

Meta-compilation of Baseline JIT Compilers with Druid

Programming '25

PALUMBO Nahuel, POLITO Guillermo, TESONE Pablo, DUCASSE Stephane



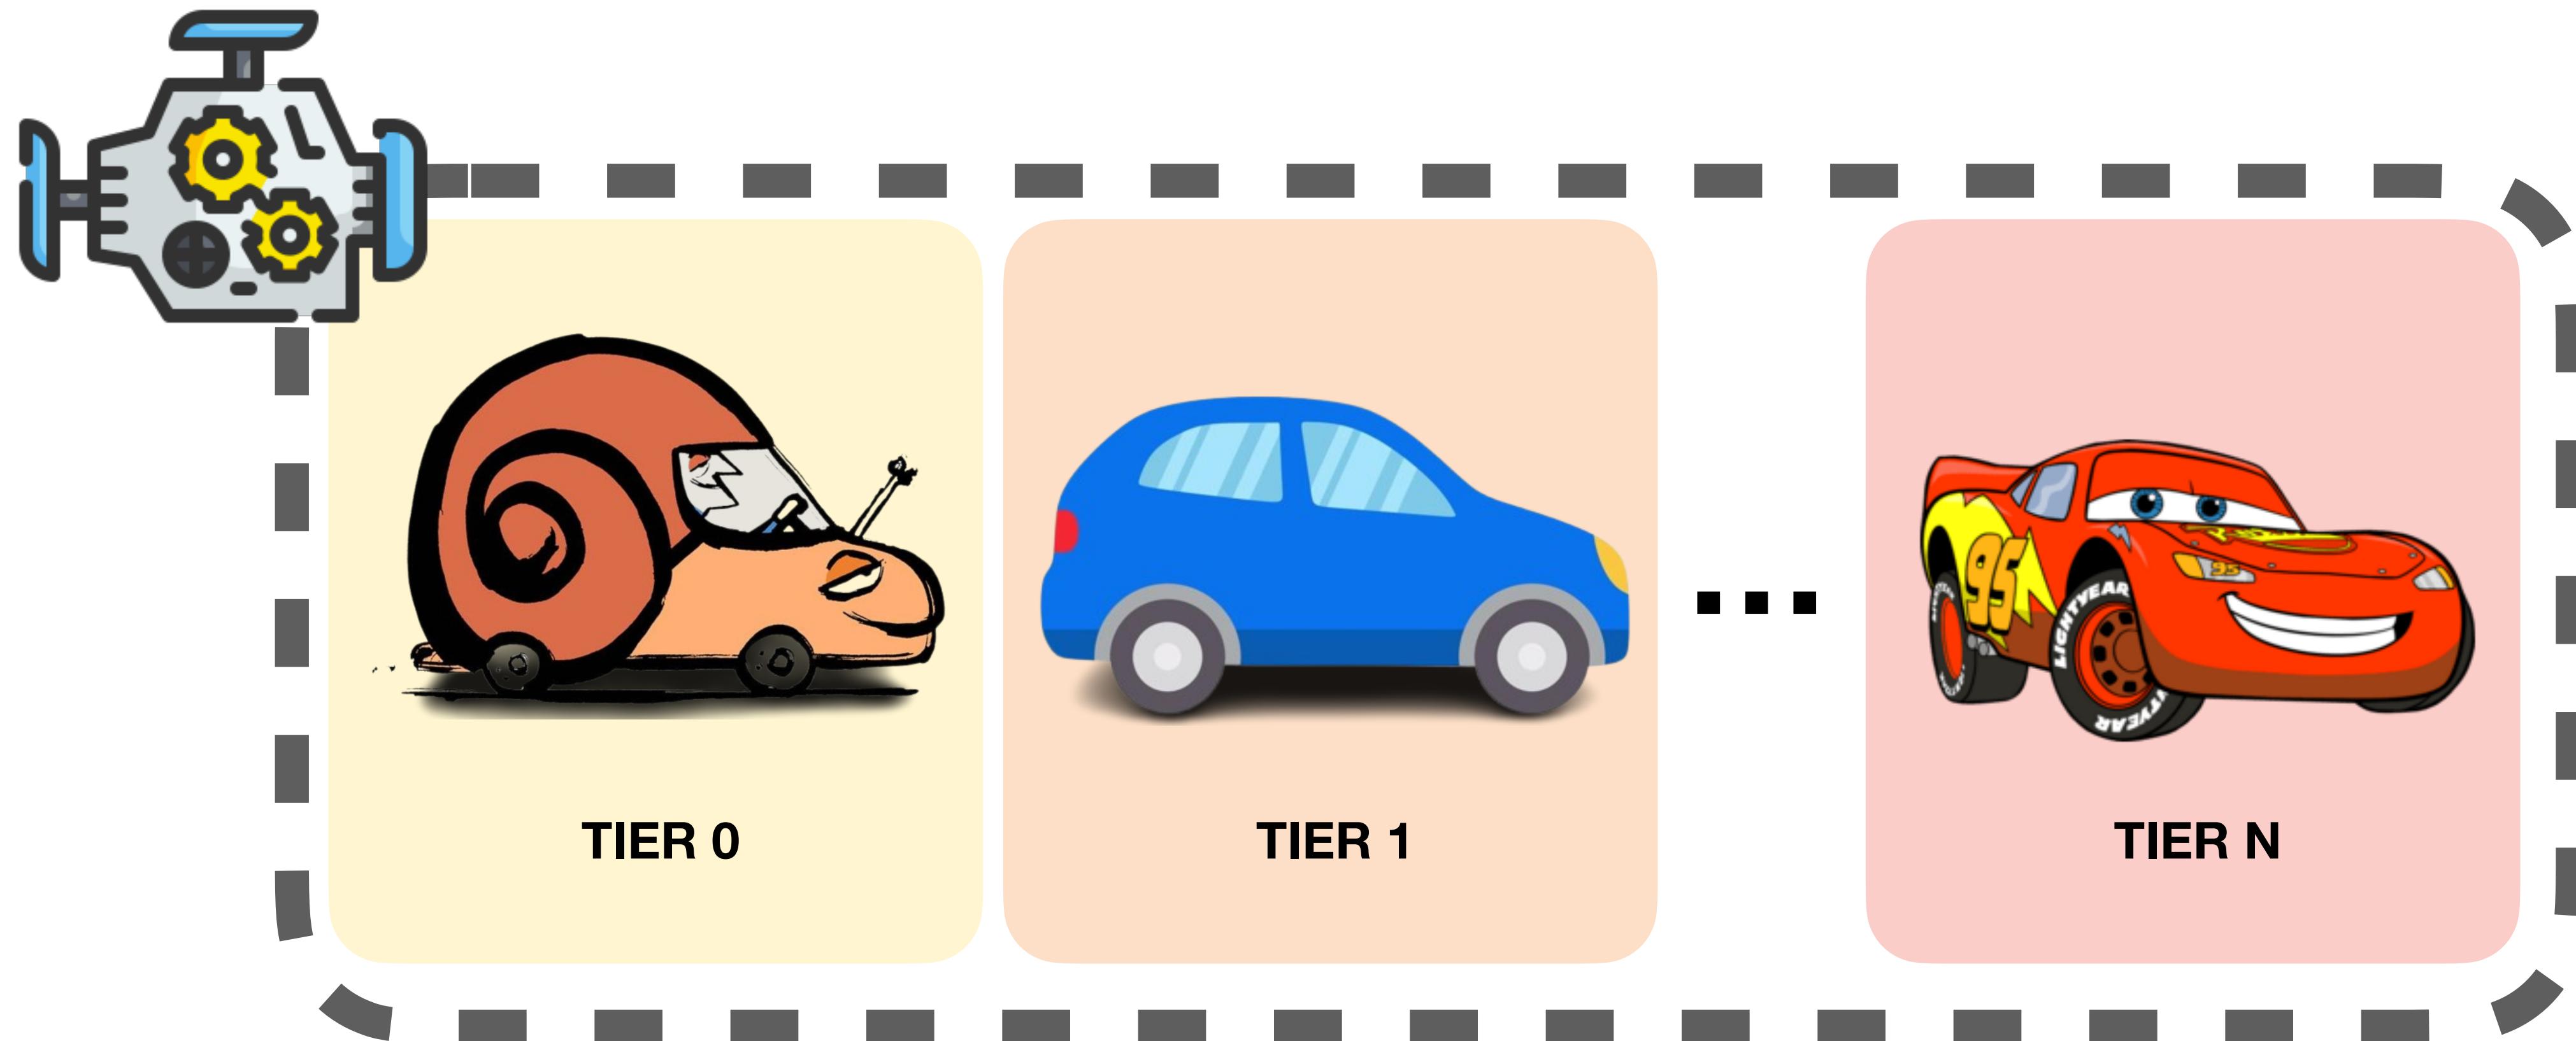
