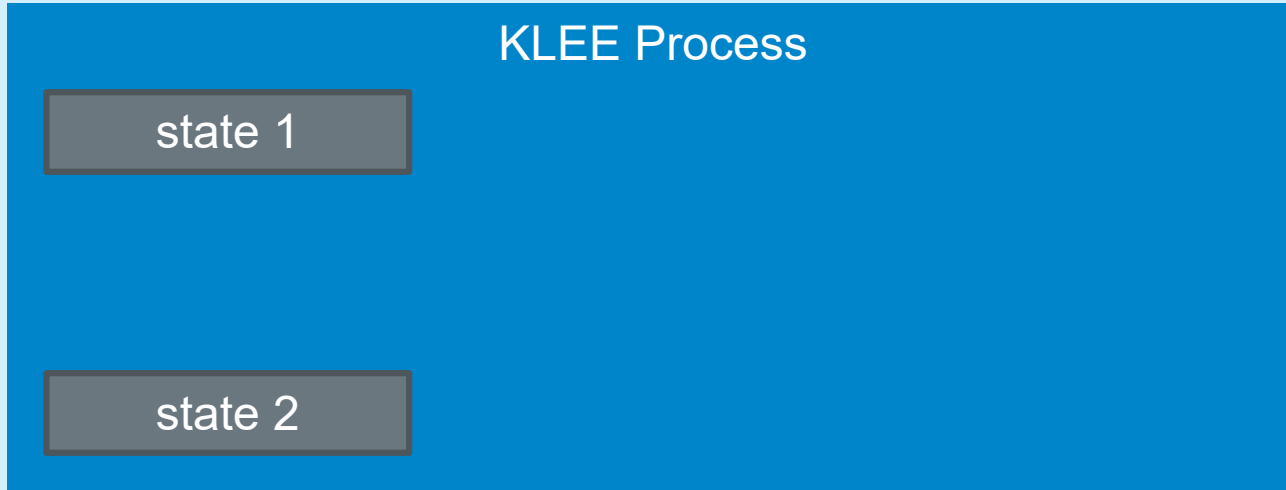
- What is the address of x ?



KLEE Process

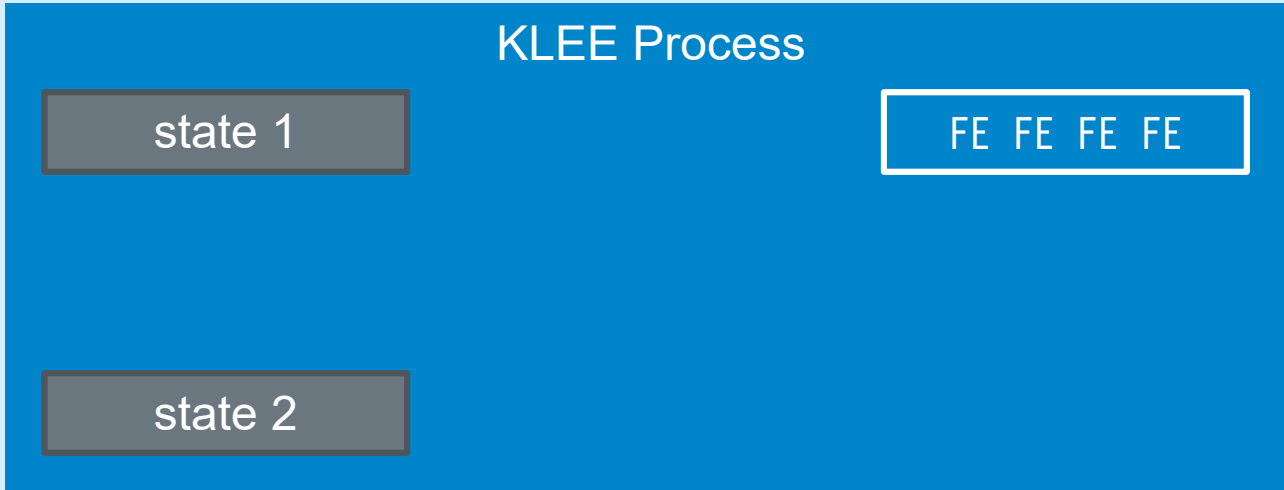
Address Translation in KLEE

- What is the address of x ?



