




# Pharo Virtual Machine

## News from the Front

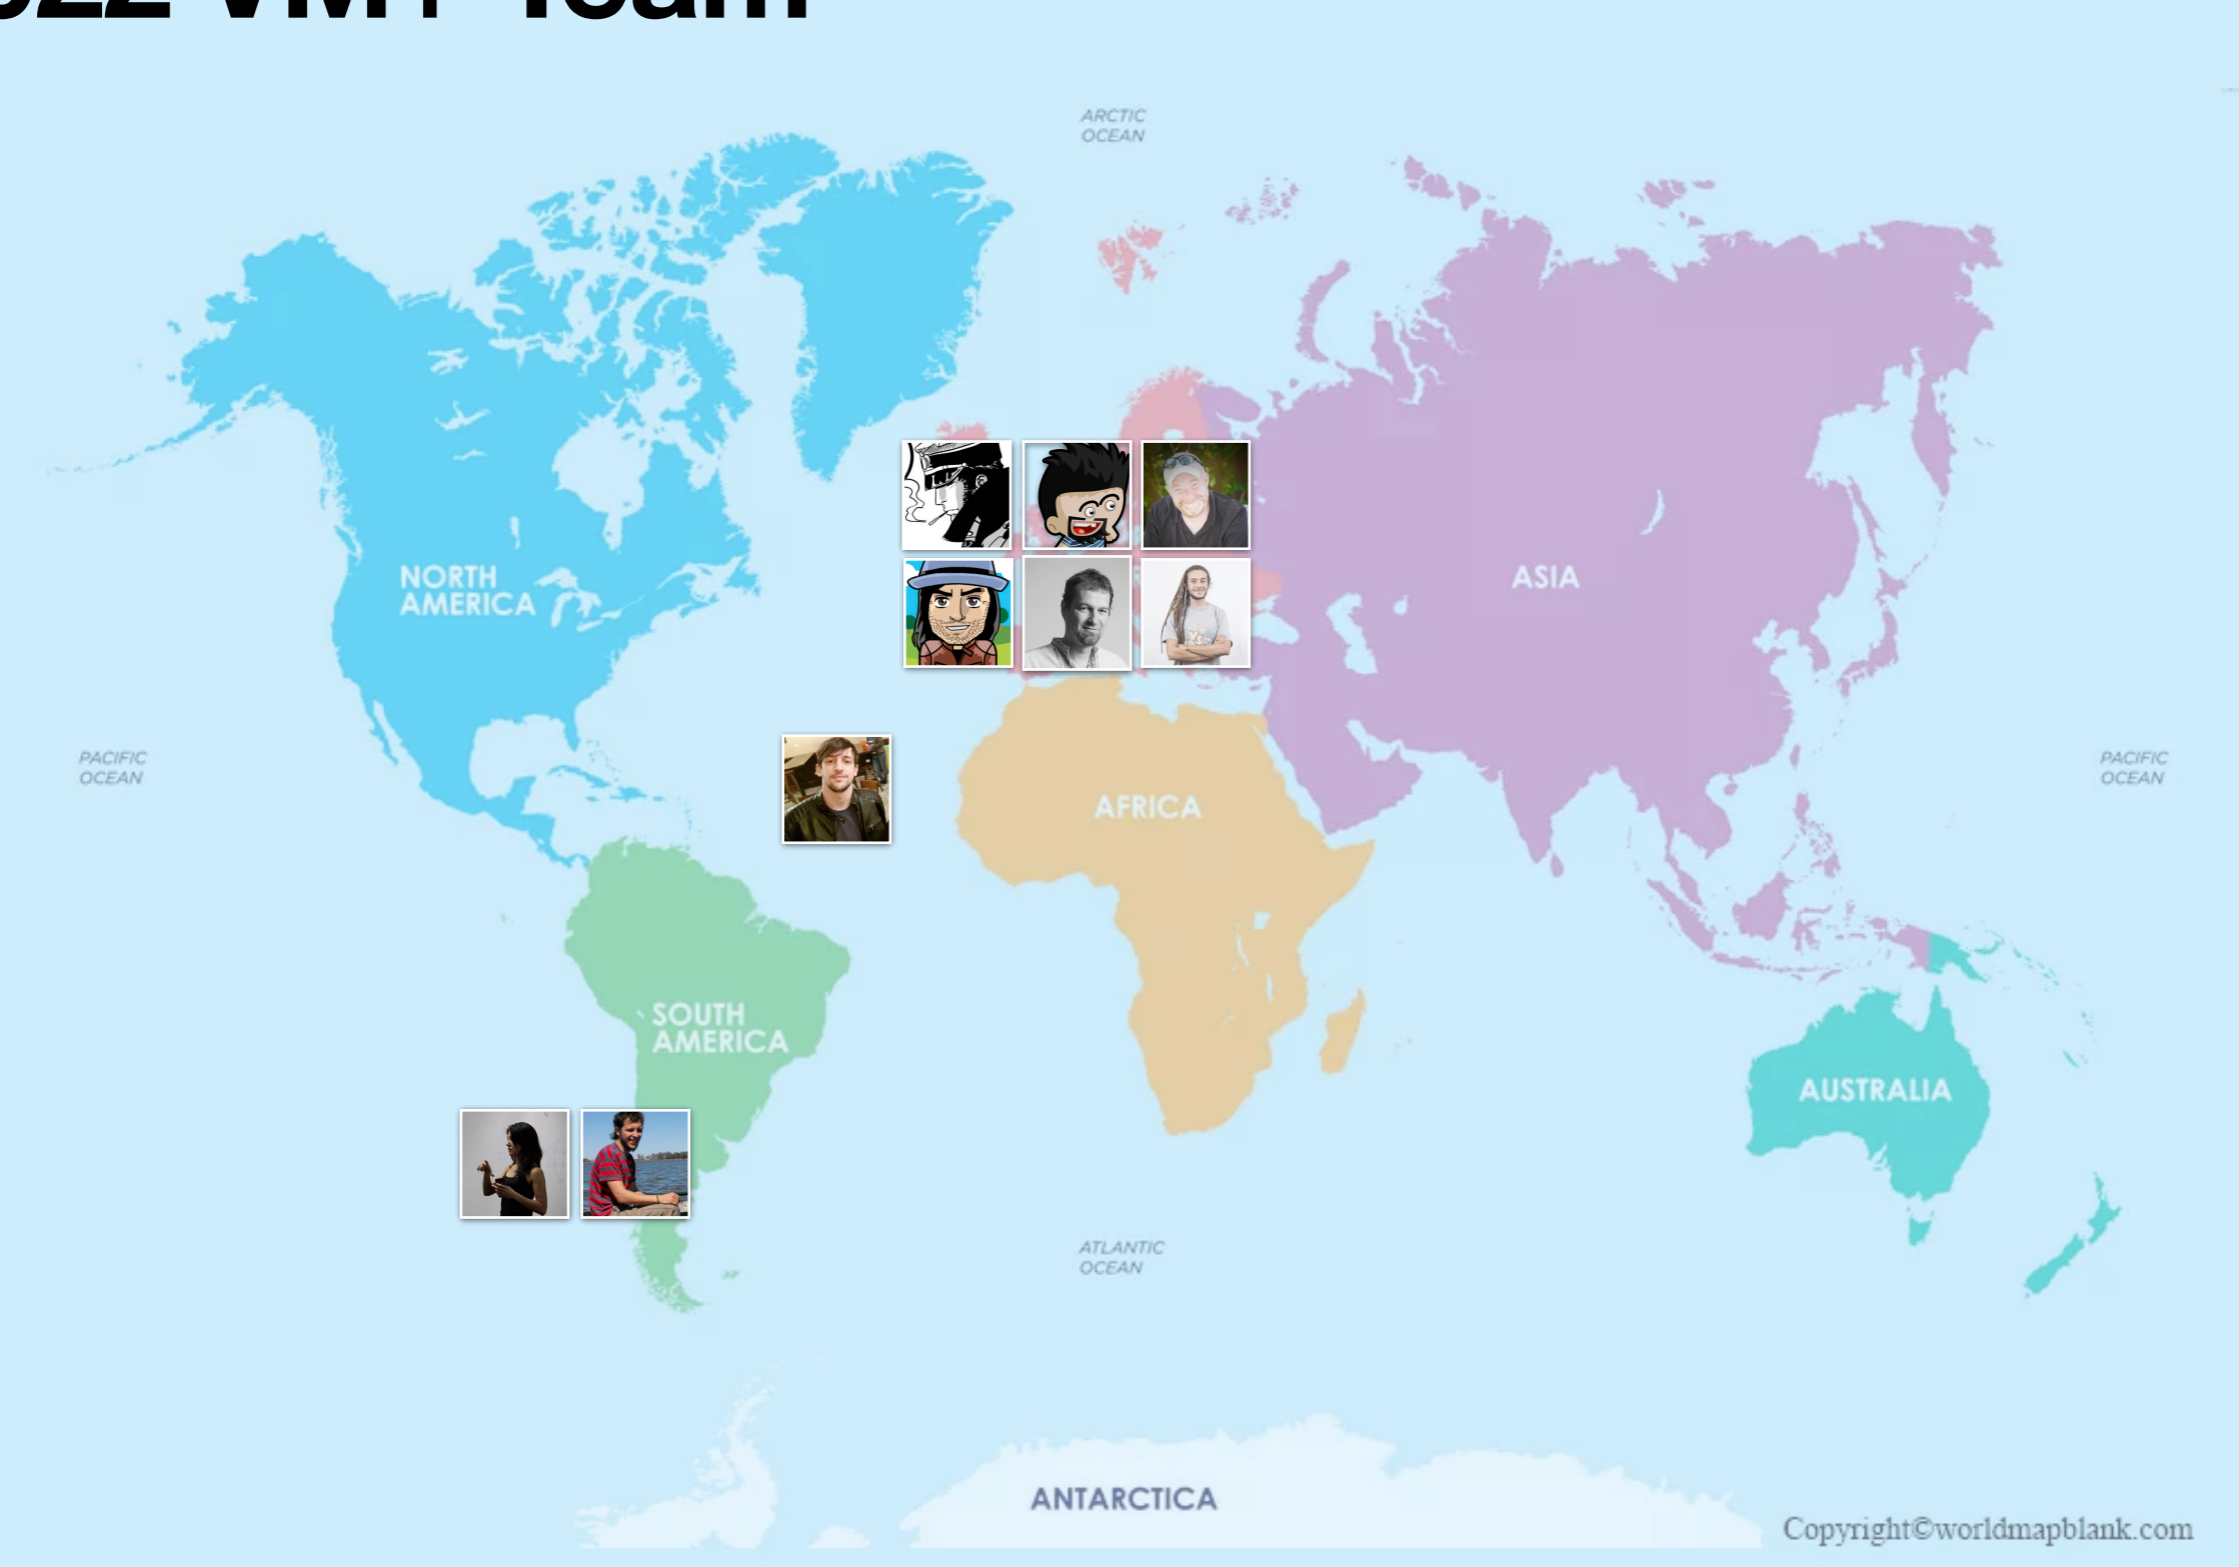
P. Tesone - G. Polito - ESUG'22  
[@tesonep pablo.tesone@inria.fr](mailto:pablo.tesone@inria.fr)  
[@guillep guillermo.polito@univ-lille.fr](mailto:guillermo.polito@univ-lille.fr)

*Inria*





 Université  
de Lille

# 2022 VM+ Team



# Virtual Machine Execution Engine




# ARM64 Backend

- ARM64 is now pervasive:
  - New Apple M1
  - Raspberry Pi 4
  - Microsoft Surface Pro X
  - PineBook Pro
  - ...

```
move r1 #1  
move r2 #17  
checkSmallInt  
checkSmallInt  
add r3 r1 r2  
checkSmallInt  
move r1 r3  
ret
```

JIT compiler IR



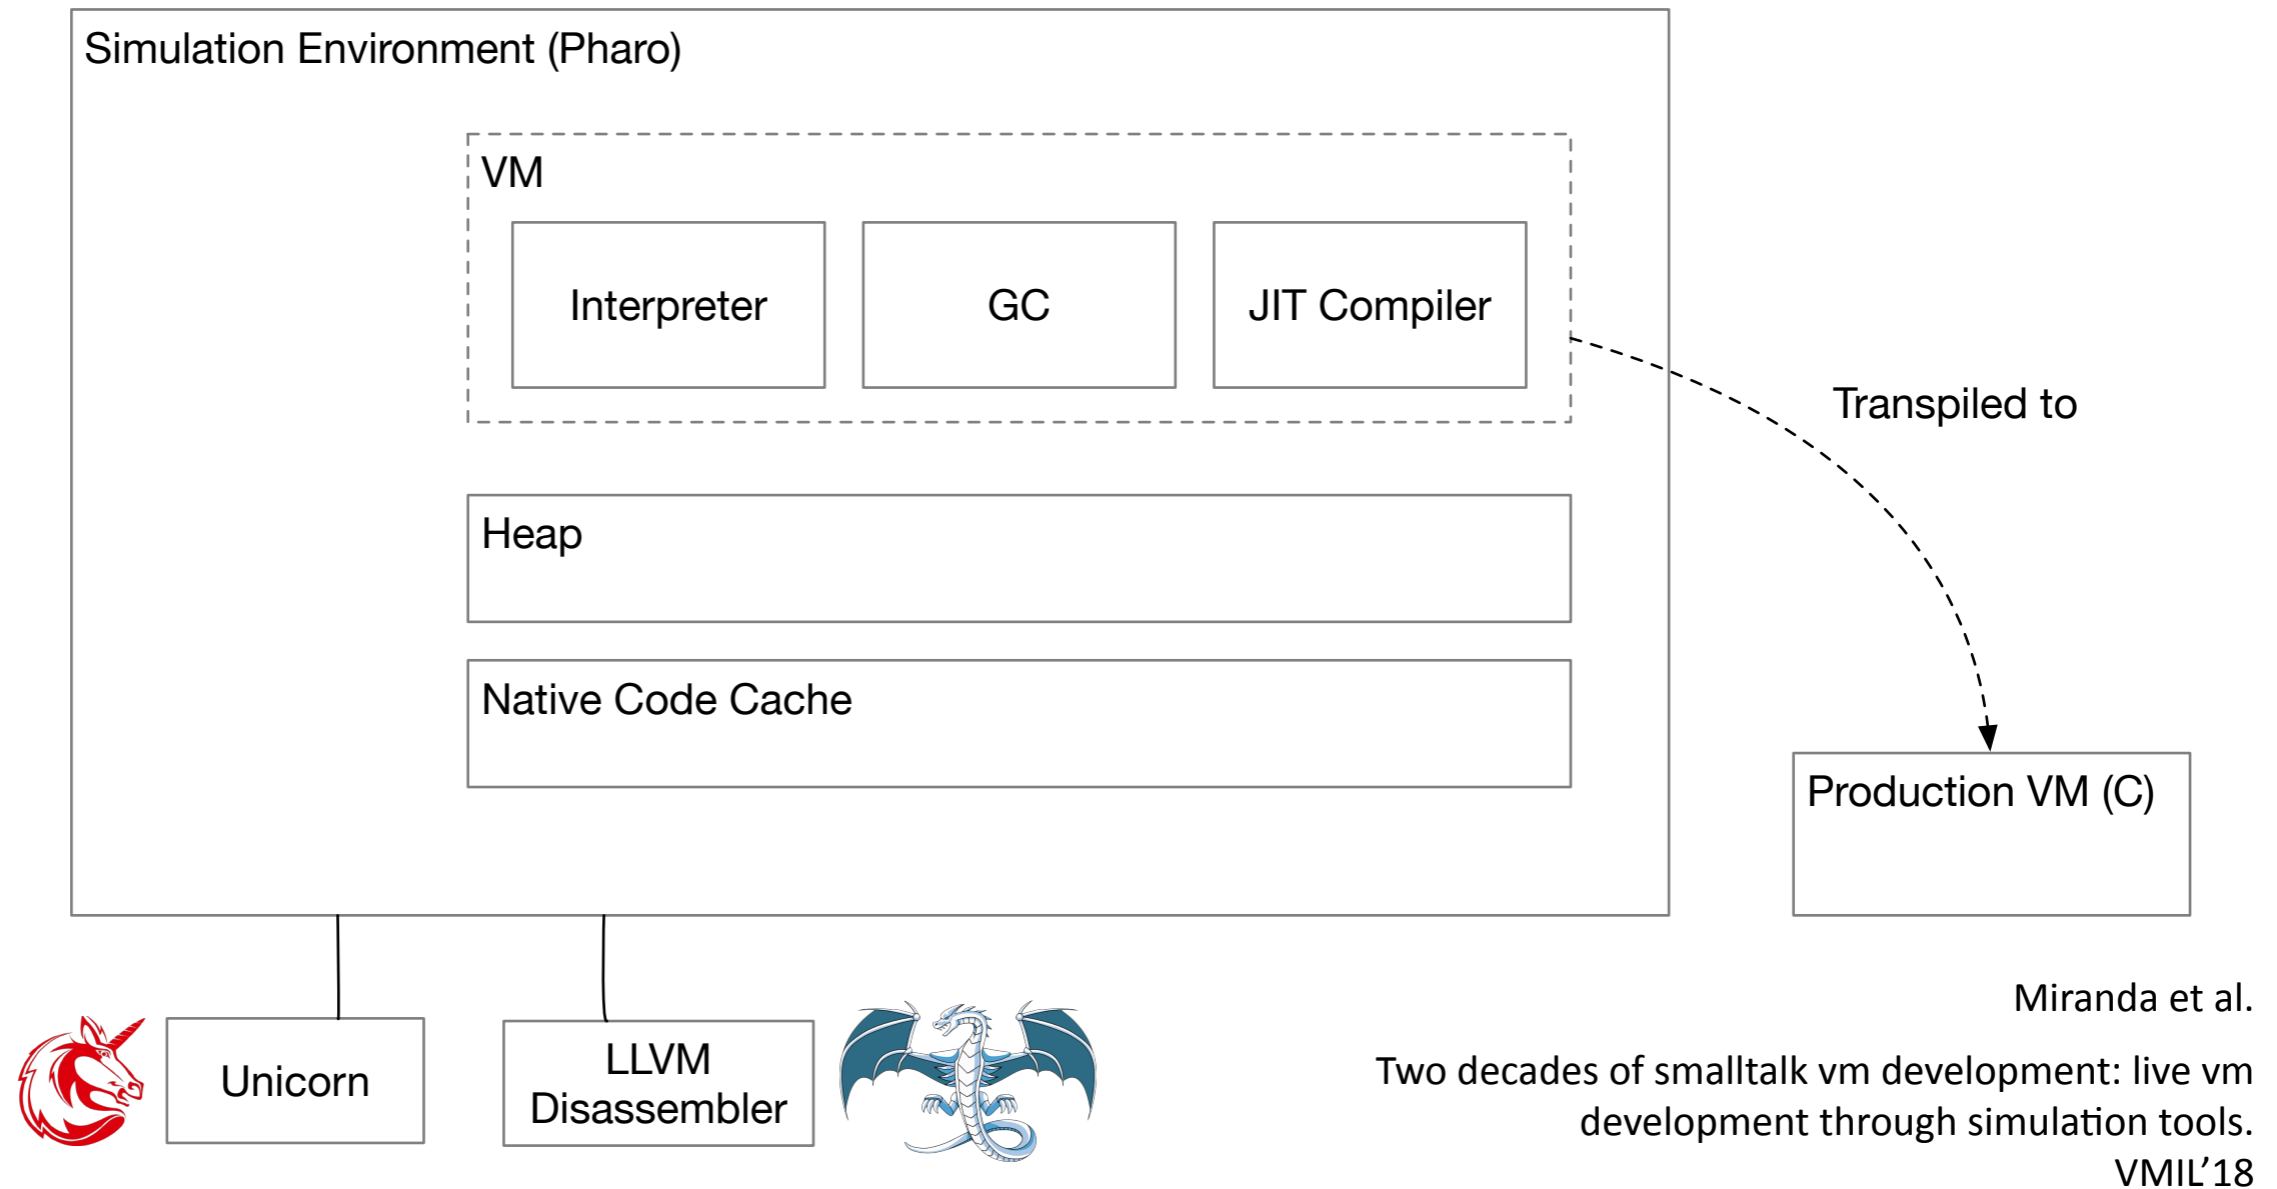


# Working Directly on Real Hardware

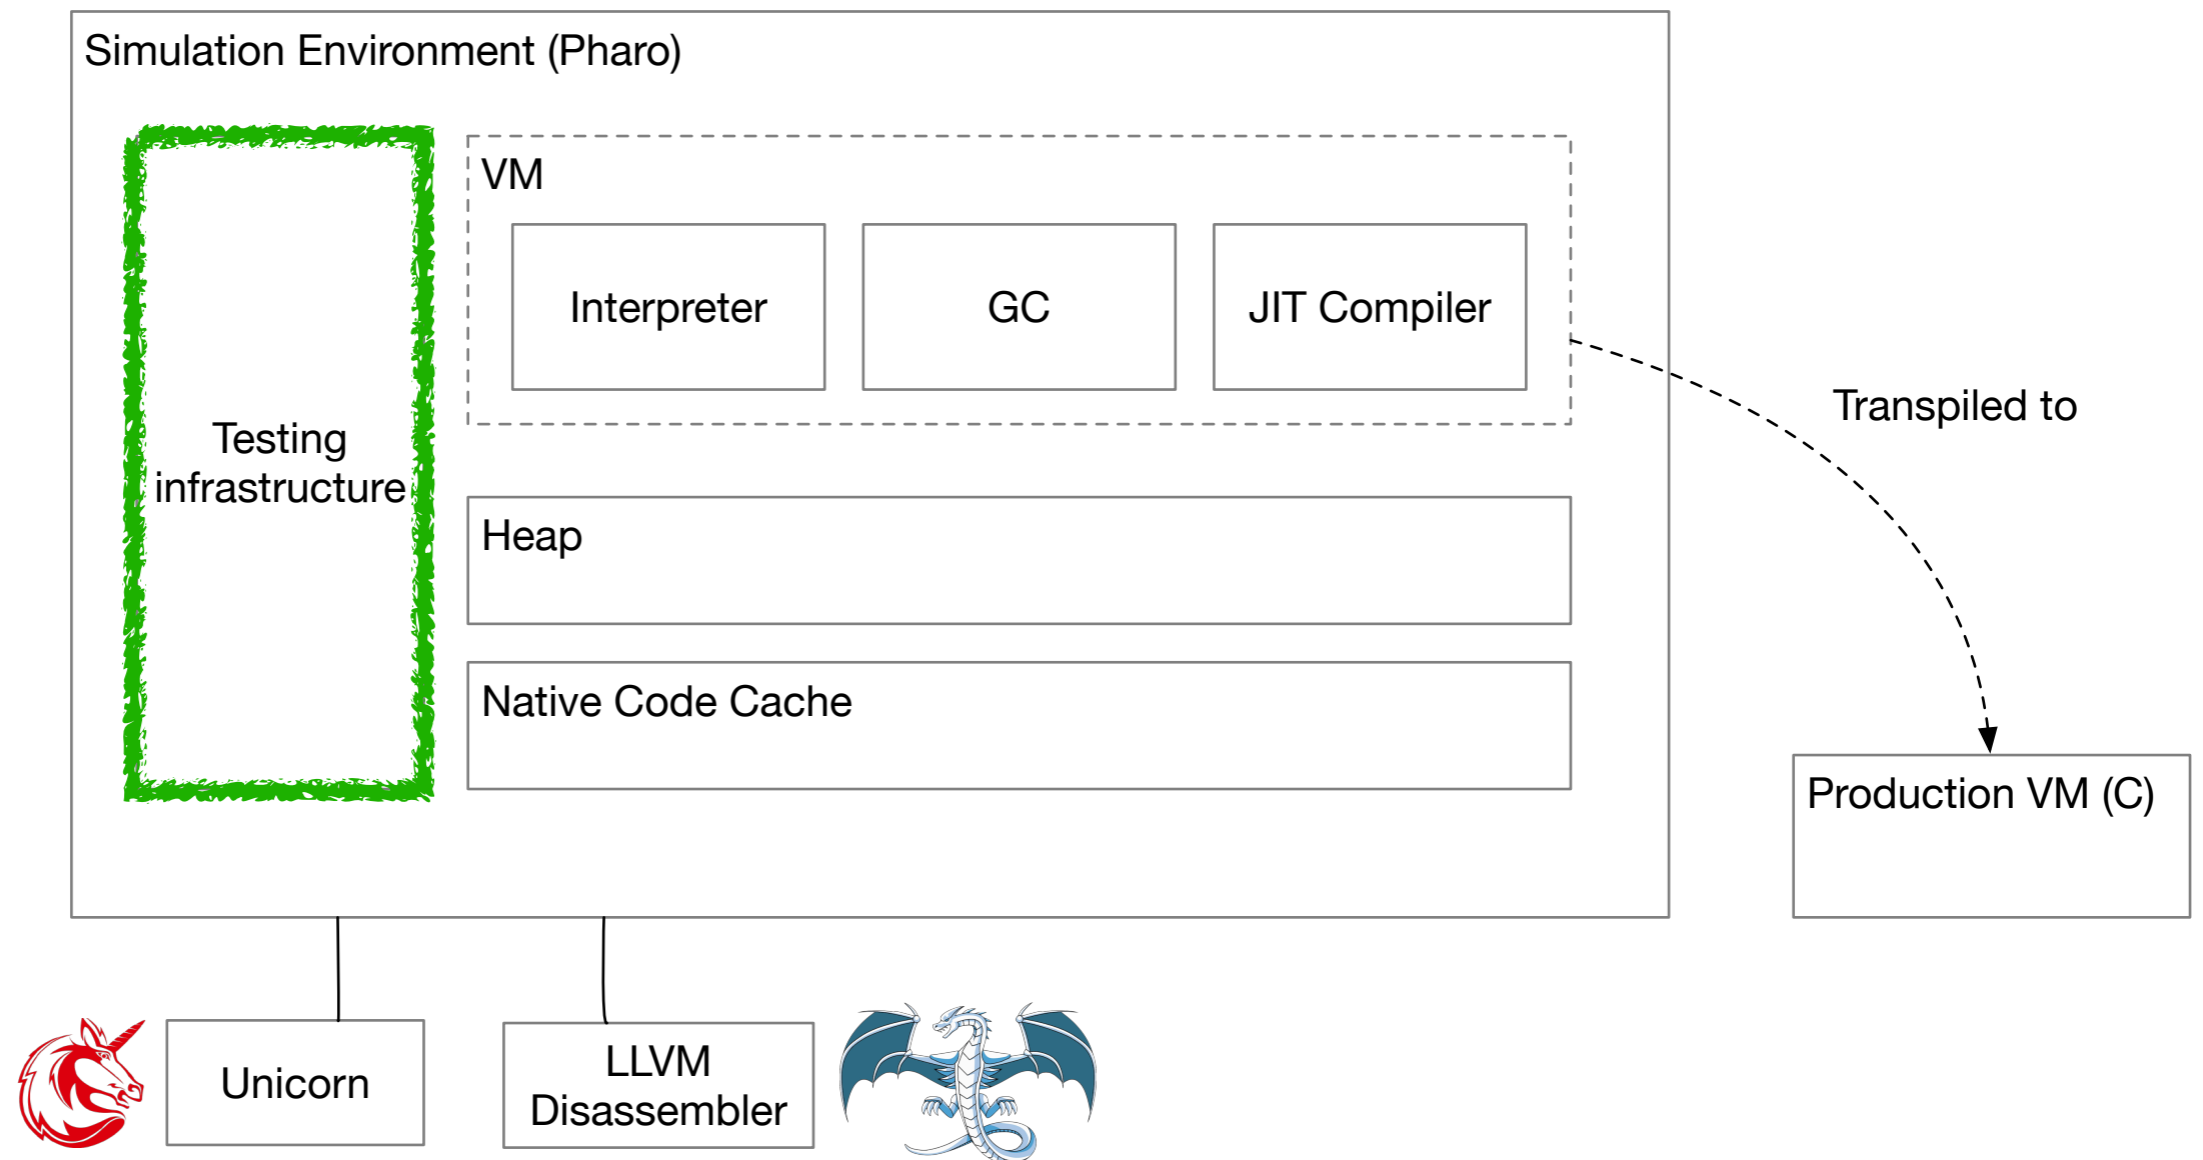
- How to do a **partial** implementation, in an iterative way?
- **Hardware availability**: did not have access to an Apple M1 yet
- **Slow** Change-Compile-Test **cycle**
- **Bug reproduction** is a demanding task



# Simulation Environment



# Extending Simulation with Unit Tests




# Our testing infrastructure by example

## testPushConstantZeroBytecodePushesASmallIntegerZero

```
self compile: [ compiler genPushConstantZeroBytecode ].  
self runGeneratedCode.
```

```
self assert: self popAddress equals: (memory integerObjectOf: 0)
```




# Our testing infrastructure by example

Reusable test fixtures covering e.g.,  
- trampoline and stub compilation  
- heap initialization

testPushConstantZeroBytecodePushesASmallIntegerZero

```
self compile: [ compiler genPushConstantZeroBytecode ].  
self runGeneratedCode.
```

```
self assert: self popAddress equals: (memory integerObjectOf: 0)
```



# Our testing infrastructure by example


Reusable test fixtures covering e.g.,  
- trampoline and stub compilation  
- heap initialization

testPushConstantZeroBytecodePushesASmallIntegerZero

Compiler internal  
DSL

```
self compile: [ compiler genPushConstantZeroBytecode  
self runGeneratedCode.
```

```
self assert: self popAddress equals: (memory integerObjectOf: 0)
```



# Our testing infrastructure by example

Reusable test fixtures covering e.g.,  
- trampoline and stub compilation  
- heap initialization


testPushConstantZeroBytecodePushesASmallIntegerZero

Compiler internal  
DSL

```
self compile: [ compiler genPushConstantZeroBytecode  
self runGeneratedCode.
```

```
self assert: self popAddress equals: (memory integerObjectOf: 0)
```

JIT Execution helpers such as e.g.,  
- run all code between two addresses  
- run until the PC hits an address



