

Automated testing of atomic instructions (lr/sc) implementations in selfie

Luis Thiele

Table of Contents

- Revisiting `selfie`
 - The `selfie` System
 - Bump-Pointer Allocator
 - Processes vs. Threads
 - Load Reserved & Store Conditional
- Treiber-Stack Assignment
 - Treiber-Stack
 - Thread-Safe `malloc`
 - Old Assignment vs. New Assignment
 - Automated Tests
- Changes to `selfie`

Revisiting selfie

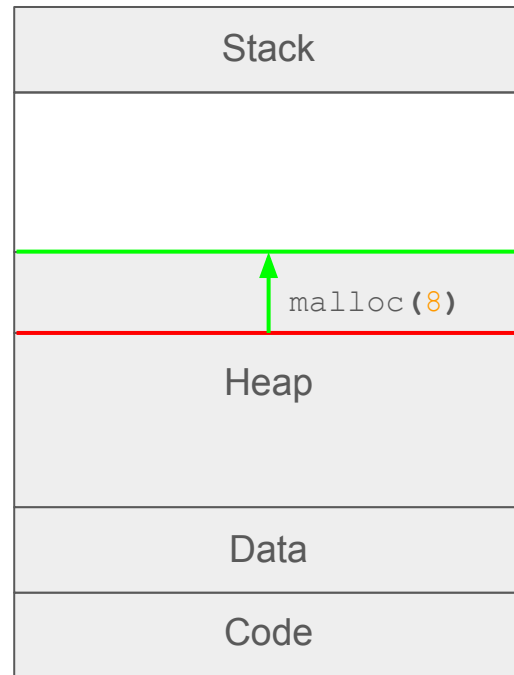
The selfie System



github.com/cksystemsteaching/selfie

Bump-Pointer Allocator

- Variable stores last allocated address
 - Addresses below value are in use
 - Addresses above value are free
- On `malloc(8)` Call:
 - Load value of bump-pointer
 - Add 8 to value
 - Syscall `BRK` to increase program break
 - Store new value of bump-pointer
 - Return new value



Processes vs. Threads

Processes:

- Independent Memory

System Calls:

- `fork()`
- `wait(uint64_t* wstatus)`
- `exit(uint64_t exitcode)`

Threads:

- Shared Memory (except Stack)

System Calls:

- `pthread_create()`
- `pthread_join(uint64_t* wstatus)`
- `pthread_exit(uint64_t exitcode)`

Load Reserved & Store Conditional (1)

- “Extended” Load & Store Instructions
- Load Reserved:
 - Load
 - Reservation on Address
- Store Conditional:
 - Condition: Reservation on Address
 - True: Store
 - False: No Store, Mark Unsuccessful

Load Reserved & Store Conditional (2)

Load:

- `ld rd,imm(rs1)`
 - `rd = memory[rs1 + imm]`

Load Reserved:

- `lr.d rd,(rs1)`
 - `rd = memory[rs1]`

Store:

- `sd rs2,imm(rs1)`
 - `memory[rs1 + imm] = rs2`

Store Conditional:

- `sc.d rd,rs2,(rs1)`
 - SUCCESS: `memory[rs1] = rs2`
 - SUCCESS: `rd = 0`
 - FAILURE: `rd = 1`

Load Reserved & Store Conditional (3)

“LR-SC Loop”:

```
1.  do {  
2.    value = lr(address);  
3.    // edit value here  
4.  } while (sc(address, value));  
5.  // sc returns 1 if unsuccessful
```

- Shared Memory Affected
 - Threads!
- Make Code Thread-Safe

- Threads:
 - Thread A
 - Thread B
- Execution:
 - lr
 - sc
 - lr
 - sc
- Interleaved Execution:
 - lr
 - lr
 - sc
 - sc
 - lr
 - sc

Treiber-Stack Assignment

Treiber-Stack

- Thread-Shared Stack
- Machine Instructions Only
 - No System Calls!
- Heap Memory
 - Uses `malloc`
 - Thus uses system call `BRK`
- Macros:
 - `void init_stack ()`
 - `void push (uint64_t value)`
 - `uint64_t pop ()`

Thread-Safe `malloc` (1)

