

# KDALLOC: The KLEE Deterministic Allocator

## Deterministic Memory Allocation during Symbolic Execution and Test Case Replay

Daniel Schemmel<sup>\*</sup>, Julian Büning<sup>†</sup>, Frank Busse<sup>\*</sup>, Martin  
Nowack<sup>\*</sup>, Cristian Cadar<sup>\*</sup>

<sup>\*</sup>Imperial College London, <sup>†</sup>RWTH Aachen University

2023-07-18




# Summary & Conclusion

- Memory addresses can impact program execution



# Summary & Conclusion

- Memory addresses can impact program execution
- `KDALLOC`: Cross-run and cross-path deterministic allocation




# Summary & Conclusion

- Memory addresses can impact program execution
- `KDALLOC`: Cross-run and cross-path deterministic allocation
- Not presented *today*: Spatial and temporal distancing, memory and time efficiency, asan integration demo, ...



# Summary & Conclusion

- Memory addresses can impact program execution
- `KDALLOC`: Cross-run and cross-path deterministic allocation
- Not presented *today*: Spatial and temporal distancing, memory and time efficiency, asan integration demo, ...
- Visit me tomorrow for more examples and everything not covered today!



# A Treap with Concrete Inputs

```
1 #include <assert.h>
2 #include <stddef.h>
3 #include <stdint.h>
4 #include <stdlib.h>
5 #include <string.h>
6
7 struct node {
8     struct node *lhs, *rhs;
9     char const *key;
10 } * root = NULL;
11
12 static uint64_t priority(void *p) {
13     uint64_t h = 0xcbf29ce48422325;
14     for (size_t i = 0; i < sizeof(p); ++i)
15         h = (h ^ ((unsigned char *)&p)[i]) * 0x100000001b3;
16     return h;
17 }
18
19 void insert(struct node **n, char const *str) {
20     if (!*n) {
21         *n = calloc(1, sizeof(struct node));
22         (*n)->key = str;
23     } else {
24         int cmp = strcmp(str, (*n)->key);
25         if (cmp < 0) {
26             insert(&(*n)->lhs, str);
27         } else if (priority((*n)->lhs) < priority(*n)) {
28             struct node *lhs = (*n)->lhs;
29             (*n)->lhs = lhs->rhs;
30             lhs->rhs = (*n);
31             *n = lhs;
32         } else if (cmp > 0) {
33             insert(&(*n)->rhs, str);
34         } else if (priority((*n)->rhs) < priority(*n)) {
35             struct node *rhs = (*n)->rhs;
36             (*n)->rhs = rhs->lhs;
37             rhs->lhs = (*n);
38             *n = rhs;
39         }
40     }
41 }
42
43 int main() {
44     insert(&root, strdup("1"));
45     insert(&root, strdup("2"));
46     insert(&root, strdup("3"));
47     assert(strcmp(root->key, "1") == 0);
48 }
```

# Running Concretely

```
$ clang treap-conc.c -g3 -o treap-conc.exe  
$ ./treap-conc.exe  
$ echo $?  
0
```

# Running Concretely

```
$ clang treap-conc.c -g3 -o treap-conc.exe
$ ./treap-conc.exe
$ echo $?
0
$ ./treap-conc.exe
treap-conc.exe: treap-conc.c:49: int main(): Assertion
`strcmp(root->key, "1") == 0' failed.
(core dumped) ./treap.exe
$ echo $?
134
```



# Running in KLEE

```
$ clang treap-conc.c -g3 -c -emit-llvm
$ klee --posix-runtime --libc=uclibc treap-conc.bc
[...]
```

KLEE: WARNING ONCE: Alignment of memory from call "calloc" is not modelled.  
Using alignment of 8.

KLEE: done: total instructions = 14801  
KLEE: done: completed paths = 1  
KLEE: done: partially completed paths = 0  
KLEE: done: generated tests = 1

# Running in KLEE

```
$ clang treap-conc.c -g3 -c -emit-llvm
$ klee --posix-runtime --libc=uclibc treap-conc.bc
[...]
KLEE: WARNING ONCE: Alignment of memory from call "calloc" is not modelled.
Using alignment of 8.

KLEE: done: total instructions = 14801
KLEE: done: completed paths = 1
KLEE: done: partially completed paths = 0
KLEE: done: generated tests = 1

$ klee --posix-runtime --libc=uclibc treap-conc.bc
[...]
KLEE: WARNING ONCE: Alignment of memory from call "calloc" is not modelled.
Using alignment of 8.
KLEE: ERROR: treap-conc.c:49: ASSERTION FAIL: strcmp(root->key, "1") == 0
KLEE: NOTE: now ignoring this error at this location

KLEE: done: total instructions = 14340
KLEE: done: completed paths = 0
KLEE: done: partially completed paths = 1
KLEE: done: generated tests = 1
```