Meta-compilation of Baseline JIT Compilers with Druid

Programming '25

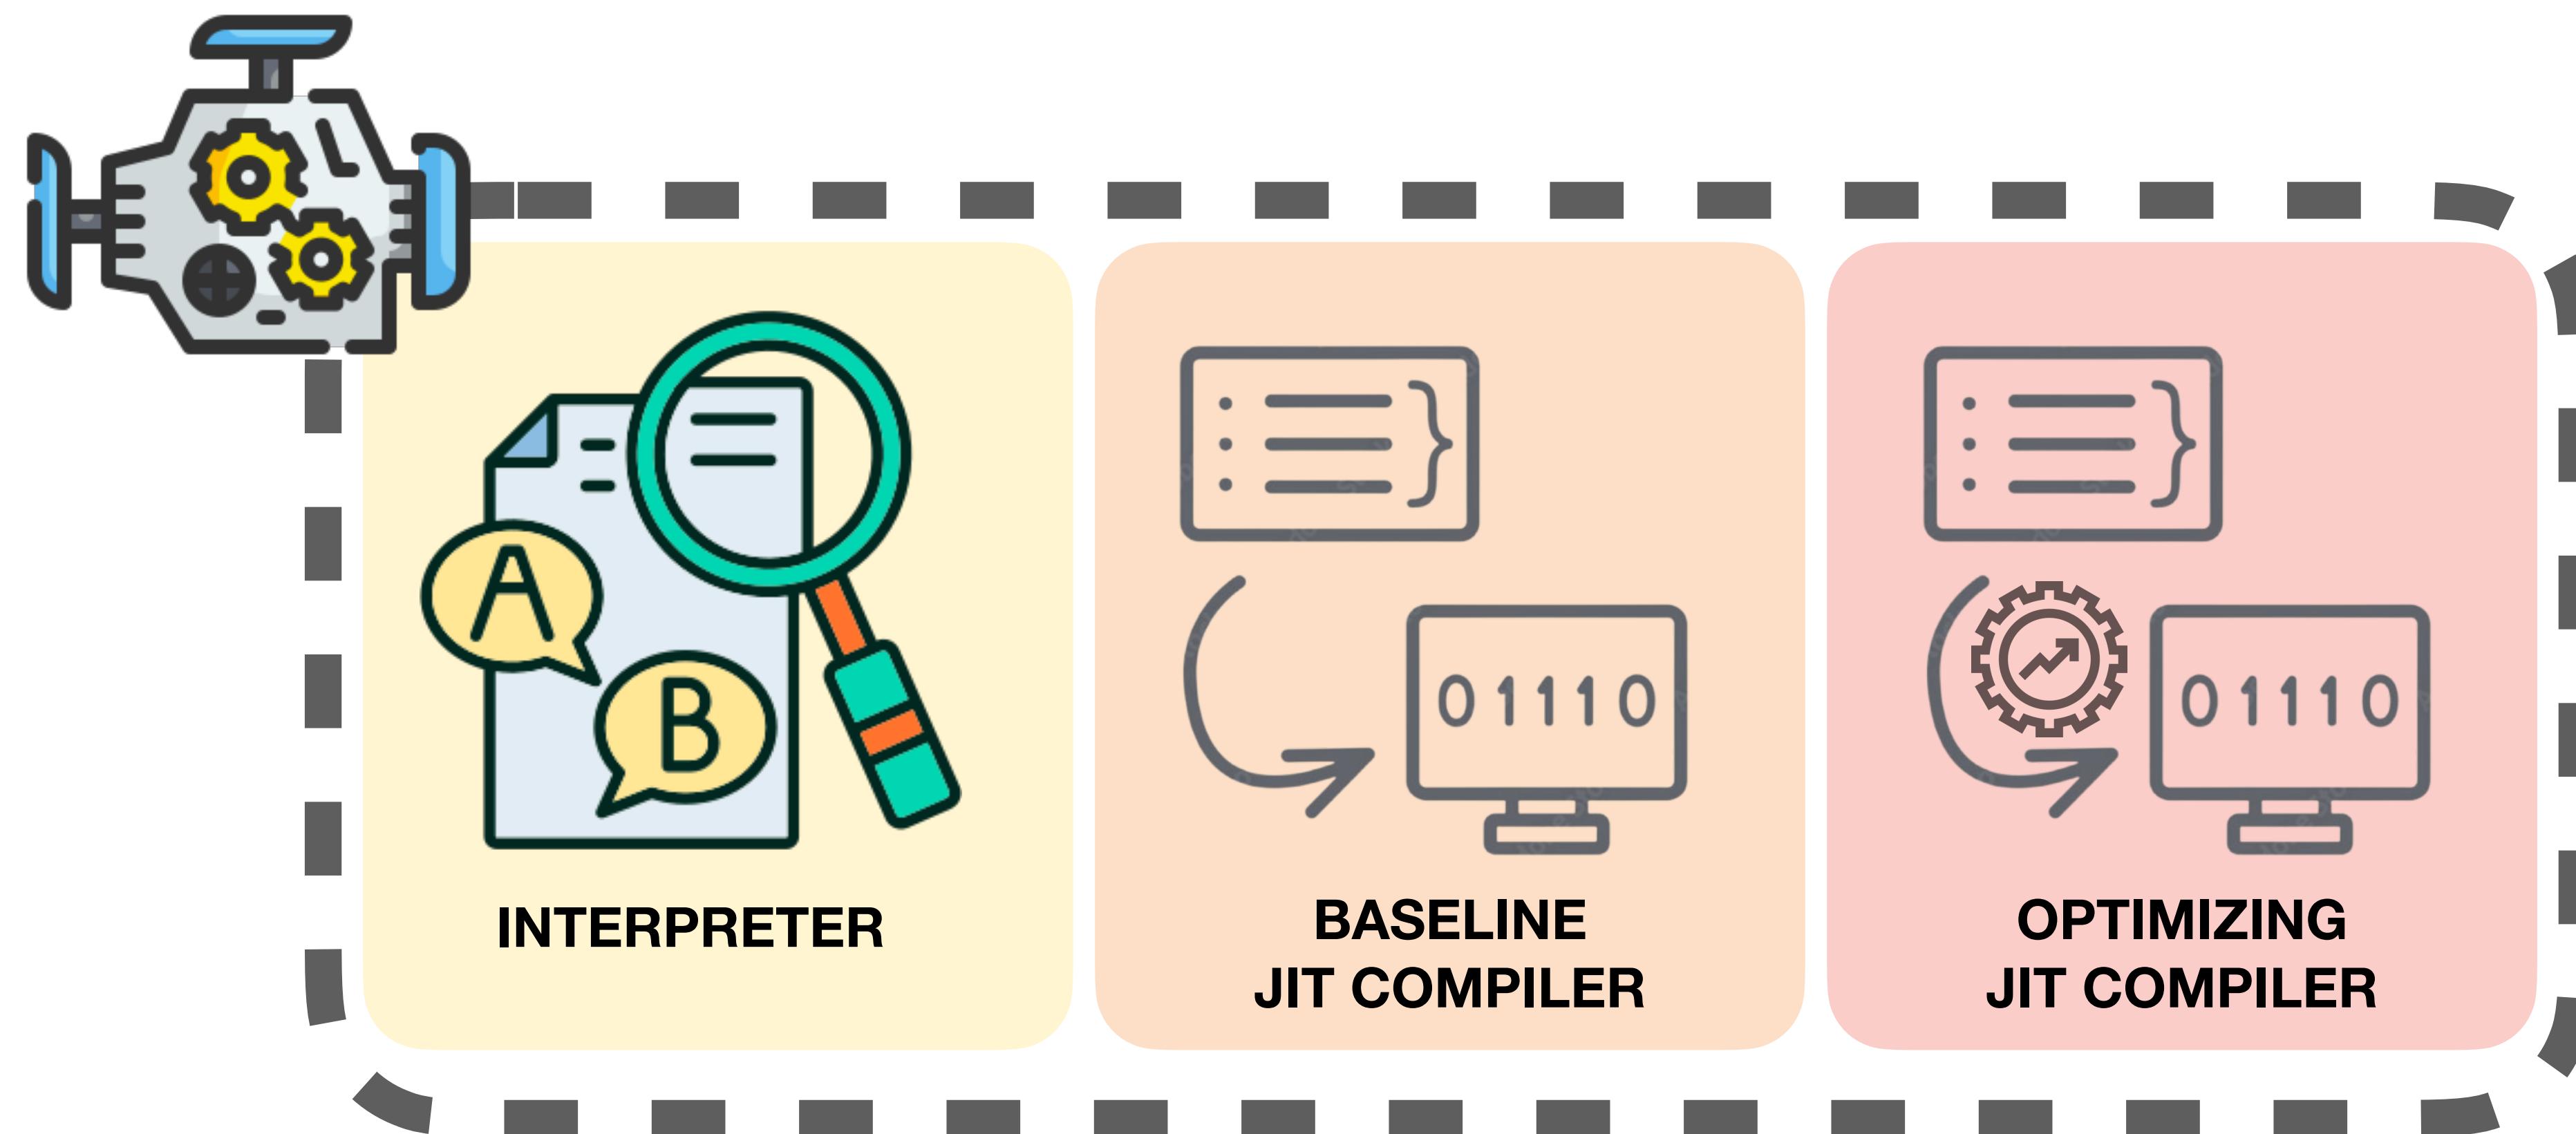