- “Old” `malloc(8)` Call:
 - `LD` value of bump-pointer
 - Add 8 to value
 - Syscall `BRK` to increase program break
 - `SD` new value of bump-pointer
 - Return new value
- Issues:
 - Not thread-safe
 - Syscalls force context switches!
- “New” `malloc(8)` Call:
 - `LR` value of bump-pointer
 - Add 8 to value
 - Syscall `BRK` to increase program break
 - `SC` new value of bump-pointer
 - SUCCESS: Return new value
 - FAILURE: Jump back to `LR`
- Short version:
 - `LD / LR`
 - `BRK`
 - `SD / SC`

Thread-Safe malloc (2)

Code:

```
1. pthread_create();  
2. malloc(8);
```

- Old malloc:

- LD
- BRK
- LD
- BRK
- SD
- SD

- New malloc:

- LR
- BRK
- LR
- BRK
- SC
- LR
- BRK
- SC
- LR
- BRK
- ...

Old Assignment vs. New Assignment

Old Assignment:

- `treiber-stack`
 - Implement `lr` & `sc`
 - Implement `treiber-stack`

New Assignment(s):

- `threadsafe-malloc`
 - Implement `lr` & `sc`
 - Make `malloc` thread-safe
 - No context switches on `malloc`
- `treiber-stack`
 - Implement `treiber-stack`
 - Make `treiber-stack` thread-safe

Automated Tests (1)

No-Context-Switch malloc Test:

- Easy solution:
 - Make sure Thread A runs first
 - Thread A calls `malloc`
 - Thread A prints eg. "Hello"
 - Thread B prints eg. "World"
- Success:
 - `malloc` did not force a switch
 - "Hello World"
- Failure:
 - "World Hello"

```
1.  pid = pthread_create ();
2.  if (pid == 0) {
3.      // child
4.      child = 1;
5.      malloc(8);
6.      write(1, "Hello  ", 8);
7.  } else {
8.      // parent
9.      while (child == 0)
10.         wait((uint64_t*) 0);
11.      write(1, "World  ", 8);
12.  }
```

Automated Tests (2)

LR & SC Semantics Test:

- Requirements:
 - LR coroutine (returns value)
 - SC coroutine (returns 1 on FAILURE)
- Solution:
 - Interleaved execution
 - 2nd SC must mark failure
 - 1st SC decides final value

```
1.  uint64_t lr(uint64_t address);
2.  uint64_t sc(uint64_t address,
3.              uint64_t value);
```

```
1.  address = malloc(8);
2.  *address = 7;
3.  lr(address);
4.  // force switch
5.  lr(address);
6.  if (sc(address, 42))
7.      return 7;
8.  pthread_wait(status);
9.  // switch
10. c = sc(address, 7);
11. pthread_exit(c);
12. // switch
13. return *status * *address;
```


Automated Tests (3)

Thread-Safe `malloc` Test:

- Force context switch between **LR** & **SC**
 - Context switch by timeout
- Idea:
 - Repeat useless loop
 - `malloc` just before switch by timeout
- Solution:
 - Thread A measures endless loop
 - Thread B stops endless loop
 - proceed as explained...
 - Child may force switch by `pthread_wait`
- Success:
 - Different addresses by `malloc`

```
1.  zero = 0;
2.  loop = 1;
3.  while (zero < loop)
4.      counter = counter + 1;
5.  // switch
6.  loop = 0;
7.  // force switch
8.  i = 2;
9.  while (i < counter);
10.   i = i + 1;
11.  malloc(8);
12.  // switch
13.  malloc(8);
```

Automated Tests (4)

Thread-Safe Treiber-Stack Test:

- “
 - “
- “
 - “
 - push/pop just before switch by timeout
- “
 - “
 - “
 - “
- “
 - {push'd} = {pop'd}
(overwrites, detached head...)

```
1.  zero = 0;
2.  loop = 1;
3.  while (zero < loop)
4.      counter = counter + 1;
5.  // switch
6.  loop = 0;
7.  // force switch
8.  i = 2;
9.  while (i < counter);
10.   i = i + 1;
11.  push(8);
12.  // switch
13.  push(8);
```

Changes to selfie

Changes to selfie

- Improve github actions
 - No more running out of quota
 - Private repo: Only run on `main` branch and only linux
 - Online dispatcher
- New assignments
 - `logical-and-or-not` (boolean)
 - `lazy-evaluation`
 - `(threadsafe-malloc, treiber-stack)`
- Restructure code
 - Array & Struct assignments a lot easier
 - Grammar also restructured