# Blackbox testing

testPushConstantZeroBytecodePushesASmallIntegerZero


```
self compile: [ compiler genPushConstantZeroBytecode ].  
self runGeneratedCode.
```

```
self assert: self popAddress equals: (memory integerObjectOf: 0)
```

**Depend only on  
observable  
behaviour**

**Reusable on  
different  
backends /**

**Resistant to  
changes in the  
implementation**





# Cross-Compilation, Cross-Execution



<http://www.unicorn-engine.org>

testPushConstantZeroBytecodePushesASmallIntegerZero


```
self compile: [ compiler genPushConstantZeroBytecode ].
```

```
self runGeneratedCode.
```

```
self assert: self popAddress equals: (memory integerObjectOf: 0)
```


**Hardware  
independent**

**Parametrizable  
tests**



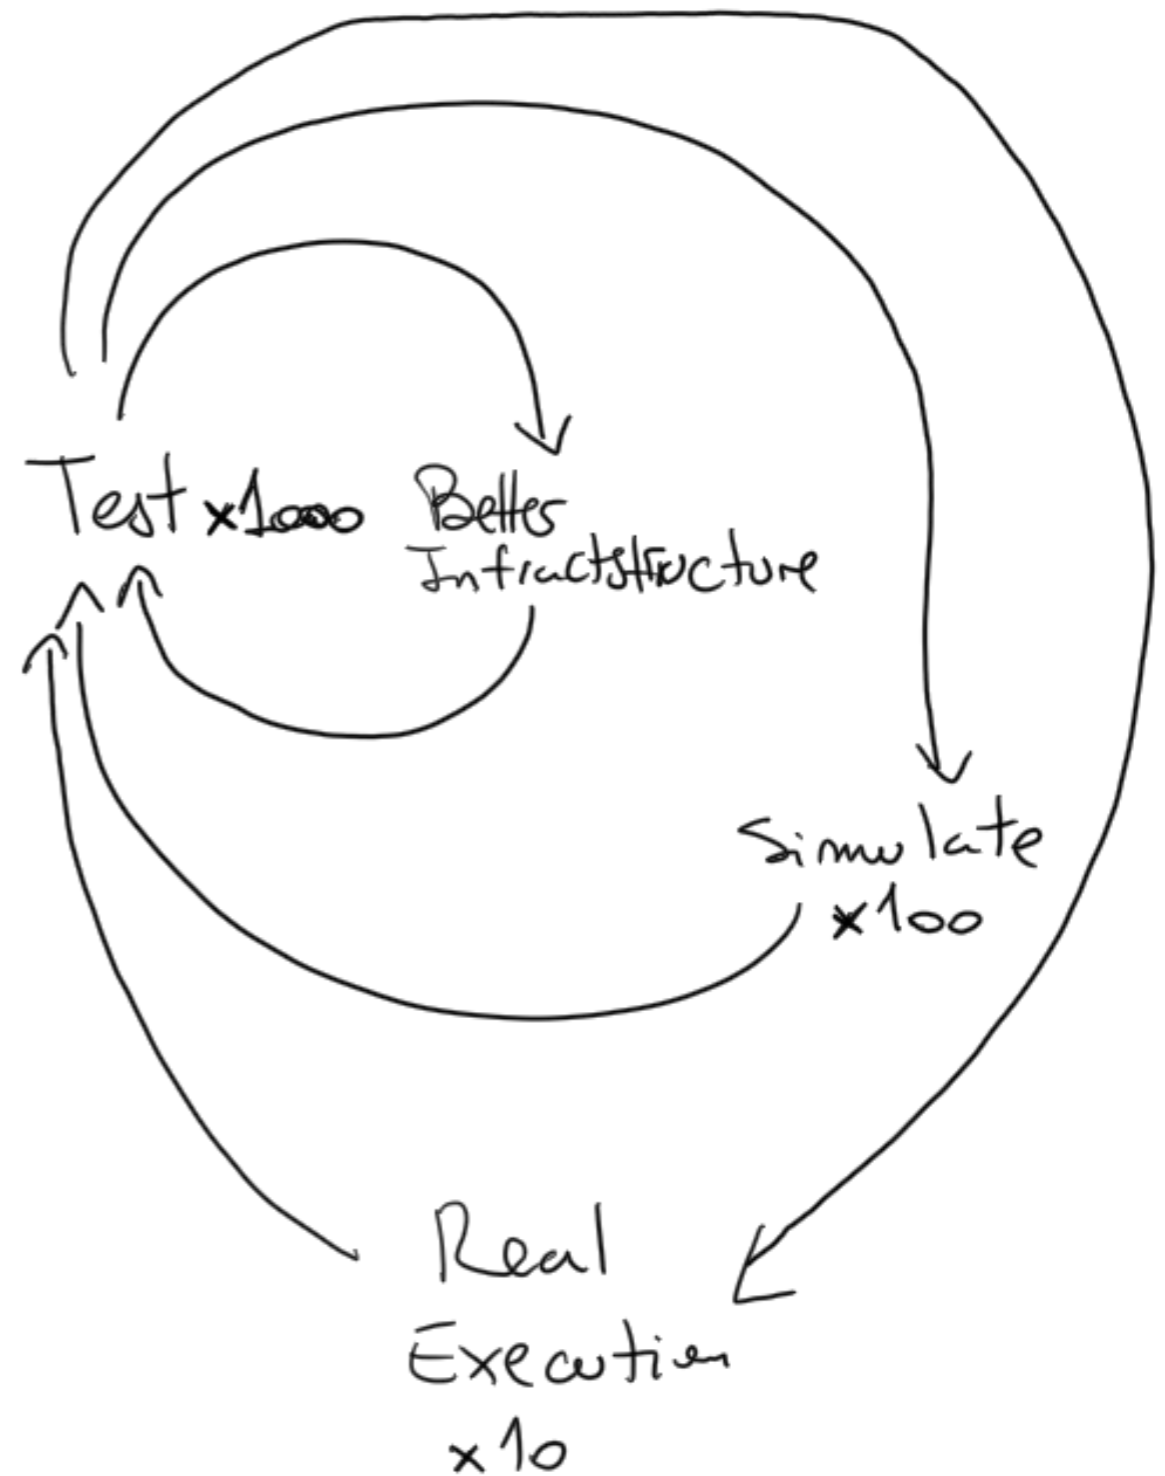
# There is no silver bullet

- Simulators are cheap, but not 100% trustworthy
- Full execution (simulated or on real HW)
  - more expensive to run
  - cannot unit-test it (less controllable)
- Unit tests only exercise specific scenarios
- Full executions exercise not yet covered scenarios



# Our testing Workflow

- Simulate the execution, less than you run tests
- Run the real app, less than you simulate
- Go back and forth:
  - Turn full execution failures into tests
  - **Fix with the aid of the test:**
    - => unit test are faster to run
    - => easier to debug
    - => detect regressions



# Testing & TDDing the VM

- No useful unit tests by ~06/2020
- Large manual testing effort during 2020 while porting to ARM64bits
  - Extended VM simulation with a (TDD compatible) unit testing infrastructure
  - **450+** written tests on the interpreter and the garbage collector\*
  - **580+** written tests on the JIT compiler\*
  - Parametrisable for 32 and 64bits, ARM32, ARM64, x86, x86-64

\* Numbers by 05/2021  
MPLR'21



# Testing & TDDing the VM


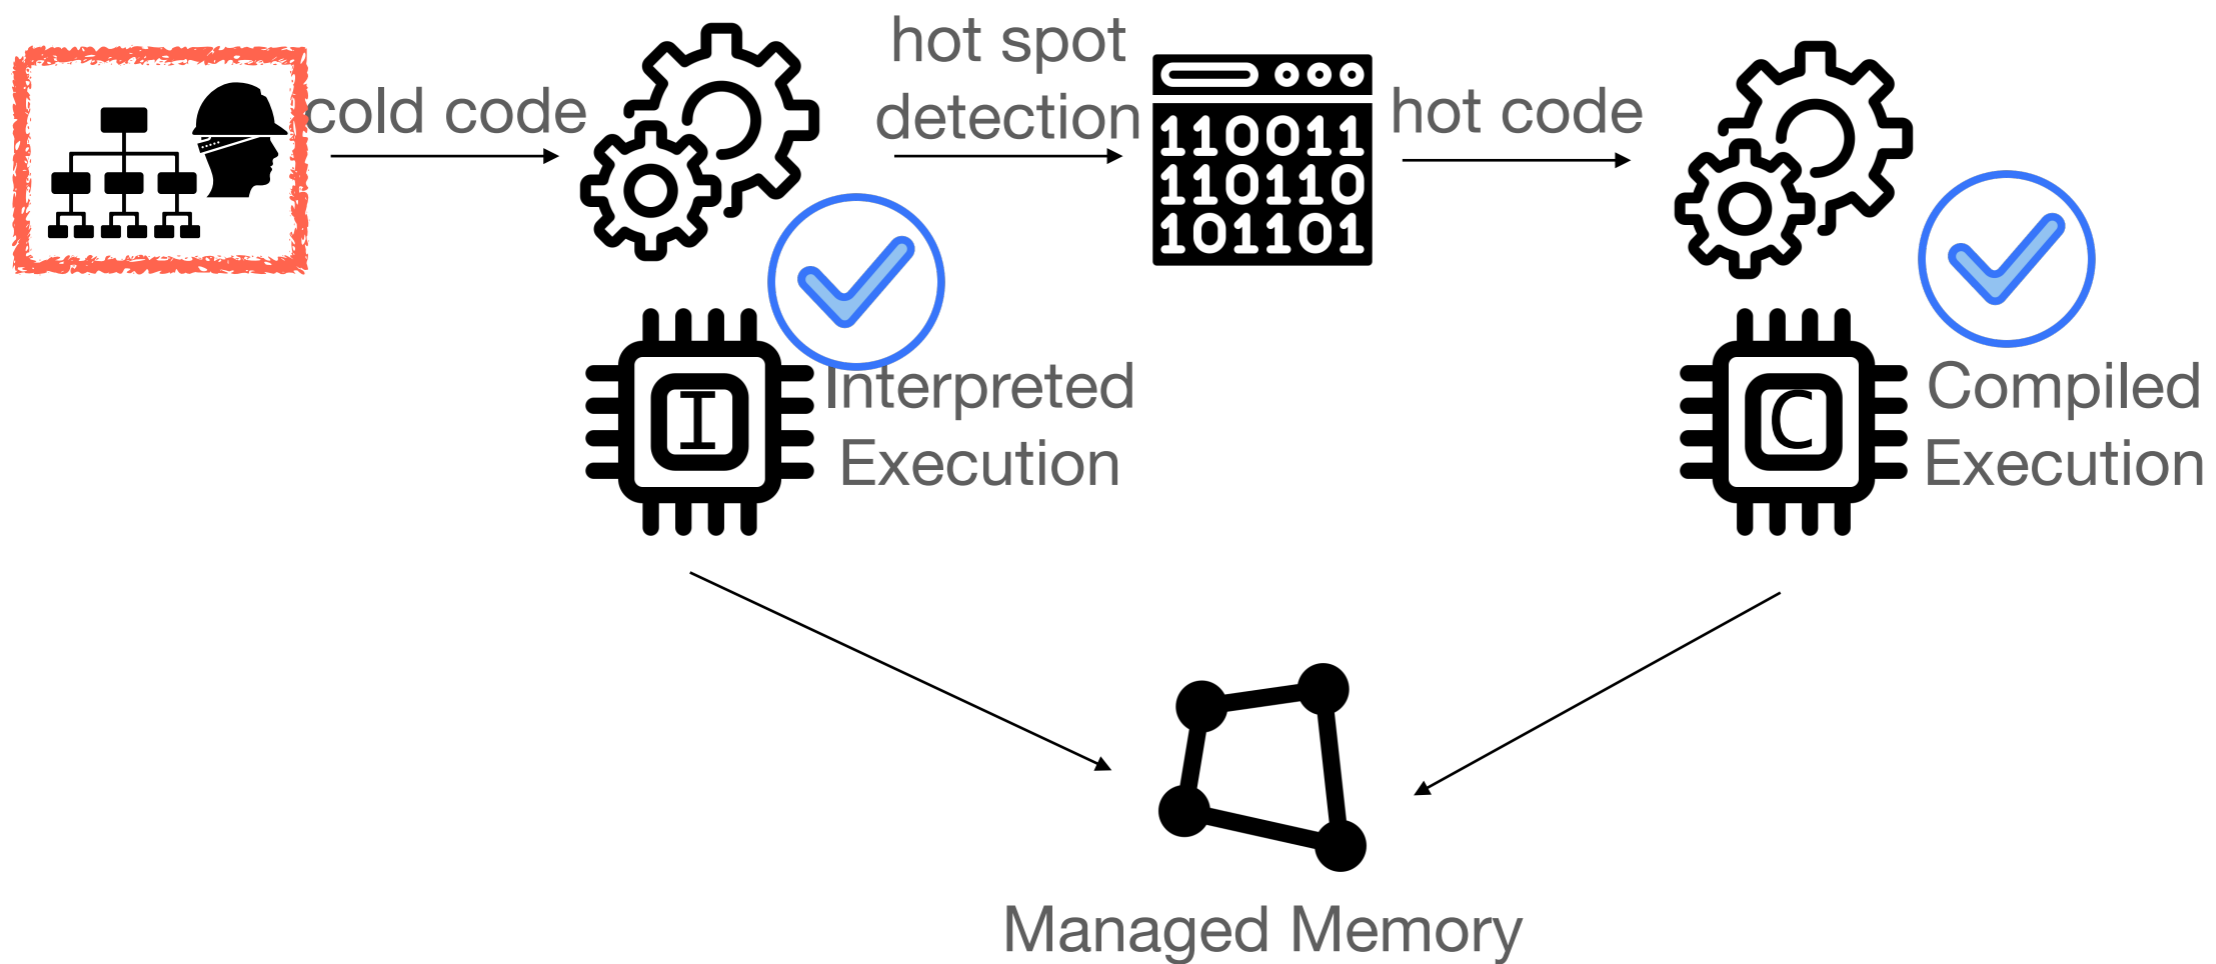
1040+ tests, are they enough?

- No useful unit tests by ~06/2020
- Large manual testing effort during 2020 while porting to ARM64bits
  - Extended VM simulation with a (TDD compatible) unit testing infrastructure
  - **450+** written tests on the interpreter and the garbage collector\*
  - **580+** written tests on the JIT compiler\*
  - Parametrisable for 32 and 64bits, ARM32, ARM64, x86, x86-64

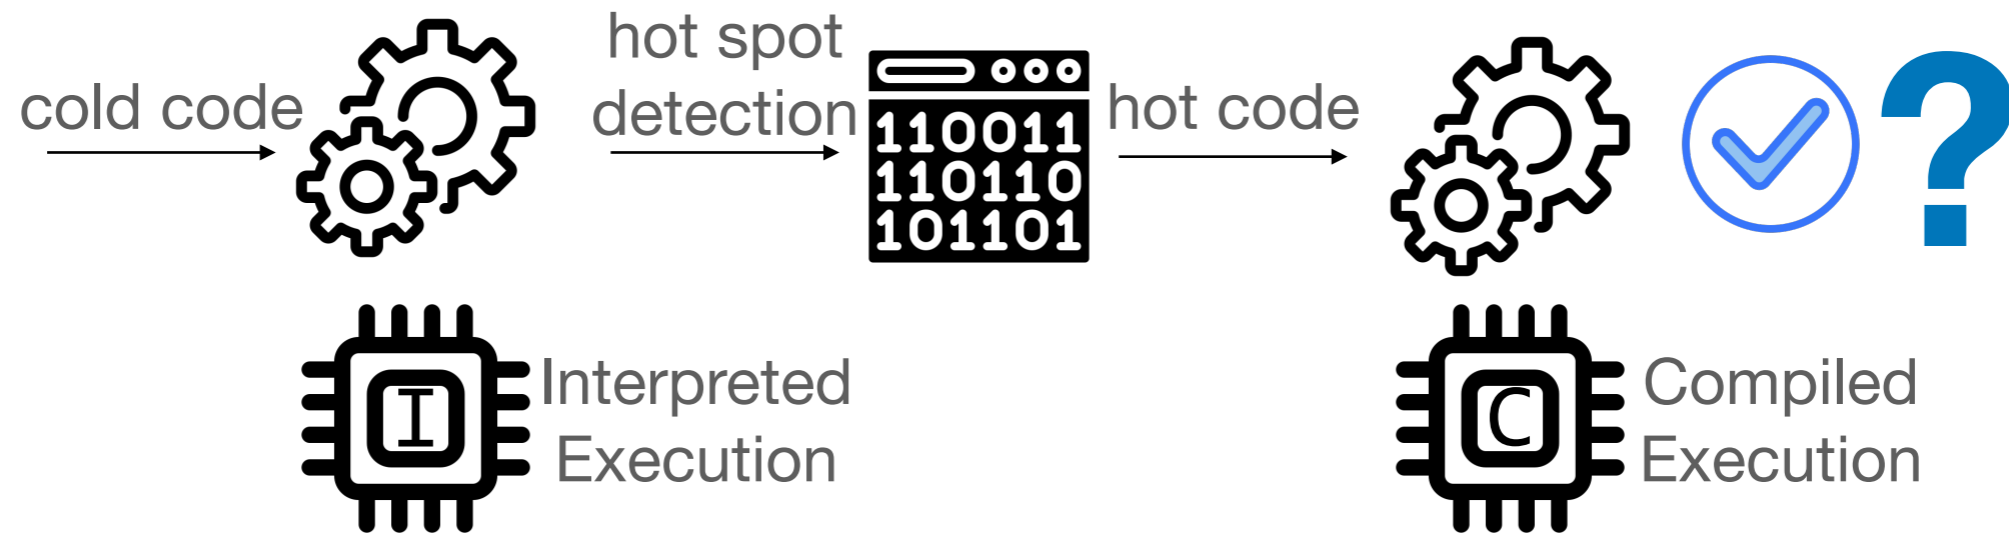
\* Numbers by 05/2021  
MPLR'21



# How can we automatically test VMs?



# Challenges of VM Test Generation




**Challenge 1: Test**

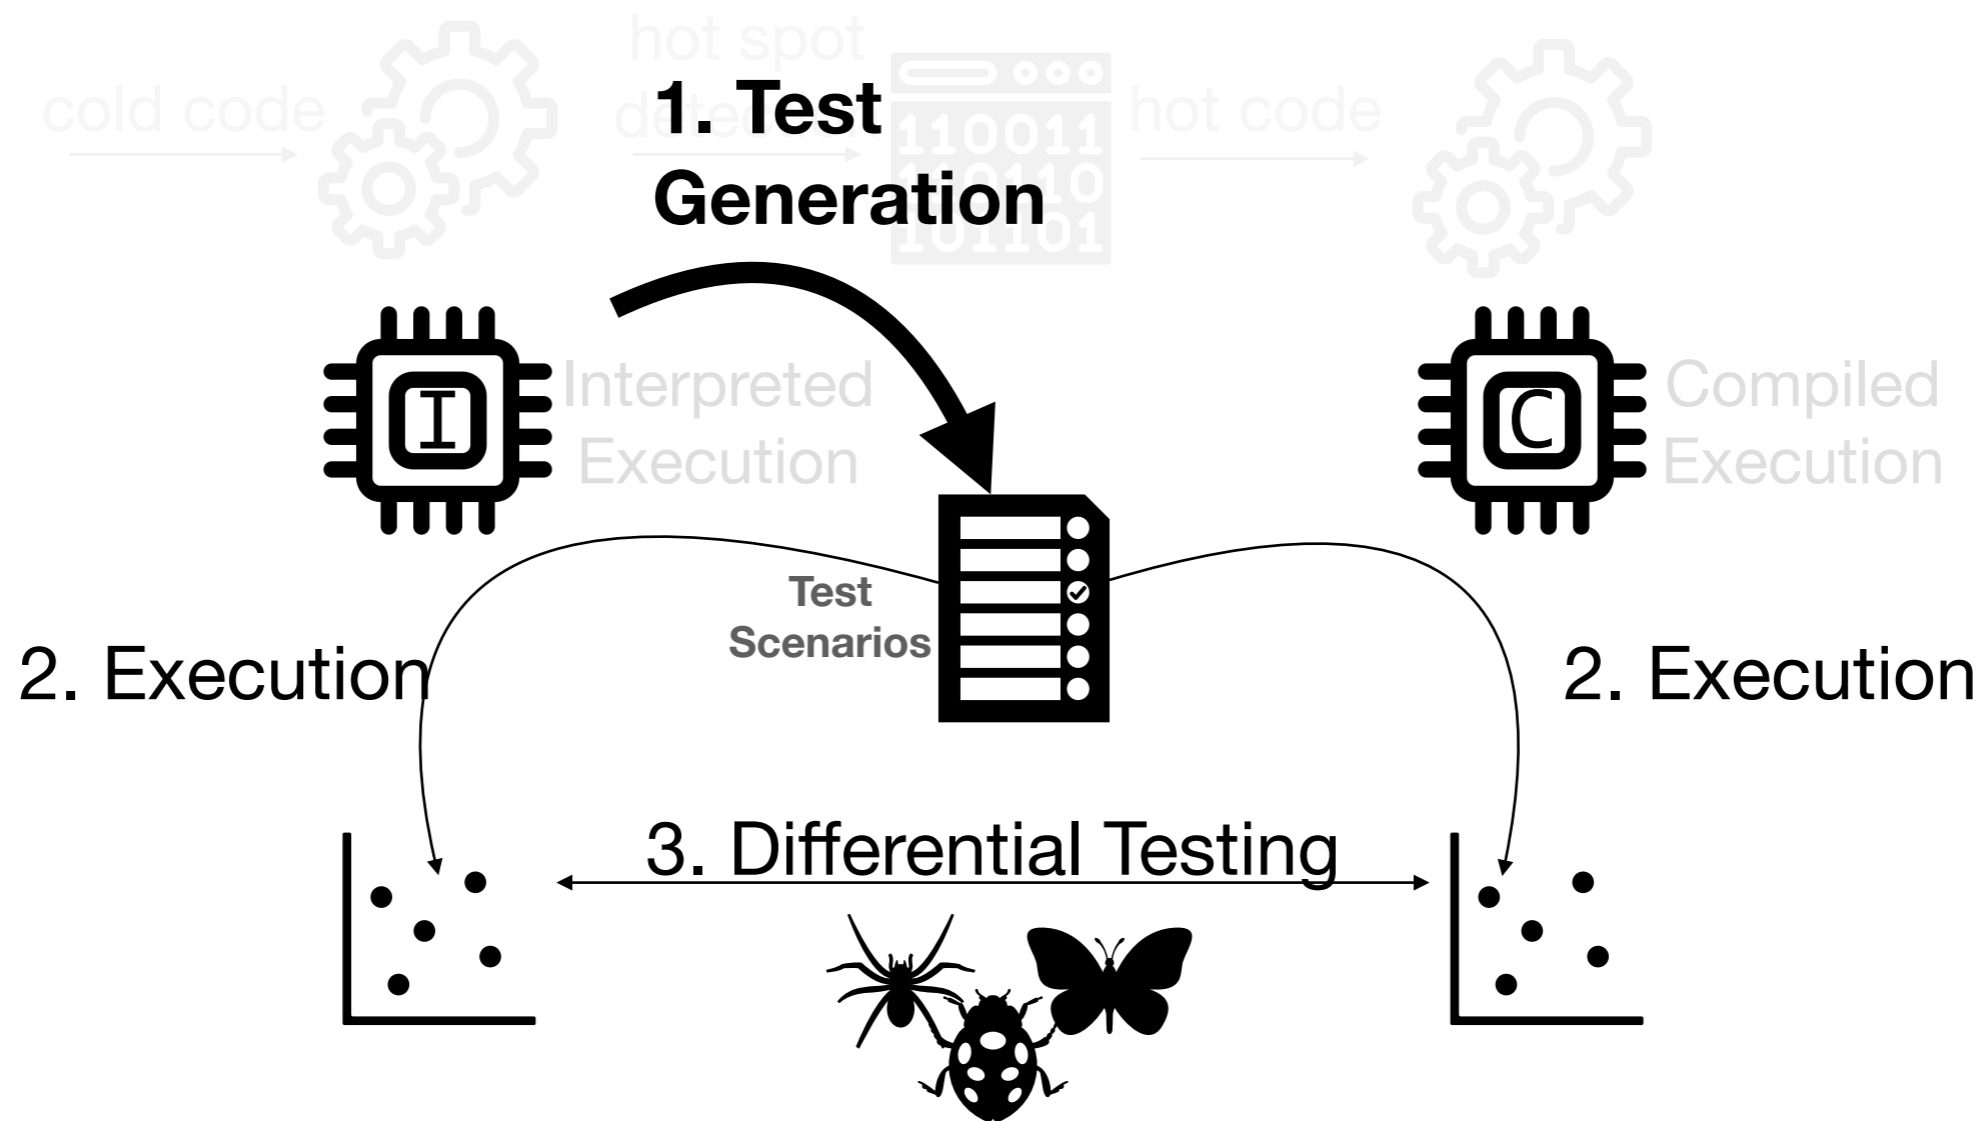
- Do they cover different code *regions/branches/paths*?

**Challenge 2: Test**

- How do we determine what is the *expected output* of a generated test?

