

A Theory of C-Style Memory Allocation

Stephan Falke, Florian Merz, and Carsten Sinz

Research Group “Verification meets Algorithm Engineering”

```
struct list_node {
    int data;
    struct list_node *tail;
};
typedef struct list_node list;

list *reverse(list *l) {
    list *r = l, *p = NULL;
    while (r != NULL) {
        list *q = r;
        r = r->tail;
        q->tail = p;
        p = q;
    }
    return p;
}
```



Introduction

- Modeling the **memory content** is important in **program analysis**
 - **Formalized** using McCarthy's theory of arrays

Introduction

- Modeling the **memory content** is important in **program analysis**
 - **Formalized** using McCarthy's theory of arrays
- Equally important: **memory protection**
 - Often only handled in an **ad-hoc** and **informal** way

- Modeling the **memory content** is important in **program analysis**
 - **Formalized** using McCarthy's theory of arrays
- Equally important: **memory protection**
 - Often only handled in an **ad-hoc** and **informal** way
- **Solution:**

An **SMT theory** for reasoning about **memory access correctness**

- **Precise** formalization of `malloc` and `free`
- Opportunities for **simplification** of formulas using rewrite rules
- Modular and local reasoning

- LLBMC = Low-Level (Software) Bounded Model Checking
 - **Low-Level**: Embedded devices, systems code, . . .
 - **Software**: Programs written in C/C++
 - **Bounded**: restricted number of nested function calls and loop iterations
 - **Model Checking**: bit-precise static analysis

- LLBMC = Low-Level (Software) Bounded Model Checking
 - **Low-Level**: Embedded devices, systems code, ...
 - **Software**: Programs written in C/C++
 - **Bounded**: restricted number of nested function calls and loop iterations
 - **Model Checking**: bit-precise static analysis
- Properties checked:
 - **Built-in properties**: invalid memory access, use-after-free, double free, arithmetic overflow, division by zero, ...
 - **User-supplied properties**: `assert` statements

Software Bounded Model Checking

- Property checking of programs is **undecidable** in general

Software Bounded Model Checking

- Property checking of programs is **undecidable** in general
- Bugs manifest themselves in **finite runs** of the program

Software Bounded Model Checking

- Property checking of programs is **undecidable** in general
- Bugs manifest themselves in **finite runs** of the program
- Software bounded model checking:
 - Analyze only **bounded** program runs
 - Restrict number of nested **function calls** and inline functions
 - Restrict number of **loop iterations** and unroll loops
 - Property checking becomes **decidable** using an SMT solver

The LLBMC Approach

- Fully supporting real-life programming languages is **cumbersome**

The LLBMC Approach


- Fully supporting real-life programming languages is **cumbersome**
- Particularly true for **C/C++** due to their complex semantics

The LLBMC Approach

- Fully supporting real-life programming languages is **cumbersome**
- Particularly true for **C/C++** due to their complex semantics
- **Solution:** Do not operate on the source code, but use a **compiler IR**
 - Well-defined, **simple** semantics
 - **Closer** to the program that is actually run on the machine
 - Support for wide range of programming languages comes “**for free**”

The LLBMC Approach

- Fully supporting real-life programming languages is **cumbersome**
- Particularly true for **C/C++** due to their complex semantics
- **Solution:** Do not operate on the source code, but use a **compiler IR**
 - Well-defined, **simple** semantics
 - **Closer** to the program that is actually run on the machine
 - Support for wide range of programming languages comes “**for free**”



- Theory of bitvectors (\mathcal{T}_{BV}) for register operations (ALU):
 - Mostly straightforward, not discussed in this talk

- Theory of **bitvectors** (\mathcal{T}_{BV}) for **register** operations (ALU):
 - Mostly straightforward, not discussed in this talk
- Theory of **arrays** (\mathcal{T}_A) for **memory** operations:
 - Memory is one big array of bytes
 - Translate `load` and `store` instructions
 - Introduce explicit memory states: a, a', \dots

- Theory of **bitvectors** (\mathcal{T}_{BV}) for **register** operations (ALU):
 - Mostly straightforward, not discussed in this talk
- Theory of **arrays** (\mathcal{T}_A) for **memory** operations:
 - Memory is one big array of bytes
 - Translate load and store instructions
 - Introduce explicit memory states: a, a', \dots