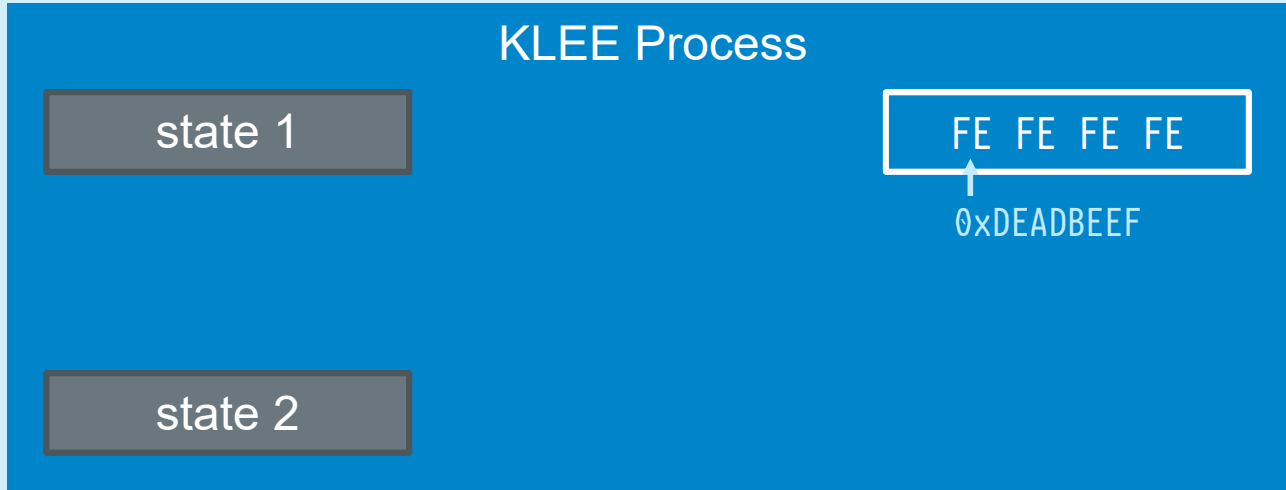
Address Translation in KLEE

- What is the address of x ?



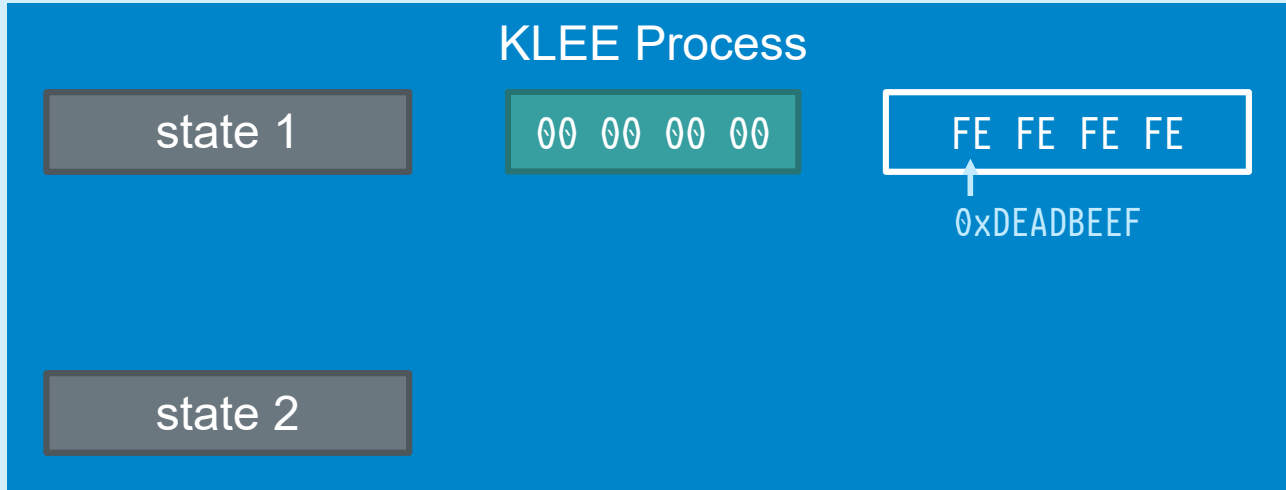
Address Translation in KLEE

- What is the address of `x`?