PALUMBO Nahuel, POLITO Guillermo, TESONE Pablo, DUCASSE Stephane



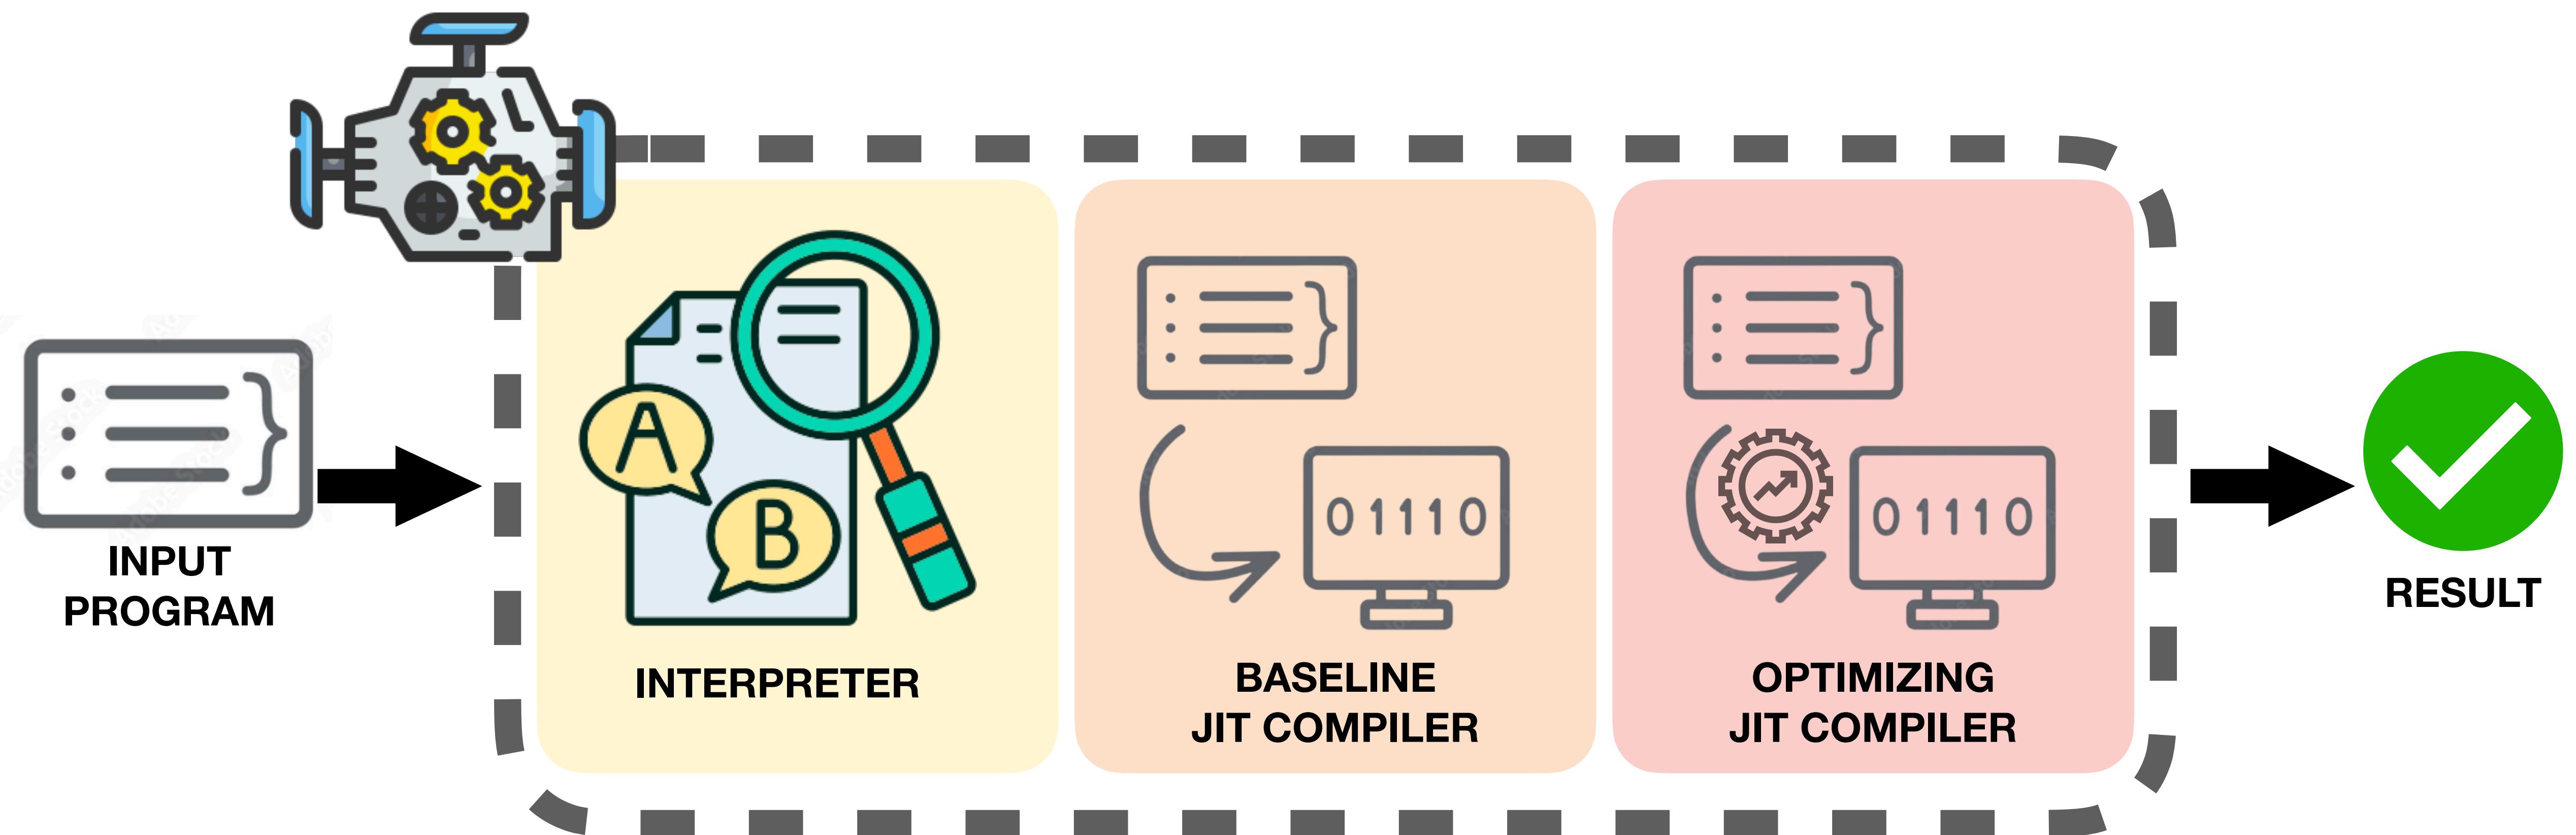
