

## AddressMonitor (AMon): Low-overhead Spatial and Temporal Memory Safety for C

Farzam Dorostkar Supervisor: Prof. Michel Dagenais Co-supervisor: Prof. Heng Li

Polytechnique Montreal

DORSAL Laboratory

Dec. 5th 2024

#### AddressMonitor (AMon): Introduction

- AddressMonitor (AMon) is a dynamic tool designed to detect heap spatial and temporal violations
  - □ Out-of-bounds accesses (buffer overflows)
  - □ Use-after-frees
  - Double-frees
  - □ Memory leaks
- Currently targets C programs running on X86-64
- Pointer tainting
- Adds untainting code at compile-time
- Designed to have minimal overhead
  - Compared to AddressSanitizer (ASan): Much lower memory overhead + More precision in detecting buffer overflows

## ASan: High Memory Overhead and Limited Overflow Detection



Red zones and shadow memory are key contributors to ASan's high memory overhead (2x-4x).



Shadow memory encodes a binary indicator (addressable or not) without object-specific details, leading to potential false negatives.

#### **ASan: Limited Overflow Detection**



Shadow memory encodes a binary indicator (addressable or not) without object-specific details, leading to potential false negatives.

#### **ASan: Limited Overflow Detection**



Shadow memory encodes a binary indicator (addressable or not) without object-specific details, leading to potential false negatives.

## **Pointer Tainting**

Polytechnique Montreal – Farzam Dorostkar

#### **Pointer Tainting**

#### When Allocating Heap Objects

- Assign a unique taint (ID) to each allocated object
- Build and maintain an object table [taint, base address, size, staus]
- On most 64 bits architecture the first 2 bytes are unused
- Embed the taint into the 2 MS bytes of the returned address (**pointer tainting**)

#### When Accessing Memory

- Retrieve the taint
- Use the taint to verify the access against the object table
- Unlike shadow-based techniques, pointer tainting provides object-specific analysis.

#### But...

• Dereferencing a tainted pointer causes segmentation fault!

## What Distinguishes AMon?

Polytechnique Montreal – Farzam Dorostkar

#### What Distinguishes AMon?

#### AMon instruments code with untainting logic at compile-time.

Similar tools, such as DataWatch, rely on costly dynamic techniques

- Let memory accesses trigger a SIGSEGV signal
- Perform dynamic patching
- Typically also require to re-taint
  - For instance, when modifying a register containing a tainted pointer, which might be shared across multiple memory-accessing instructions.
- Usually used for targeted protection, not full program verification

AMon instruments code at compile-time to call untainting

- No re-tainting logic is required as it operates on IR registers rather than physical registers
- Full program coverage

## AMon: Implementation (Runtime Library & Compile-time Transformation)

#### **AMon: Implementation**

#### Runtime Library (libamon.so)

- Intercepts standard heap allocation and de-allocation functions
  - To add taint to returned pointers, etc.
- Intercepts other standard C functions as well
  - To untaint possibly tainted arguments
- Maintains the object table
- Defines the verification logic
- Defines environment variables to control the behavior of AMon
  - Supported object sizes
- It can be either preloaded using LDPRELOAD or added as a dependency using patchelf

#### **Compile-time Transformation**

Implemented as an LLVM module pass

- 1. Identifies all IR instructions that perform pointer dereferencing
  - All <u>Memory Access and Addressing Operations</u> (such as load, store, and cmpxchg)
  - Memory-related <u>Standard C Library Intrinsics</u> (such as 11vm.memcpy and 11vm.memset)
- 2. Untaints the dereferenced pointer in each identified instruction
  - Extracts the pointer operand
  - Generates an untainted version of the pointer using the llvm.ptrmask intrinsic
  - Replaces the pointer in the current instruction with the untainted version using the setOperand method
- 3. Instruments each identified instruction for runtime verification
  - amon\_access(tainted pointer, access size, is\_write)

#### AMon: Implementation

#### AMon is easy to use.

- Compile your source code to LLVM IR clang -S -emit-llvm test.c -o test.ll
- Apply the 'amon' LLVM pass to the generated IR opt -passes='amon' -S test.ll -o test\_amon.ll
- Compile the transformed IR to executable clang test\_amon.ll -o test\_amon
- Run the executable while preloading libamon.so LD\_PRELOAD=./libamon.so ./test\_amon

```
1 #include <stdlib.h>
2
3 int main(int argc, char *argv[])
4 {
5 int tmp;
6 int *ptr_taint = (int*)malloc(sizeof(int));
7 *ptr_taint = 2024;
8 tmp = *(ptr_taint + 1);
9 free(ptr_taint);
10 *ptr_taint = 2025;
11 free(ptr_taint);
12 return 0;
13 }
```