<code>%x = load i8* %p</code>	\rightsquigarrow	$x = \text{read}(a, p)$
<code>store i8* %p, i8 %x</code>		$\hat{a} = \text{write}(a, p, x)$

- Theory of **bitvectors** (\mathcal{T}_{BV}) for **register** operations (ALU):
 - Mostly straightforward, not discussed in this talk
- Theory of **arrays** (\mathcal{T}_A) for **memory** operations:
 - Memory is one big array of bytes
 - Translate load and store instructions
 - Introduce explicit memory states: a, a', \dots

<code>store i8* %p, i8 %x</code>	\rightsquigarrow	$\hat{a} = \text{write}(a, p, x)$
<code>%x = load i8* %p</code>		$x = \text{read}(\hat{a}, p)$

Memory Related Program Bugs

- Kinds of memory related bugs:
 - load from a non-allocated region of memory
 - store to a non-allocated region of memory
 - free of a non-allocated region of memory
 - free of an already freed region of memory
 - ...

Memory Related Program Bugs

- Kinds of memory related bugs:
 - load from a non-allocated region of memory
 - store to a non-allocated region of memory
 - free of a non-allocated region of memory
 - free of an already freed region of memory
 - ...
- Not handled by \mathcal{T}_{ABV} , but should be detected by LLBMC

Memory Related Program Bugs

- Kinds of memory related bugs:
 - load from a non-allocated region of memory
 - store to a non-allocated region of memory
 - free of a non-allocated region of memory
 - free of an already freed region of memory
 - ...
- Not handled by \mathcal{T}_{ABV} , but should be detected by LLBMC
- Need to be handled atop of \mathcal{T}_{ABV}

- Current approach in LLBMC (SSV '10):
 - Introduce explicit heap states: h, h', \dots
 - Memory access correctness constraints explicitly encoded in \mathcal{T}_{ABV}
 - Creates **huge** subformulas for each memory access operation
 - Needs knowledge of the **complete** heap state history

- Current approach in LLBMC (SSV '10):
 - Introduce explicit heap states: h, h', \dots
 - Memory access correctness constraints explicitly encoded in \mathcal{T}_{ABV}
 - Creates **huge** subformulas for each memory access operation
 - Needs knowledge of the **complete** heap state history

$$\text{accessible}(h, q, t) \equiv \bigvee_{\substack{h' \preceq h \\ l: h' = \text{malloc}(\widehat{h}', p, s)}} c_{\text{exec}}(l) \wedge \neg \text{deallocated}(h', h, p) \wedge \text{contained}(p, s, q, t)$$

$$\text{deallocated}(h, h', p) \equiv \bigvee_{\substack{h \preceq h^* \preceq h' \\ l: h^* = \text{free}(\widehat{h}^*, q)}} c_{\text{exec}}(l) \wedge p = q$$

- Current approach in LLBMC (SSV '10):
 - Introduce explicit heap states: h, h', \dots
 - Memory access correctness constraints explicitly encoded in \mathcal{T}_{ABV}
 - Creates **huge** subformulas for each memory access operation
 - Needs knowledge of the **complete** heap state history
- Goals of this work:
 - **Precise** formalization of `malloc` and `free` as an **SMT theory**
 - Opportunities for **simplification** of formulas using rewrite rules
 - Modular and local reasoning: **partial** heap state history suffices

Motivation: McCarthy's Theory of Arrays

McCarthy's *read-over-write* axioms:

$$i = j \Rightarrow \text{read}(\text{write}(a, i, x), j) = x$$

$$i \neq j \Rightarrow \text{read}(\text{write}(a, i, x), j) = \text{read}(a, j)$$

Motivation: McCarthy's Theory of Arrays

McCarthy's *read-over-write* axioms:

$$i = j \Rightarrow \text{read}(\text{write}(a, i, x), j) = x$$

$$i \neq j \Rightarrow \text{read}(\text{write}(a, i, x), j) = \text{read}(a, j)$$

Reduction to equality logic:

1. Apply axioms in the form of the rewrite rule

$$\text{read}(\text{write}(a, i, x), j) \rightarrow \text{ITE}(i = j, x, \text{read}(a, j))$$