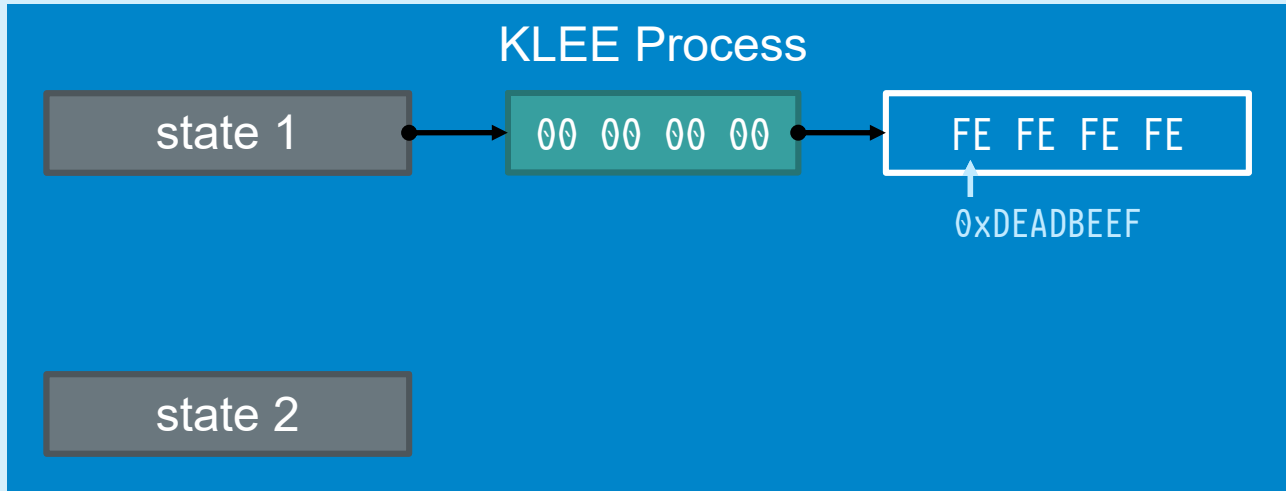
Address Translation in KLEE

- What is the address of x ?



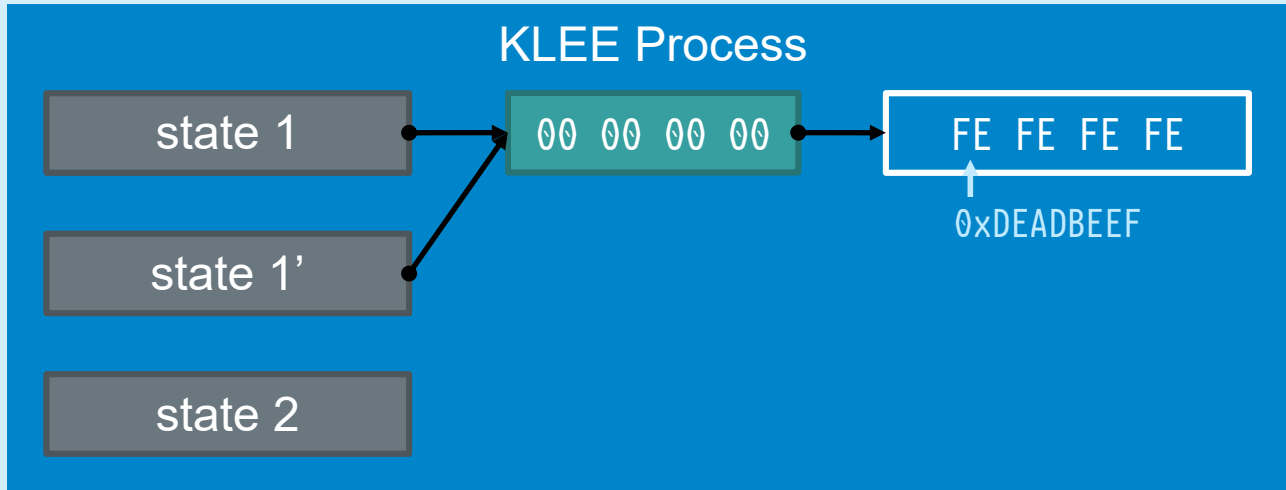
Address Translation in KLEE

- What is the address of x ?



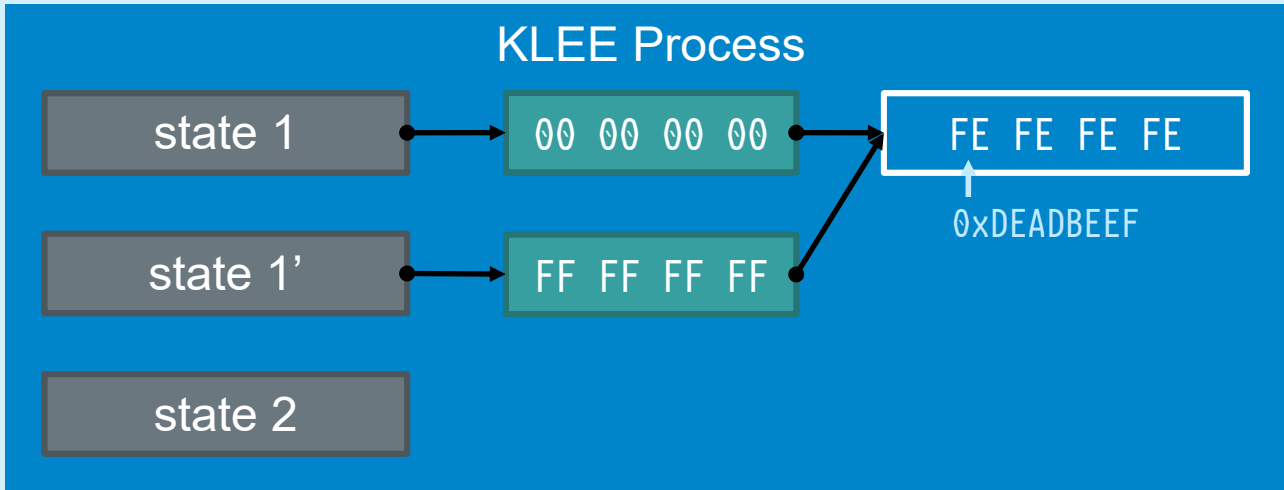
Address Translation in KLEE

- What is the address of x ?