<pre>ptr_taint = 0x03ec6231ba38b1b0 0x03ec 0 1 #include <stdlib.h> 2 3 int main(int argc, char *argv[]) 4 { 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2</stdlib.h></pre>	∂x6231ba38b1b0	 4	main.c	
<pre>ptr_taint = 0x03ec6231ba38b1b0 0x03ec 0 1 #include <stdlib.h> 2 3 int main(int argc, char *argv[]) 4 { 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2</stdlib.h></pre>	0x6231ba38b1b0	4		
<pre>1 #include <stdlib.h> 2 3 int main(int argc, char *argv[]) 4 { 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2</stdlib.h></pre>			main:6	Allocated
<pre>3 int main(int argc, char *argv[]) 1 s 4 { 2</pre>				
<pre>5 int tmp; 6 int *ptr_taint = (int*)malloc(sizeof(int)); 7 *ptr_taint = 2024; 8 tmp = *(ptr_taint + 1); 9 free(ptr_taint); 10 *ptr_taint = 2025; 11 free(ptr_taint); 12 return 0;</pre>	<pre>static void*    malloc(size_t s {     // internal     void *result =     result = amon_     objtbl_add(res     return result; }</pre>	size) l checks libc_ protect( ult, siz	malloc(size); result); // ta e); // added <sup>-</sup>	aint added to objtbl

objtbl	Taint	Base Address	Size	Call Stack	Status
			• • •		
<pre>ptr_taint = 0x03ec6231ba38b1b0</pre>	0x03ec	0x6231ba38b1b0	4	main:6	Allocated
<pre>1 #include <stdlib.h> 2 3 int main(int argc, char *argv[]) 4 { 5 int tmp; 6 int *ptr_taint = (int*)malloc(sizeof(int)); 7 *ptr_taint = 2024; 8 tmp = *(ptr_taint + 1); 9 free(ptr_taint); 10 *ptr_taint = 2025; 11 free(ptr_taint); 12 return 0; 13 }</stdlib.h></pre>	1 2 3 4 5 6 7 8 <b>te</b>	<pre>; Function Attrs: r define dso_local i3 entry:      %ptr_taint = call     store i32 2024, p  } st.ll</pre>	noinline B2 @main( l noalias otr %ptr_	nounwind uwtab i32 %argc, ptr ptr @malloc(i taint, align 4 <b>Segme</b>	le %argv) { 64 noundef 4) ntation Fault

	objtbl Ta			ldress	Size	Call Stack	Status	
					•••			
	<pre>ptr_taint = 0x03ec6231ba38b1b0</pre>	0x03ec	0x6231ba	a38b1b0	4	main:6	Allocated	
Adds untainting code at Compile-								
1 ; Funct 2 define 3 entry: 4 5 %ptr_ 6 store 7 8 }	<pre>tion Attrs: noinline nounwind uwtable   dso_local i32 @main(i32 %argc, ptr %argv) {    taint = call noalias ptr @malloc(i64 noundef 4    e i32 2024, ptr %ptr_taint, align 4         Segmentation Fau </pre>	1 ; 2 0 3 4 5 6 7 8 9 10 3	; Function At define dso_lo entry:  %ptr_taint %1 = call p call void ( store i32 2 	trs: noinl cal i32 @m call noa ptr @llvm. @amon_acce 2024, ptr 9	ine nounwi ain(i32 %a lias ptr @ ptrmask.p0 ss(ptr %pt %1, align	nd uwtable rgc, ptr %argv) { malloc(i64 nounde .i64(ptr %0, i64 r_taint, i64 4, 4	ef 4) 0x0000FFFFFFFFFFF il true)	
test.ll		tes	st_amon.	11				
				Ru	ntime v	erification ag	gainst objtbl	

objtbl	Taint	Base Address	Size	Call Stack	Status
			•••		
<pre>ptr_taint = 0x03ec6231ba38b1b0</pre>	0x03ec	0x6231ba38b1b0	4	main:6	Allocated
<pre>1 #include <stdlib.h> 2 3 int main(int argc, char *argv[]) 4 { 5 int tmp; 6 int *ptr_taint = (int*)malloc(sizeof(int)); 7 *ptr_taint = 2024; 8 tmp = *(ptr_taint + 1); 9 free(ptr_taint); 10 *ptr_taint = 2025; 11 free(ptr_taint); 12 return 0; 13 }</stdlib.h></pre>	1; 2d 3e 4 5 6 7 8} tes	<pre>Function Attrs: noin efine dso_local i32 @ ntry: %add.ptr = getelemen %2 = load i32, ptr % t.ll</pre>	line noun main(i32 tptr inbo add.ptr,	wind uwtable %argc, ptr %arg ounds i32, ptr % align 4 <b>Segme</b>	v) { ptr_taint, i64 entation Fau

