

MaxSim: A Simulation Platform for Managed Applications

Open-source: <https://github.com/beehive-lab/MaxSim>

Andrey Rodchenko, Christos Kotselidis,
Andy Nisbet, Antoniu Pop, Mikel Lujan

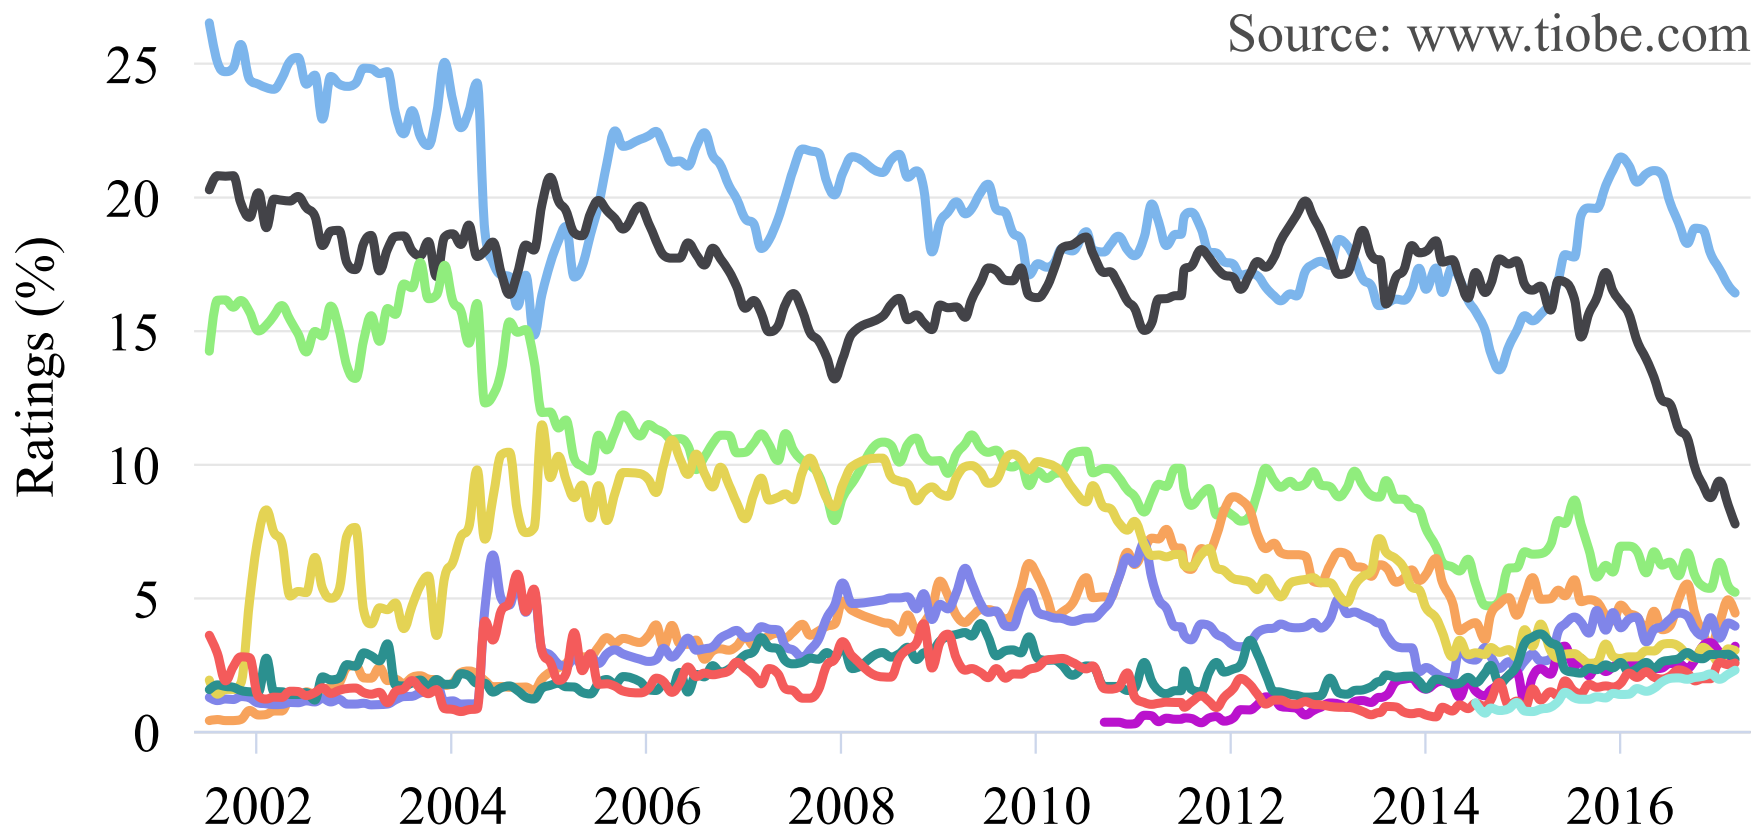
Advanced Processor Technologies Group,
School Of Computer Science,
The University of Manchester