# Running in KLEE with KDAIloc 1/2

```
$ klee --posix-runtime --libc=uclibc --kdalloc treap-conc.bc
[...]
KLEE: Deterministic allocator: Using quarantine queue size 8
KLEE: Deterministic allocator: globals (start-address=0x7f992f600000 size=10 GiB)
KLEE: Deterministic allocator: constants (start-address=0x7f96af600000 size=10 GiB)
KLEE: Deterministic allocator: heap (start-address=0x7e96af600000 size=1024 GiB)
KLEE: Deterministic allocator: stack (start-address=0x7e76af600000 size=128 GiB)
KLEE: WARNING ONCE: calling external: syscall(16, 0, 21505, 139093310701568) at
/klee-src/runtime/POSIX/fd.c:997 10
KLEE: WARNING ONCE: Alignment of memory from call "malloc" is not modelled. Using
alignment of 8.
KLEE: WARNING ONCE: calling __klee_posix_wrapped_main with extra arguments.
KLEE: WARNING ONCE: Alignment of memory from call "calloc" is not modelled. Using
alignment of 8.
KLEE: ERROR: treap-conc.c:49: ASSERTION FAIL: strcmp(root->key, "1") == 0
KLEE: NOTE: now ignoring this error at this location

KLEE: done: total instructions = 14393
KLEE: done: completed paths = 0
KLEE: done: partially completed paths = 1
KLEE: done: generated tests = 1
```

## Running in KLEE with KDAIloc 2/2

```
$ klee --posix-runtime --libc=uclibc --kdalloc treap-conc.bc
[...]  
KLEE: Deterministic allocator: Using quarantine queue size 8  
KLEE: Deterministic allocator: globals (start-address=0x7fcc56e00000 size=10 GiB)  
KLEE: Deterministic allocator: constants (start-address=0x7fc9d6e00000 size=10 GiB)  
KLEE: Deterministic allocator: heap (start-address=0x7ec9d6e00000 size=1024 GiB)  
KLEE: Deterministic allocator: stack (start-address=0x7ea9d6e00000 size=128 GiB)  
KLEE: WARNING ONCE: calling external: syscall(16, 0, 21505, 139313016733696) at  
/klee-src/runtime/POSIX/fd.c:997 10  
KLEE: WARNING ONCE: Alignment of memory from call "malloc" is not modelled. Using  
alignment of 8.  
KLEE: WARNING ONCE: calling __klee_posix_wrapped_main with extra arguments.  
KLEE: WARNING ONCE: Alignment of memory from call "calloc" is not modelled. Using  
alignment of 8.  
  
KLEE: done: total instructions = 14854  
KLEE: done: completed paths = 1  
KLEE: done: partially completed paths = 0  
KLEE: done: generated tests = 1
```

# Running in KLEE with KDAIloc and Fixed Base-Addresses

```
$ klee --posix-runtime --libc=uclibc --kdalloc \
--kdalloc-constants-start-address=0x610000000000 \
--kdalloc-globals-start-address=0x620000000000 \
--kdalloc-heap-start-address=0x640000000000 \
--kdalloc-stack-start-address=0x630000000000 \
treap-conc.bc
[...]
KLEE: Deterministic allocator: Using quarantine queue size 8
KLEE: Deterministic allocator: globals (start-address=0x620000000000 size=10 GiB)
KLEE: Deterministic allocator: constants (start-address=0x610000000000 size=10 GiB)
KLEE: Deterministic allocator: heap (start-address=0x640000000000 size=1024 GiB)
KLEE: Deterministic allocator: stack (start-address=0x630000000000 size=128 GiB)
KLEE: WARNING ONCE: calling external: syscall(16, 0, 21505, 108896748306432) at /klee-src/runtime
KLEE: WARNING ONCE: Alignment of memory from call "malloc" is not modelled. Using alignment of 8
KLEE: WARNING ONCE: calling __klee_posix_wrapped_main with extra arguments.
KLEE: WARNING ONCE: Alignment of memory from call "calloc" is not modelled. Using alignment of 8
KLEE: ERROR: treap-conc.c:49: ASSERTION FAIL: strcmp(root->key, "1") == 0
KLEE: NOTE: now ignoring this error at this location

KLEE: done: total instructions = 14415
KLEE: done: completed paths = 0
KLEE: done: partially completed paths = 1
KLEE: done: generated tests = 1
```