		objtbl	Taint	Base /	Address	Size	Call Stack	Status
						•••		
	<pre>ptr_taint = 0x03ec</pre>	6231ba38b1b0	0x03ec	0x6231	1ba38b1b0 4		main:6	Allocated
					Adds un	tainting	g code at Co	mpile-time
<pre>1; Function Attrs: noinline nounwind uwtable 2 define dso_local i32 @main(i32 %argc, ptr %argv) { 3 entry: 4 5 %add.ptr = getelementptr inbounds i32, ptr %ptr_taint, i64 1 6 %2 = load i32, ptr %add.ptr, align 4 7 8 } Segmentation Fault!</pre>				<pre>; Function / define dso_1 entry:</pre>	<pre>Attrs: noinlin local i32 @mai = getelementpt ptr @llvm.pt @amon_access i32, ptr %3,</pre>	e nounwind n(i32 %argc r inbounds rmask.p0.i6 (ptr %add.p align 4	uwtable , ptr %argv) { i32, ptr %ptr_tain 4(ptr %add.ptr, i6 tr, i64 4, i1 fals	t, i64 1 64 0x0000FFFFFFFFF Ge)
lest.11			Le	'S L_amor	Runt	ime ver	ification aga	inst obitbl
				am AM	on_acces	s(0x030 s the ou	ec6231ba38b1 ut-of-bounds	b4, 4, false

objtbl	Taint	Base Address	Size	Call Stack	Status
			• • •		
<pre>ptr_taint = 0x03ec6231ba38b1b0</pre>	0x03ec	0x6231ba38b1b0	4	main:9	Freed
<pre>1 #include <stdlib.h> 2 3 int main(int argc, char *argv[]) 4 { 5 int tmp; 6 int *ptr_taint = (int*)malloc(sizeof(int)); 7 *ptr_taint = 2024; 8 tmp = *(ptr_taint + 1); 9 free(ptr_taint); 10 *ptr_taint = 2025; 11 free(ptr_taint); 12 return 0; 13 }</stdlib.h></pre>		<pre>1 static void* 2 free(void *ptr) 3 { 4 // internal of 5 if(amon_is_proted 6 elselibc_freed 7 } .ibamon.so</pre>	checks cted(ptr); (ptr);	) amon_free_prot	ect( <b>ptr</b> );

#### **Compile-time Transformation: Optimizations**

Not all identified instructions require monitoring

- 1. Filtering out accesses to stack
  - Filters out instructions where the pointer operand originates directly or indirectly (through pointer arithmetic) from an alloca instruction.
- 2. Not instrumenting memory accesses that can be analyzed statically (out-of-bounds)
  - AMon additionally incorporates a degree of static analysis
  - When both the allocation size and the access size are known at compile time
  - No instrument (runtime check) needed, but still needed to untaint the pointer
- 3. If all accesses to a pointer can be analyzed statically, no need to taint it
  - For instance, replaces a malloc with a direct \_\_libc\_malloc

# Results & Conclusion

#### Results

- Three SPEC CPU 2017 benchmarks
- Under ASan and AMon
- Compared to native compilation

		Native		ASan			AMon		
Benchmark	Time (Sec)	MRSS (MB)	Exec. Size	Time (x)	MRSS (x)	Exec. Size	Time (x)	MRSS (x)	Exec. Size
505.mcf_r	7.8	292	83 KB	1.4 x	2.1 x	11 MB	1.3 x	≈ 1 (1.03) x	120 KB
519.lbm_r	1.3	420	64 KB	2.6 x	1.2 x	11 MB	1.6 x	≈ 1 (1.02) x	92 KB
544.nab_r	1.1	2.6	506 KB	1.4 x	5.5 x	12 MB	1.2	1.3 x	665 KB
557.xz_r	0.7	531	588 KB	1.7 x	1.2 x	12 MB	1.4	≈ 1 (1.02) x	763 KB

- AddressMonitor (AMon): A tool designed to detect heap-based spatial and temporal memory access violations in C programs targeting the x86-64 architecture
- **Core approach:** Leverages pointer tainting and introduces a novel compile-time transformation for efficient pointer untainting.
- **Compared to ASan:** Achieves significantly lower memory overhead, relatively reduced runtime overhead, and enhanced precision in identifying buffer overflow vulnerabilities
- Implementation Flexibility: Can be integrated with compilers beyond LLVM

### Thanks! Questions? Comments?

farzam.dorostkar@polymtl.ca https://github.com/farzamdorostkar https://farzamdorostkar.github.io/