- What simulation platform for managed applications is needed and why?
- VM Selection Justification: Maxine VM
- Simulator Selection Justification: ZSim
- MaxSim: Overview and Features
- Use Cases: Characterization, Profiling, and HW/SW Co-design
- Conclusion

What simulation platform for managed applications is needed and why?

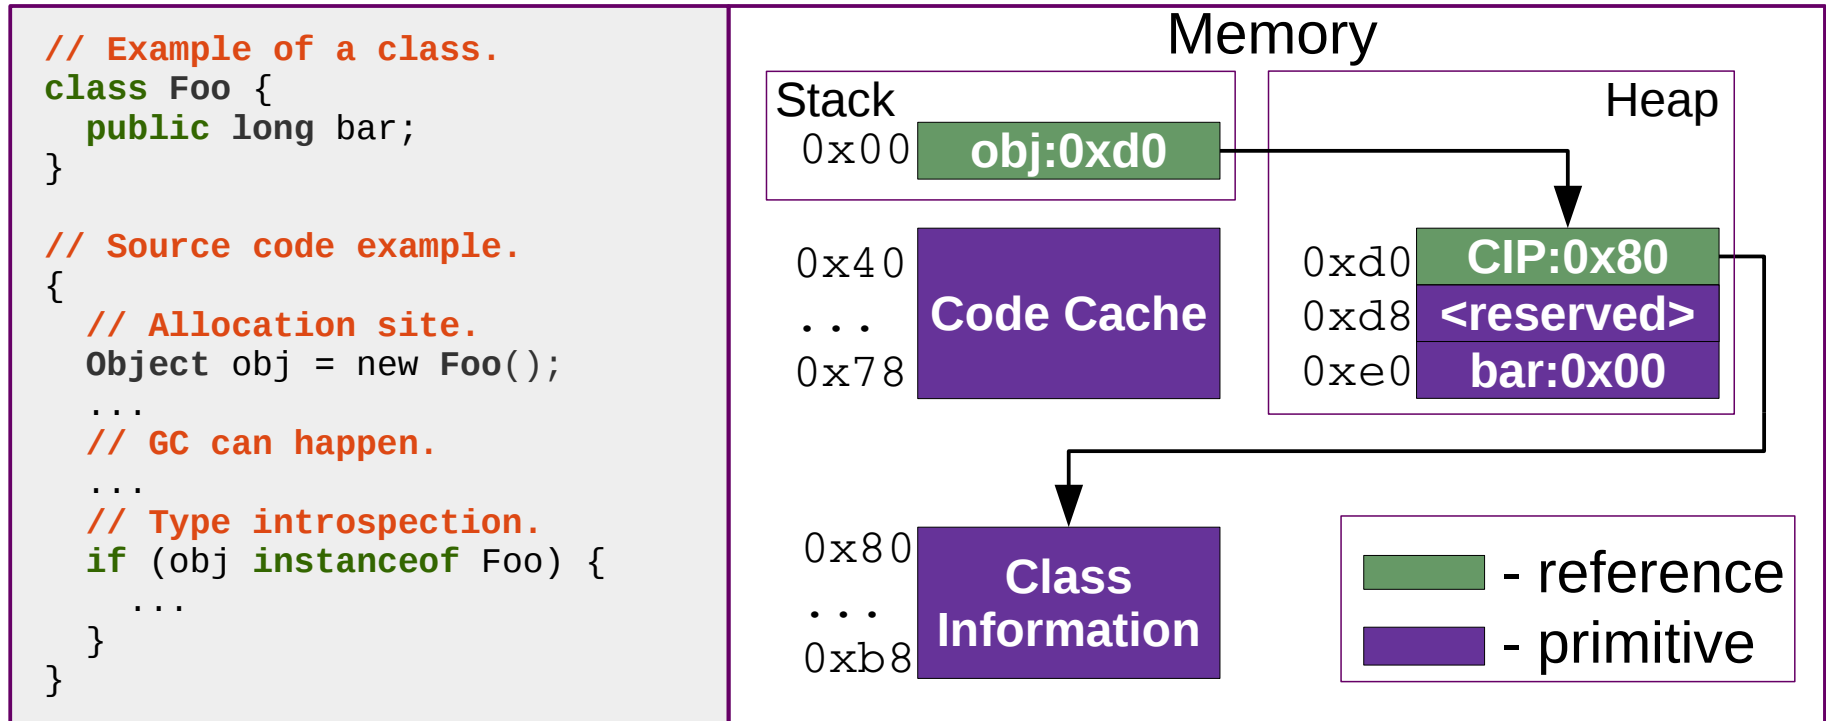
TIOBE Programming Community Index (March 2017)