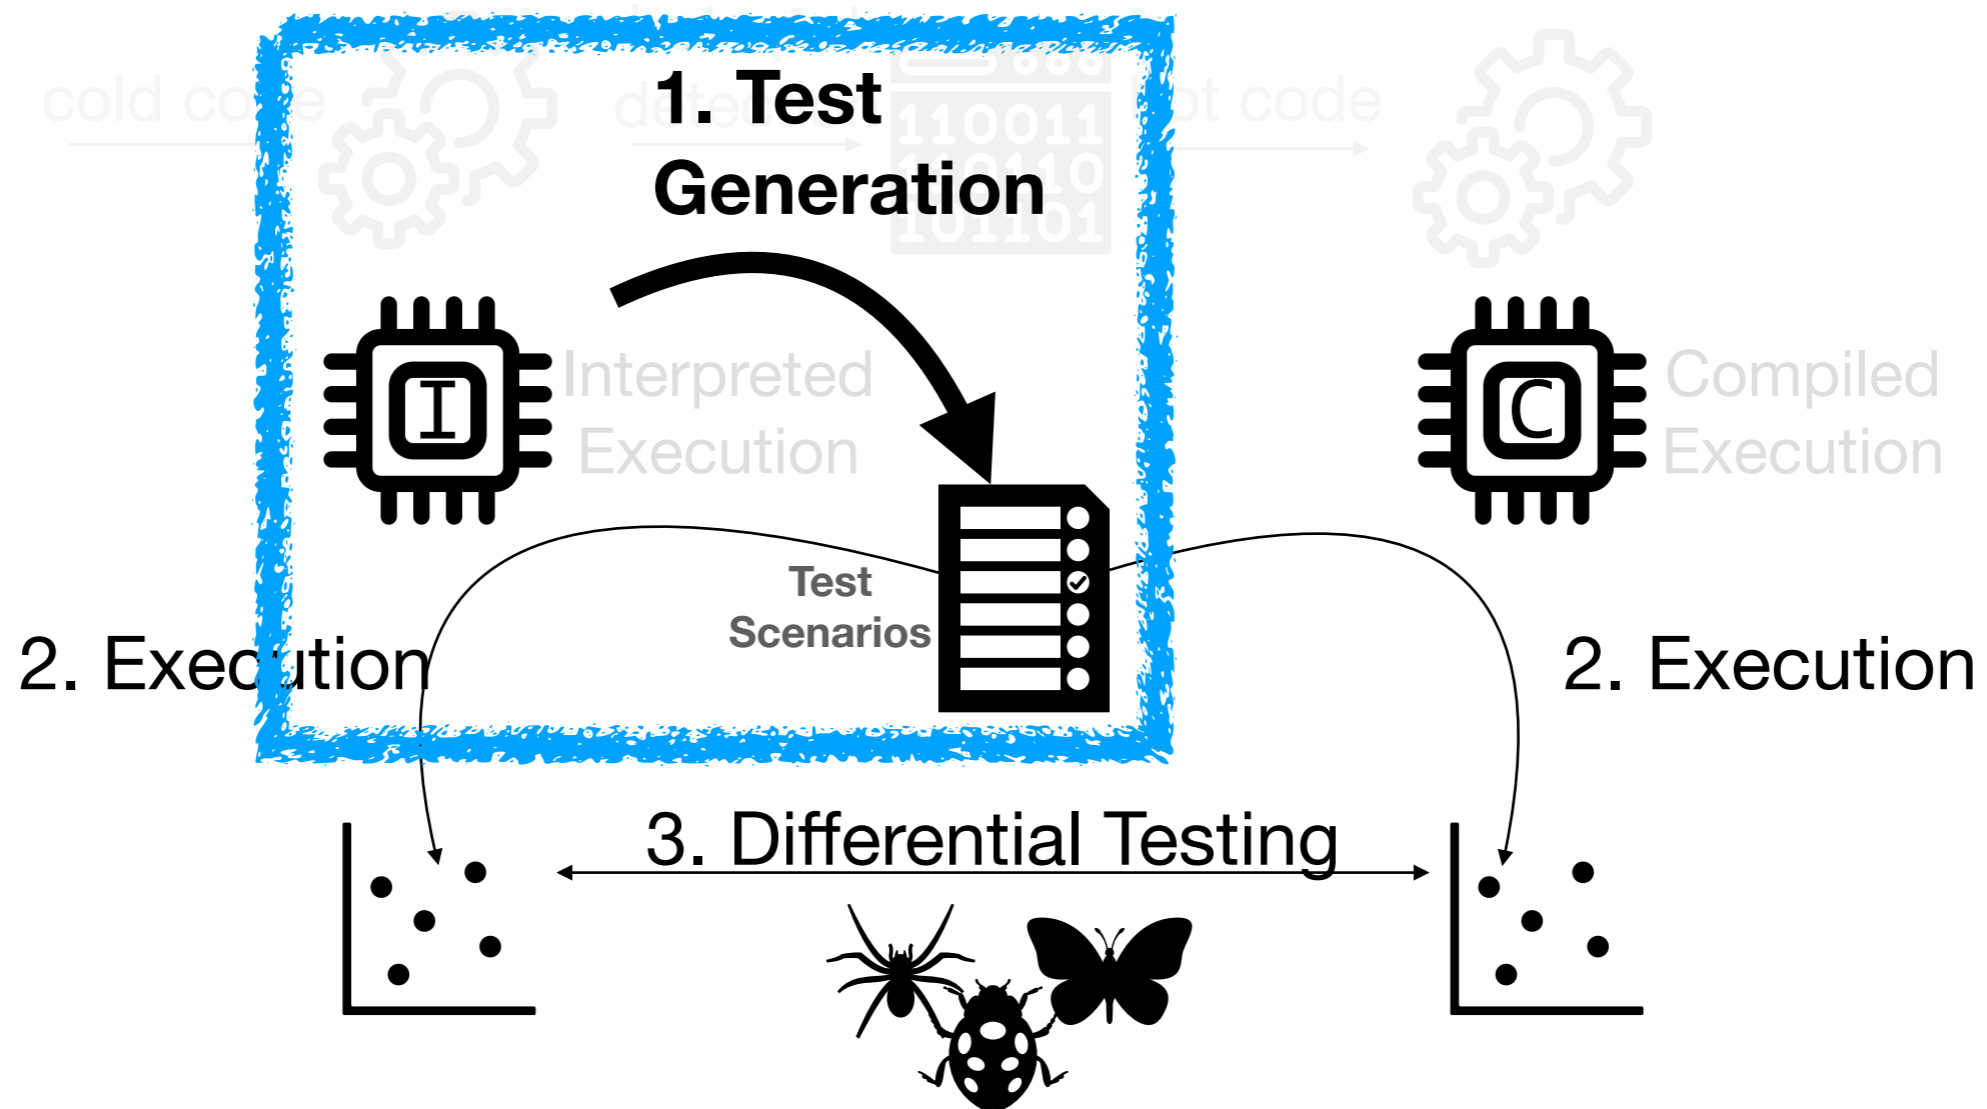


# Interpreter-Guided Automatic JIT Compiler Unit Testing


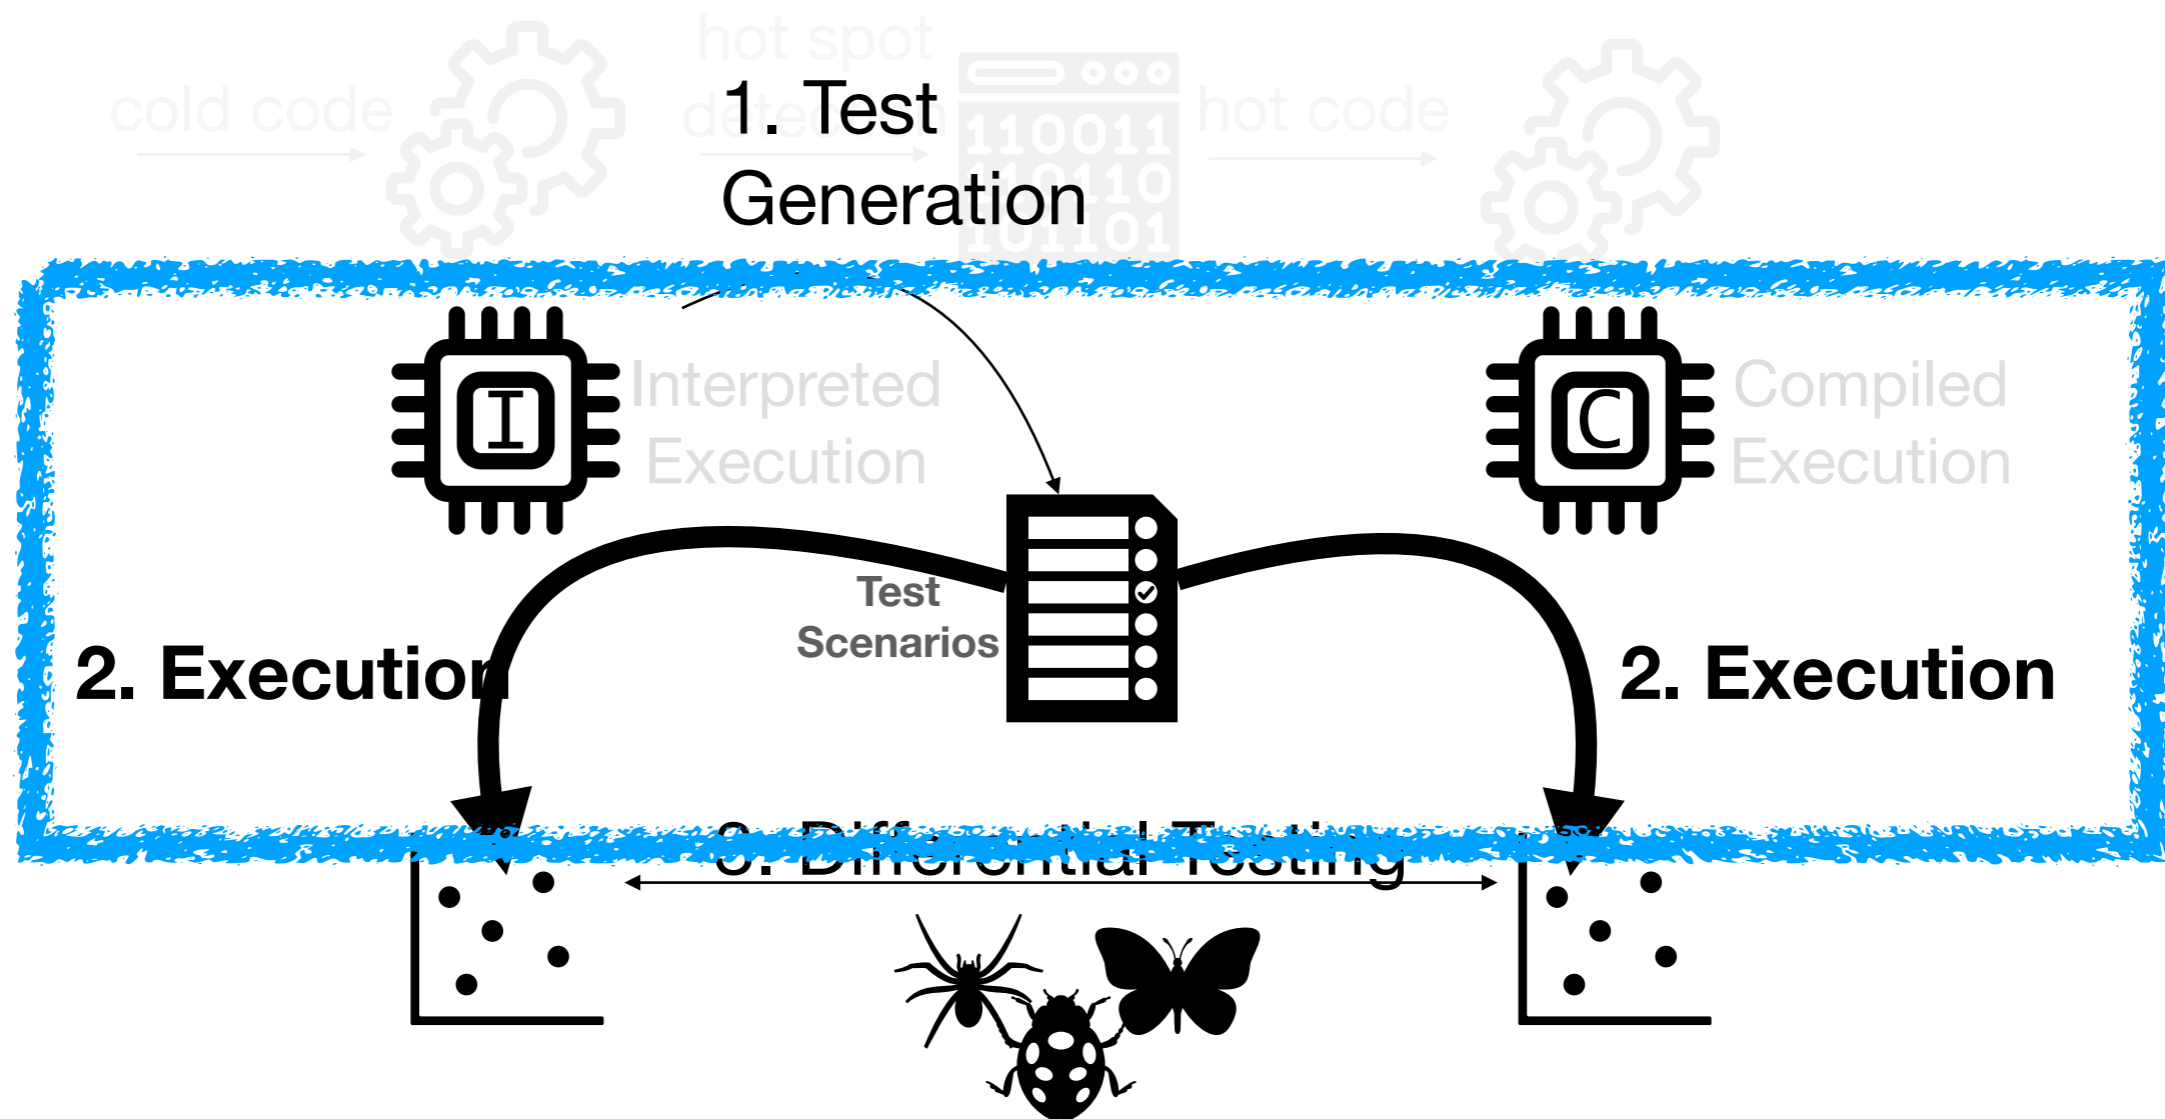




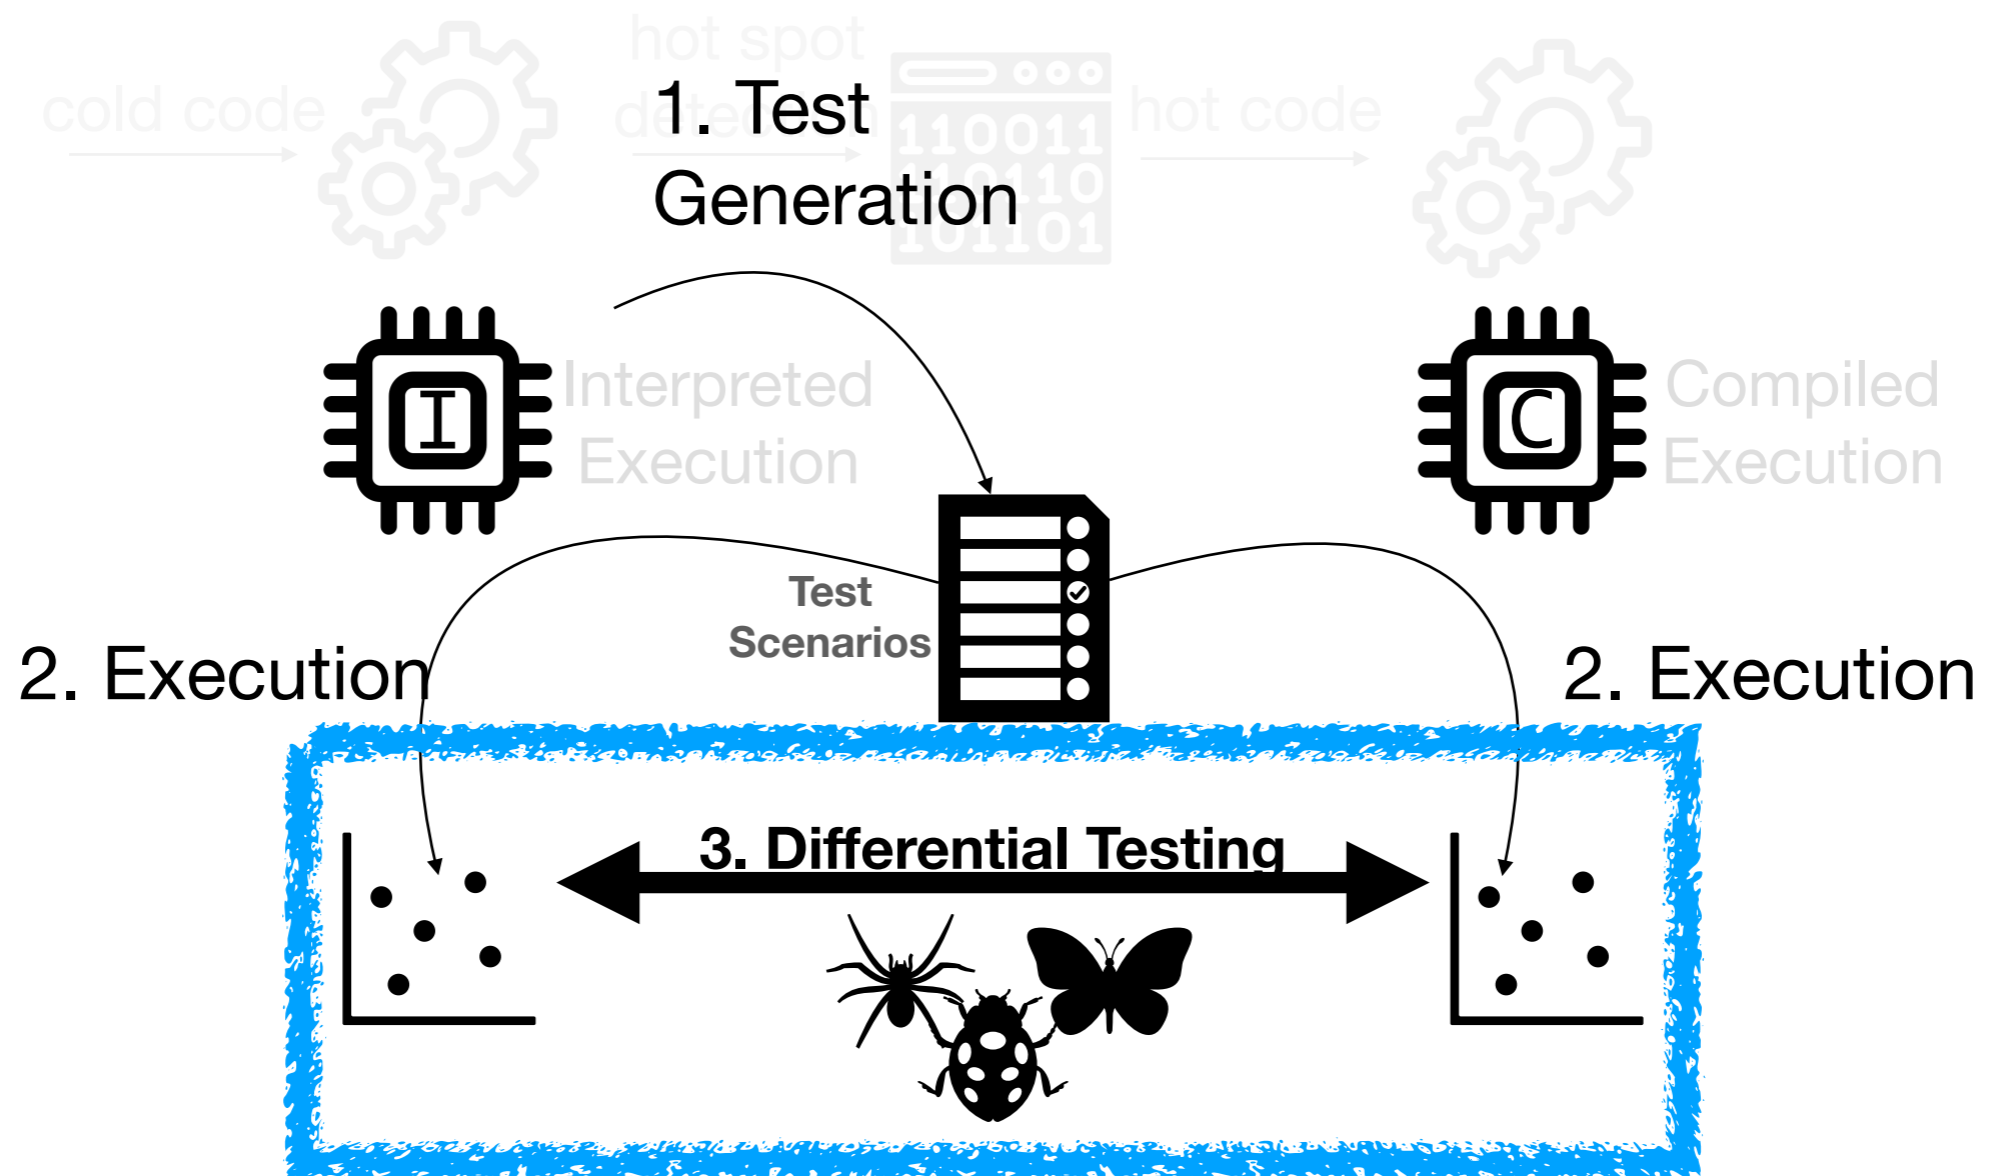
# Interpreter-Guided Automatic JIT Compiler Unit Testing



# Interpreter-Guided Automatic JIT Compiler Unit Testing




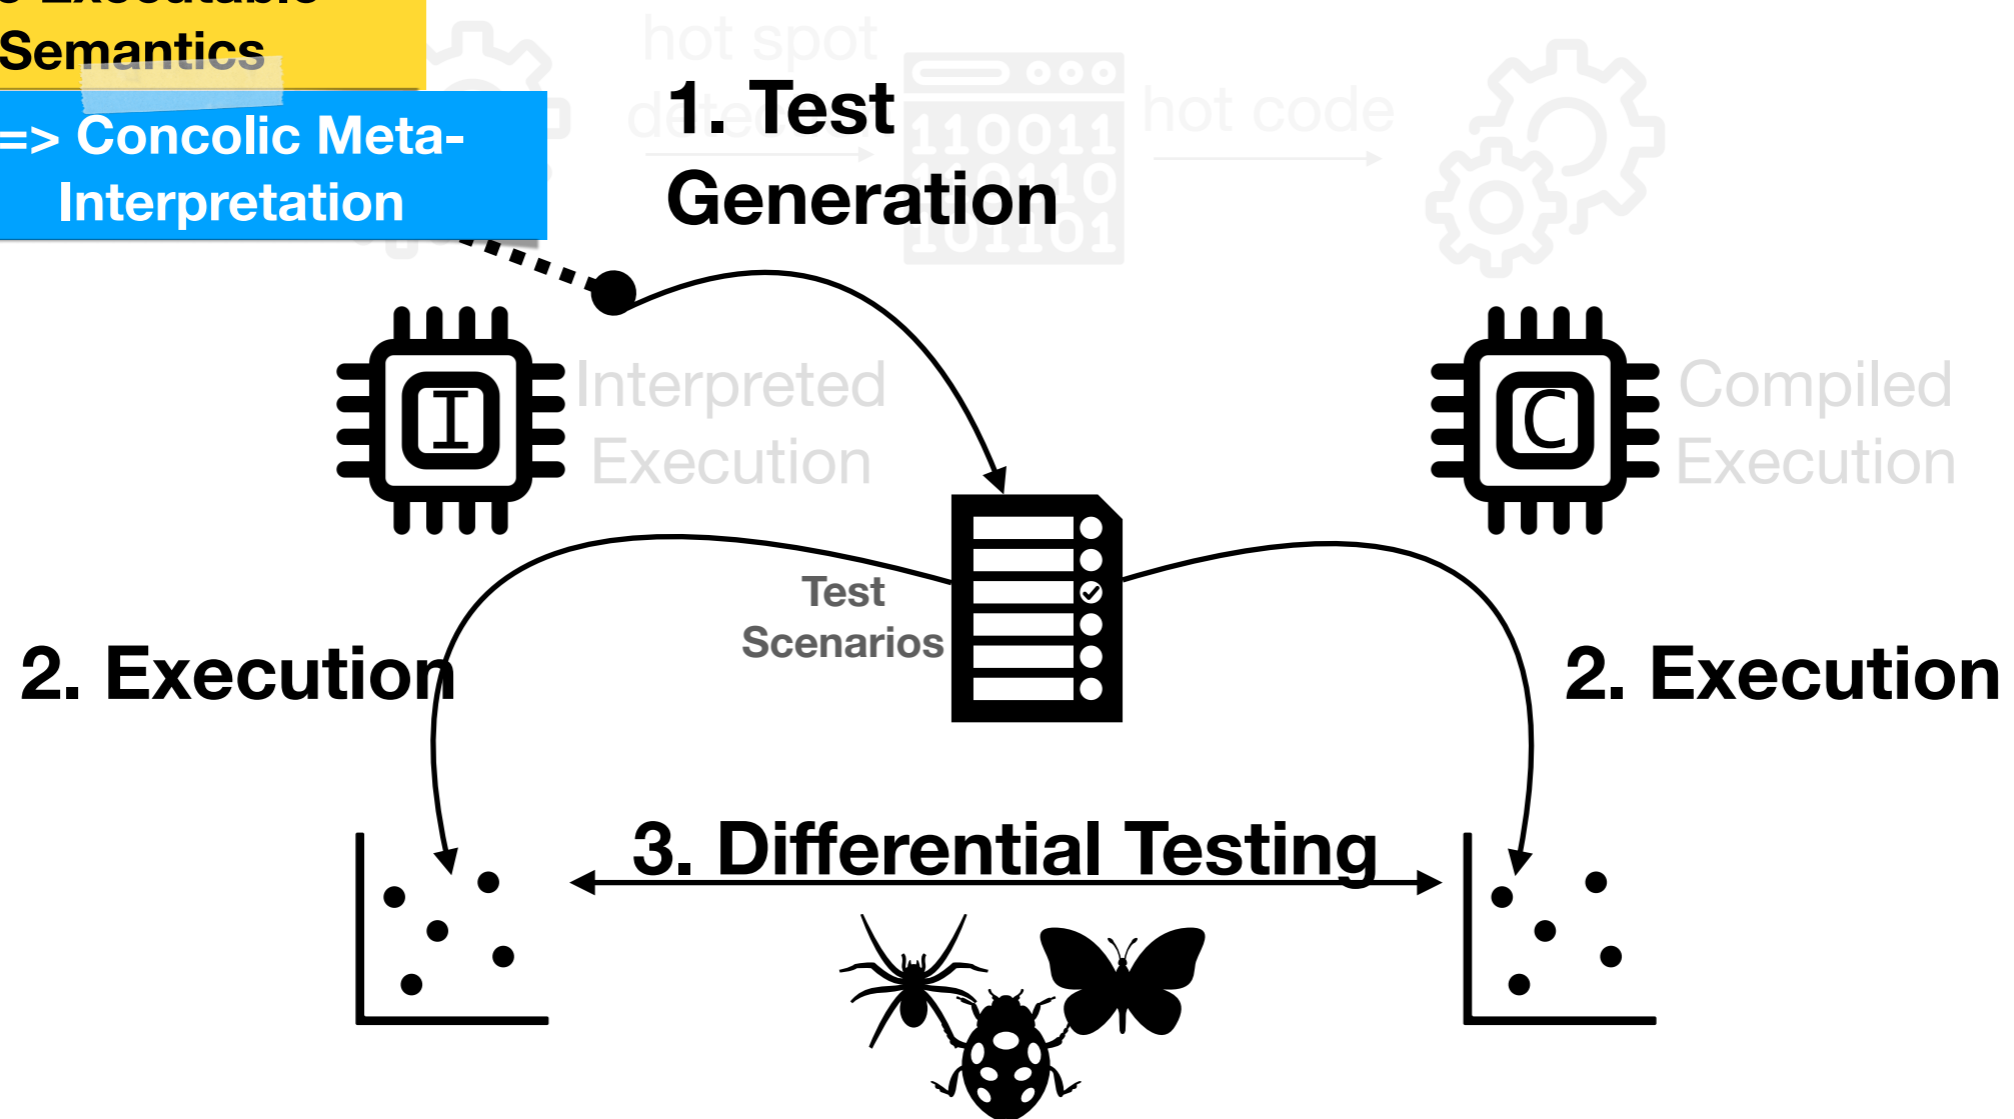
# Interpreter-Guided Automatic JIT Compiler Unit Testing



# Interpreter-Guided Automatic JIT Compiler Unit Testing

Insight 1: Interpreters are Executable Semantics


=> Concolic Meta-Interpretation



# Interpreter are Executable Semantics

## Pharo VM Example

```
1 Interpreter >> bytecodePrimAdd
2 | rcvr arg result |
3 rcvr := self internalStackValue: 1.
4 arg := self internalStackValue: 0.
5 (objectMemory areIntegers: rcvr and: arg) ifTrue: [
6   result := (objectMemory integerValueOf: rcvr) + (
7     objectMemory integerValueOf: arg).
8   "Check for overflow"
9   (objectMemory isIntegerValue: result) ifTrue: [
10    self
11    internalPop: 2
12    thenPush: (objectMemory integerObjectOf: result).
13    ^ self fetchNextBytecode "success"]].
14 "Slow path, message send"
15 self normalSend
```



# Interpreter are Executable Semantics

## Pharo VM Example

```
1 Interpreter >> bytecodePrimAdd
2 | rcvr arg result |
3 rcvr := self internalStackValue: 1.
4 arg := self internalStackValue: 0.
5 (objectMemory areIntegers: rcvr and: arg) ifTrue: [
6   result := (objectMemory integerValueOf: rcvr) + (
7     objectMemory integerValueOf: arg).
8   "Check for overflow"
9   (objectMemory isIntegerValue: result) ifTrue: [
10    self
11    internalPop: 2
12    thenPush: (objectMemory integerObjectOf: result).
13    A self fetchNextBytecode "success"]].
14 "Slow path, message send"
15 self normalSend
```