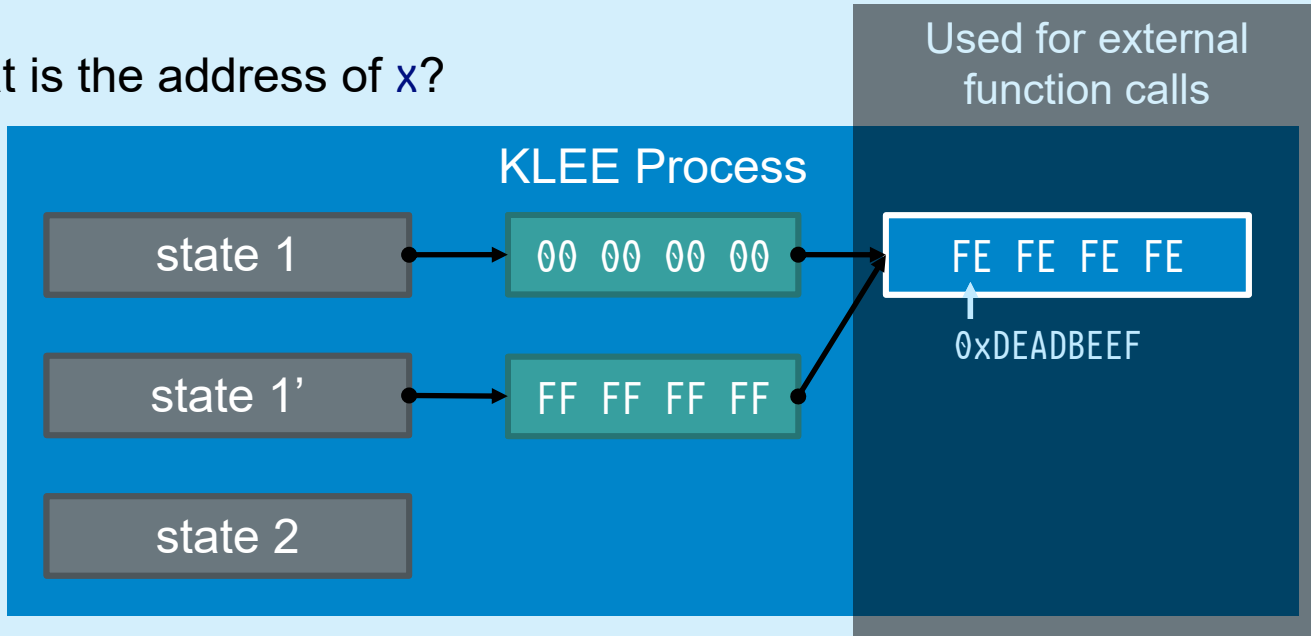
Address Translation in KLEE

- What is the address of x ?