1. Java
2. C
3. C++
4. C#
5. Python
6. Visual Basic .NET
7. PHP
8. JavaScript
9. Delphi/Object Pascal
10. Swift



What simulation platform for managed applications is needed and why?

Specific Characteristics of Managed Applications

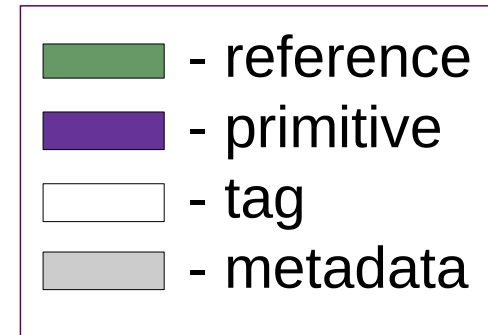



- Distributed in the verifiable bytecode format
- Automatic memory management
- JIT compilation and interpretation
- Object orientation and associated metadata

What simulation platform for managed applications is needed and why?

Support for Tagged Pointers

- An option for object metadata storage



- Support in commodity 64-bit architectures
 - AArch64: [tag:8b | pointer:48b]
 - SPARC M7: [tag:8b | pointer:48b] - [tag:32b | pointer:32b]
 - x86-64: [signExtension | pointer:(48b|57b)]

What simulation platform for managed applications is needed and why?

Design Goals

- Productivity for research
 - VM modularity and support of other languages
 - High simulation speed (DaCapo benchmarks in one day on a single PC)
- Awareness of the VM in the simulator
- Advanced features
 - Support of tagged 64-bit pointers
 - Ability to experiment with different object layouts
 - Ability to perform power and energy modeling

Maxine VM¹: A Platform for Research in VM Technology

- Mostly written in Java, with a substrate written in C
- Modular design: schemes for object layouts, object references, heap and GC, thread synchronization, etc.
- Compilers: T1X (O0), C1X (O1), Graal (O2)
 - Graal supports other languages via Truffle (JavaScript, R, Ruby, others)
- Target ISAs: x86-64, ARMv7
- Class library: JDK 7

[1] Wimmer *et al.*, “Maxine: An approachable virtual machine for, and in, Java”, TACO, 2013

VM Selection Justification: Maxine VM

Maxine Inspector: Integrated Debugging Support


The screenshot displays the Maxine Inspector interface for a VM process. The main window is titled "Maxine Inspector (mode=CREATE) VM Process Stopped". The interface is divided into several panes:

- Threads:** A list of threads with columns for ID and Name. The threads listed are: 1: main, 2: VmOperationThread, 4: Reference Handler, 5: Finalizer, and 3: Signal Dispatcher.
- Stack: main [1] (Breakpoint):** A stack trace showing the current execution path:
 - 0: HelloWorld.main0[B]
 - 1: test_output_HelloWorld\$main\$95.invoke0[B]
 - 2: OPT2BASELINE-Adapter(RRR)
 - 3: Method.invoke0[O]
 - 4: JavaRunScheme.lookupAndInvokeMain0[O]
 - 5: JavaRunScheme.run0[O]
 - 6: VmThread.executeRunnable0[O]
 - 7: VmThread.run0[O]
 - 8: thread_run
- Method: void main(String[])[B] in test.output.HelloWorld:** A window showing the machine code for the selected method. The code is displayed in a table with columns for Tag, Addr., Label, Instr., and Operands. The current instruction is highlighted in red:

Tag	Addr.	Label	Instr.	Operands
0: getstatic	7fffbb9d6013L2:		mov	rdi, <50143>StaticTuple{System}
	7fffbb9d601a		mov	esi, 0x18
	7fffbb9d601f		movsxd	rsi, esi
	7fffbb9d6022		mov	rax, rdi[rsi]
	7fffbb9d6026		subq	rsp, 0x10
	7fffbb9d602a		mov	[rsp], rax
3: ldc	7fffbb9d602e		mov	r11, <50145>"Hello World!"
	7fffbb9d6035		subq	rsp, 0x10
	7fffbb9d6039		mov	[rsp], r11
5: invokevirtual	7fffbb9d603d		mov	edi, 0x2F
	7fffbb9d6042		mov	rsi, [rsp + 16]
	7fffbb9d6047		mov	rdx, <49535>MethodProfile
	7fffbb9d604e		mov	ecx, 0x1
	7fffbb9d6053		mov	rax, rsi

VM Selection Justification: Maxine VM

Maxine VM: Performance Comparison Against Hotspot VM



- Maxine VM performance is ~59% of the highly optimized Hotspot VM
- Graal (O2) compiler delivers 8% better performance than C1X (O1)

ZSim¹: Fast and Accurate Microarchitectural Simulation

- x86-64 execution-driven timing simulator based on Pin
- Bound-weave technique for scalable simulation
- Lightweight user-level virtualization
- Comparison with open simulators supporting managed applications


Simulator	Engine	Full-System	Simulation Speed
gem5	Emulation	yes	~100-300 KIPS
Sniper *	DBT	no	~1-3 MIPS
ZSim	DBT	no	~7-20 MIPS

[1] Sanchez *et al.* "ZSim: Fast and Accurate Microarchitectural Simulation of Thousand-Core Systems", ISCA, 2013

* Sniper can simulate DaCapo benchmarks on 32-bit Jikes RVM only.