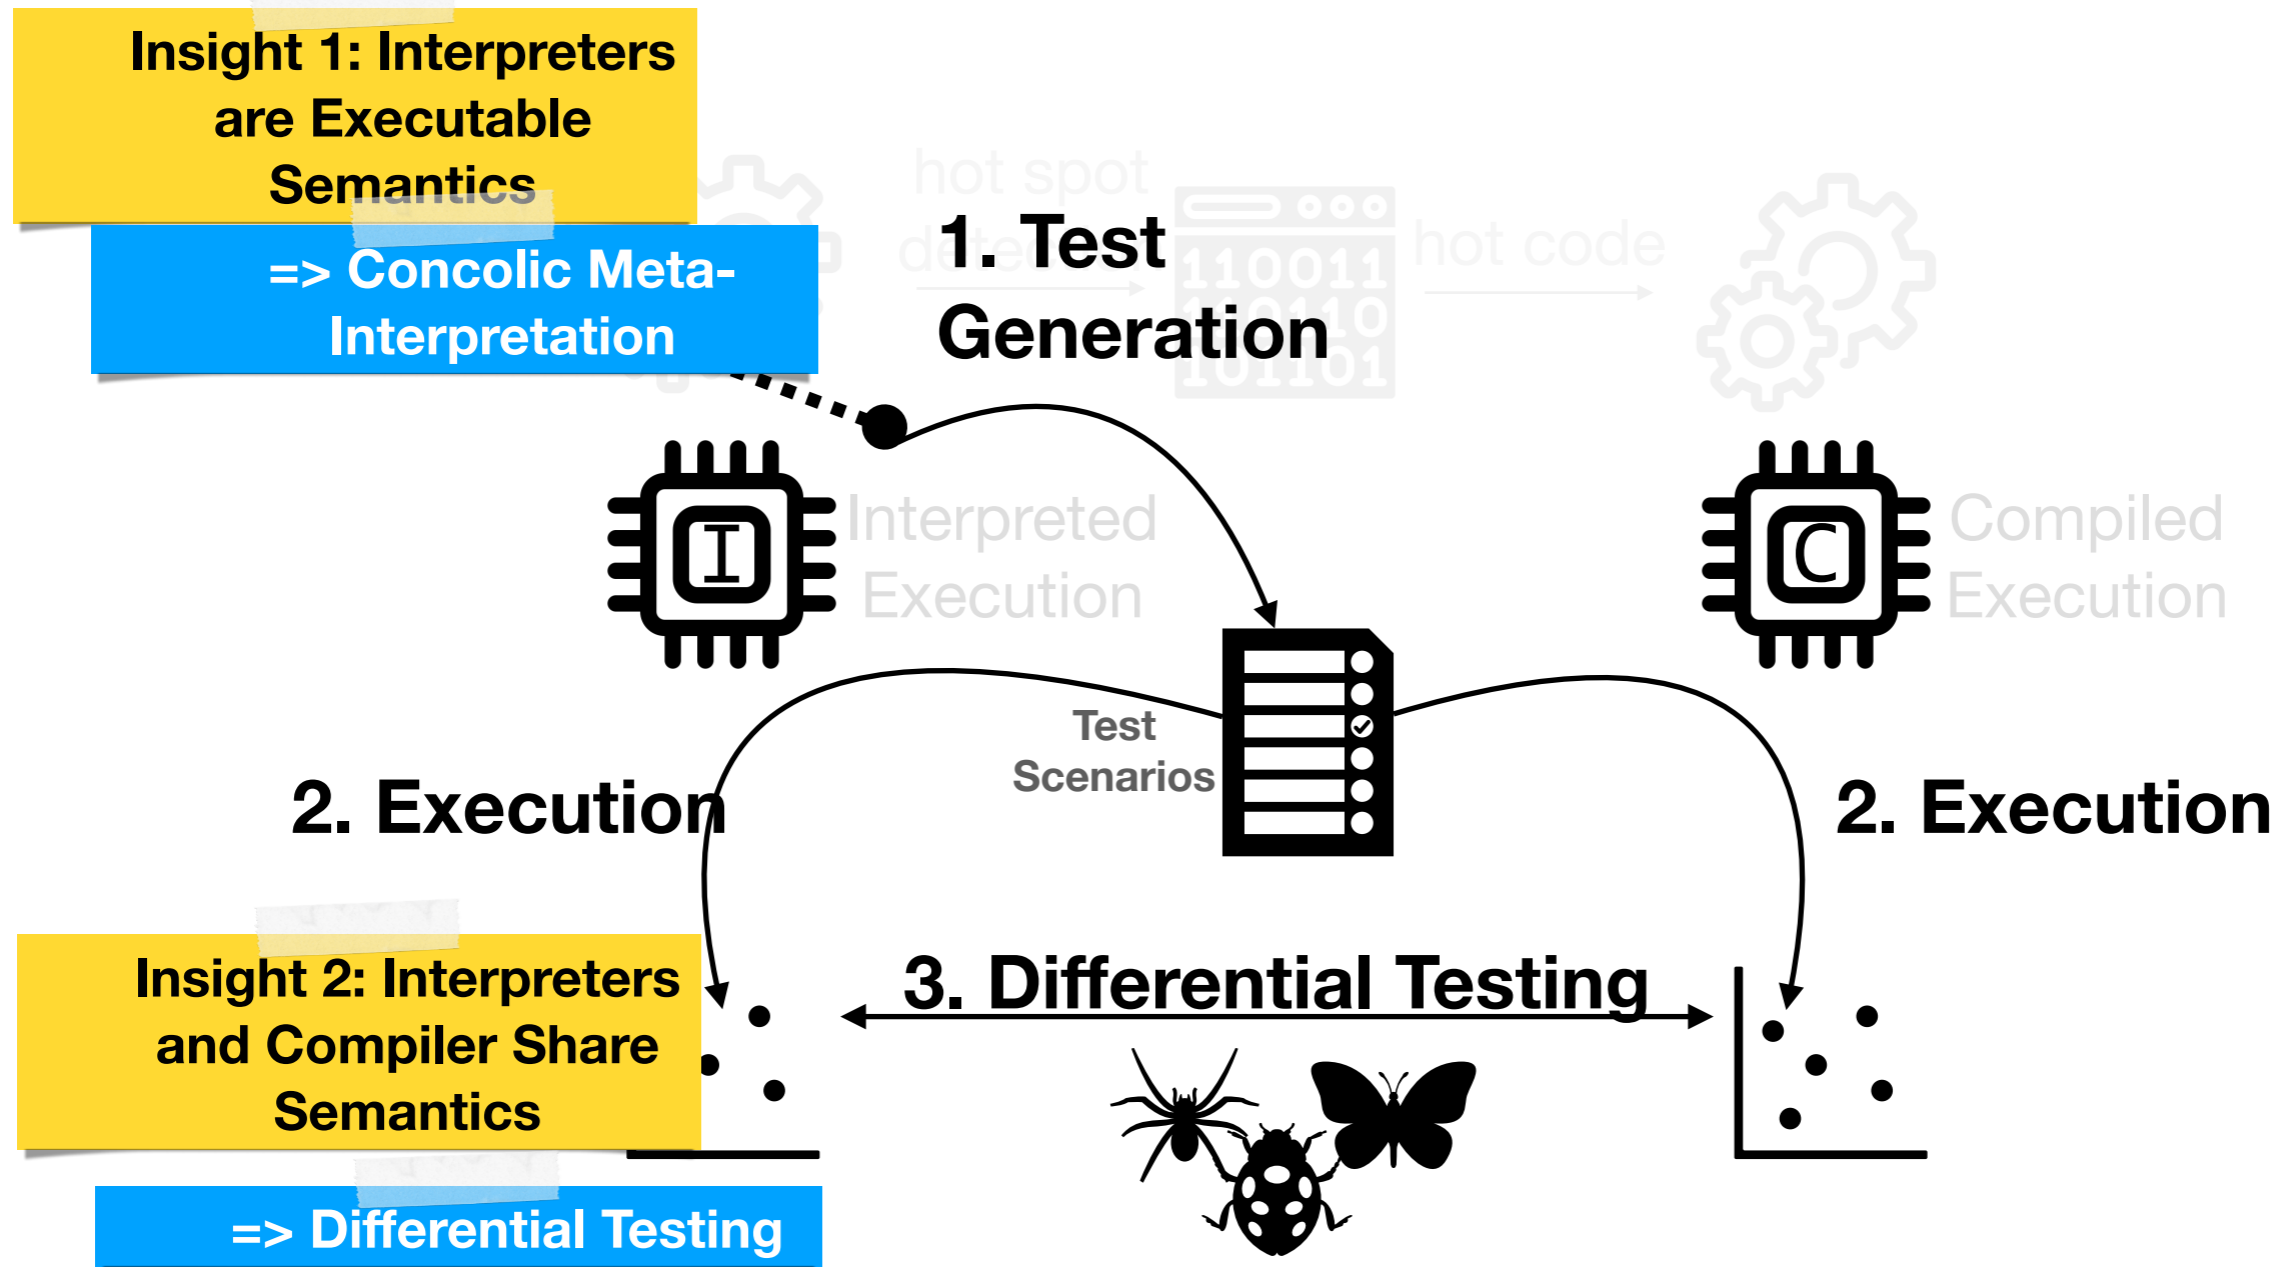
If both operands are integers

If their sum does not overflow

Else, slow path => message send



# Interpreter-Guided Automatic JIT Compiler Unit Testing




# Interpreter VS Compiled Code

## Pharo VM Example

```
1 Interpreter >> bytecodePrimAdd
2 | rcvr arg result |
3 rcvr := self internalStackValue: 1.
4 arg := self internalStackValue: 0.
5 (objectMemory areIntegers: rcvr and: arg) ifTrue: [
6   result := (objectMemory integerValueOf: rcvr) + (
7     objectMemory integerValueOf: arg).
8   "Check for overflow"
9   (objectMemory isIntegerValue: result) ifTrue: [
10    self
11    internalPop: 2
12    thenPush: (objectMemory integerObjectOf: result).
13    ^ self fetchNextBytecode "success" ].
14 "Slow path, message send"
15 self normalSend
```


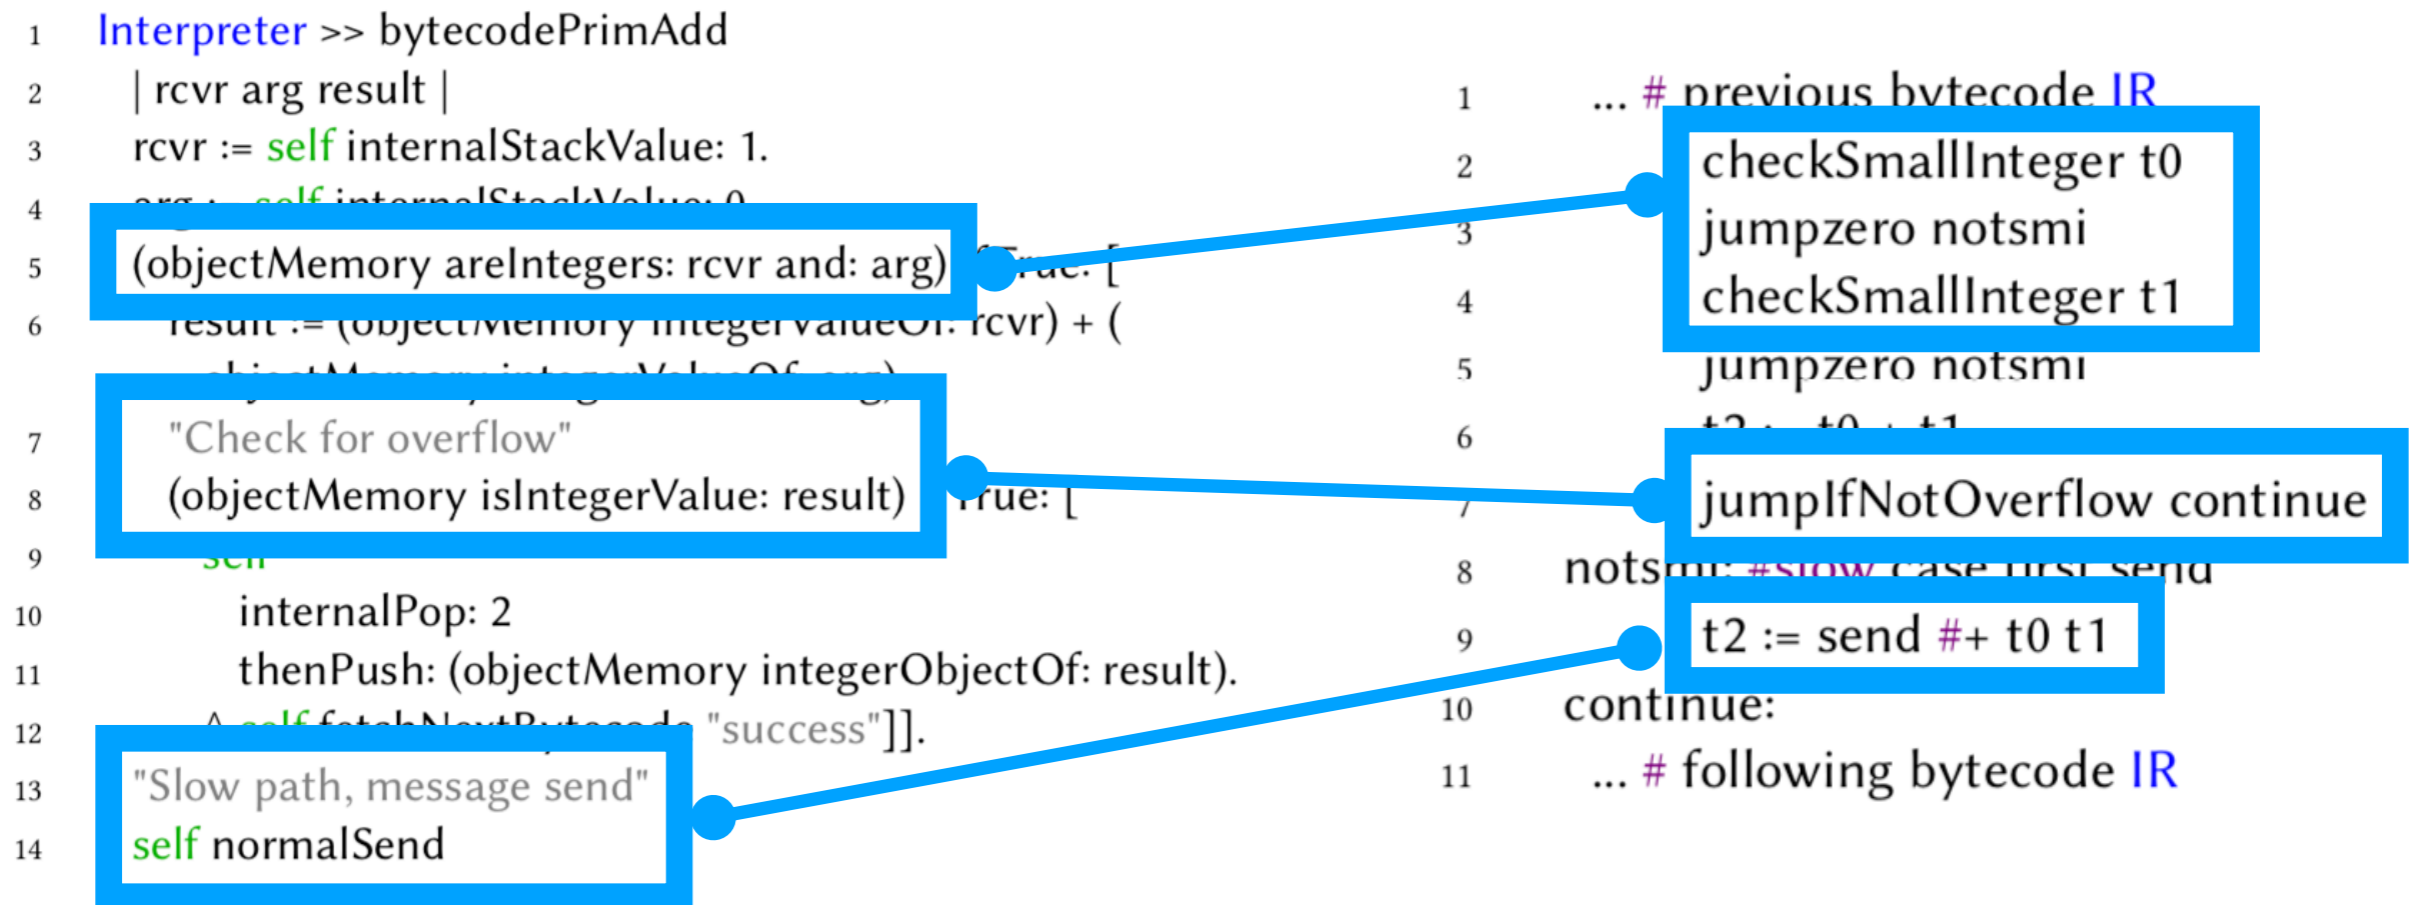
```
1 ... # previous bytecode IR
2   checkSmallInteger t0
3   jumpzero notsmi
4   checkSmallInteger t1
5   jumpzero notsmi
6   t2 := t0 + t1
7   jumpIfNotOverflow continue
8 notsmi: #slow case first send
9   t2 := send #+ t0 t1
10 continue:
11 ... # following bytecode IR
```





# Interpreter VS Compiled Code

## Pharo VM Example



# Concolic Testing through Meta-interpretation

- Idea: Guide test generation by looking at the implementation


```
int f(int x, int y){  
  if (x > 100){  
    if (y == 1023){  
      segfault(!!)  
    }  
  }  
}
```

Different cases  
if  $x > 100$  or  $\leq 100$ !!

Different cases  
if  $x = 1023$  or  $\neq 1023$

Godefroid et al. DART: Directed Automated Random Testing.  
PLDI'05

Set et al. CUTE: a concolic unit testing engine for C. FSE'05




# Concolic Testing by Example

- **Concrete + Symbolic** execution
- Goal: automatically discover *all* execution *paths*

```
int f(int x, int y){  
  if (x > 100){  
    if (y == 1023){  
      segfault(!!)  
    }  
  }  
}
```

x	y	constraints	next?

Godefroid et al. DART: Directed Automated Random Testing. PLDI' 05  
Set et al. CUTE: a concolic unit testing engine for C. FSE'05




# Concolic Testing by Example

- **Concrete + Symbolic** execution
- Goal: automatically discover *all* execution *paths*

```
int f(int x, int y){  
    if (x > 100){  
        if (y == 1023){  
            segfault(!!)  
        }  
    }  
}
```

x	y	constraints	next?
0	0	$x \leq 100$	

Godefroid et al. DART: Directed Automated Random Testing. PLDI' 05  
Set et al. CUTE: a concolic unit testing engine for C. FSE'05




# Concolic Testing by Example

- **Concrete + Symbolic** execution
- Goal: automatically discover *all* execution *paths*

```
int f(int x, int y){  
    if (x > 100){  
        if (y == 1023){  
            segfault(!!)  
        }  
    }  
}
```

x	y	constraints	next?
0	0	$x \leq 100$	$x > 100$

Godefroid et al. DART: Directed Automated Random Testing. PLDI' 05  
Set et al. CUTE: a concolic unit testing engine for C. FSE'05




# Concolic Testing by Example

- **Concrete + Symbolic** execution
- Goal: automatically discover *all* execution *paths*

```
int f(int x, int y){  
    if (x > 100){  
        if (y == 1023){  
            segfault(!!)  
        }  
    }  
}
```

x	y	constraints	next?
0	0	$x \leq 100$	$x > 100$
101	0		

Godefroid et al. DART: Directed Automated Random Testing. PLDI' 05  
Set et al. CUTE: a concolic unit testing engine for C. FSE'05




# Concolic Testing by Example

- **Concrete + Symbolic** execution
- Goal: automatically discover *all* execution *paths*

```
int f(int x, int y){  
    if (x > 100){  
        if (y == 1023){  
            segfault(!!)  
        }  
    }  
}
```

x	y	constraints	next?
0	0	$x \leq 100$	$x > 100$
101	0	$x > 100, y \neq 1023$	

Godefroid et al. DART: Directed Automated Random Testing. PLDI' 05  
Set et al. CUTE: a concolic unit testing engine for C. FSE'05



# Concolic Testing by Example

- **Concrete + Symbolic** execution
- Goal: automatically discover *all* execution *paths*

```
int f(int x, int y){  
    if (x > 100){  
        if (y == 1023){  
            segfault(!!)  
        }  
    }  
}
```

x	y	constraints	next?
0	0	$x \leq 100$	$x > 100$
101	0	$x > 100, y \neq 1023$	$x > 100, y == 1023$

Godefroid et al. DART: Directed Automated Random Testing. PLDI' 05  
Set et al. CUTE: a concolic unit testing engine for C. FSE'05






# Concolic Testing by Example

- **Concrete + Symbolic** execution
- Goal: automatically discover *all* execution *paths*

```
int f(int x, int y){  
    if (x > 100){  
        if (y == 1023){  
            segfault(!!)  
        }  
    }  
}
```

x	y	constraints	next?
0	0	$x \leq 100$	$x > 100$
101	0	$x > 100, y \neq 1023$	$x > 100, y == 1023$
101	1023		

Godefroid et al. DART: Directed Automated Random Testing. PLDI' 05  
Set et al. CUTE: a concolic unit testing engine for C. FSE'05




# Concolic Testing by Example

- **Concrete + Symbolic** execution
- Goal: automatically discover *all* execution *paths*

```
int f(int x, int y){  
    if (x > 100){  
        if (y == 1023){  
            segfault(!!)  
        }  
    }  
}
```

x	y	constraints	next?
0	0	$x \leq 100$	$x > 100$
101	0	$x > 100, y \neq 1023$	$x > 100, y == 1023$
101	1023	$x > 100, y \neq 1023$	finished!

Godefroid et al. DART: Directed Automated Random Testing. PLDI' 05  
Set et al. CUTE: a concolic unit testing engine for C. FSE'05



# Some Numbers

- 3 bytecode compilers + 1 native method compiler
- 4928 tests generated
- **478 differences**


Compiler	# Tested Instructions	# Interpreter Paths	# Curated Paths	# Differences (%)
Native Methods (primitives)	112	2024	1520	440 (28,95%)
Simple Stack BC Compiler	175	1308	1136	18 (1,59%)
Stack-to-Register BC Compiler	175	1308	1136	10 (0,88%)
Linear-Scan Allocator BC Compiler	175	1308	1136	10 (0,88%)
<b>Total</b>	<b>637</b>	<b>5948</b>	<b>4928</b>	<b>478 (9,7%)</b>



# Analysis of Differences through Manual Inspection


- 91 causes, *6 different categories*
- Errors both in the interpreter AND the compilers
- 14 causes of ***segmentation faults!***

<u>Family</u>	<u># Cases</u>
Missing interpreter type check	1
Missing compiled type check	13
Optimisation difference	10
Behavioral difference	5
Missing Functionality	60
Simulation Error	2



# Practical and Cheap

- Test generation ~5 minutes
- Total run time of ~10 seconds
  - Avg 30ms per instruction



# More in the PLDI article!

- Discovered Bugs
- Concolic Model
- Testing Infrastructure

PLDI'22

## Interpreter-Guided Differential JIT Compiler Unit Testing

Guillermo Polito  
Univ. Lille, CNRS, Inria, Centrale Lille,  
UMR 9189 CRIStAL, F-59000 Lille  
France  
guillermo.polito@univ-lille.fr

Stéphane Ducasse  
Univ. Lille, Inria, CNRS, Centrale Lille,  
UMR 9189 CRIStAL  
France  
stephane.ducasse@inria.fr

Pablo Tesone  
Pharo Consortium  
Univ. Lille, Inria, CNRS, Centrale Lille,  
UMR 9189 CRIStAL  
France  
pablo.tesone@inria.fr

### Abstract

Modern language implementations using Virtual Machines feature diverse execution engines such as byte-code interpreters and machine-code dynamic translators, a.k.a. JIT compilers. Validating such engines requires not only validating each in isolation, but also that they are functionally equivalent. Tests should be duplicated for each execution engine, exercising the same execution paths on each of them.

In this paper, we present a novel automated testing approach for virtual machines featuring byte-code interpreters. Our solution uses concolic meta-interpretation: it applies

San Diego, CA, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3519939.3523457>

### 1 Introduction

Modern Virtual Machines support code generation for compilation and dynamic code patching for techniques such as inline caching. They are often structured around a baseline code interpreter, a baseline JIT compiler, and a speculative inliner. This complexity is aggravated when the VM supports multiple target architectures [1]. Validating



**Is that all?**

# Ongoing RISC-V64 Port

- Currently under development: **Real HW testing** stage
- Taking advantage of our harness test suite
- Improving tests and scenarios
  
- Collaboration with Q. Ducasse, P. Cortret, L. Lagadec from ENSTA Bretagne
- Future work on: *Hardware-based security enforcement*



RISC-V

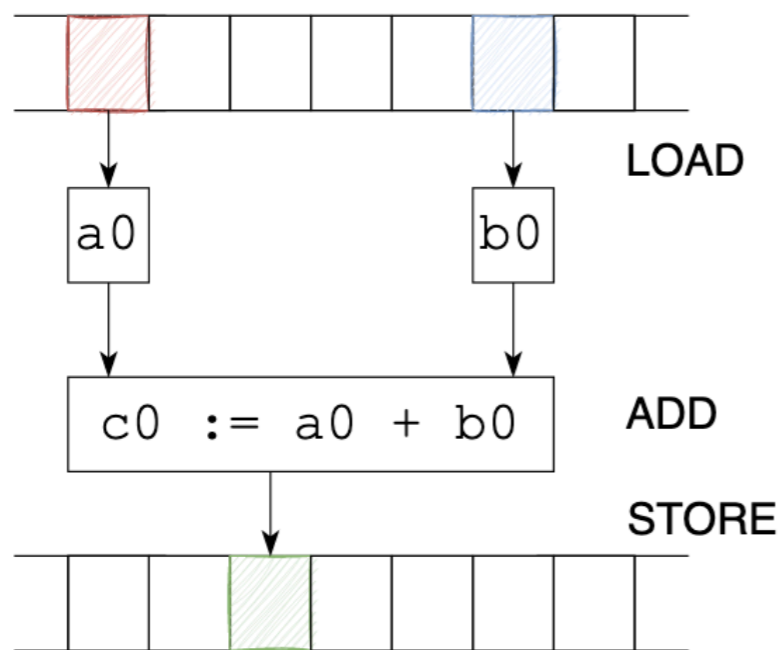


ENSTA  
BRETAGNE

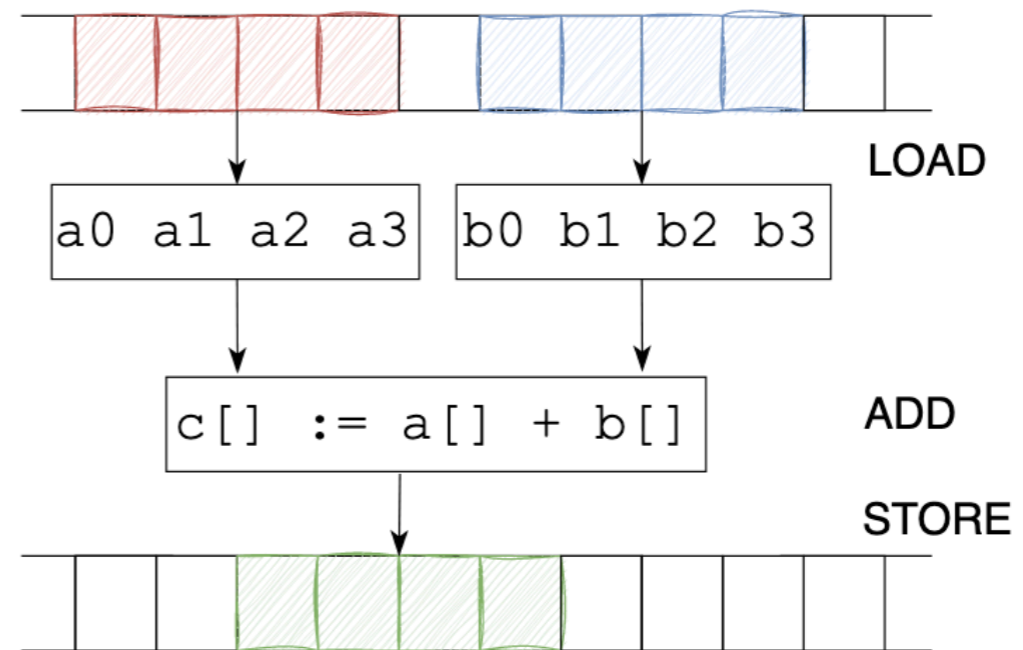


# Single Instruction Multiple Data Extensions

Scalar




Vectorial





# SIMD Design Space

- VM Primitives
  - **Specialised**
  - Faster, less checks
- Vectorised Bytecode
  - **Composable**
  - Safe at the expense of speed

Throughput Scalability \*



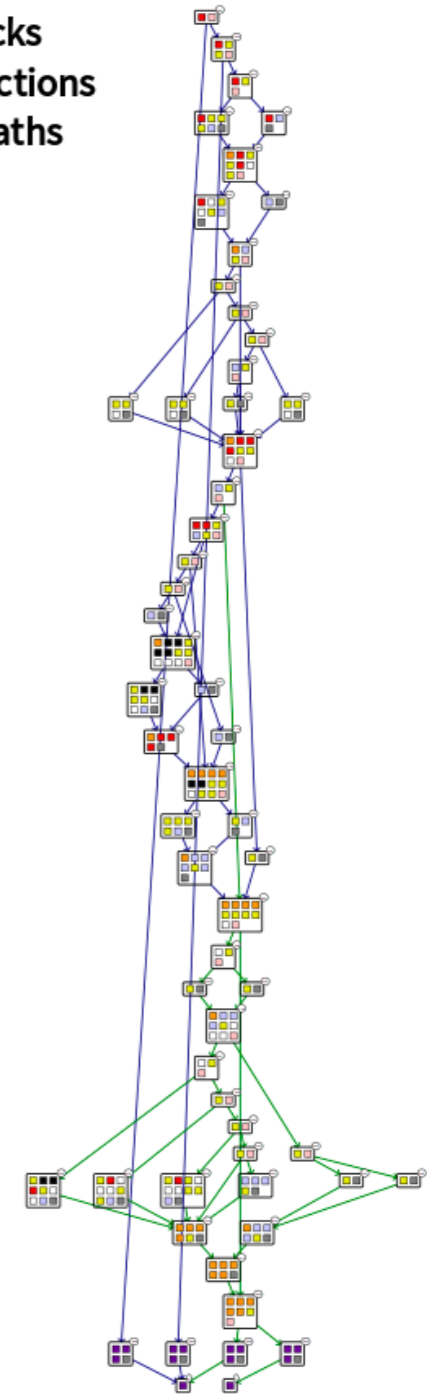
Speedup ratio (1x = Scalar performance) \*



# Tools for Debugging

- Machine Code Debugger
- Compiler IR Visualisations
- Disassembler DSL
- ...