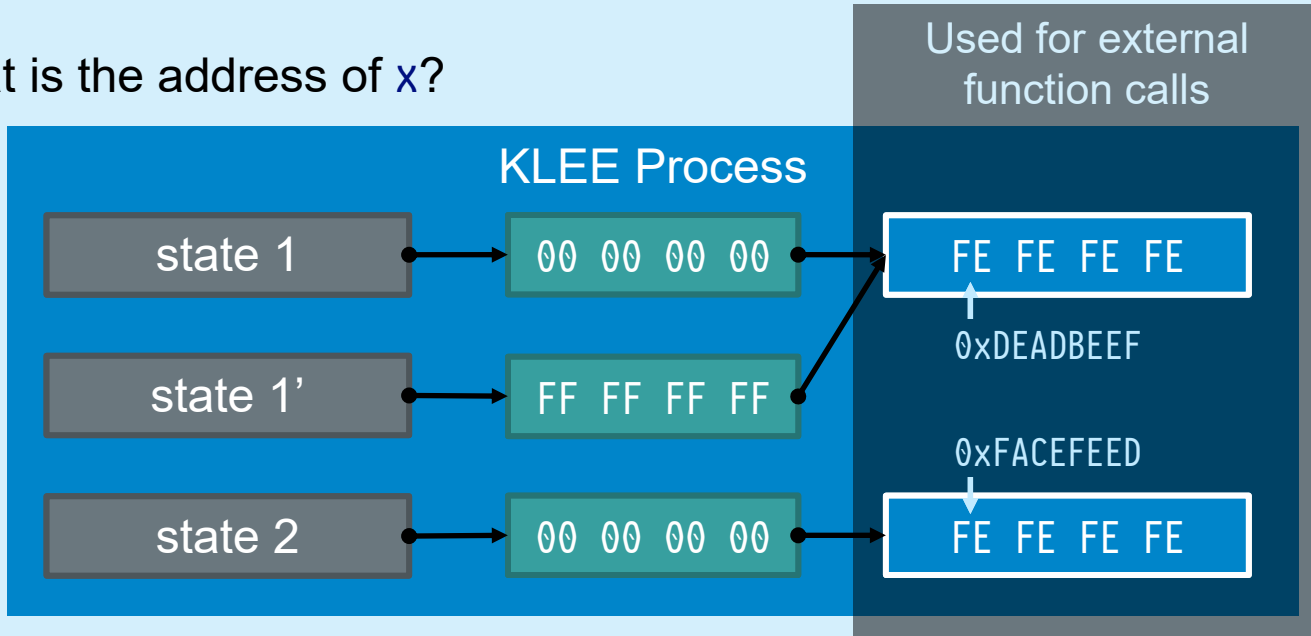
Address Translation in KLEE

- What is the address of x ?



Address Translation in KLEE

- What is the address of x ?



The Need for Deterministic Memory Allocation

- For experiments to be repeatable, memory allocation must be repeatable

The Need for Deterministic Memory Allocation

- For experiments to be repeatable, memory allocation must be repeatable
- Advanced symbolic execution techniques benefit from or outright require deterministic execution

The Need for Deterministic Memory Allocation

- For experiments to be repeatable, memory allocation must be repeatable
- Advanced symbolic execution techniques benefit from or outright require deterministic execution
 - POR-SE [Symbolic partial-order execution for testing multi-threaded programs. Schemmel et al. CAV 2020]

The Need for Deterministic Memory Allocation

- For experiments to be repeatable, memory allocation must be repeatable
- Advanced symbolic execution techniques benefit from or outright require deterministic execution
 - POR-SE [Symbolic partial-order execution for testing multi-threaded programs. Schemmel et al. CAV 2020]
 - SYMLIVE [Symbolic liveness analysis of real-world software. Schemmel et al. CAV 2018]

The Need for Deterministic Memory Allocation

- For experiments to be repeatable, memory allocation must be repeatable
- Advanced symbolic execution techniques benefit from or outright require deterministic execution
 - POR-SE [Symbolic partial-order execution for testing multi-threaded programs. Schemmel et al. CAV 2020]
 - SYMLIVE [Symbolic liveness analysis of real-world software. Schemmel et al. CAV 2018]
 - MoKLEE [Running symbolic execution forever. Busse et al. ISSTA 2020]

KDALLOC

- An allocator specifically for dynamic symbolic execution can do better!

KDALLOC

- An allocator specifically for dynamic symbolic execution can do better!
- Important properties:

KDALLOC

- An allocator specifically for dynamic symbolic execution can do better!
- Important properties:
 1. Support for external calls (addresses valid in host process)

KDALLOC

- An allocator specifically for dynamic symbolic execution can do better!
- Important properties:
 1. Support for external calls (addresses valid in host process)
 2. Cross-run determinism (multiple runs should behave the same)

KDALLOC

- An allocator specifically for dynamic symbolic execution can do better!
- Important properties:
 1. Support for external calls (addresses valid in host process)
 2. Cross-run determinism (multiple runs should behave the same)
 3. Cross-path determinism (multiple paths should behave the same)

KDALLOC

- An allocator specifically for dynamic symbolic execution can do better!
- Important properties:
 1. Support for external calls (addresses valid in host process)
 2. Cross-run determinism (multiple runs should behave the same)
 3. Cross-path determinism (multiple paths should behave the same)
 4. Spatially distanced allocations (misindexing an array should trap)

KDALLOC

- An allocator specifically for dynamic symbolic execution can do better!
- Important properties:
 1. Support for external calls (addresses valid in host process)
 2. Cross-run determinism (multiple runs should behave the same)
 3. Cross-path determinism (multiple paths should behave the same)
 4. Spatially distanced allocations (misindexing an array should trap)
 5. Temporally distanced allocations (use-after-free should trap)

KDALLOC

- An allocator specifically for dynamic symbolic execution can do better!
- Important properties:
 1. Support for external calls (addresses valid in host process)
 2. Cross-run determinism (multiple runs should behave the same)
 3. Cross-path determinism (multiple paths should behave the same)
 4. Spatially distanced allocations (misindexing an array should trap)
 5. Temporally distanced allocations (use-after-free should trap)
 6. Stability (minor changes should not snowball)

General Architecture

- `mmap` one large region and attach forkable metadata to the initial state

General Architecture

- `mmap` one large region and attach forkable metadata to the initial state
 - This region is only used to provide addresses and for external calls