Multi-tier Virtual Machine



Multi-tier Virtual Machine

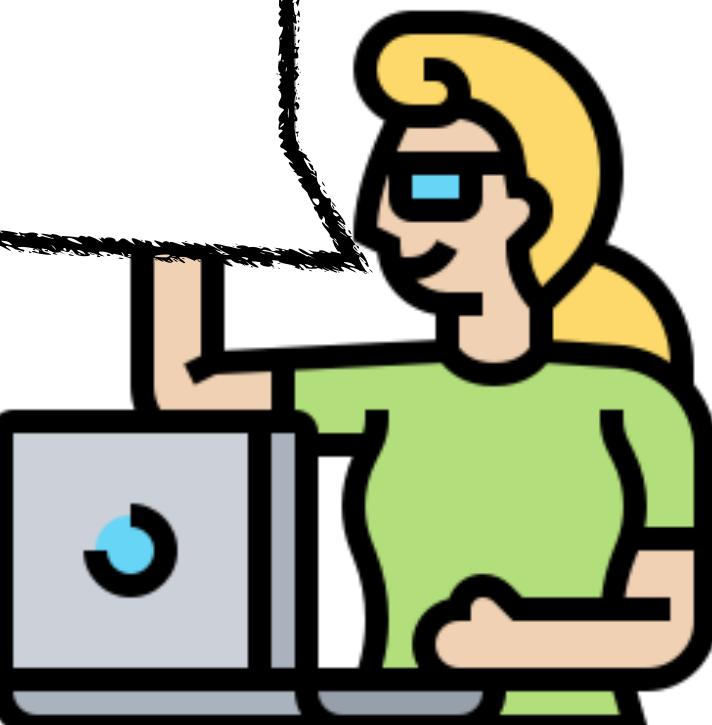