Simulator Selection Justification: ZSim

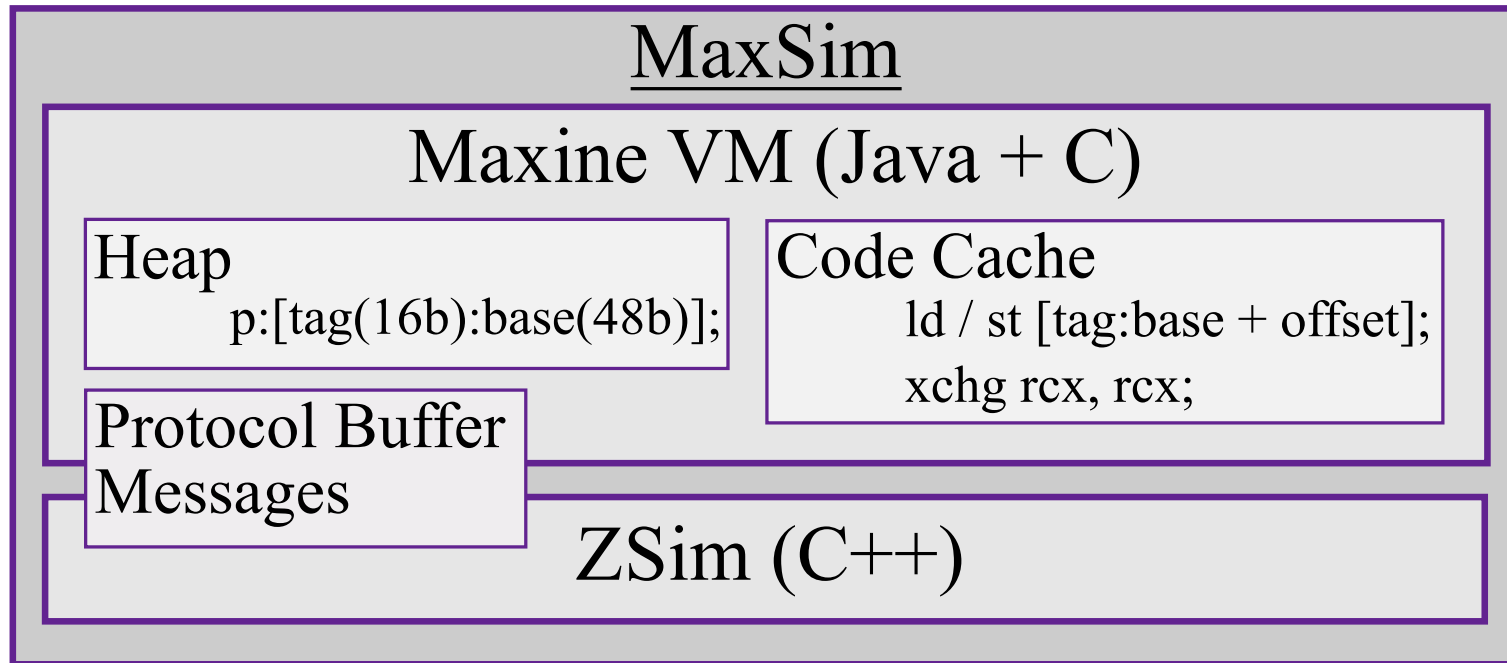
ZSim Validation: DaCapo on Maxine VM



- 100% pass rate and ~10% geomean simulation error at ~12 MIPS
- Inconsistencies:
 - eclipse, tradesoap (1C-*): Round Robin vs CFS scheduling
 - avrora: spends more than 50% of execution in the kernel

MaxSim: Overview and Features

Maxine-ZSim Integration Scheme



- Protocol Buffer Messages
 - Interface definition
 - Configuration
 - Profile serialization
- Magic NOPs
 - Simulation control
 - VM awareness
 - Sending/receiving protocol buffer messages
- Tagged Pointers
 - VM awareness
 - Profiling

VM Awareness in the Simulator

- VM memory regions
 - Stack
 - TLS
 - Heap
 - Code cache
 - Native code
 - Others
- VM operations
 - Garbage collection
 - Object allocation
- Object binding
 - To its class
 - To its allocation site

Pointer Tagging

- Two types of pointer tagging are supported

- Class ID tagging

- Allocation site ID tagging

```
// Example of a class.  
class Foo {  
    public long bar;  
}  
  
// Source code example.  
{  
    // Allocation site.  
    Foo obj = new Foo();  
    obj.bar = 42;  
}
```

- Tagging/untagging of all pointers at arbitrary places of execution

- Enables simulation fast-forwarding

- After tagging the following properties are preserved:

- Pointers to the same object are tagged with the same tag

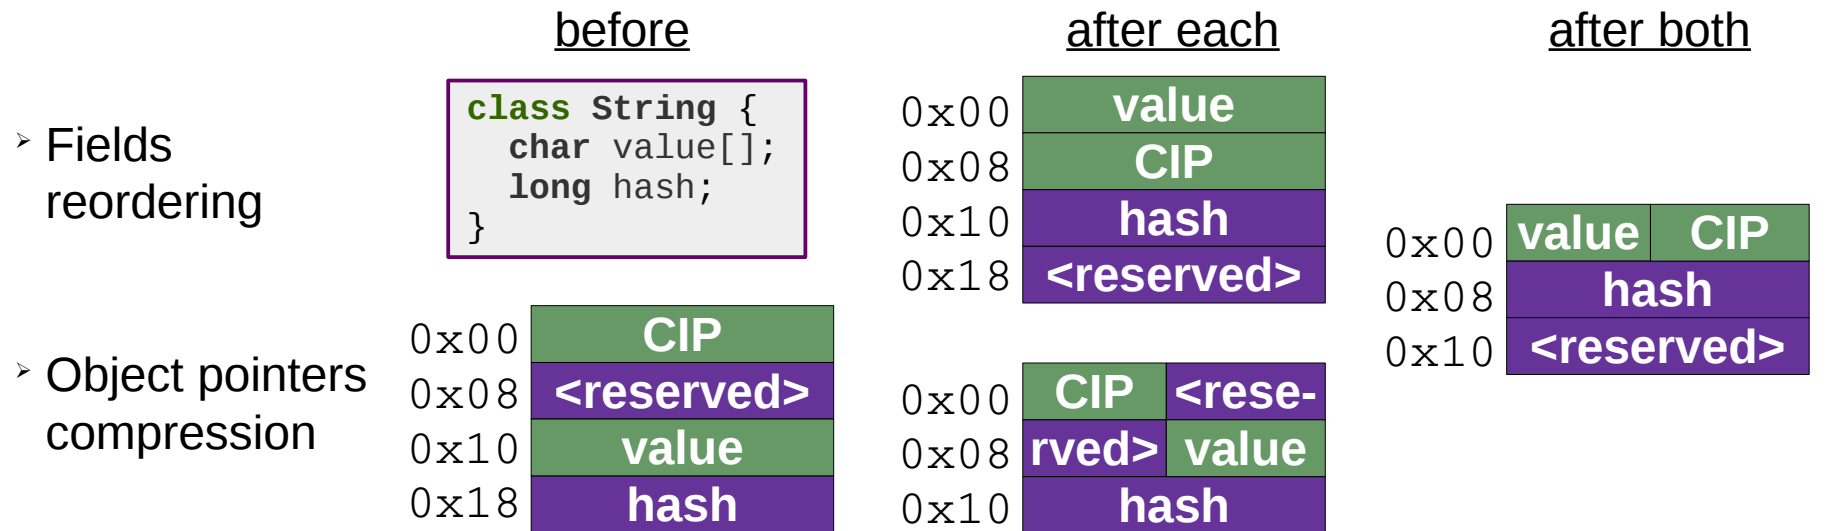
- Tags are immutable between an allocation and a garbage collection

- Objects are accessed using `[tag:base + offset]` addressing mode

MaxSim: Overview and Features

Address Space Morphing

- Motivation: easy experimentation with object layouts without adding extra complexity or breaking modularity of Maxine VM
- Supports two object layout transformations



- Makes use of two properties of MaxSim
 - › Flexibility of Maxine VM to expand object fields
 - › Ability of ZSim to remap memory addresses