59 blocks  
262 instructions  
10082 paths



VM Debugger									
IR Instructions	Address	ASM	Bytes						
'(PopR 10 13503 810113)'	16r1000000	ld a0, 0(sp) #[3 53 1 0]		lr		'16r1001000'	SP	16r1002FE8	16r1013400
'(Label 1)'	16r1000004	addi sp, sp, 8 #[19 1 129 0]		pc		'16r1000'		16r1002FF0	16r1013400
'(TstCqR 7 10 757D93)'	16r1000008	andi s11, a0, #[147 125 11		sp		'16r1002FE8'		16r1002FF8	16r1013400
'(JumpNonZero (Label 2) 20D9063)'	16r100000C	bnez s11, 32 #[99 144 13		fp		'16r1003000'	FP	16r1003000	16r0
'(MoveMwrR 0 10 22/16 53B03)'	16r1000010	ld s6, 0(a0) #[3 59 5 0]		x0	zero	'16r0'		16r1003008	16r0
'(AndCqR 4194295/3FFFF7 22/16 no mcode)'	16r1000014	lui t0, 1024 #[183 2 64 0]		x1	ra	'16r1001000'		16r1003010	16r0
'(JumpNonZero (Label 2) D9663)'	16r1000018	addiw t0, t0, #[155 130 11		x2	sp   sp	'16r1002FE8'		16r1003018	16r0
'(MoveMwrR 8 10 10 853503)'	16r100001C	and s6, s6, t0 #[51 123 91		x3	gp	'16r0'		16r1003020	16r0
'(Jump (Label 1) FE1FF06F)'	16r1000020	bnez s11, 12 #[99 150 13		x4	tp	'16r0'		16r1003028	16r0
'(Label 2)'	16r1000024	ld a0, 8(a0) #[3 53 133 0]		x5	t0   ip1	'16r0'		16r1003030	16r0
'(MoveMwrR 0 2 23/17 13B83)'	16r1000028	j -32 #[111 240 31		x6	t1   ip2	'16r0'		16r1003038	16r0
'(Label 3)'	16r100002C	ld s7, 0(sp) #[131 59 1 0]		x7	t2	'16r0'		16r1003040	16r0
'(TstCqR 7 23/17 7BFD93)'	16r1000030	andi s11, s7, i#[147 253 12		x8	s0(fp)   fp	'16r1003000'		16r1003048	16r0
'(JumpNonZero (Label 4) 20D9063)'	16r1000034	bnez s11, 32 #[99 144 13		x9	s1	'16r0'		16r1003050	16r0
'(MoveMwrR 0 23/17 22/16 BBB03)'	16r1000038	ld s6, 0(s7) #[3 187 11 0]		x10	a0   arg0	'16r0'		16r1003058	16r0
'(AndCqR 4194295/3FFFF7 22/16 no mcode)'	16r100003C	lui t0, 1024 #[183 2 64 0]		x11	a1   arg1	'16r0'		16r1003060	16r0
'(JumpNonZero (Label 4) D9663)'	16r1000040	addiw t0, t0, #[155 130 11		x12	a2   carg0	'16r0'		16r1003068	16r0
'(MoveMwrR 8 23/17 23/17 8BBB83)'	16r1000044	and s6, s6, t0 #[51 123 91		x13	a3   carg1	'16r0'		16r1003070	16r0
'(Jump (Label 3) FE1FF06F)'	16r1000048	bnez s11, 12 #[99 150 13		x14	a4   carg2	'16r0'		16r1003078	16r0
'(Label 4)'	16r100004C	ld s7, 8(s7) #[131 187 13		x15	a5   carg3	'16r0'		16r1003080	16r0
'(CmpRR 10 23/17 41750DB3)'	16r1000050	j -32 #[111 240 31		x16	a6	'16r0'		16r1003088	16r0
'(JumpNonZero (MoveCqR 16856080/1013410 16	16r1000054	sub s11, a0, s #[179 13 117		x19	s3   extra1	'16r0'		16r1003090	16r0
'(MoveCqR 16856096/1013420 10 1013537 42050	16r1000058	bnez s11, 16 #[99 152 13		x20	s4   extra2	'16r0'		16r1003098	16r0
'(Jump (MoveRMwr 10 0 2 A13023) C0006F)'	16r100005C	lui a0, 4115 #[55 53 1 1]		x22	s6   temp	'16r0'		16r10030A0	16r0



# Pharo VM Manual Variable Localisation

**interpret**

```
self fetchNextBytecode.  
[ true ] whileTrue: [  
  self  
    dispatchOn: currentBytecode  
    in: BytecodeTable ].
```

**pushReceiverBytecode**

```
self fetchNextBytecode.  
self internalPush: self receiver
```

**pushBool: trueOrFalse**

```
<inline: true>  
self push: (objectMemory booleanObjectOf: trueOrFalse)
```

**internalAboutToReturn: resultOop through: aContext**

```
<inline: true>  
[...]  
self internalPush: resultOop  
[...]
```


**internalPush: aValue**

```
localSP := localSP - bytesPerWord.  
self longAt: localSP put: aValue
```

**push: aValue**

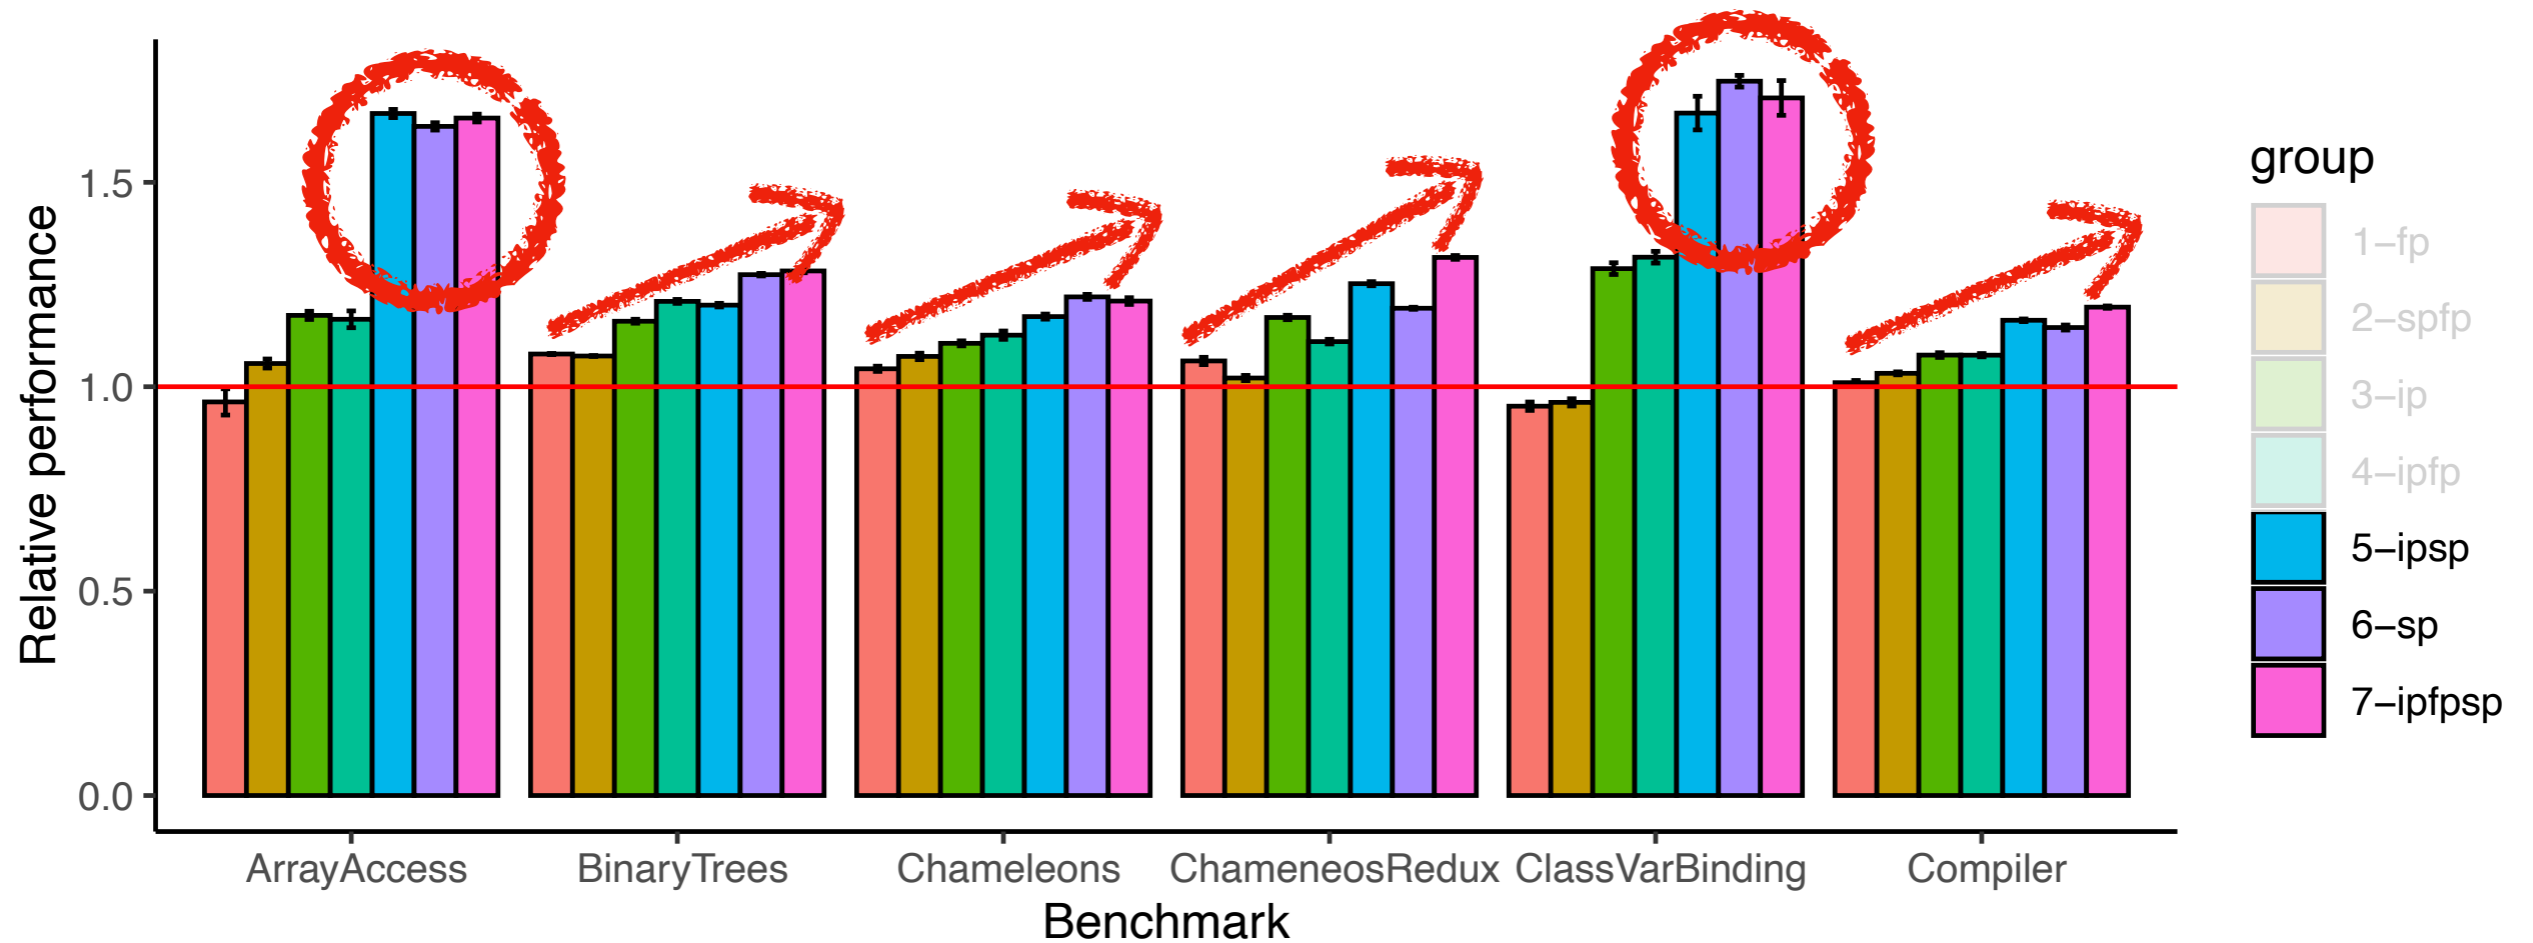
```
stackPointer := stackPointer - byte  
self longAt: stackPointer put: aValue
```

**Developer should know  
C generation semantics**



# Automatic Variable localisation!

Intel x86-64




Averages of 100 iterations + stdev. Relative to baseline (no optimisation). Higher



# Analysing Code Cache Behavior

Occupation  
Rate

Red: Compaction Events




Analysing Events  
We see trashing  
in the code

We need to  
increase the size  
of the code



# Code Cache Unexpected results




Young Space 1MB 10 MB 100 MB




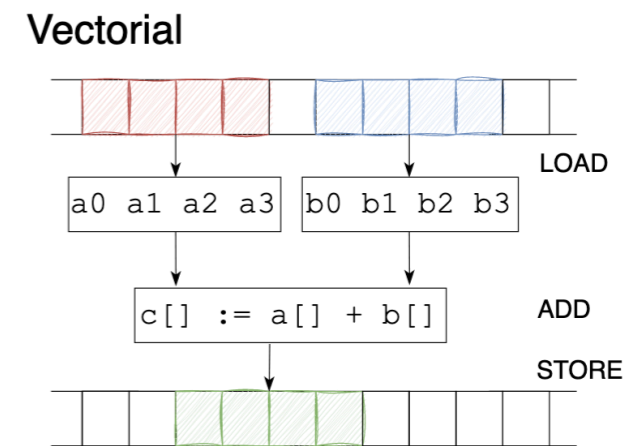
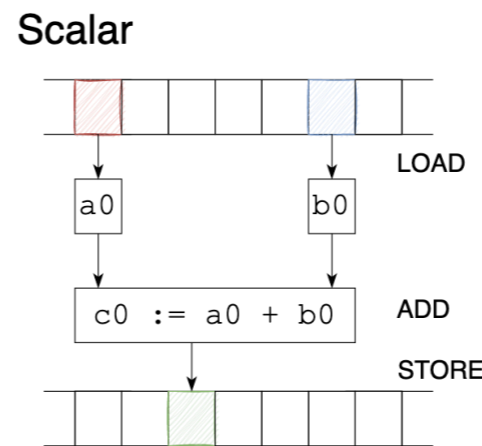
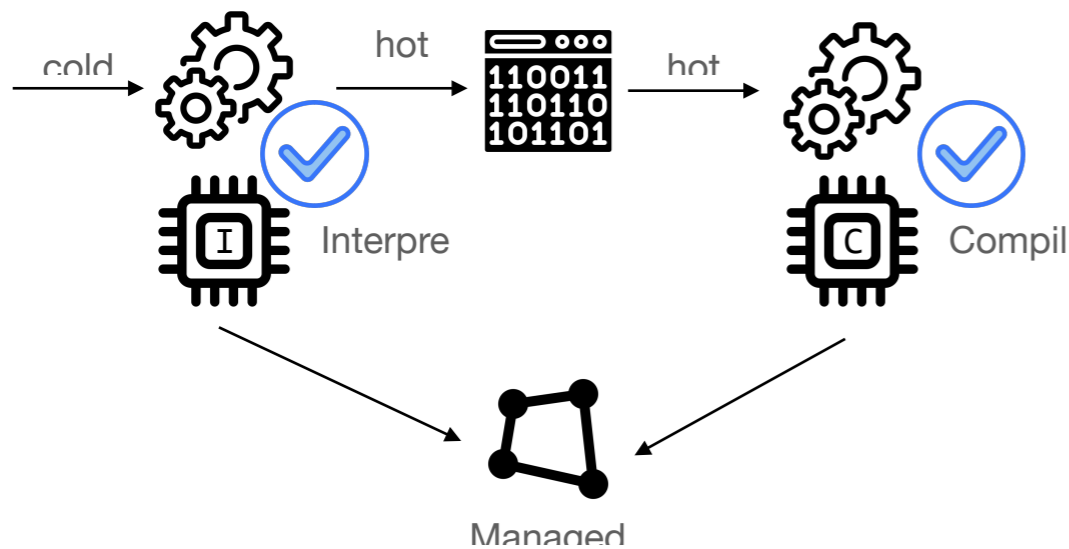
# We are hiring!

- We have
  - Engineer Positions
  - Phd Positions
- Keywords: *Compilers, Interpreters, Memory Management, Security*
- **Come talk to us!**





# Pharo VM - News from the Front



## Stay in Sync!

Permanent Space

Ephemérons

New Image Format

► Lifeware

Faster Startup / Saving

Speculative  
Compilation

[guillermo.polito@univ-lille.fr](mailto:guillermo.polito@univ-lille.fr)

@pharoproject

pharo.org


consortium-adm@pharo.org

discord.gg/QewZMZA

thepharo.dev

# Conclusion

- 478 differences found, 91 causes, 6 categories
- Practical:
  - 4928 tests generated in ~8 minutes



Guillermo Polito - Pablo Tesone - Stéphane Ducasse

[guillermo.polito@univ-lille.fr](mailto:guillermo.polito@univ-lille.fr)

@guillep



# Improvements: Clean Up

- V3 Support
- Old Memory Format
- Old Block Closures
- Dead Code
- ~ 65KLOC



# Improvements: Sockets

- Unified Implementation in all Platforms
- Better Async Support
- Unix Sockets (Under Work)
- IPv6 Addresses (Under Work)



# Improvements: Serial Port FFI

- Pure FFI implementation
- Working in all Platforms (Unix / Windows / OSX)
- Migrating Plugins to FFI

# Improvements: RISCv64

## Ongoing Port

- Currently under development: Real HW testing stage
- Taking advantage of our harness test suite.
- Improving tests and scenarios
- Collaboration with Q. Ducasse, P. Cortret, L. Lagadec from ENSTA Bretagne
- Future work on: Hardware-based security enforcement





# Improvements: Open Build Service

## Better Support for Linux Distributions

	Arch	Debian_10	Debian_9.0	Debian_Testing	Fedora_31	Fedora_32	Fedora_33	Raspbian_10		Raspbian_9.0	
	↑↓  x86_64↓	x86_64↓	x86_64↓	x86_64 ↑↓	x86_64↓	x86_64↓	x86_64↓	aarch64↓	x86_64↓	aarch64↓	x86_64↓
libffi7		succeeded	succeeded		succeeded	succeeded	succeeded	succeeded	succeeded	succeeded	succeeded
libgit2-1		succeeded		failed							
pharo9	failed	succeeded	failed	failed	failed	failed	failed	succeeded	succeeded	failed	failed
pharo9-ui	succeeded	succeeded	succeeded	failed	succeeded	succeeded	succeeded		succeeded		succeeded

	Raspbian_9.0		openSUSE_Leap_15.1	openSUSE_Leap_15.2	openSUSE_Tumbleweed	xUbuntu_18.04	xUbuntu_19.04	xUbuntu_20.04	
	↑↓ arch64↓	x86_64↓	x86_64 ↑↓	x86_64 ↑↓	x86_64 ↑↓	x86_64 ↑↓	x86_64 ↑↓	aarch64↓	x86_64↓
libffi7	succeeded	succeeded	succeeded	succeeded	succeeded	succeeded	succeeded	succeeded	succeeded
libgit2-1			succeeded	succeeded		succeeded	succeeded		succeeded
pharo9	failed	failed	failed	failed	failed	failed	succeeded	succeeded	succeeded
pharo9-ui		succeeded	succeeded	succeeded	succeeded	succeeded	succeeded		succeeded

### Initial targets:

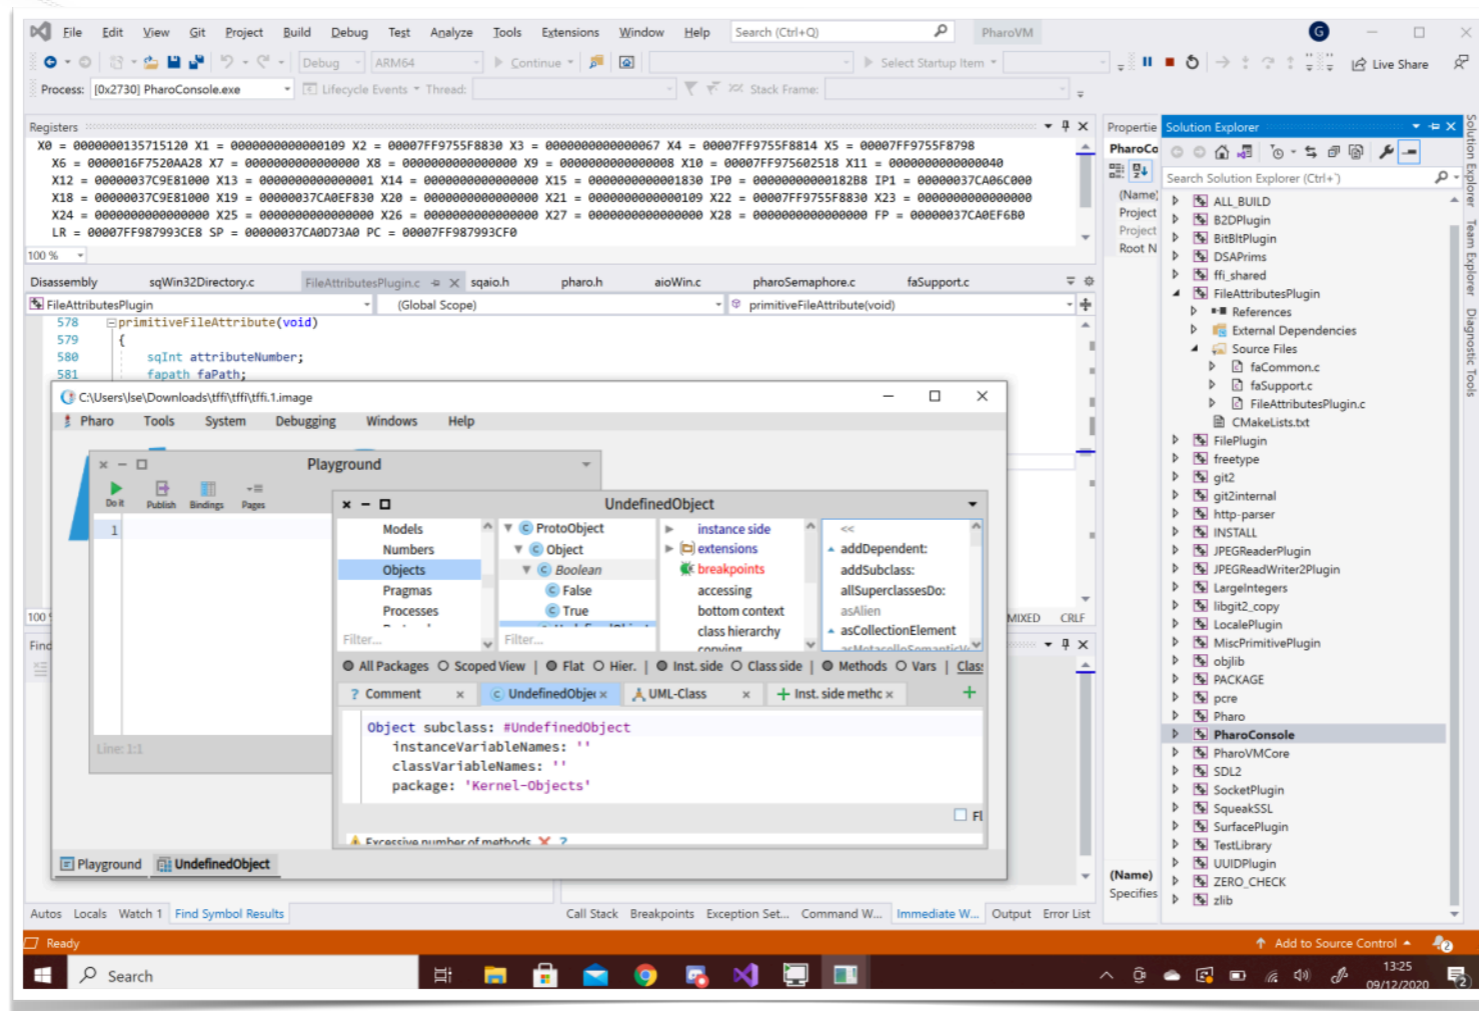
- Arch / Manjaro
- Debian
- Fedora
- Raspbian
- Ubuntu
- openSuse

Multiple Architectures

Supporting system

Building using existing system

# Improvements: Visual Studio Support Building & Debugging



MSVC - No  
cygwin



# Improvements: Windows ARM



The screenshot displays the Visual Studio Code IDE interface on a Windows ARM system. The main window is titled "Playground" and contains the following code:

```
1 canvas := RSCanvas new.  
2 shapes := #(20 10 5 30 24 32) collect: [ :e | RSEllipse model: e ]  
3 canvas addAll: shapes.  
4 RSNormalizer size shapes: shapes; normalize: #yourself.  
5 RSFlowLayout new alignCenter; on: shapes.  
6 canvas @ RSCanvasController
```

The Test Runner window shows the following test results:

```
274 ran, 270 passed, 0 skipped, 3 expected failures, 0 failures, 1 error, 0 passed unexpected
```

The Inspector window shows a canvas with several gray circles of different sizes and positions.

The Class Explorer window shows the project structure, including the following packages and classes:

- BaselineOfSpec2
- Commander-Spec2-Compatil
- Spec2-Adapters-Morphic
- Spec2-Adapters-Morphic-Test
- Spec2-Adapters-Stub
- Spec2-Backend-Tests
- Spec2-Code
- spec2

The System Reporter window shows the following system information:

```
Operating System/Hardware  
Win32 Windows-ARM64 ARM64
```

MSVC - No  
cygwin





# Back to the Future

## Objectives for 2022



# Permanent Space

## Problem

- Many permanent objects
- They have references from/to other objects
- We are traversing them to GC
- E.g., Classes, Methods, Literals, Resources



**Generates  
Long Pauses**



# Permanent Space

## Our Solution



- New Object Space for permanent Objects
- Minimise or Eliminate GC passes
- Persisting them through executions



# Permanent Space

## Our Solution



- New Object Space for permanent Objects
- Minimise or Eliminate GC passes
- Persisting them through executions

**We need to  
put them in a**





# New Image Format



## Problem

- Current Image format only support a single object space
- No extensible: not new metadata nor new data
- Cannot be Memory Mapped (it is modified before save/load)
- Requires to discard all state of the running VM (slow saves)



# New Image Format



## Problem

- Current Image format only support a single object space
- No extensible: not new metadata nor new data
- Cannot be Memory Mapped (it is modified before save/load)
- Requires to discard all state of the running vmi (slow saves)

**Slow and Restricting**





# New Image Format



## Our Solution

- New Image format based in directories / bundles
- Many Elements of data and metadata
- Metadata en User & Machine readable format (STON?)
- Extensible format



# **Fast Snapshots / Loading**

## **Based on PermSpace & Image Format**

- Memory Mapped Image
- Shareable State
- Saving / Loading Warm State of the VM

## Next Objectives

Permanent Space      Ephemeron  
New Image Format      Speculative  
Faster Startup / Saving      Compilation

- ARM64, RISCV64, Slang...
- Lots of Tests!
- Integration: Sockets, serial      @pharoproject  
   pharo.org
- Visual Studio, Open Build Service @pharo.org  
   discord.gg/QewZMZA  
   thepharo.dev




# Cross-ISA Testing of the Pharo VM

Lessons learned while porting to ARMv8 64bits  
Tool Paper – MPLR'21

**Guille Polito**, Stéphane Ducasse, Pablo Tesone,  
Théo Rogliano, Pierre Misse-Chanabier, Carolina Hernandez, Luc Fabresse  
**RMoD Team – Inria Lille Nord Europe – UMR9189 CRISTAL – CNRS**



*Inria*

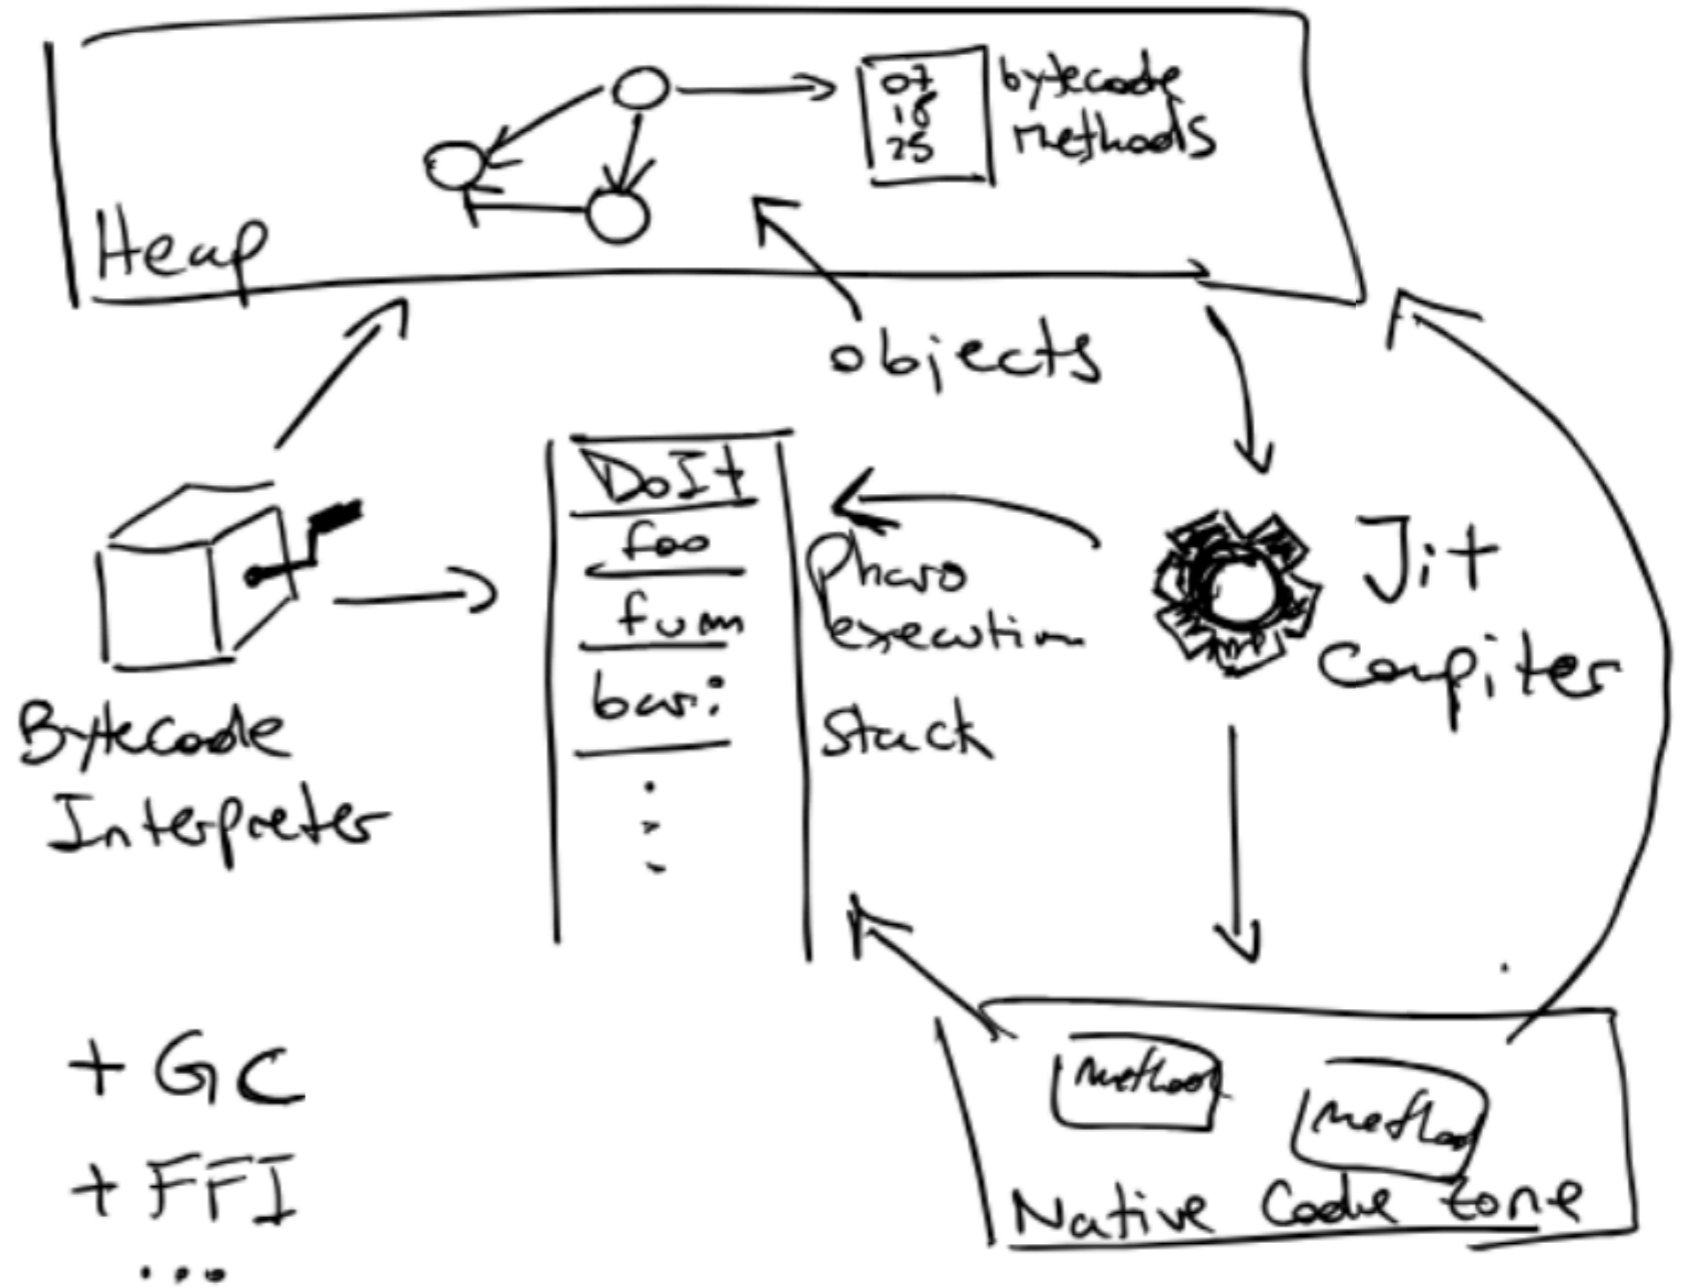


72



# Context

## The Pharo VM



# Some Numbers

- 255 stack based bytecodes (77 different) + ~340 primitives/native methods
- 146 different IR instructions
- polymorphic inline caches
- threaded code interpreter
- generational scavenger GC

*Lots of combinations!*




# Objective: Implementing an ARM64 Backend

- ARM64 is now pervasive
  - New Apple M1
  - Raspberry Pi 4
  - Microsoft Surface Pro
  - PineBook Pro
  - ...

```
move r1 #1  
move r2 #17  
checkSmallInt  
checkSmallInt  
add r3 r1 r2  
checkSmallInt  
move r1 r3  
ret
```


JIT compiler IR



# Case Study 1

## Porting the Cogit JIT Compiler to ARM64

- Started with no tests and no hardware (main target Apple M1)
- Incremental test development: bytecode, native methods, PICs, code patching
- All tests run from the beginning on our four targets: x86, x86-64, ARM32 and ARM64
- Test allowed *safe* modifications in the IR to support e.g., ARM64 Multiplication overflow
- ARM64 specific tests covered stack alignment, W+X ...