Multi-tier Virtual Machine



Multi-tier Virtual Machine

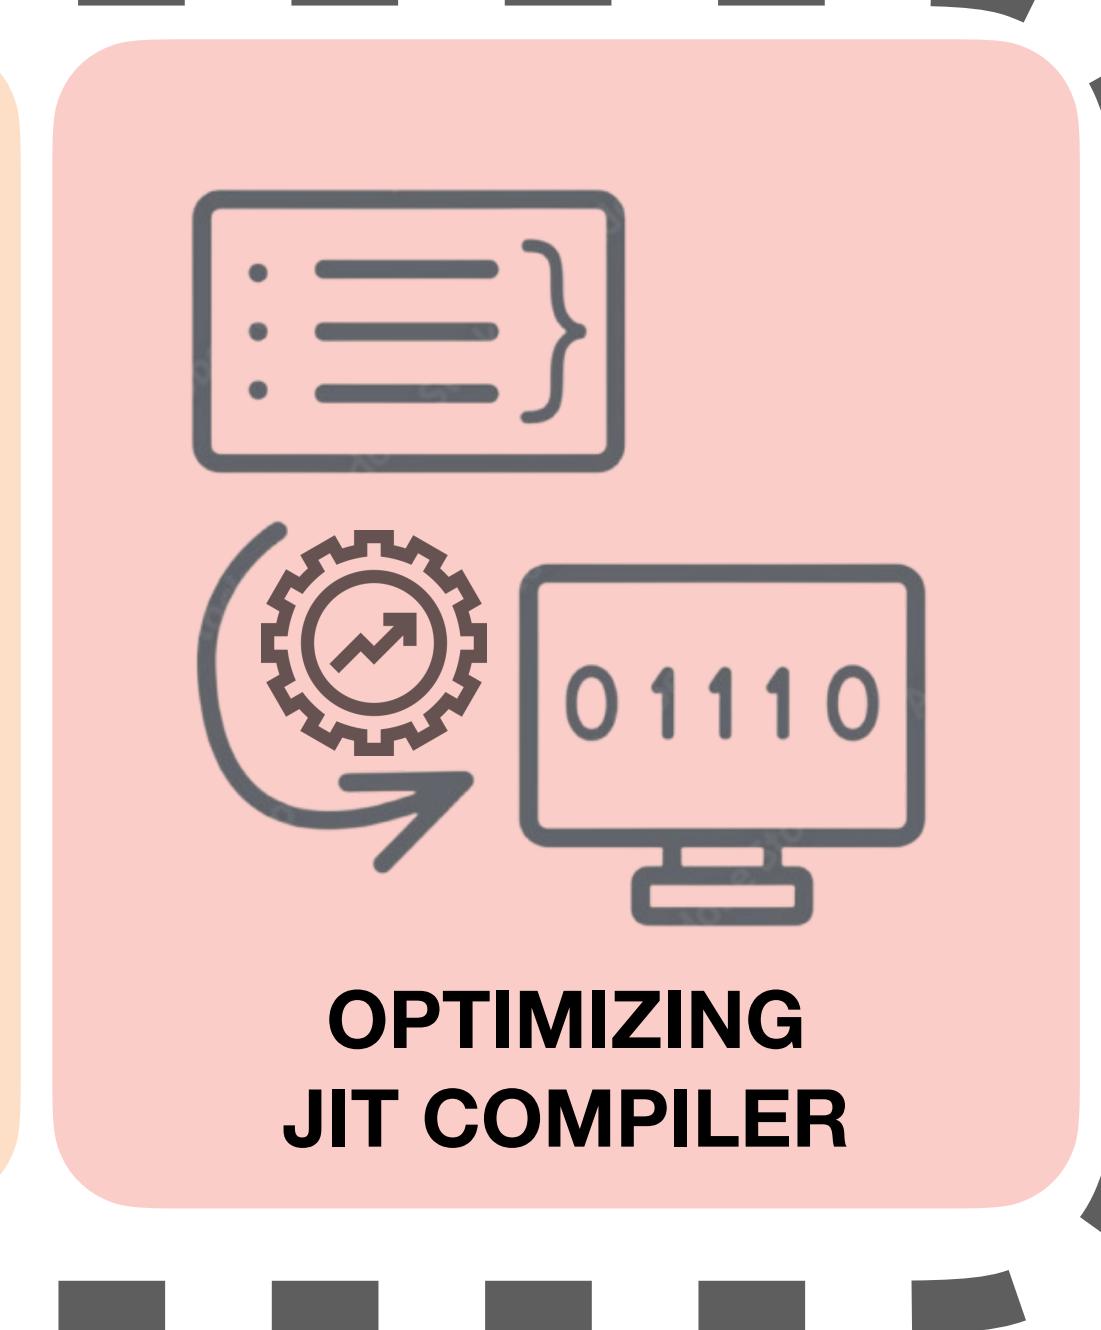