MaxSim: Overview and Features

Stages of Address Space Morphing

$f_e(1,2)$ - expansion
in Maxine VM

$f_c(2)$ - contraction
in ZSim

$f_r(m_c)$ - reordering
in ZSim

Layout

0x00	ref.0
0x08	prim.1
0x10	ref.2
0x18	prim.3

0x00	ref.0
0x08	prim.1
0x10	
0x18	ref.2
0x20	prim.3
0x28	

0x00	ref.0	pri-
0x08	m.1	ref.2
0x10	prim.3	

0x00	ref.2	ref.0
0x08	prim.3	
0x10	prim.1	

Addressing

$$[b_o + o_o]$$

$$[f_e(b_o) + f_e(o_o)]$$

$$[b_e/2 + o_e/2]$$

$$[b_c + m_c(o_c)]$$

Fields Reordering Map

0x00 → 0x08
0x08 → 0x18
0x10 → 0x00
0x18 → 0x10


m_o

0x00 → 0x08
0x08 → 0x20
0x18 → 0x00
0x20 → 0x10

m_e

0x00 → 0x04
0x04 → 0x10
0x0C → 0x00
0x18 → 0x08

m_c

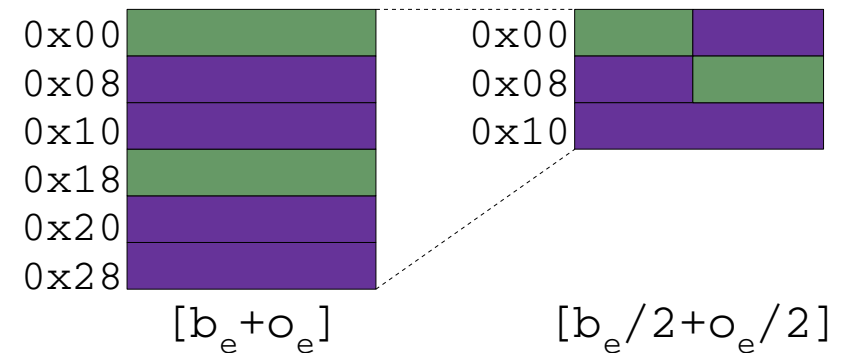


Address Space Morphing: Special Cases and Validation

- Simulation filtering of copying and initialization

```
// Loop used for initialization.  
void setWords(Pointer p, int n) {  
    ZSIM_MAGIC_NOP(BEGIN_LOOP_FILTERING);  
    for (int i = 0; i < n; i++) {  
        p.writeWord(i, 0);  
    }  
    ZSIM_MAGIC_NOP(END_LOOP_FILTERING);  
}
```

$f_c(2)$ – contraction
in ZSim



- Special cases for fast simulation

- Array of primitives and code cache objects are handled differently

- Validation

- References and primitives were expanded twice in Maxine VM and contracted twice in ZSim
- Less than 1% difference in comparison with the original layout


MaxSim: Use Cases

DaCapo Tomcat Characterization


1 Core 2MB LLC: 

4 Cores 8MB LLC: 


GC part: 




Instructions per Clock



L3 Load Cache Misses per Kilo Instruction



Consumed Power



L2 Load Cache Misses per Kilo Instruction

Analysis of L2 Cache Misses via Profiling

MaxSim output of class profiling information

```
char[]([C)(i:43 mf:57163720 (s:56(200337) ... r2m:722499 w2m:158200 r3m:108784 w3m:7723):  
...  
(o:16 f:.35 r:18602074 w:759093 r2m:596449 w2m:62251 r3m:80211 w3m:161)  
...
```

Maxsim output of cache miss site profiling information

```
...  
[java.lang.String.equals(Object)+108(k:I bci:23)](m:539629 i:43 ol:16 oh:16)  
...
```

String.class bytecode

bci	Instr.	Line
20	<i>getfield</i>	980
23	<i>arraylength</i>	980
24	<i>istore_3</i>	980


String.java source code

```
974 public boolean equals(Object anObject) {  
975     if (this == anObject) {  
976         return true;  
977     }  
978     if (anObject instanceof String) {  
979         String anotherString = (String) anObject;  
980         int n = value.length;  
981         if (n != anotherString.value.length)  
982             return false;  
983         ...
```

MaxSim: Use Cases

Storing Array Length in a Pointer Tag

- Having 16-bit-tagged pointers it is possible to store a range of array lengths [0;0xFFFFE], when 0xFFFF is Not an Array Length (NaAL) indicator
- Array length retrieval in software