# Case Study 2

## Ongoing Port to RISCV64

- Currently under development
- Is our harness test suite enough to develop a new backend?
- Are our tests general enough?
  
- Collaboration with Q. Ducasse, P. Cortret, L. Lagadec from ENSTA Bretagne
- Future work on: Hardware-based security enforcement



# Case Study 3

## Debugging and Testing Memory Corruptions

- Bug report using Ephemeron  
<https://github.com/pharo-project/pharo/issues/8153>
- Starting the other way around
  - First reproducing the bug in real-hardware
    - => long to execute (even longer in simulation)
    - => required manual developer intervention
  - Then building a unit test from observations
  - Test becomes a part of the regression test suite

# Future Perspectives



## Automatic VM Validation

- Automatic (Unit?) Test Case Generation
- Interpreter vs Compiler Differential Testing
- VM Tailored Multi-level Debugging

# Cross-ISA Testing of the Pharo VM

Lessons learned while porting to ARMv8 64bits

	Real Hardware Execution	Full-System Simulation	Unit-Testing
Feedback-cycle speed	Very low	Low	High
Availability	Low	High	High
Reproducibility	Low	Low	High
Precision	High	Low	Low
Debuggability	Low	High	High



# Debugging a compiler

Insights: build your own tools, based on needs, not desires

The screenshot shows a VM Debugger window with the following components:

- IR Instructions:** A list of instructions with their addresses and IR representations, such as '(PopR 10 13503 810113)', '(Label 1)', '(TstCqR 7 10 757D93)', etc.
- Address:** A column showing the memory addresses for each instruction, ranging from 16r1000000 to 16r100005C.
- ASM:** A column showing the assembly code for each instruction, such as 'ld a0, 0(sp) #[3 53 1 0]', 'addi sp, sp, 8 #[19 1 129 0]', etc.
- Bytes:** A column showing the byte sequences for each instruction, such as 'lr', 'pc', 'sp', 'fp', 'x0', 'x1', etc.
- Registers:** A table on the right showing the current values of various registers, including SP, FP, x0-x22, s0-s7, a0-a6, and extra1-extra2.
- Control Panel:** At the bottom, there are buttons for 'Jump to', 'Step', and 'Disassemble at PC', with the current step number '81' displayed.

Examples:

- Machine code debugger
- Bytecode-IR visualization
- Disassembler



# Interpreter-Guided JIT Compiler Test Generation


Validating the Pharo JIT compiler through  
**concolic** execution and **differential** testing

Guille Polito - Pablo Tesone - Stéphane Ducasse  
[guillermo.polito@univ-lille.fr](mailto:guillermo.polito@univ-lille.fr)  
@guillep

PLDI'22 — San Diego



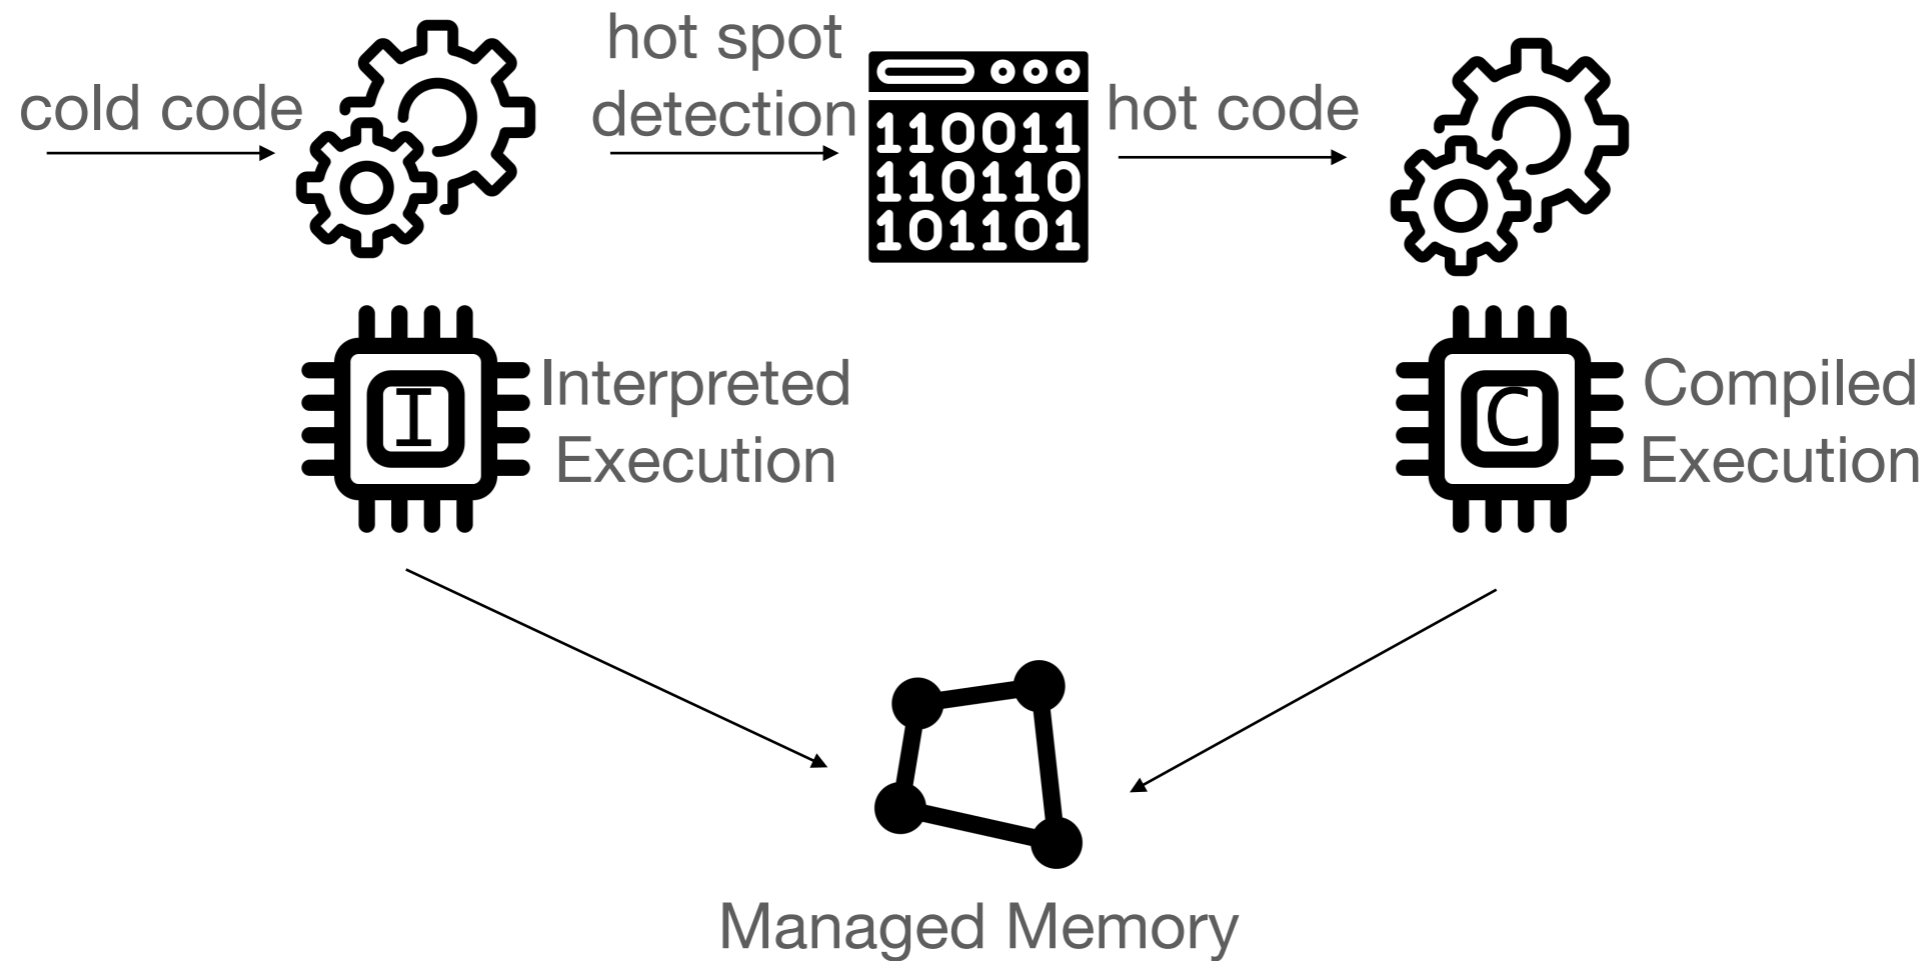
*Inria*




82



# Virtual Machine Execution Engine




# Virtual Machine Execution Engine






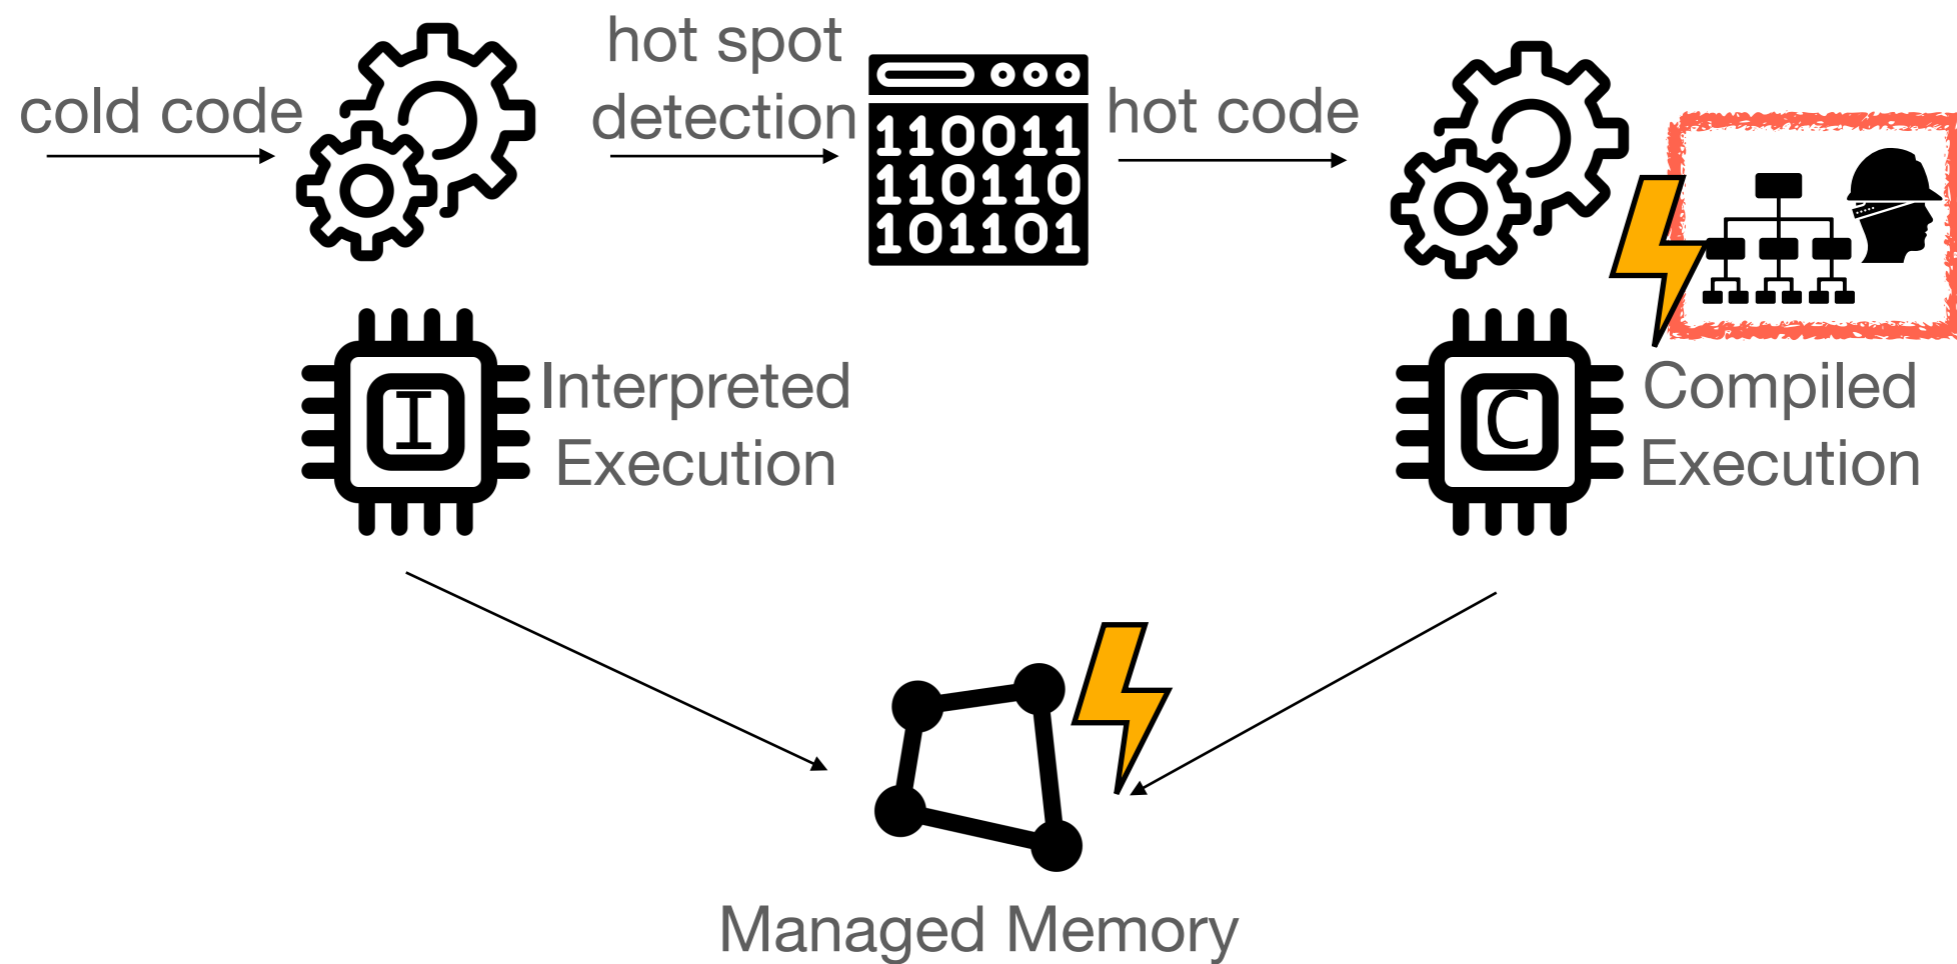
# Virtual Machine Execution Engine



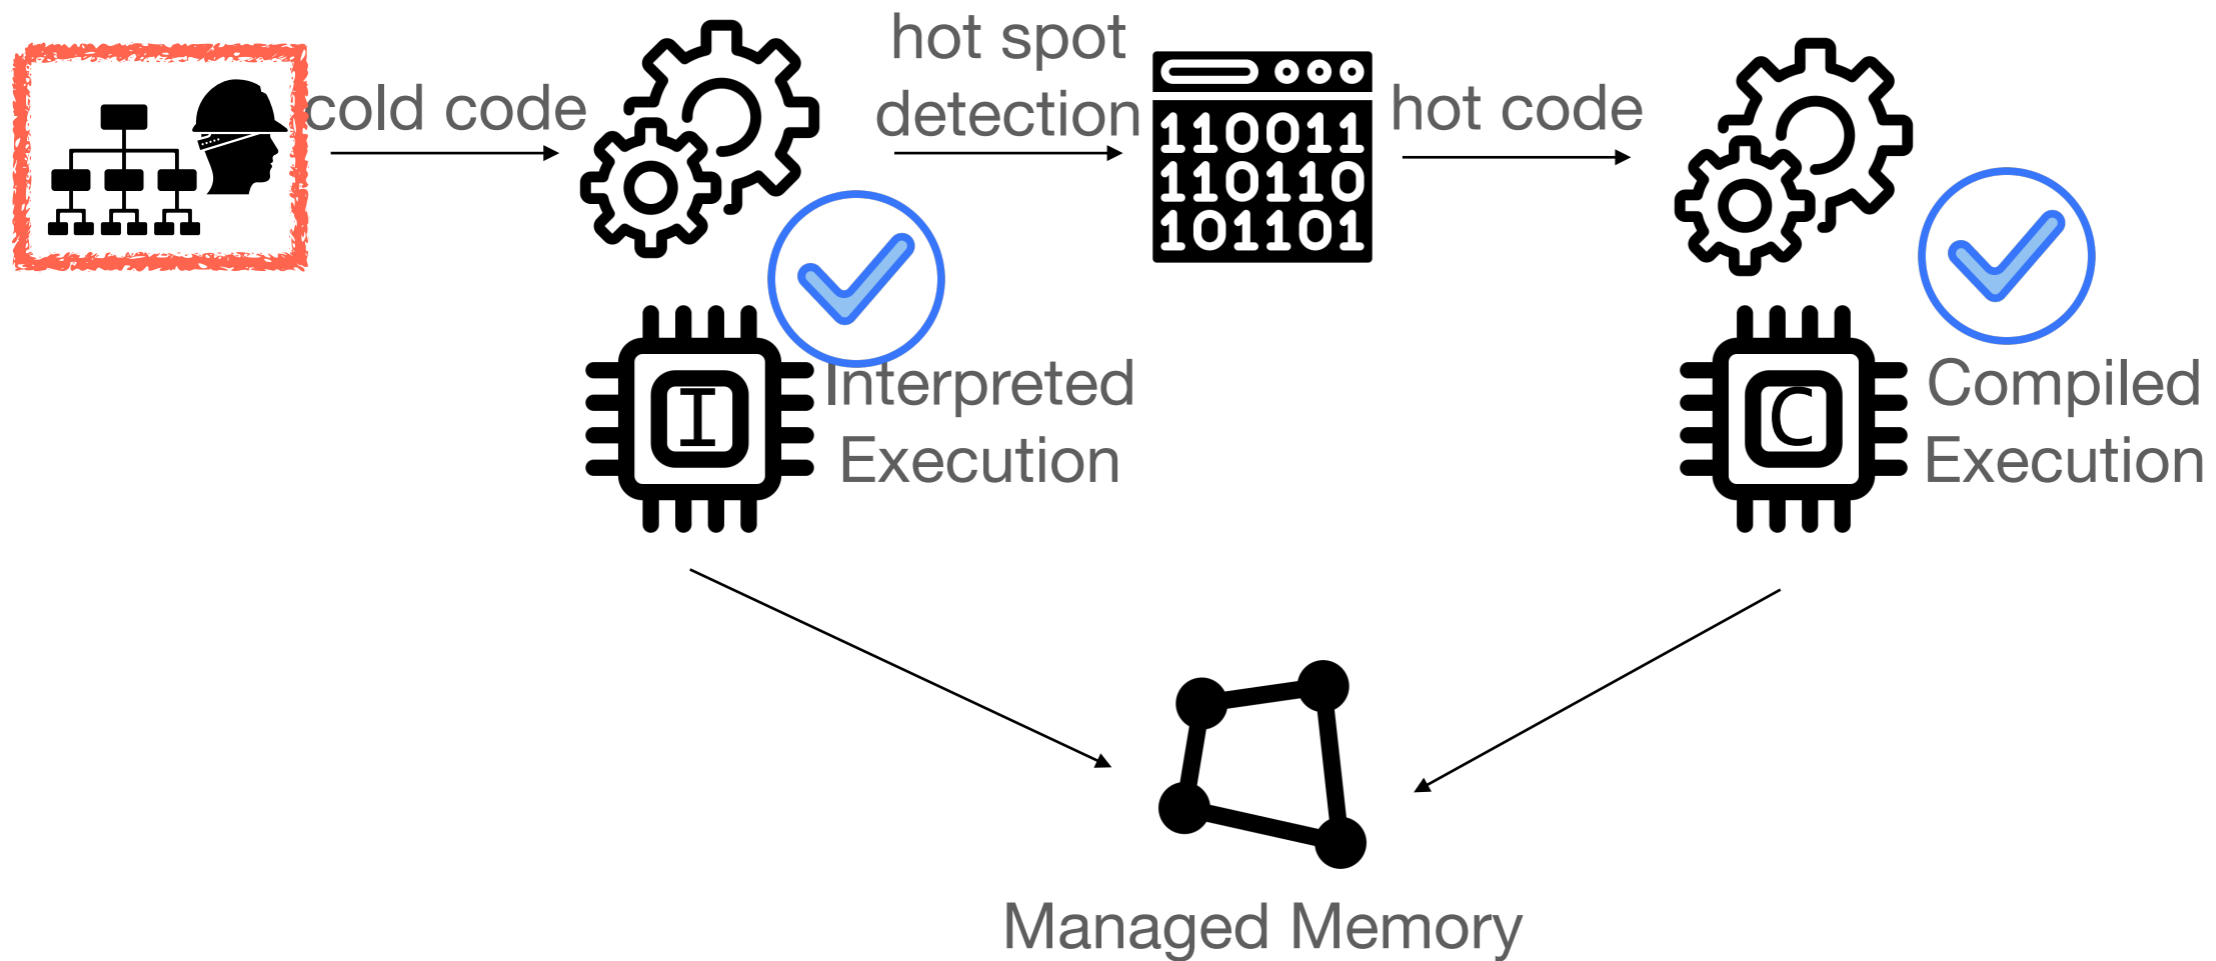
# Virtual Machine Execution Engine



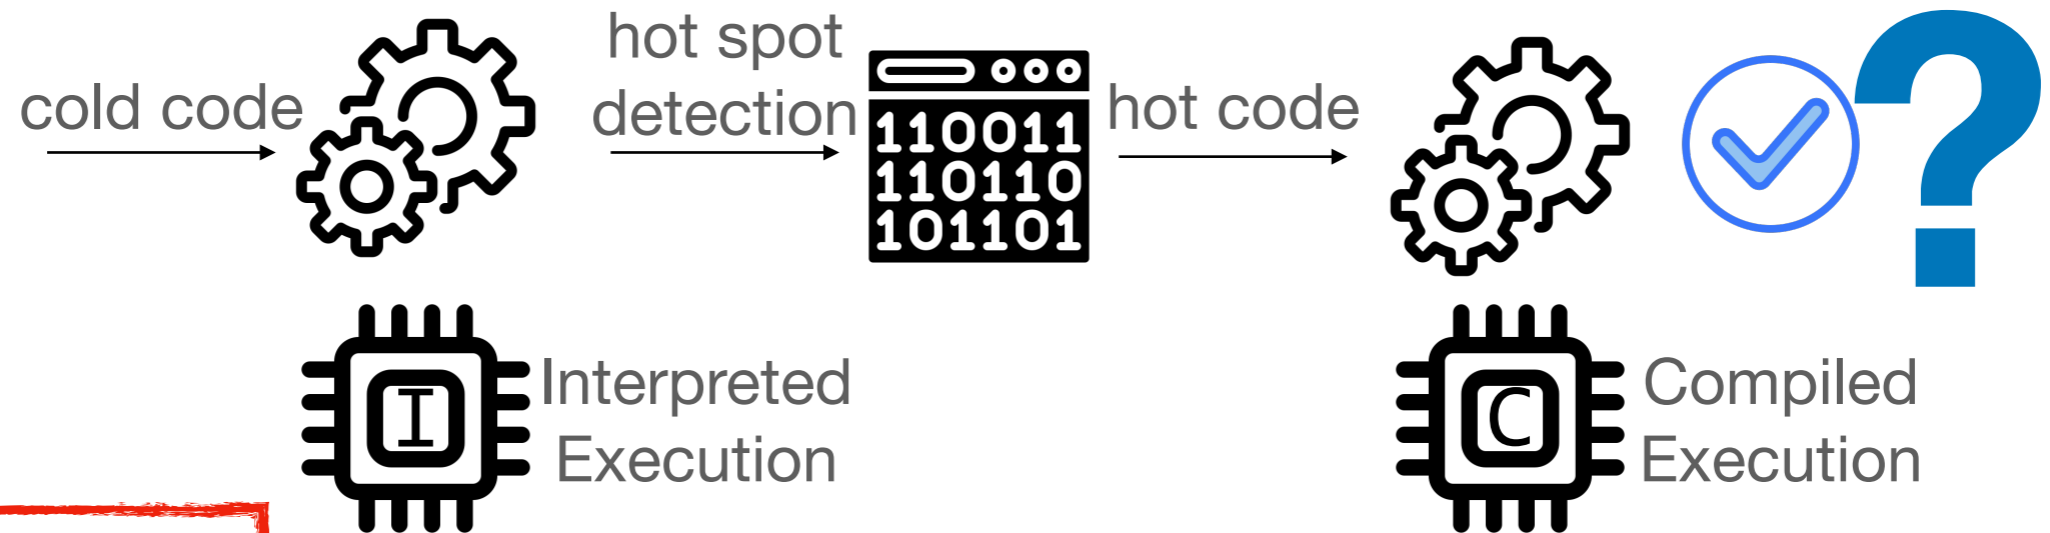
# Virtual Machine Execution Engine



# How can we automatically test VMs?



# Challenges of VM Test Generation




**Challenge 1: Test**


**Challenge 2: Test**

- Do they cover different code *regions/branches/paths*?
- How do we determine what is the *expected output* of a generated test?