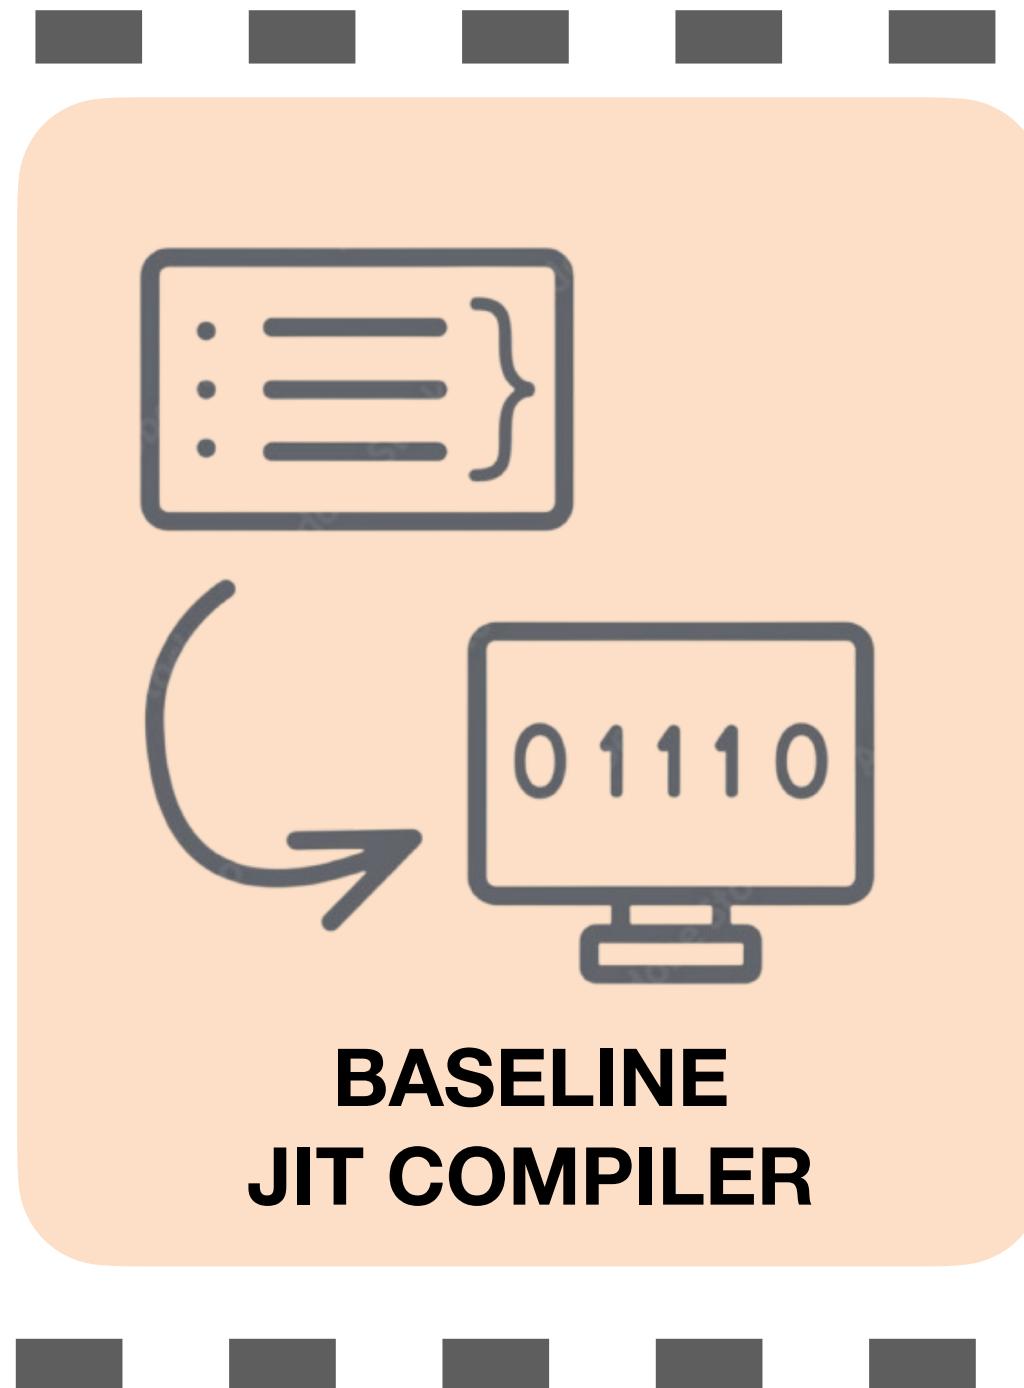
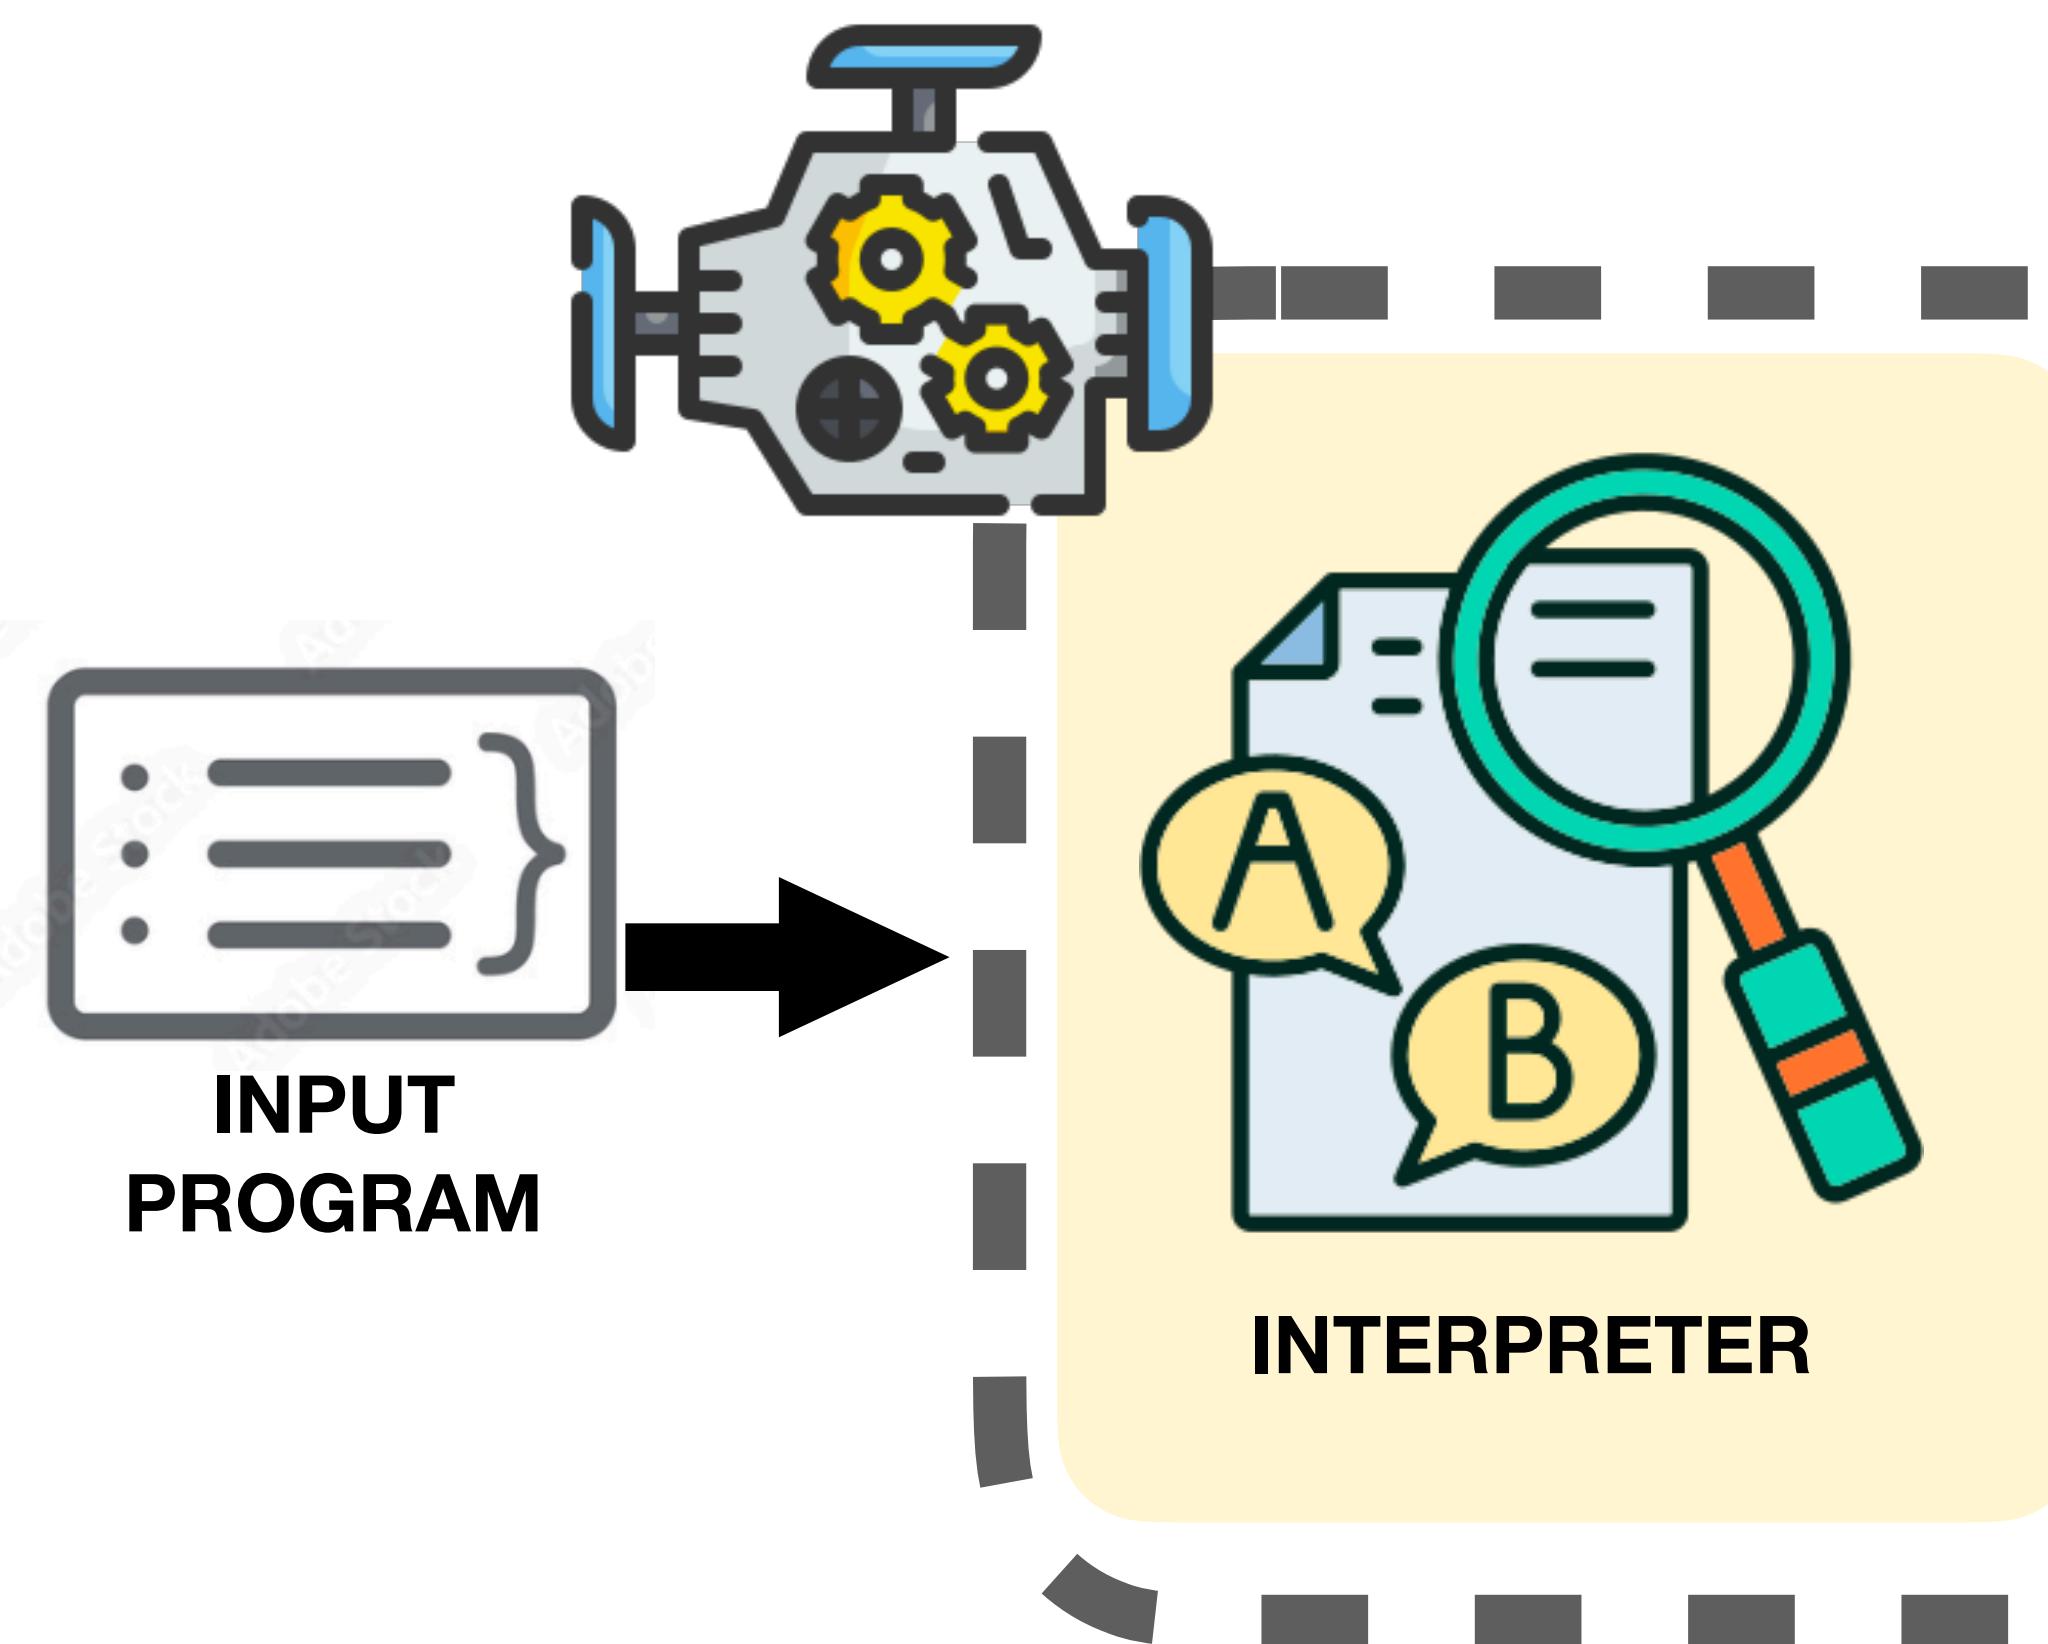