General Architecture

- `mmap` one large region and attach forkable metadata to the initial state
 - This region is only used to provide addresses and for external calls
 - Object data is already state-dependent

General Architecture

- `mmap` one large region and attach forkable metadata to the initial state
 - This region is only used to provide addresses and for external calls
 - Object data is already state-dependent
- Categorize allocations to reduce snowball effect


General Architecture

- `mmap` one large region and attach forkable metadata to the initial state
 - This region is only used to provide addresses and for external calls
 - Object data is already state-dependent
- Categorize allocations to reduce snowball effect
 - Multiple allocators, especially to disconnect stack and heap

General Architecture

- `mmap` one large region and attach forkable metadata to the initial state
 - This region is only used to provide addresses and for external calls
 - Object data is already state-dependent
- Categorize allocations to reduce snowball effect
 - Multiple allocators, especially to disconnect stack and heap
 - Binned allocations


Memory Layout for KDALLOC




Slot Allocator for Sized Bins: Spatially Distanced



Slot Allocator for Sized Bins: Spatially Distanced

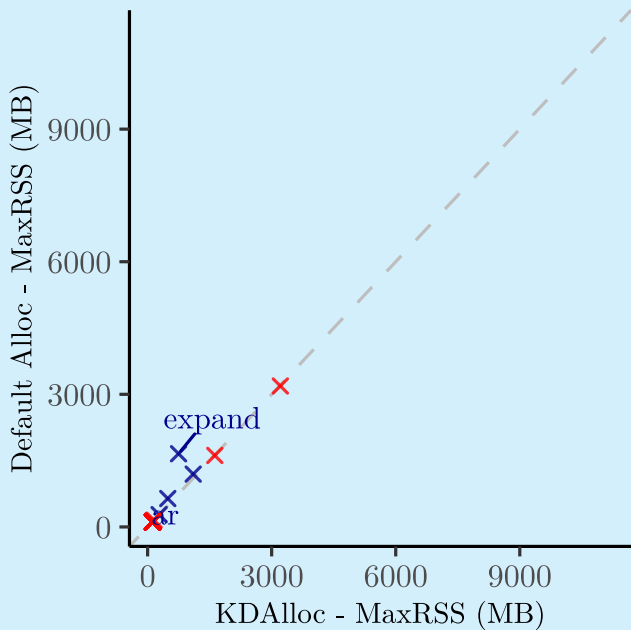


Slot Allocator for Sized Bins: Spatially Distanced

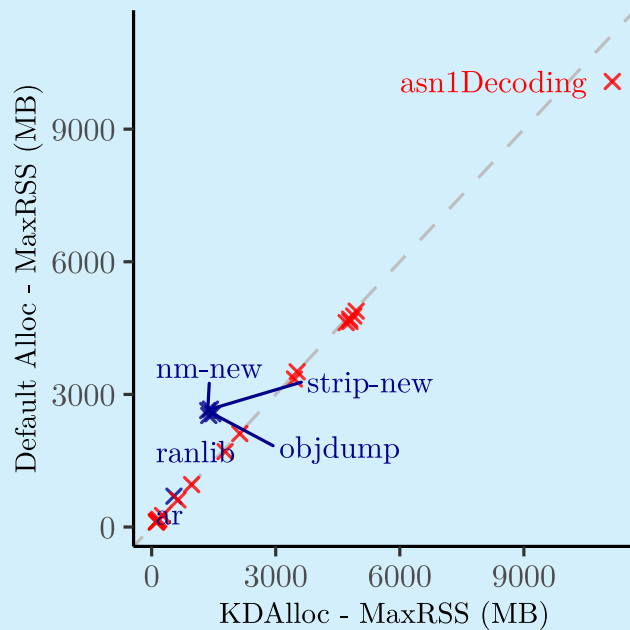


Memory Consumption

DFS

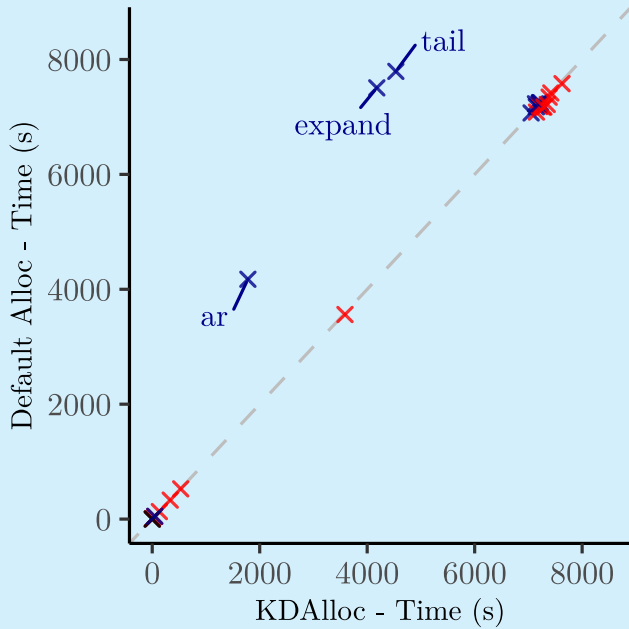


RNDCOV

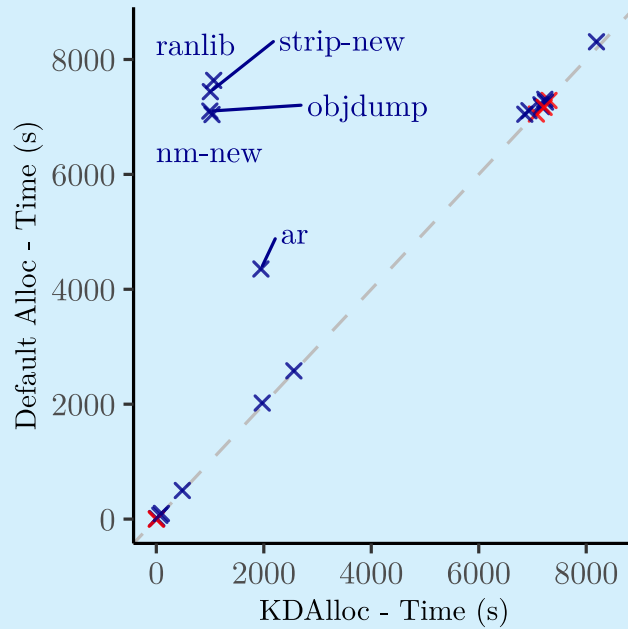


Performance

DFS

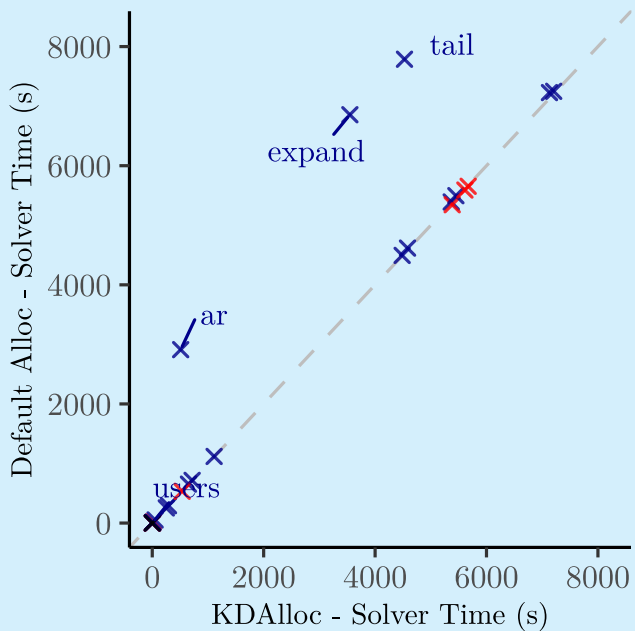


RNDCOV

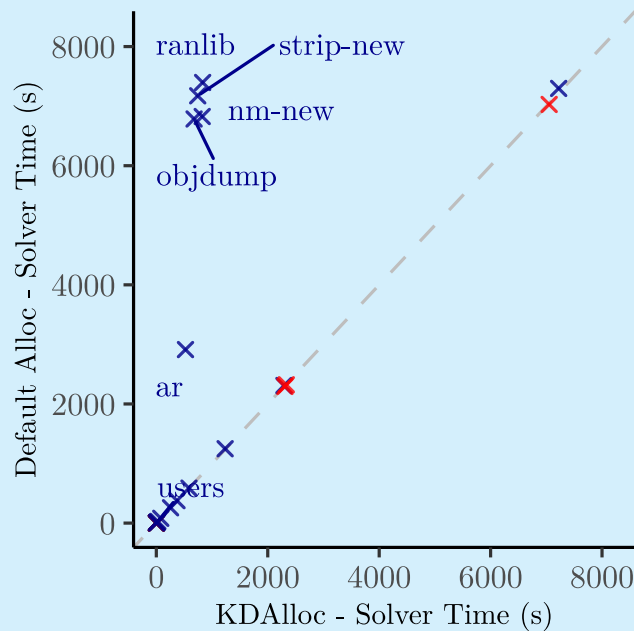


Solver Time

DFS



RNDCOV



MoKLEE: Fewer Diverging Locations

Suite	DFS		RNDCOV	
	MoKLEE	KDALLOC	MoKLEE	KDALLOC
Coreutils	22	12	42	32
Findutils	1	0	1	1
Libspng	0	0	1	0
Binutils	0	0	0	0
Diffutils	0	0	0	0
Grep	0	0	1	1
Tcpdump	0	0	0	0