# Black Box Testing + Fuzzing




cold code



# Black Box Testing + Fuzzing



cold code



**Slow**



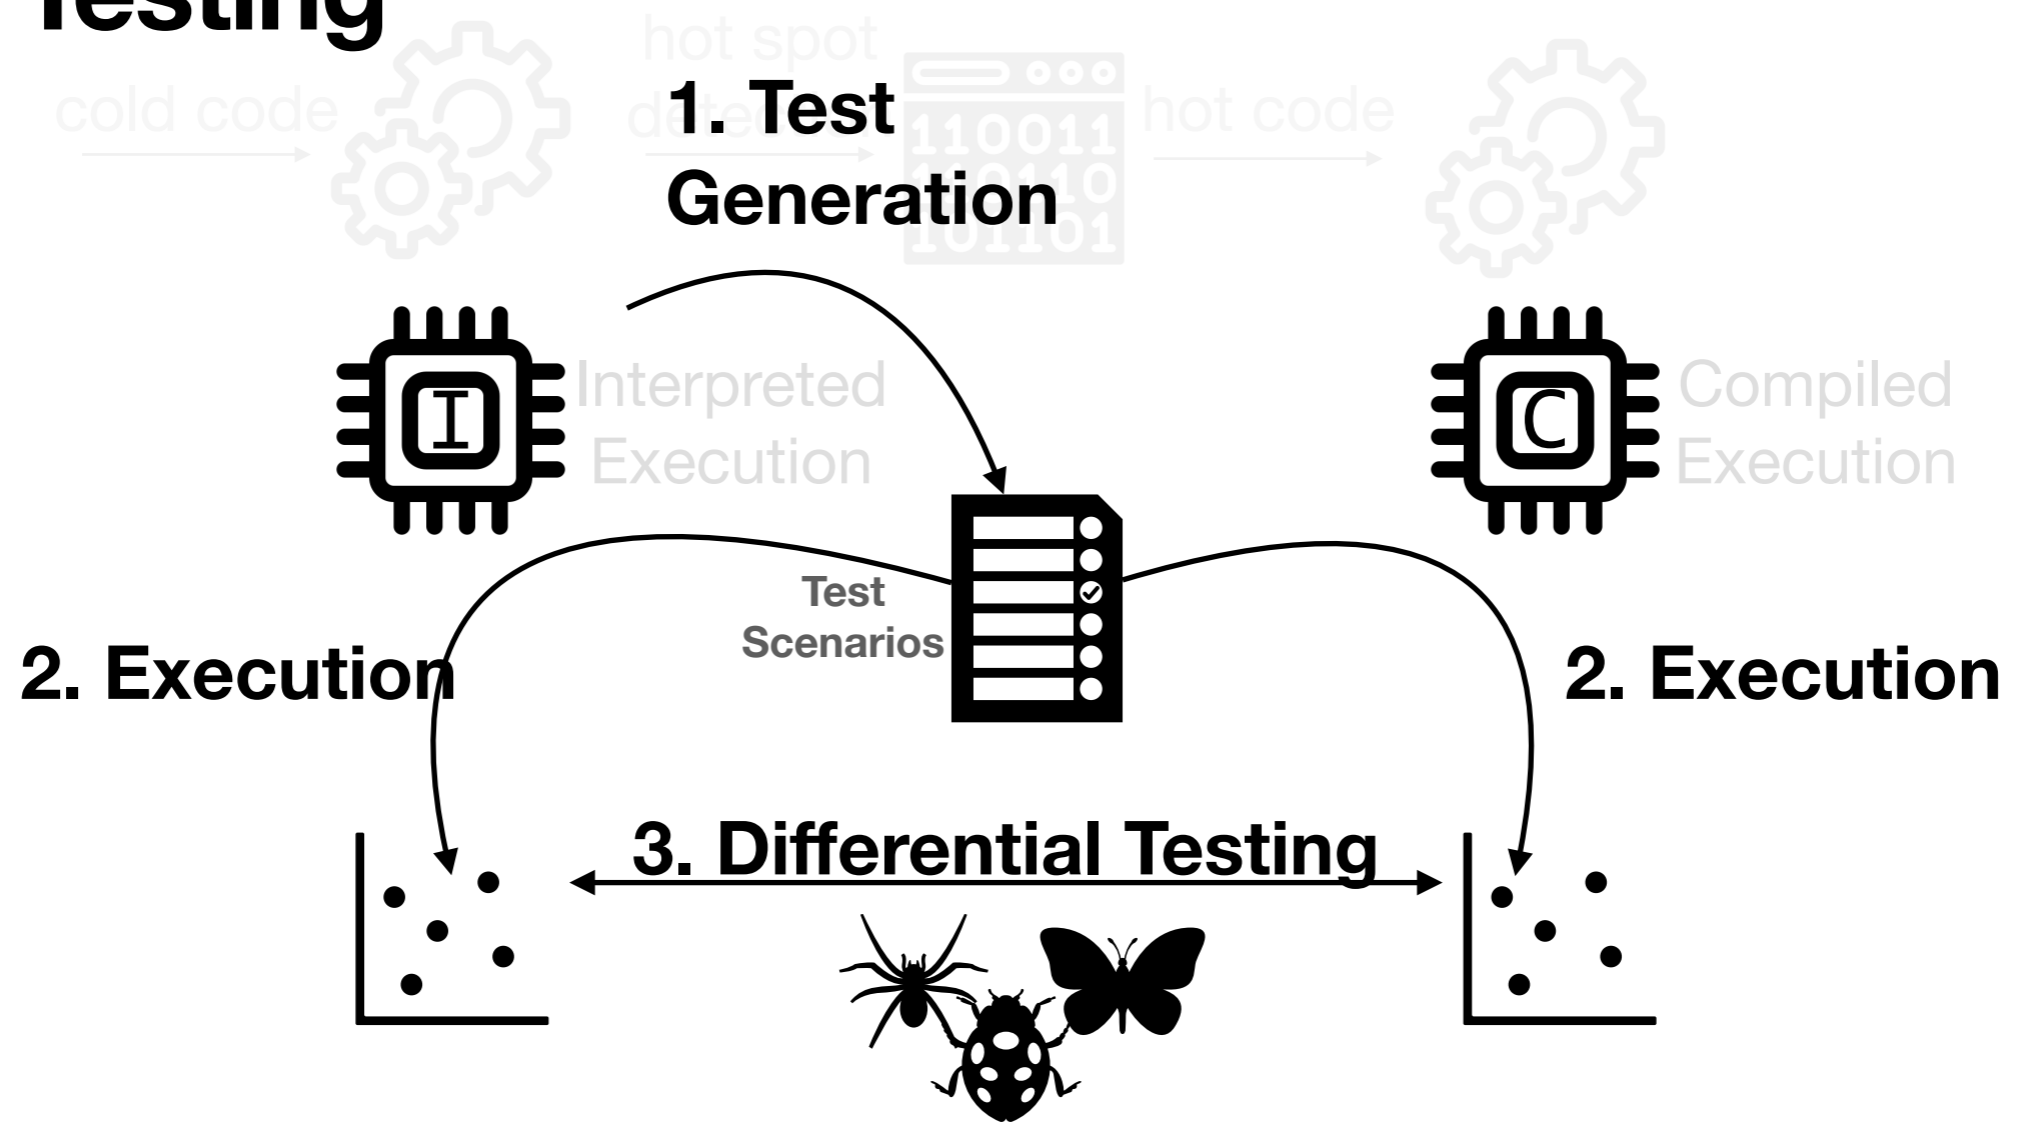
**Coarse Grained**

**Non Determinism**

**Require Multiple Reference**

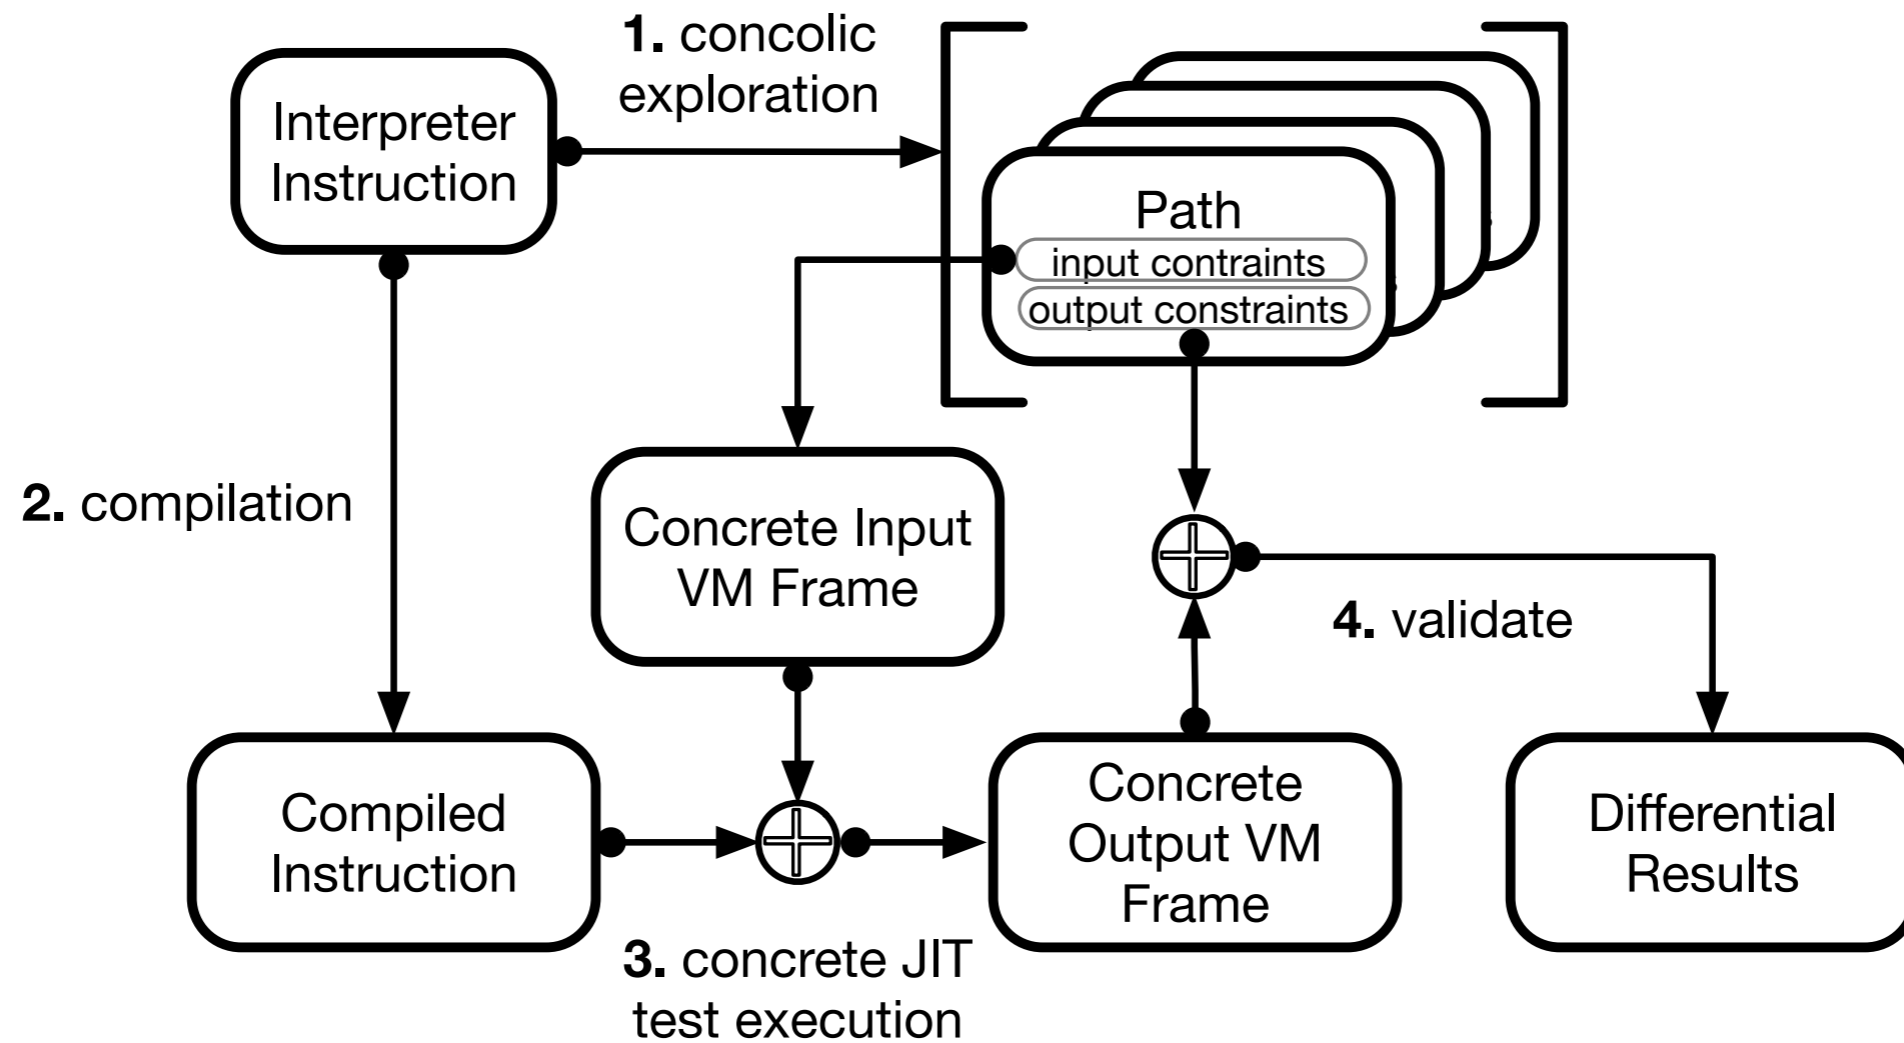
**Implementations**

# Interpreter-Guided Automatic JIT Compiler Unit Testing

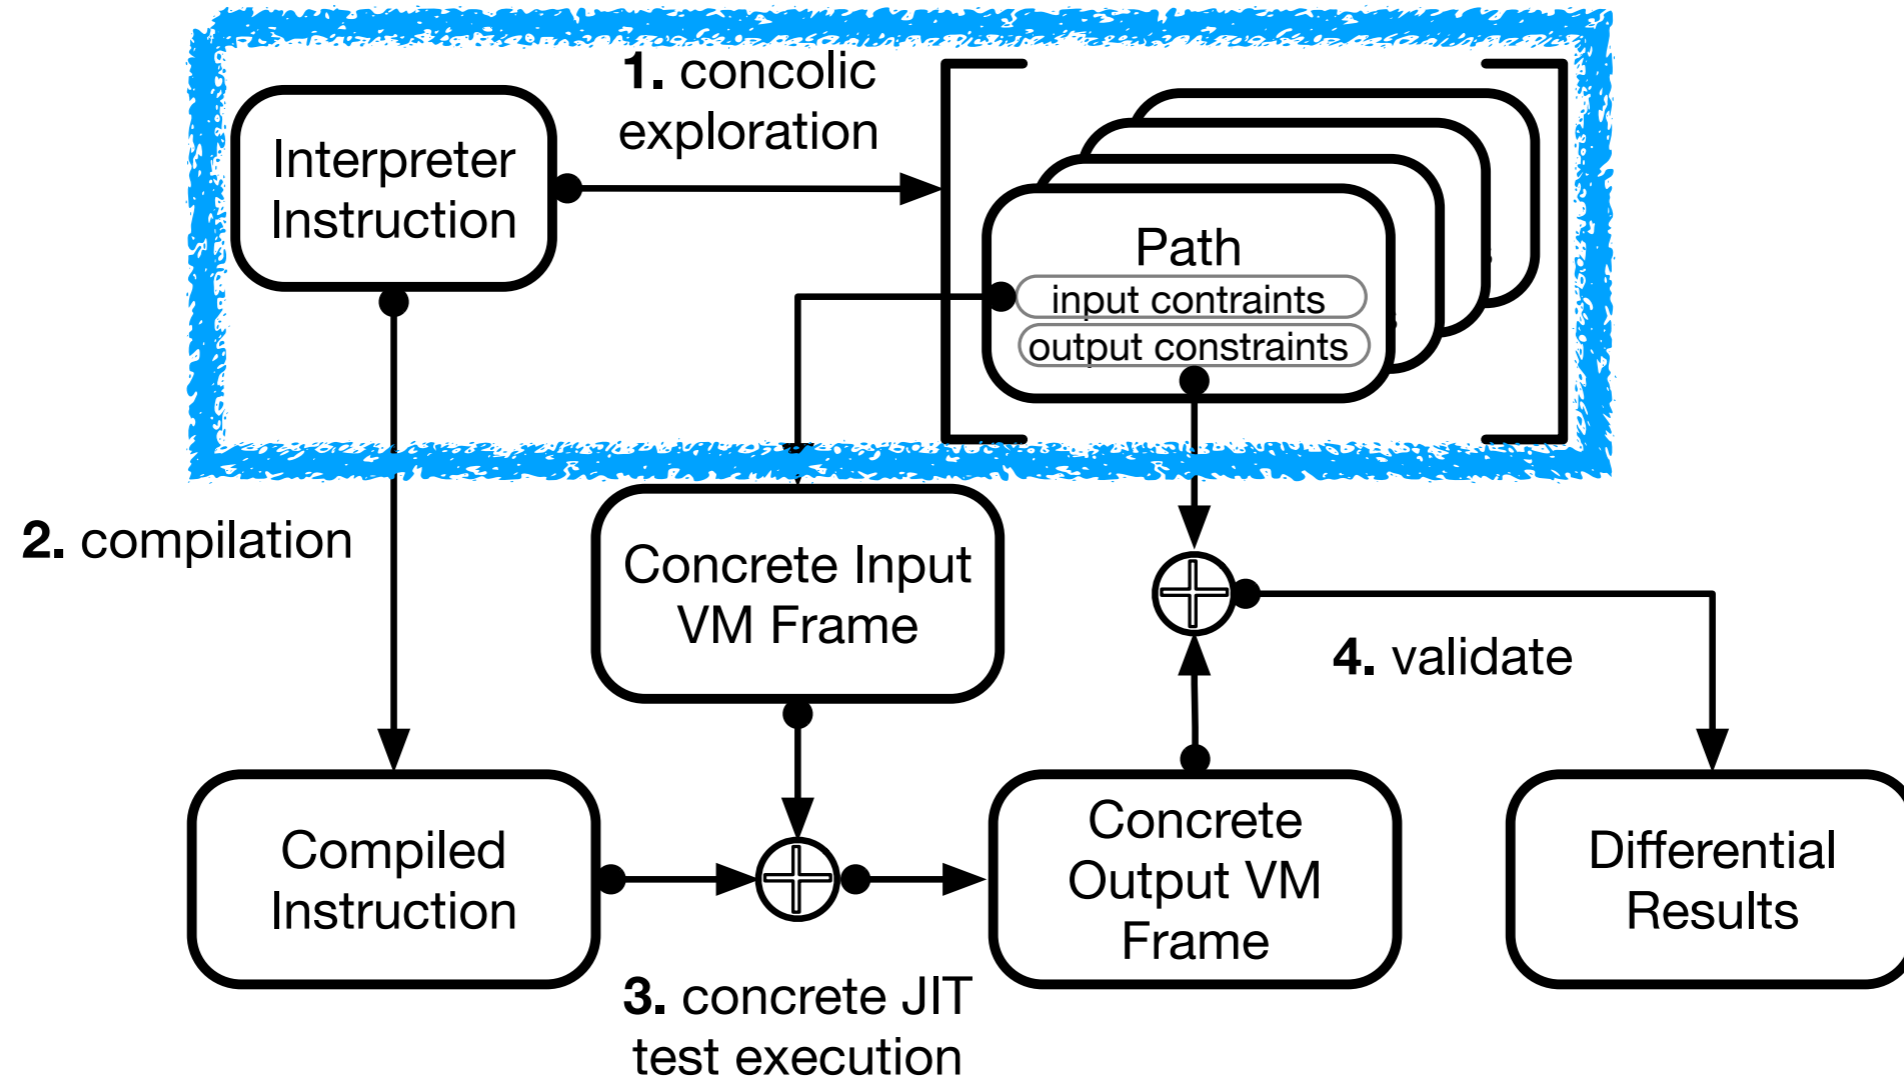




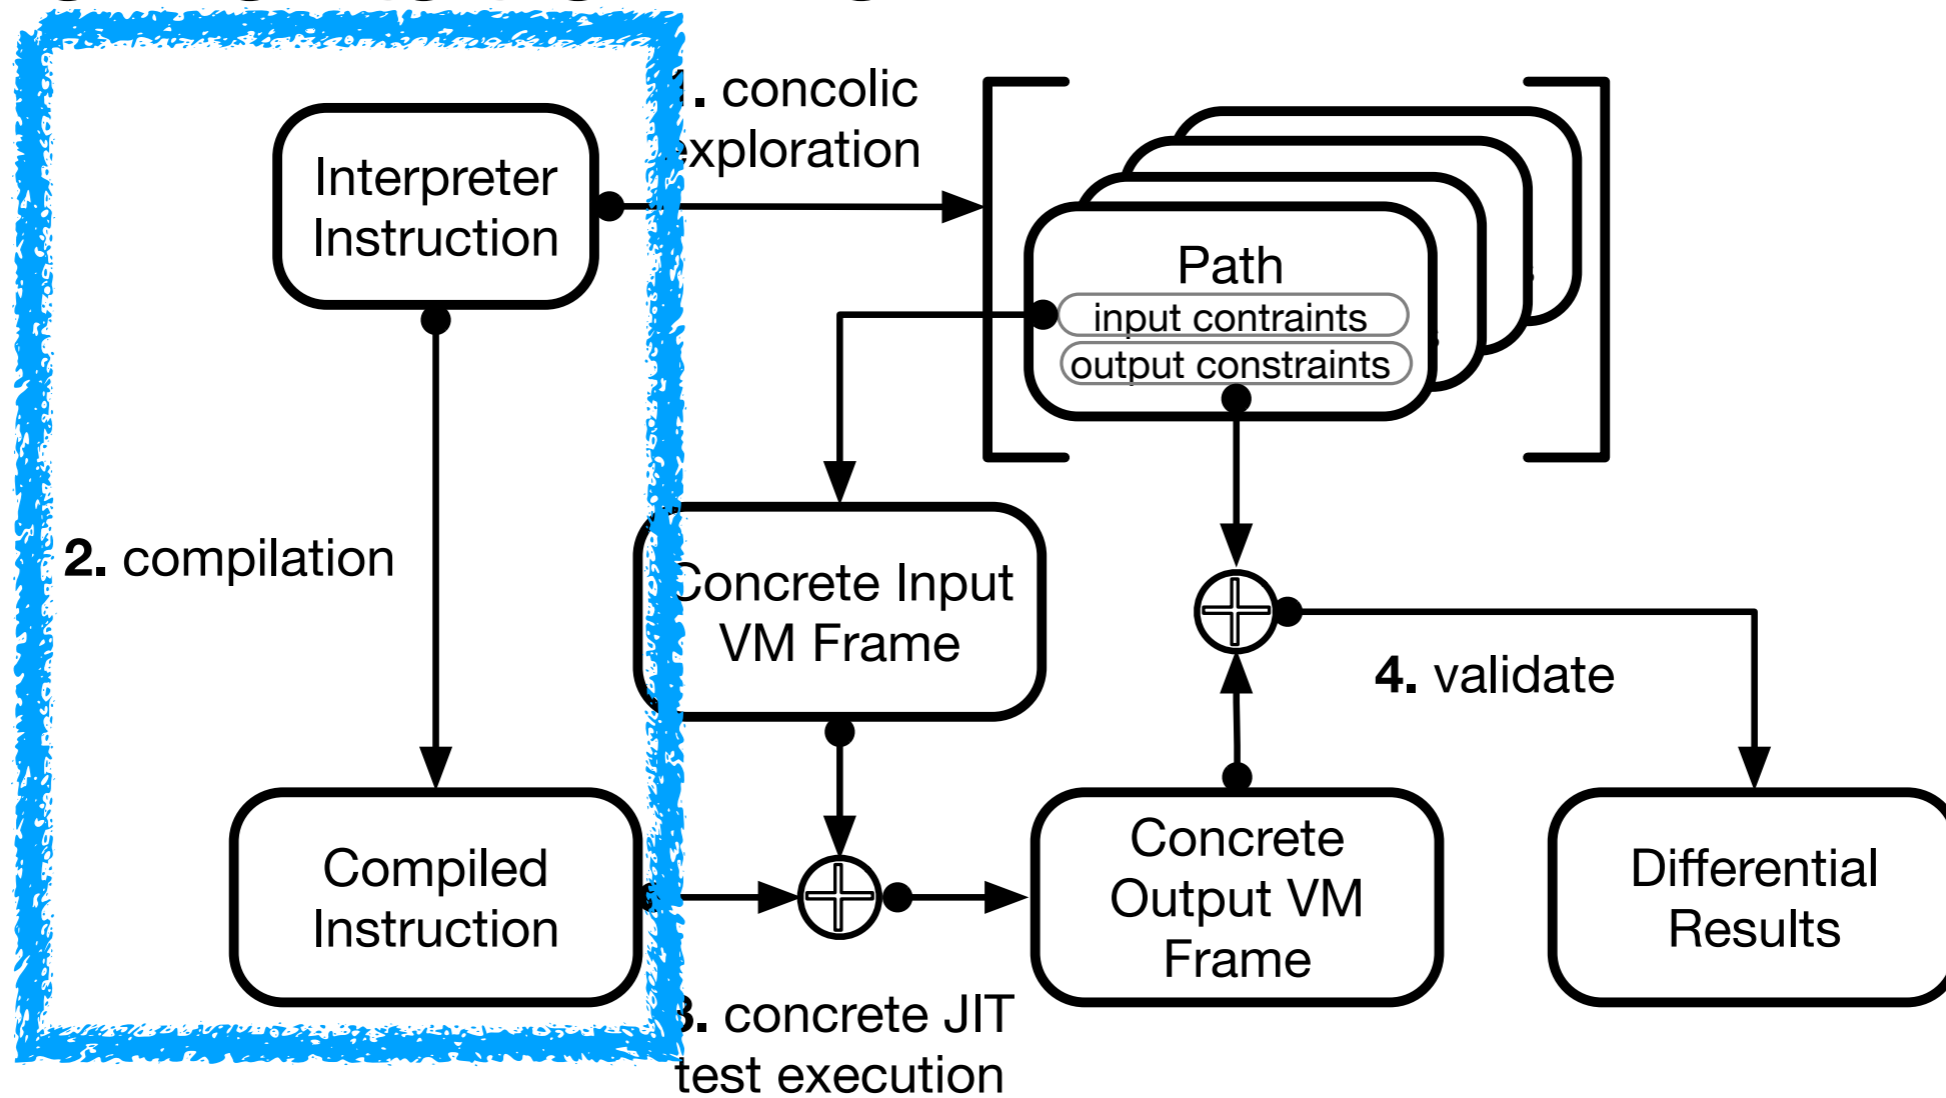
# Implementation View



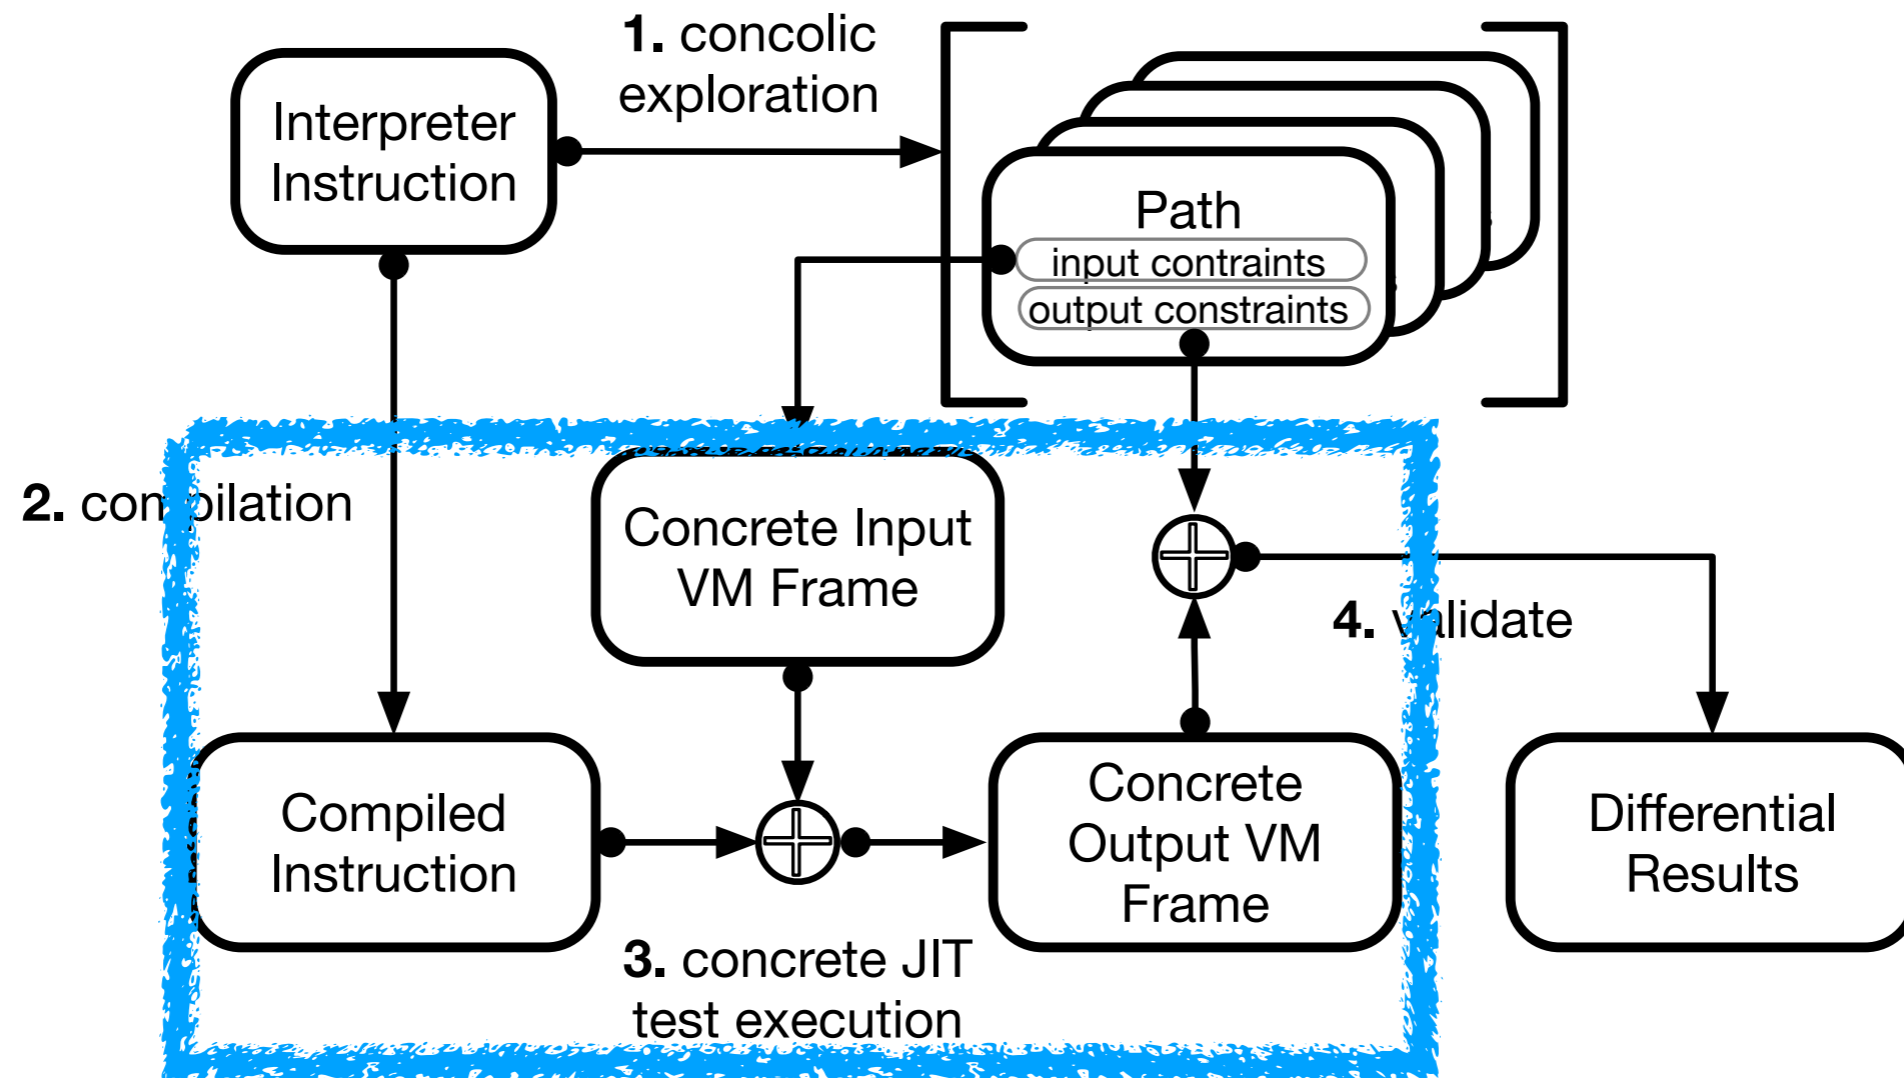
# Implementation View



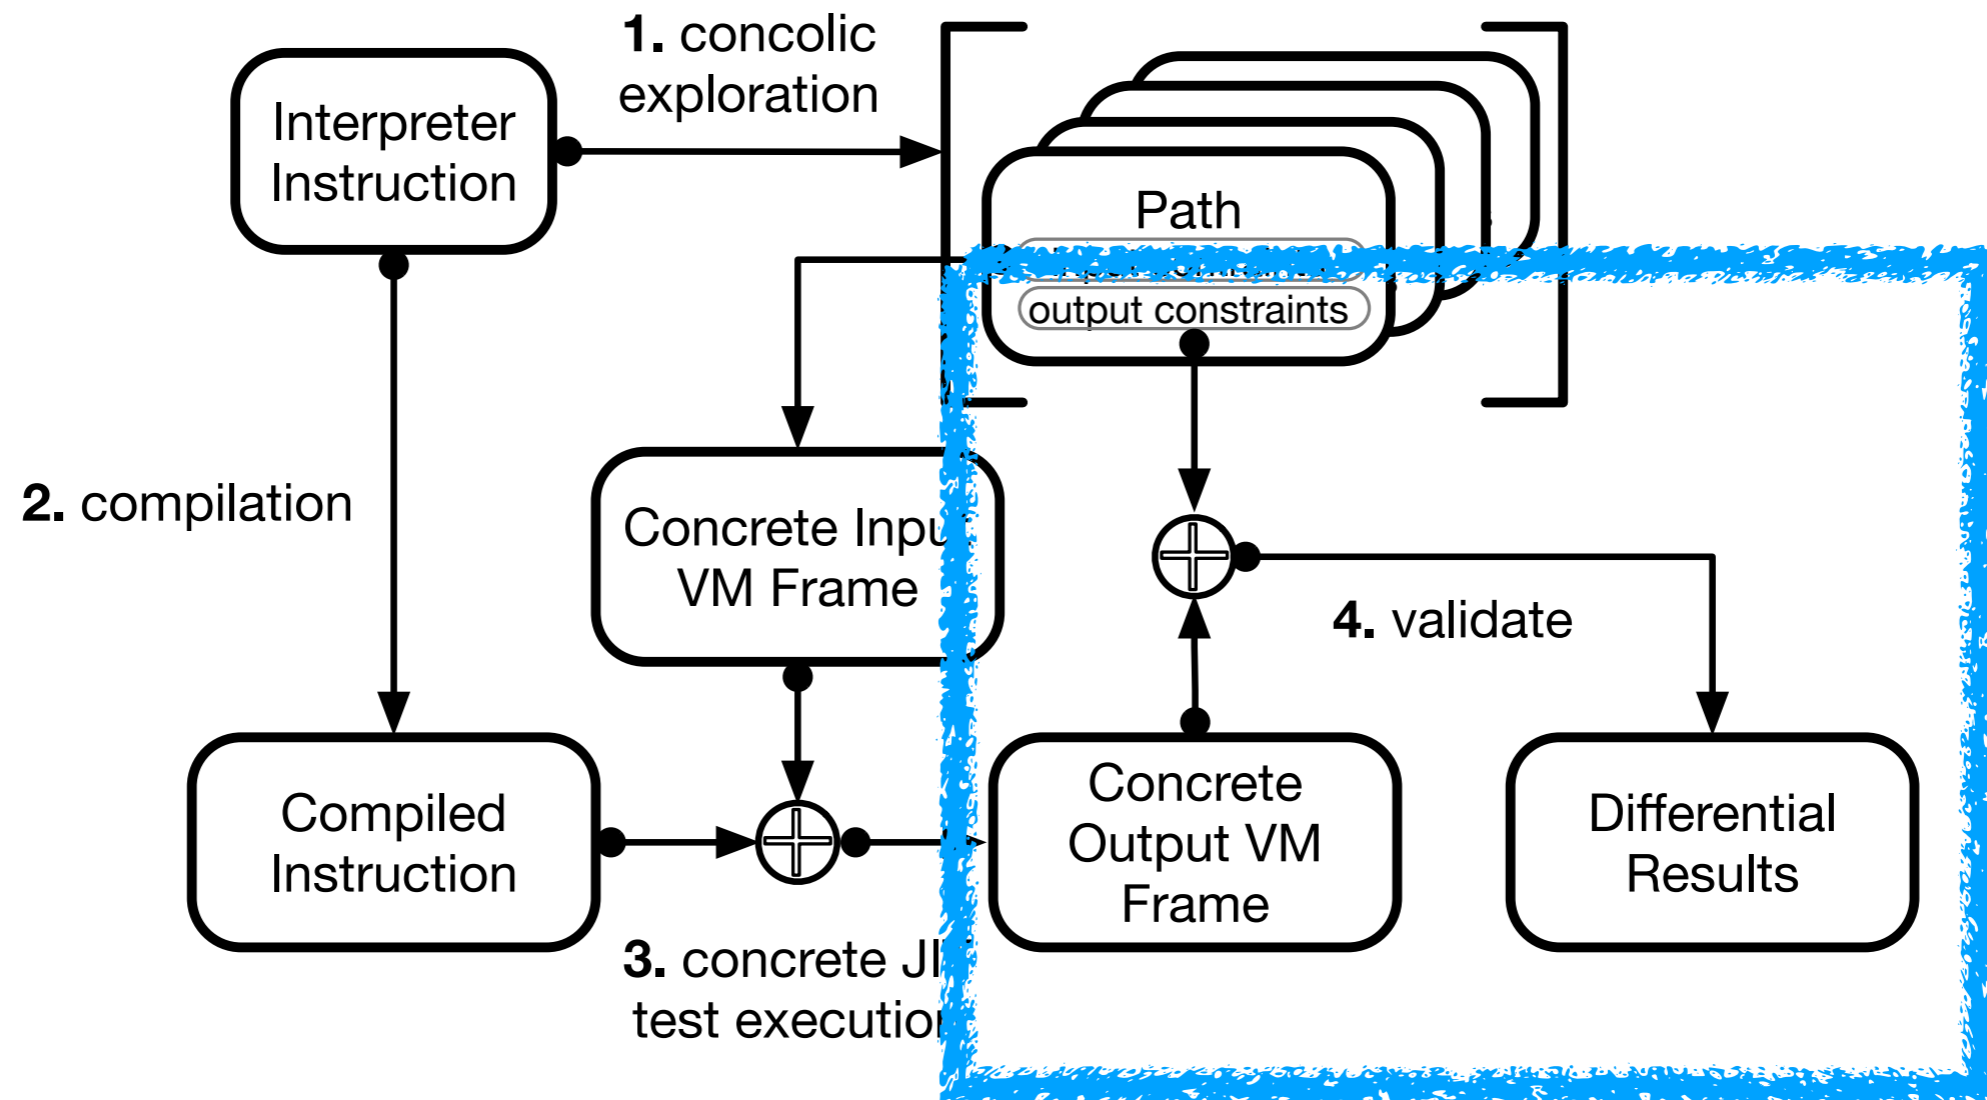
# Implementation View



# Implementation View

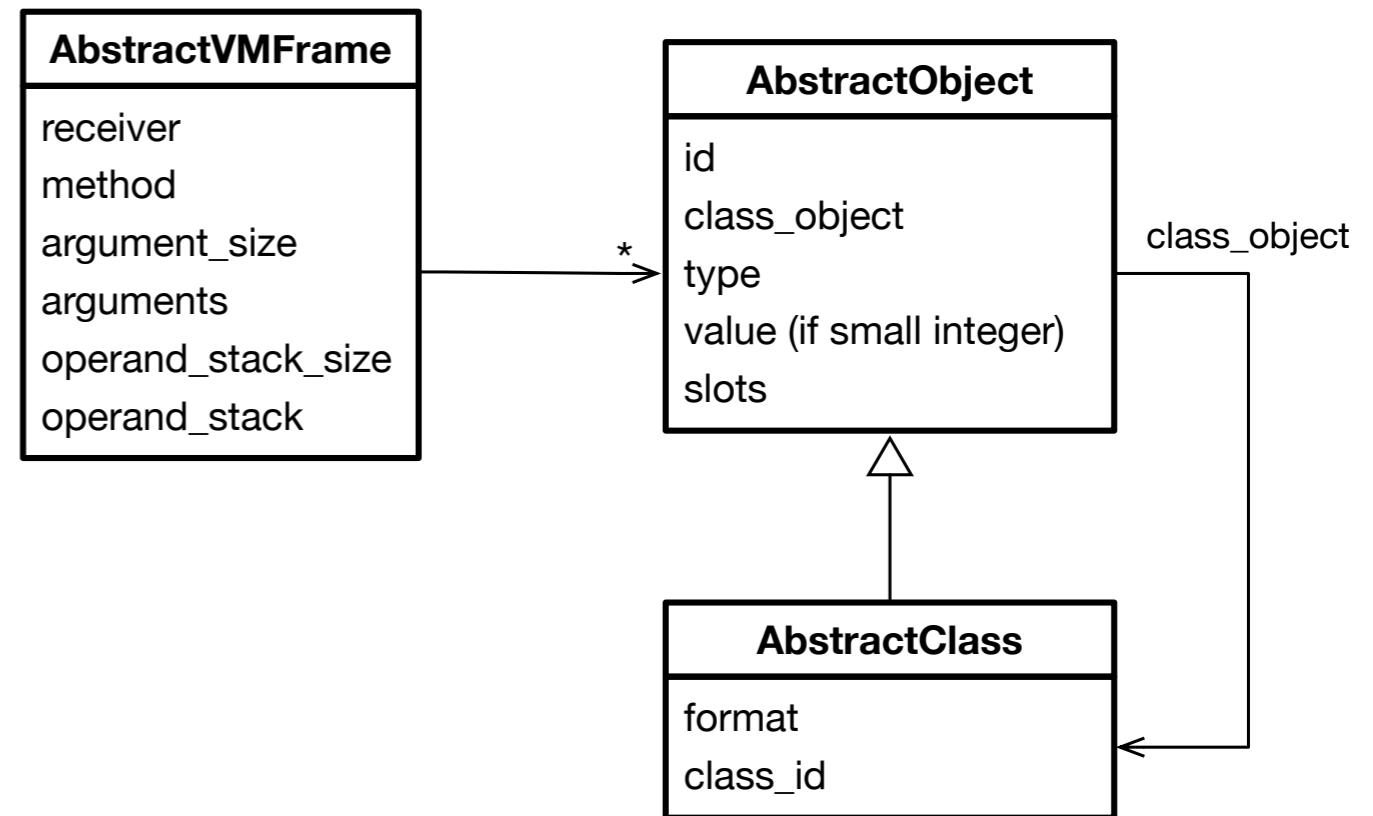


# Implementation View



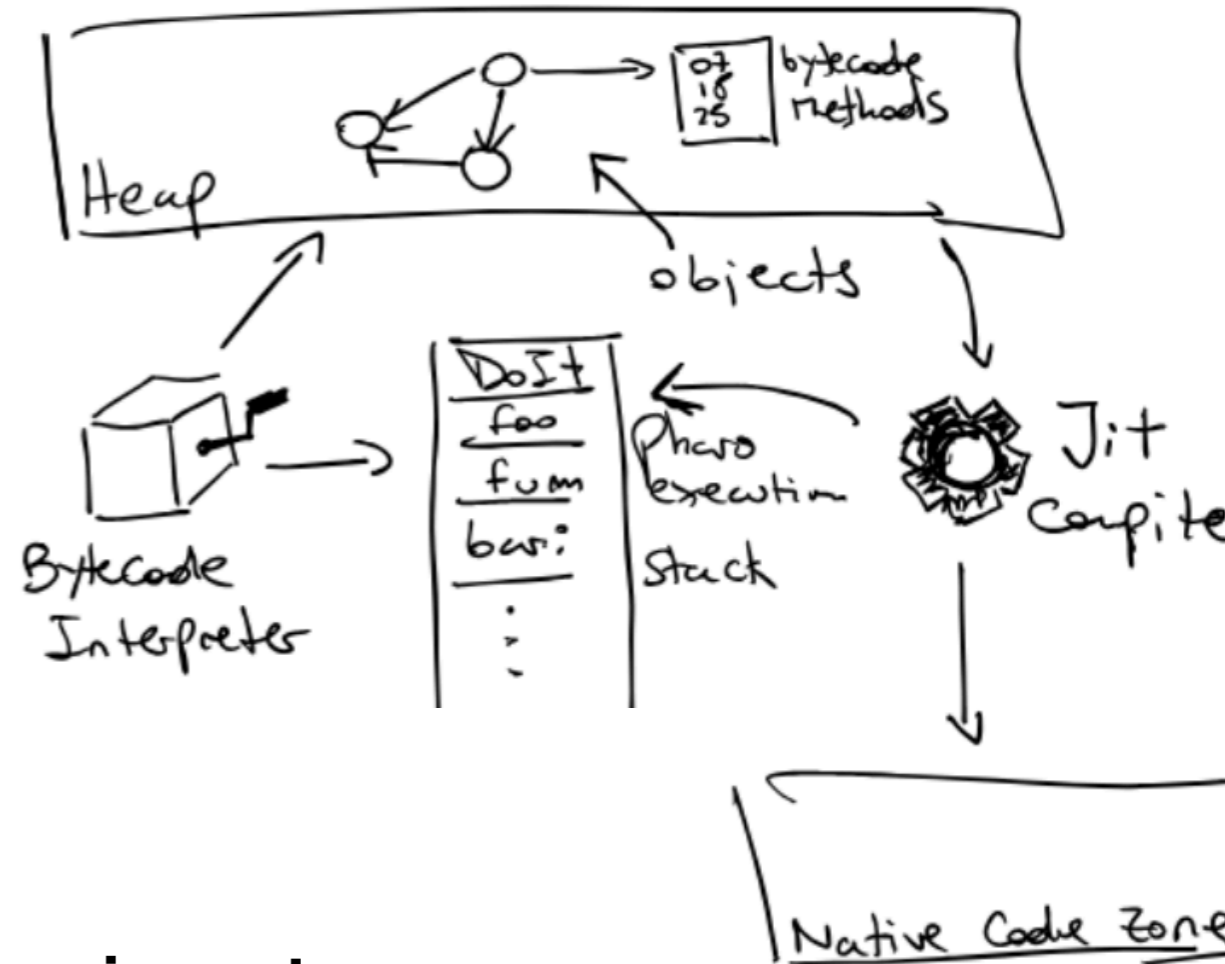
# Concolic Meta-Interpretation Model

- Models VM behaviour during concolic execution
  - Frame
  - Objects + types
  - Classes
- Then flattened into SAT solver equations



# Experimental Context: The Pharo VM

- Interpreted-compiled mixed execution
- Some numbers:
  - 255 stack based bytecodes
  - ~340 primitives/native methods
  - 146 different IR instructions
  - x86, x86-64, ARMv7, ARMv8, RISC-V
- Industrial consortium:
  - **28 International companies, 26 academic partners**



# Previous Manual Testing Effort

- No useful unit tests by ~06/2020
  - Large manual testing effort during 2020 while porting to ARM64bits
    - Extended VM simulation with a (TDD compatible) unit testing infrastructure
    - **450+** written tests on the interpreter and the garbage collector\*
    - **580+** written tests on the JIT compiler\*
    - Parametrisable for 32 and 64bits, **ARM32, ARM64, x86, x86\_64** numbers by 05/2020
- Cross-ISA Testing of the Pharo VM. Lessons learned while porting to ARMv8 64bits.





# Evaluation

- 3 bytecode compilers + 1 native method compiler
- 4928 tests generated

Compiler	# Tested Instructions	# Interpreter Paths	# Curated Paths	# Differences (%)
Native Methods (primitives)	112	2024	1520	440 (28,95%)
Simple Stack BC Compiler	175	1308	1136	18 (1,59%)
Stack-to-Register BC Compiler	175	1308	1136	10 (0,88%)
Linear-Scan Allocator BC Compiler	175	1308	1136	10 (0,88%)
<b>Total</b>	<b>637</b>	<b>5948</b>	<b>4928</b>	<b>478 (9,7%)</b>



# Analysis of Differences through Manual Inspection

- 91 causes, *6 different categories*
- Errors both in the interpreter AND the compilers

- 14 causes of ***segmentation***

<u>Family</u>	<u># Cases</u>
Missing interpreter type check	1
Missing compiled type check	13
Optimisation difference	10
Behavioral difference	5
Missing Functionality	60
Simulation Error	2




# Characterising Concolic Execution

## Paths per instruction

- Native methods present in average more paths than bytecode instructions

=> longer time to explore

=> potentially more bugs




**Paths per  
Type of  
Instruction**