Motivation: McCarthy's Theory of Arrays

McCarthy's *read-over-write* axioms:

$$i = j \Rightarrow \text{read}(\text{write}(a, i, x), j) = x$$

$$i \neq j \Rightarrow \text{read}(\text{write}(a, i, x), j) = \text{read}(a, j)$$

Reduction to equality logic:

1. Apply axioms in the form of the rewrite rule

$$\text{read}(\text{write}(a, i, x), j) \rightarrow \text{ITE}(i = j, x, \text{read}(a, j))$$

2. Replace *reads* by fresh variables, add Ackermann constraints

$$i = j \Rightarrow v_{\text{read}(a,i)} = v_{\text{read}(a,j)}$$

Motivation: McCarthy's Theory of Arrays

McCarthy's *read-over-write* axioms:

$$i = j \Rightarrow \text{read}(\text{write}(a, i, x), j) = x$$

$$i \neq j \Rightarrow \text{read}(\text{write}(a, i, x), j) = \text{read}(a, j)$$

Reduction to equality logic:

1. Apply axioms in the form of the rewrite rule

$$\text{read}(\text{write}(a, i, x), j) \rightarrow \text{ITE}(i = j, x, \text{read}(a, j))$$

2. Replace *reads* by fresh variables, add Ackermann constraints

$$i = j \Rightarrow v_{\text{read}(a,i)} = v_{\text{read}(a,j)}$$

More *efficient* approaches exist, e.g., Boolector's *lemmas-on-demand*

Introducing \mathcal{T}_H

- Objects of type H encode information about the **heap state**
 - Contain the whole history of `malloc` and `free` operations
 - Do not encode information about the data stored in the heap
 - Are used to determine memory access correctness

Introducing \mathcal{T}_H

- Objects of type H encode information about the **heap state**
 - Contain the whole history of `malloc` and `free` operations
 - Do not encode information about the data stored in the heap
 - Are used to determine memory access correctness

Functions of type H :

- | | |
|---------------------------------------------|----------------------------------------------------------------|
| ■ ε | Heap without any allocation |
| ■ <code>malloc(h, p, s)</code> | Allocate region $[p, p + s)$ (if possible) |
| ■ <code>free(h, p)</code> | De-allocate region starting at p
(if currently allocated) |

Introducing \mathcal{T}_H

- Objects of type H encode information about the **heap state**
 - Contain the whole history of `malloc` and `free` operations
 - Do not encode information about the data stored in the heap
 - Are used to determine memory access correctness

Functions of type H :

- ε Heap without any allocation
- `malloc(h, p, s)` Allocate region $[p, p + s)$ (if possible)
- `free(h, p)` De-allocate region starting at p
(if currently allocated)

Auxilliary functions:

- `mallocsize(h, p)` Size of the memory region starting at p
(if currently allocated)

Predicates of \mathcal{T}_H

Predicates:

- $\text{accessible}(h, p, s)$ Is $[p, p + s)$ contained in an allocated region?
- $\text{freeable}(h, p)$ Is p the start of an allocated region?
- $\text{mallocable}(h, p, s)$ Can $[p, p + s)$ be allocated?

Predicates of \mathcal{T}_H

Predicates:

- $\text{accessible}(h, p, s)$ Is $[p, p + s)$ contained in an allocated region?
- $\text{freeable}(h, p)$ Is p the start of an allocated region?
- $\text{mallocable}(h, p, s)$ Can $[p, p + s)$ be allocated?

The semantics of the predicates is specified using axioms

Axioms for mallocable

`mallocable(h, p, s)` Can $[p, p + s)$ be allocated?

Axioms for mallocable

$\text{mallocable}(h, p, s)$ Can $[p, p + s)$ be allocated?

- $[p, p + s)$ can be allocated if it **does not overlap** with any allocated memory region

Axioms for mallocable

$\text{mallocable}(h, p, s)$ Can $[p, p + s)$ be allocated?

- $[p, p + s)$ can be allocated if it **does not overlap** with any allocated memory region
- Various possible axiomatizations:
 - **Stack-like**: p needs to be “on the right” of all regions that have been allocated before
 - Axiomatize non-overlap **precisely**
 - ...

Axioms for freeable

$\text{freeable}(h, q)$

Is q the start of an allocated region?