# A Treap with Symbolic Inputs

```
1  #include <assert.h>
2  #include <stddef.h>
3  #include <stdint.h>
4  #include <stdlib.h>
5  #include <string.h>
6  #include <sys/types.h>
7  #include <sys/unistd.h>
8  #include <unistd.h>
9
10 struct node {
11     struct node *lhs, *rhs;
12     char const *key;
13 }* root = NULL;
14
15 static uint64_t priority(void *p) {
16     uint64_t h = 0xcbf29ce484222325;
17     for (size_t i = 0; i < sizeof(p); ++i)
18         h = (h ^ ((unsigned char *)p)[i]) * 0x100000001B3;
19     return h;
20 }
21
22 void insert(struct node **n, char const *str) {
23     if (!*n) {
24         *n = calloc(1, sizeof(struct node));
25         (*n)->key = str;
26     } else {
27         int cmp = strcmp(str, (*n)->key);
28         if (cmp < 0) {
29             insert(&(*n)->lhs, str);
30             if (priority((*n)->lhs) < priority(*n)) {
31                 struct node *lhs = (*n)->lhs;
32                 (*n)->lhs = lhs->rhs;
33                 lhs->rhs = (*n);
34                 *n = lhs;
35             }
36         } else if (cmp > 0) {
37             insert(&(*n)->rhs, str);
38             if (priority((*n)->rhs) < priority(*n)) {
39                 struct node *rhs = (*n)->rhs;
40                 (*n)->rhs = rhs->lhs;
41                 rhs->lhs = (*n);
42                 *n = rhs;
43             }
44         }
45     }
46 }
47
48 int main() {
49     char sym2[2], sym3[2];
50     sym2[1] = '\0';
51     sym3[1] = '\0';
52     read(0, sym2, 1);
53     read(0, sym3, 1);
54
55     insert(&root, strdup("1"));
56     insert(&root, sym2);
57     insert(&root, sym3);
58
59     assert(strcmp(root->key, "1") == 0);
60 }
```

# Replaying Symbolic Execution with KDAIloc

```
$ clang treap-sym.c -g3 -c -emit-llvm
$ klee --posix-runtime --libc=uclibc --emit-all-errors \
  treap-sym.bc --sym-stdin 2
[...]
$ clang treap-sym.c -g3 -o treap-sym.exe
$ klee-replay klee-last/test000002.ktest ./treap-sym.exe
[...]
KLEE-REPLAY: NOTE: EXIT STATUS: NORMAL (0 seconds)
[...]
```

# Replaying Symbolic Execution with KDAIloc

```
$ clang treap-sym.c -g3 -c -emit-llvm
$ klee --posix-runtime --libc=uclibc --emit-all-errors \
  treap-sym.bc --sym-stdin 2
[...]
$ clang treap-sym.c -g3 -o treap-sym.exe
$ klee-replay klee-last/test000002.ktest ./treap-sym.exe
[...]
KLEE-REPLAY: NOTE: EXIT STATUS: NORMAL (0 seconds)
[...]
$ klee-replay klee-last/test000002.ktest ./treap-sym.exe
[...]
treap-sym.exe: treap-sym.c:59: int main(): Assertion `strcmp(root->key, "1") == 0' failed.
KLEE-REPLAY: NOTE: EXIT STATUS: CRASHED signal 6 (0 seconds)
[...]
```



# Replaying Symbolic Execution with KDAIloc

```
$ klee --posix-runtime --libc=uclibc --emit-all-errors \  
--kdalloc \  
--kdalloc-constants-start-address=0x610000000000 \  
--kdalloc-globals-start-address=0x620000000000 \  
--kdalloc-heap-start-address=0x640000000000 \  
--kdalloc-stack-start-address=0x630000000000 \  
treap-sym.bc --sym-stdin 2 \  
[...]  
$ klee-replay klee-last/test000002.ktest ./treap-sym.exe  
[...]  
KDAIloc initialized at 0x640000000000 with 1024GiB and quarantine 8  
treap-sym.exe: treap-sym.c:59: int main(): Assertion `strcmp(root->key, "1") == 0' failed.  
KLEE-REPLAY: NOTE: EXIT STATUS: CRASHED signal 6 (0 seconds)  
[...]
```

# Summary & Conclusion

- Memory addresses can impact program execution



# Summary & Conclusion

- Memory addresses can impact program execution
- `KDALLOC`: Cross-run and cross-path deterministic allocation



# Summary & Conclusion

- Memory addresses can impact program execution
- `KDALLOC`: Cross-run and cross-path deterministic allocation
- Not presented *today*: Spatial and temporal distancing, memory and time efficiency, asan integration demo, ...



# Summary & Conclusion

- Memory addresses can impact program execution
- `KDALLOC`: Cross-run and cross-path deterministic allocation
- Not presented *today*: Spatial and temporal distancing, memory and time efficiency, asan integration demo, ...
- Visit me tomorrow for more examples and everything not covered today!