Source code

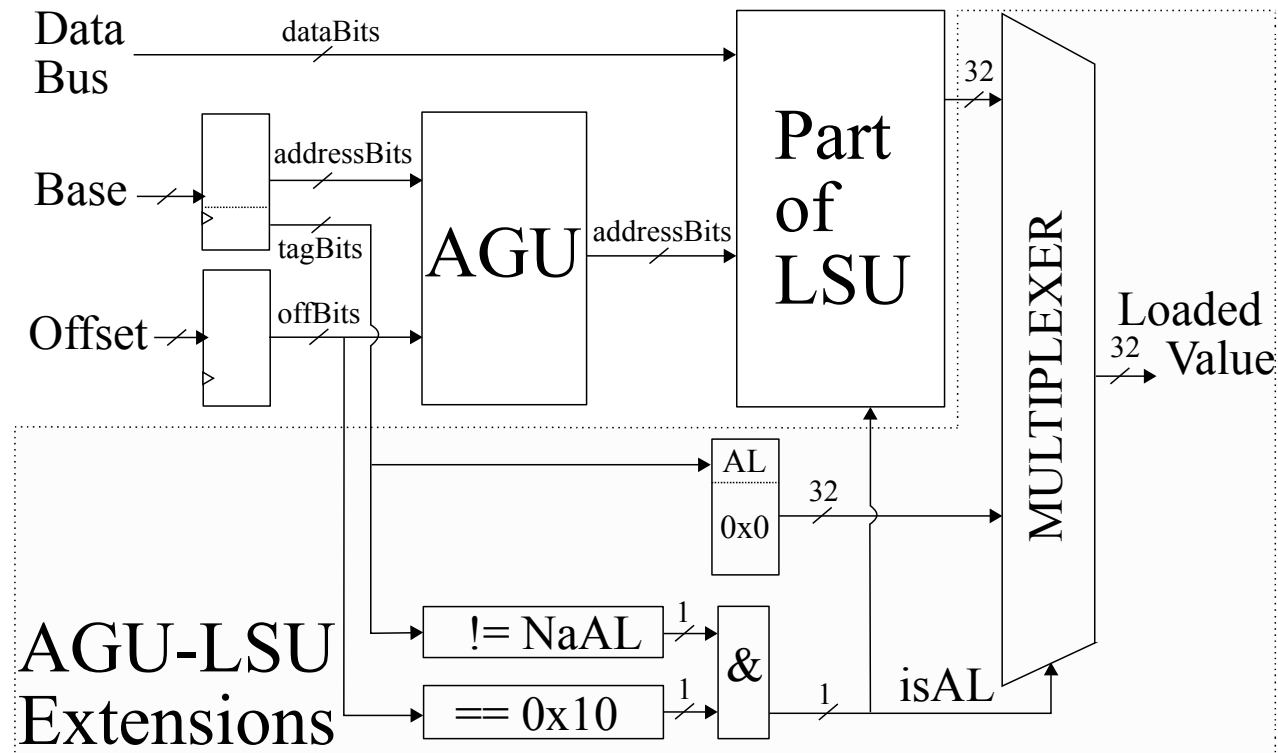
```
inline int retrieveArrayLength(Pointer_t objectPointer)
{
    TAG_t tag = extractTAG(objectPointer);
    if (tag != NaAL) {
        return (int) tag;
    }
    return * ((int *) (objectPointer + 0x10));
}
```

x86-64 assembler

```
// objectAddress in %rdi
movq %rdi, %rax
shrq $48, %rax
cmpq $65535, %rax
jne .L1
movq 16(%rdi), %rax
.L1:
// array length in %rax
```

- Dynamic execution height of 4.5 instructions of 19 bytes
- Originally 1 instruction of 4 bytes

HW-Assisted Array Length Retrieval from Tagged Pointers





- Array length retrieval in one instruction

```
inline int retrieveArrayLength(Address_t objectAddress) {  
    return * (( CIP_t *) (objectAddress + 0x10));  
}
```

```
// objectAddress in %rdi  
movq 16(%rdi), %rax  
// array length in %rax
```

MaxSim: Use Cases

Evaluation of HW-Assisted Array Length Retrieval



- Novel simulation platform for managed applications
 - Based of the state-of-the art VM and simulator
 - Awareness of the VM in the simulator
 - Simulation of 16-bit tagged pointers on x86-64
 - Low-overhead memory access profiling
 - Address-space morphing technique
- Use cases
 - Workload characterization and profiling
 - HW/SW co-design and exploration of architectural specialization for managed applications
 - Easy experimentation with object layout transformations
- Open-source platform is available at:
<https://github.com/beehive-lab/MaxSim>