# Practical and Cheap

- Test generation ~5 minutes
- Total run time of ~10 seconds
  - Avg 30ms per instruction



# More in the article!

- Discovered Bugs
- Concolic Model
- Testing Infrastructure

## Interpreter-Guided Differential JIT Compiler Unit Testing

Guillermo Polito  
Univ. Lille, CNRS, Inria, Centrale Lille,  
UMR 9189 CRISTAL, F-59000 Lille  
France  
guillermo.polito@univ-lille.fr

Stéphane Ducasse  
Univ. Lille, Inria, CNRS, Centrale Lille,  
UMR 9189 CRISTAL  
France  
stephane.ducasse@inria.fr

Pablo Tesone  
Pharo Consortium  
Univ. Lille, Inria, CNRS, Centrale Lille,  
UMR 9189 CRISTAL  
France  
pablo.tesone@inria.fr

### Abstract

Modern language implementations using Virtual Machines feature diverse execution engines such as byte-code interpreters and machine-code dynamic translators, a.k.a. JIT compilers. Validating such engines requires not only validating each in isolation, but also that they are functionally equivalent. Tests should be duplicated for each execution engine, exercising the same execution paths on each of them.

In this paper, we present a novel automated testing approach for virtual machines featuring byte-code interpreters. Our solution uses concolic meta-interpretation: it applies concolic testing to a byte-code interpreter to explore all possible execution interpreter paths and obtain a list of concrete values that explore such paths. We then use such values to enable differential testing on the VM interpreter and JIT

San Diego, CA, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3519939.3523457>

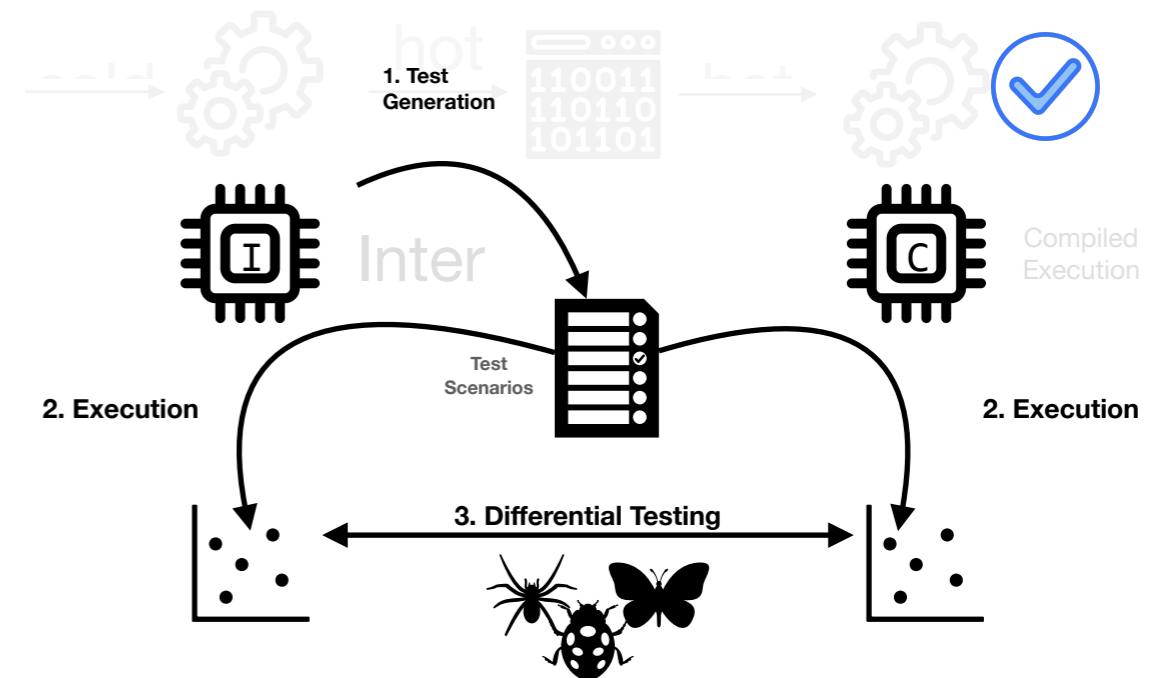
### 1 Introduction

Modern Virtual Machines support code generation for compilation and dynamic code patching for techniques such as inline caching. They are often structured around a baseline code interpreter, a baseline JIT compiler, and a speculative inliner. This complexity is aggravated when the VM supports and runs on multiple target architectures [1]. Validating the execution of interpreted code and its compiled counterparts is challenging.

Several solutions have been proposed to aid in VM testing tasks. Traditionally, VM simulation environments have

# Conclusion

- 478 differences found, 91 causes, 6 categories
- Practical:
  - 4928 tests generated in ~8 minutes

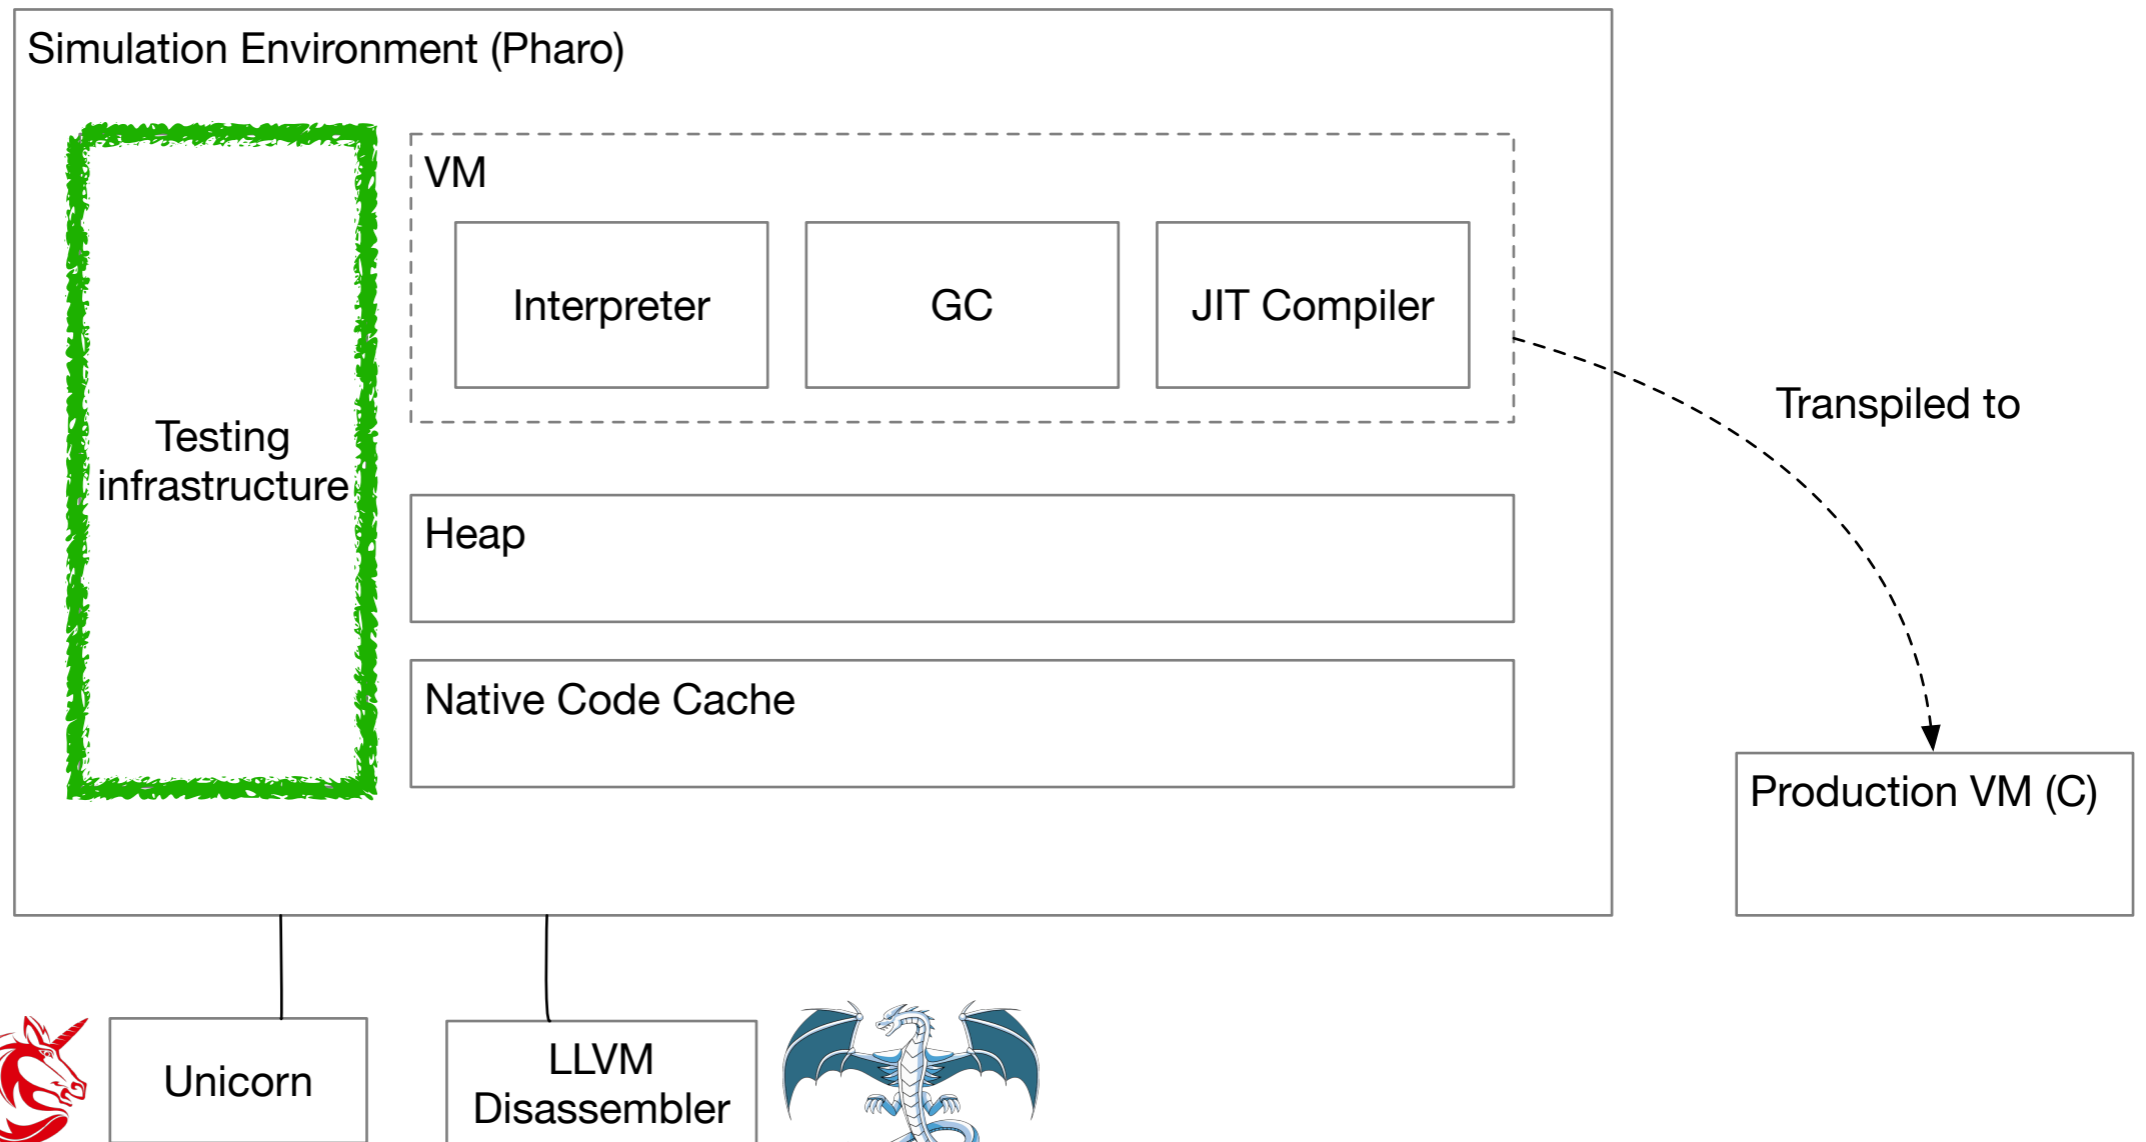


1928 tests run in ~140 seconds - Stéphane Ducasse  
**Guillermo Polito**  
[guillermo.polito@univ-lille.fr](mailto:guillermo.polito@univ-lille.fr)  
 @guillep

# Extras



# Simulation + Testing Environment





# Unit Testing Infrastructure Comparison

	Real Hardware Execution	Full-System Simulation	Unit-Testing
Feedback-cycle speed	Very low	Low	High
Availability	Low	High	High
Reproducibility	Low	Low	High
Precision	High	Low	Low
Debuggability	Low	High	High