Axioms for freeable

$\text{freeable}(h, q)$

Is q the start of an allocated region?

$$\text{freeable}(\varepsilon, q) \Leftrightarrow \perp$$

Axioms for freeable

$\text{freeable}(h, q)$ Is q the start of an allocated region?

$$\text{freeable}(\varepsilon, q) \Leftrightarrow \perp$$

$$\text{mallocable}(h, p, s) \wedge p = q \Rightarrow \text{freeable}(\text{malloc}(h, p, s), q) \Leftrightarrow \top$$

$$\neg \text{mallocable}(h, p, s) \vee p \neq q \Rightarrow \text{freeable}(\text{malloc}(h, p, s), q) \Leftrightarrow \text{freeable}(h, q)$$

Axioms for freeable

$\text{freeable}(h, q)$

Is q the start of an allocated region?

$$\text{freeable}(\varepsilon, q) \Leftrightarrow \perp$$

$$\text{mallocable}(h, p, s) \wedge p = q \Rightarrow \text{freeable}(\text{malloc}(h, p, s), q) \Leftrightarrow \top$$

$$\neg \text{mallocable}(h, p, s) \vee p \neq q \Rightarrow \text{freeable}(\text{malloc}(h, p, s), q) \Leftrightarrow \text{freeable}(h, q)$$

$$\text{freeable}(h, p) \wedge p = q \Rightarrow \text{freeable}(\text{free}(h, p), q) \Leftrightarrow \perp$$

$$\neg \text{freeable}(h, p) \vee p \neq q \Rightarrow \text{freeable}(\text{free}(h, p), q) \Leftrightarrow \text{freeable}(h, q)$$

Axioms for mallocsize

`mallocsize(h, q)`

Size of the memory region starting at q
(if currently allocated)

Axioms for mallocsize

`mallocsize(h, q)`

Size of the memory region starting at q
(if currently allocated)

$$\text{mallocsize}(\varepsilon, q) = 0$$

Axioms for accessible

$\text{accessible}(h, q, t)$ Is $[q, q + t)$ contained in an allocated region?

Axioms for accessible

$\text{accessible}(h, q, t)$ Is $[q, q + t)$ contained in an allocated region?

$\text{accessible}(\varepsilon, q, t) \Leftrightarrow \perp$

Axioms for accessible

$\text{accessible}(h, q, t)$ Is $[q, q + t)$ contained in an allocated region?

$$\text{accessible}(\varepsilon, q, t) \Leftrightarrow \perp$$

$$\begin{aligned} \text{mallocable}(h, p, s) \wedge \text{contained}(p, s, q, t) &\Rightarrow \text{accessible}(\text{malloc}(h, p, s), q, t) \Leftrightarrow \top \\ \neg \text{mallocable}(h, p, s) \vee \neg \text{contained}(p, s, q, t) &\Rightarrow \text{accessible}(\text{malloc}(h, p, s), q, t) \\ &\Leftrightarrow \text{accessible}(h, q, t) \end{aligned}$$

Axioms for accessible

$\text{accessible}(h, q, t)$ Is $[q, q + t)$ contained in an allocated region?

$$\text{accessible}(\varepsilon, q, t) \Leftrightarrow \perp$$

$$\text{mallocable}(h, p, s) \wedge \text{contained}(p, s, q, t) \Rightarrow \text{accessible}(\text{malloc}(h, p, s), q, t) \Leftrightarrow \top$$

$$\neg \text{mallocable}(h, p, s) \vee \neg \text{contained}(p, s, q, t) \Rightarrow \text{accessible}(\text{malloc}(h, p, s), q, t) \\ \Leftrightarrow \text{accessible}(h, q, t)$$

$$\neg \text{freeable}(h, p) \Rightarrow \text{accessible}(\text{free}(h, p), q, t) \\ \Leftrightarrow \text{accessible}(h, q, t)$$

$$\text{freeable}(h, p) \wedge \text{disjoint}(p, \text{mallocsize}(h, p), q, t) \Rightarrow \text{accessible}(\text{free}(h, p), q, t) \\ \Leftrightarrow \text{accessible}(h, q, t)$$

$$\text{freeable}(h, p) \wedge \neg \text{disjoint}(p, \text{mallocsize}(h, p), q, t) \Rightarrow \text{accessible}(\text{free}(h, p), q, t) \Leftrightarrow \perp$$

Implementation

- \mathcal{T}_H is **not supported** by off-the-shelf SMT solvers (yet?)