MoKLEE: Fewer Divergences in memmove

```
char* s = (char*)dest;
const char* p = (const char*)src;
if (p >= s) {
    while (n) {
        *s++ = *p++;
        --n;
    }
} else {
    while (n) {
        --n;
        s[n] = p[n];
    }
}
return dest;
```

MoKLEE: Fewer Divergences in memmove

```
char* s = (char*)dest;
const char* p = (const char*)src;
if (p >= s) {
    while (n) {
        *s++ = *p++;
        --n;
    }
} else {
    while (n) {
        --n;
        s[n] = p[n];
    }
}
return dest;
```

MoKLEE: Fewer Divergences in memmove

```
char* s = (char*)dest;
const char* p = (const char*)src;
if (p >= s) {
    while (n) {
        *s++ = *p++;
        --n;
    }
} else {
    while (n) {
        --n;
        s[n] = p[n];
    }
}
return dest;
```

- uClibc's memmove is sensitive to memory layout

Summary & Conclusion

Summary & Conclusion

- The memory allocator has significant impact on dynamic symbolic execution

Summary & Conclusion

- The memory allocator has significant impact on dynamic symbolic execution
- We implemented KDALLOC in KLEE and show:

Summary & Conclusion

- The memory allocator has significant impact on dynamic symbolic execution
- We implemented KDALLOC in KLEE and show:
 - Performance and memory consumption are not impacted negatively

Summary & Conclusion

- The memory allocator has significant impact on dynamic symbolic execution
- We implemented KDALLOC in KLEE and show:
 - Performance and memory consumption are not impacted negatively
 - Use-after-free detection is improved (general benefit)

Summary & Conclusion

- The memory allocator has significant impact on dynamic symbolic execution
- We implemented KDALLOC in KLEE and show:
 - Performance and memory consumption are not impacted negatively
 - Use-after-free detection is improved (general benefit)
 - Specific benefits for multiple DSE-based techniques

Summary & Conclusion

- The memory allocator has significant impact on dynamic symbolic execution
- We implemented KDALLOC in KLEE and show:
 - Performance and memory consumption are not impacted negatively
 - Use-after-free detection is improved (general benefit)
 - Specific benefits for multiple DSE-based techniques
- KDALLOC is becoming part of mainline KLEE!

Guaranteed Use-After-Free Behavior

```
char *mallocfree() {  
    char *s = strdup("A");  
    free(s);  
    char *t = strdup("B");  
    return s;  
}  
  
int main(void) {  
    char *s = mallocfree();  
    puts(s);  
    return 0;  
}
```

- KDALLOC guarantees detection when quarantine is enabled

Query Structure with KDAAlloc

```
(Extract w32 0
  (Add w64 0xFFFFDDBC00000000 (Select w64 C 0x0000000000000000 0x0000224400000000)))
----- ↓ Extract(Add): (Extract (Add x y)) → (Add (Extract x) (Extract y)) -----
(Add w32 (Extract w32 0 0xFFFFDDBC00000000)
  (Extract w32 0 (Select w64 C 0x0000000000000000 0x0000224400000000)))
----- (Extract w32 0 0xFFFFDDBC00000000) → 0x00000000 and ↓ Extract(Select) -----
(Add w32 0x00000000
  (Select w32 C (Extract w32 0 0x0000000000000000) (Extract w32 0 0x0000224400000000)))
----- (Extract w32 0 0x0000000000000000) → 0x00000000 -----
----- (Extract w32 0 0x0000224400000000) → 0x00000000 -----
(Add w32 0x00000000 (Select w32 C 0x00000000 0x00000000))
----- (Select w32 C 0x00000000 0x00000000) → 0x00000000 -----
(Add w32 0x00000000 0x00000000) = 0x00000000
```

Query Structure without KDAIloc

```
(Extract w32 0  
  (Add w64 0xFFFFFAAAA7290C00 (Select w64 C 0x0000000000000000 0x0000555558D6F400)))
```

————— ↓ Extract(Add): (Extract (Add x y)) → (Add (Extract x) (Extract y)) —————

```
(Add w32 (Extract w32 0 0xFFFFFAAAA7290C00)  
  (Extract w32 0 (Select w64 C 0x0000000000000000 0x0000555558D6F400)))
```

————— (Extract w32 0 0xFFFFFAAAA7290C00) → 0xA7290C00 and ↓ Extract(Select) —————

```
(Add w32 0xA7290C00  
  (Select w32 C (Extract w32 0 0x0000000000000000) (Extract w32 0 0x0000555558D6F400)))
```

————— (Extract w32 0 0x0000000000000000) → 0x00000000 —————

————— (Extract w32 0 0x0000555558D6F400) → 0x58D6F400 —————

```
(Add w32 0xA7290C00 (Select w32 C 0x00000000 0x58D6F400))
```