Implementation

- $\mathcal{T}_{\mathcal{H}}$ is **not supported** by off-the-shelf SMT solvers (yet?)
- **Reduction** to $\mathcal{T}_{\mathcal{B}\gamma}$: Apply axioms in the form of rewrite rules

Implementation


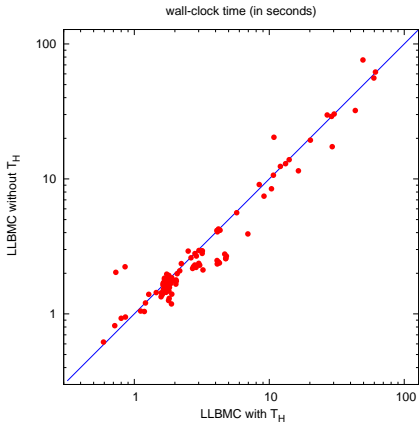
- $\mathcal{T}_{\mathcal{H}}$ is **not supported** by off-the-shelf SMT solvers (yet?)
- **Reduction** to $\mathcal{T}_{\mathcal{B}\mathcal{V}}$: Apply axioms in the form of rewrite rules

$$\begin{aligned} \text{mallocable}(h, p, s) \wedge \text{contained}(p, s, q, t) &\Rightarrow \text{accessible}(\text{malloc}(h, p, s), q, t) \Leftrightarrow \top \\ \neg \text{mallocable}(h, p, s) \vee \neg \text{contained}(p, s, q, t) &\Rightarrow \text{accessible}(\text{malloc}(h, p, s), q, t) \\ &\Leftrightarrow \text{accessible}(h, q, t) \end{aligned}$$

\rightsquigarrow

$$\begin{aligned} \text{accessible}(\text{malloc}(h, p, s), q, t) &\rightarrow \text{ITE}(\\ &\quad \text{mallocable}(h, p, s) \wedge \text{contained}(p, s, q, t), \\ &\quad \top, \\ &\quad \text{accessible}(h, q, t) \\ &\quad) \end{aligned}$$

Comparison of the current approach (SSV '10) and \mathcal{T}_H in LLBMC on 97 C-programs



	#mallocs / #frees	#accessible	time		memory	
			SSV '10	$\mathcal{T}_{\mathcal{H}}$	SSV '10	$\mathcal{T}_{\mathcal{H}}$
sparsemem	129 / 51	8374	76.2	49.5	861	404
binary-tree	127 / 127	3048	7.5	9.1	150	94
inplace-reverse	100 / 100	1800	20.4	10.8	260	119
wcet-bsort100	3 / 0	120204	12.4	12.1	246	246
wcet-statemate	106 / 0	2816	2.2	0.9	35	9

Conclusions and Future Work

- We have presented an **SMT theory** for reasoning about **memory access correctness**
 - **Precise** formalization of `malloc` and `free`
 - Opportunities for **simplification** of formulas using rewrite rules
 - Modular and local reasoning

Conclusions and Future Work

- We have presented an **SMT theory** for reasoning about **memory access correctness**
 - **Precise** formalization of `malloc` and `free`
 - Opportunities for **simplification** of formulas using rewrite rules
 - Modular and local reasoning
- Performance is roughly **comparable** to the current approach in LLBMC

Conclusions and Future Work

- We have presented an **SMT theory** for reasoning about **memory access correctness**
 - **Precise** formalization of `malloc` and `free`
 - Opportunities for **simplification** of formulas using rewrite rules
 - Modular and local reasoning
- Performance is roughly **comparable** to the current approach in LLBMC
- Potential **performance increases**:
 - Improved rewrite engine/strategy in the reduction
 - Use of rewrite rules that are **derived** from the axioms

- We have presented an **SMT theory** for reasoning about **memory access correctness**
 - **Precise** formalization of `malloc` and `free`
 - Opportunities for **simplification** of formulas using rewrite rules
 - Modular and local reasoning
- Performance is roughly **comparable** to the current approach in LLBMC
- Potential **performance increases**:
 - Improved rewrite engine/strategy in the reduction
 - Use of rewrite rules that are **derived** from the axioms
- **Open Questions**:

1. Can \mathcal{T}_H be **integrated into SMT solvers**?
 2. Is the **lemmas-on-demand** approach **applicable**?