Oh no... yet another tier to maintain... 😡



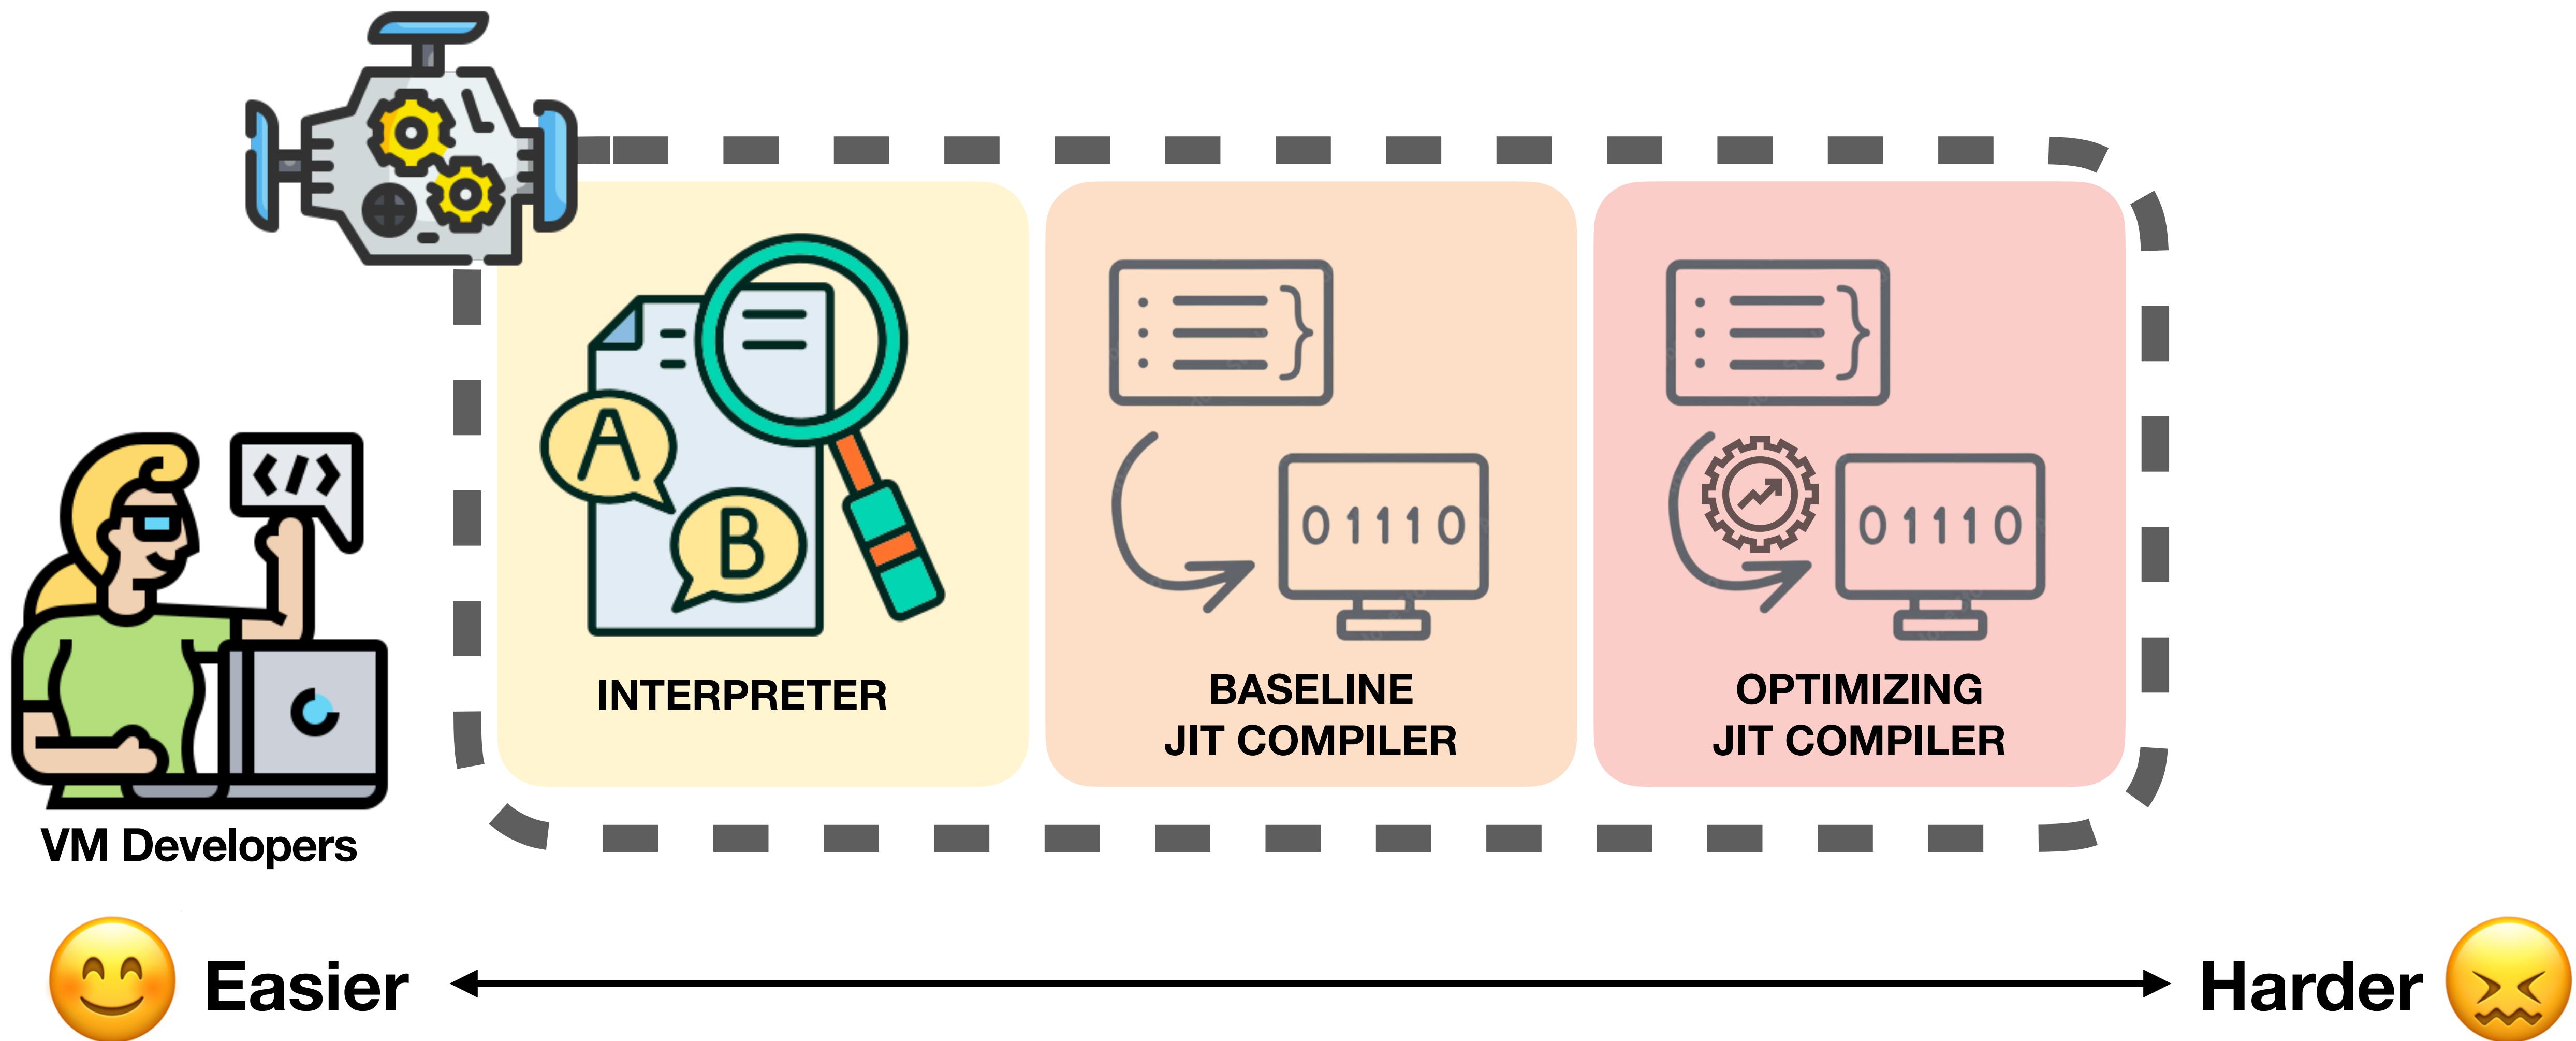
VM Developers

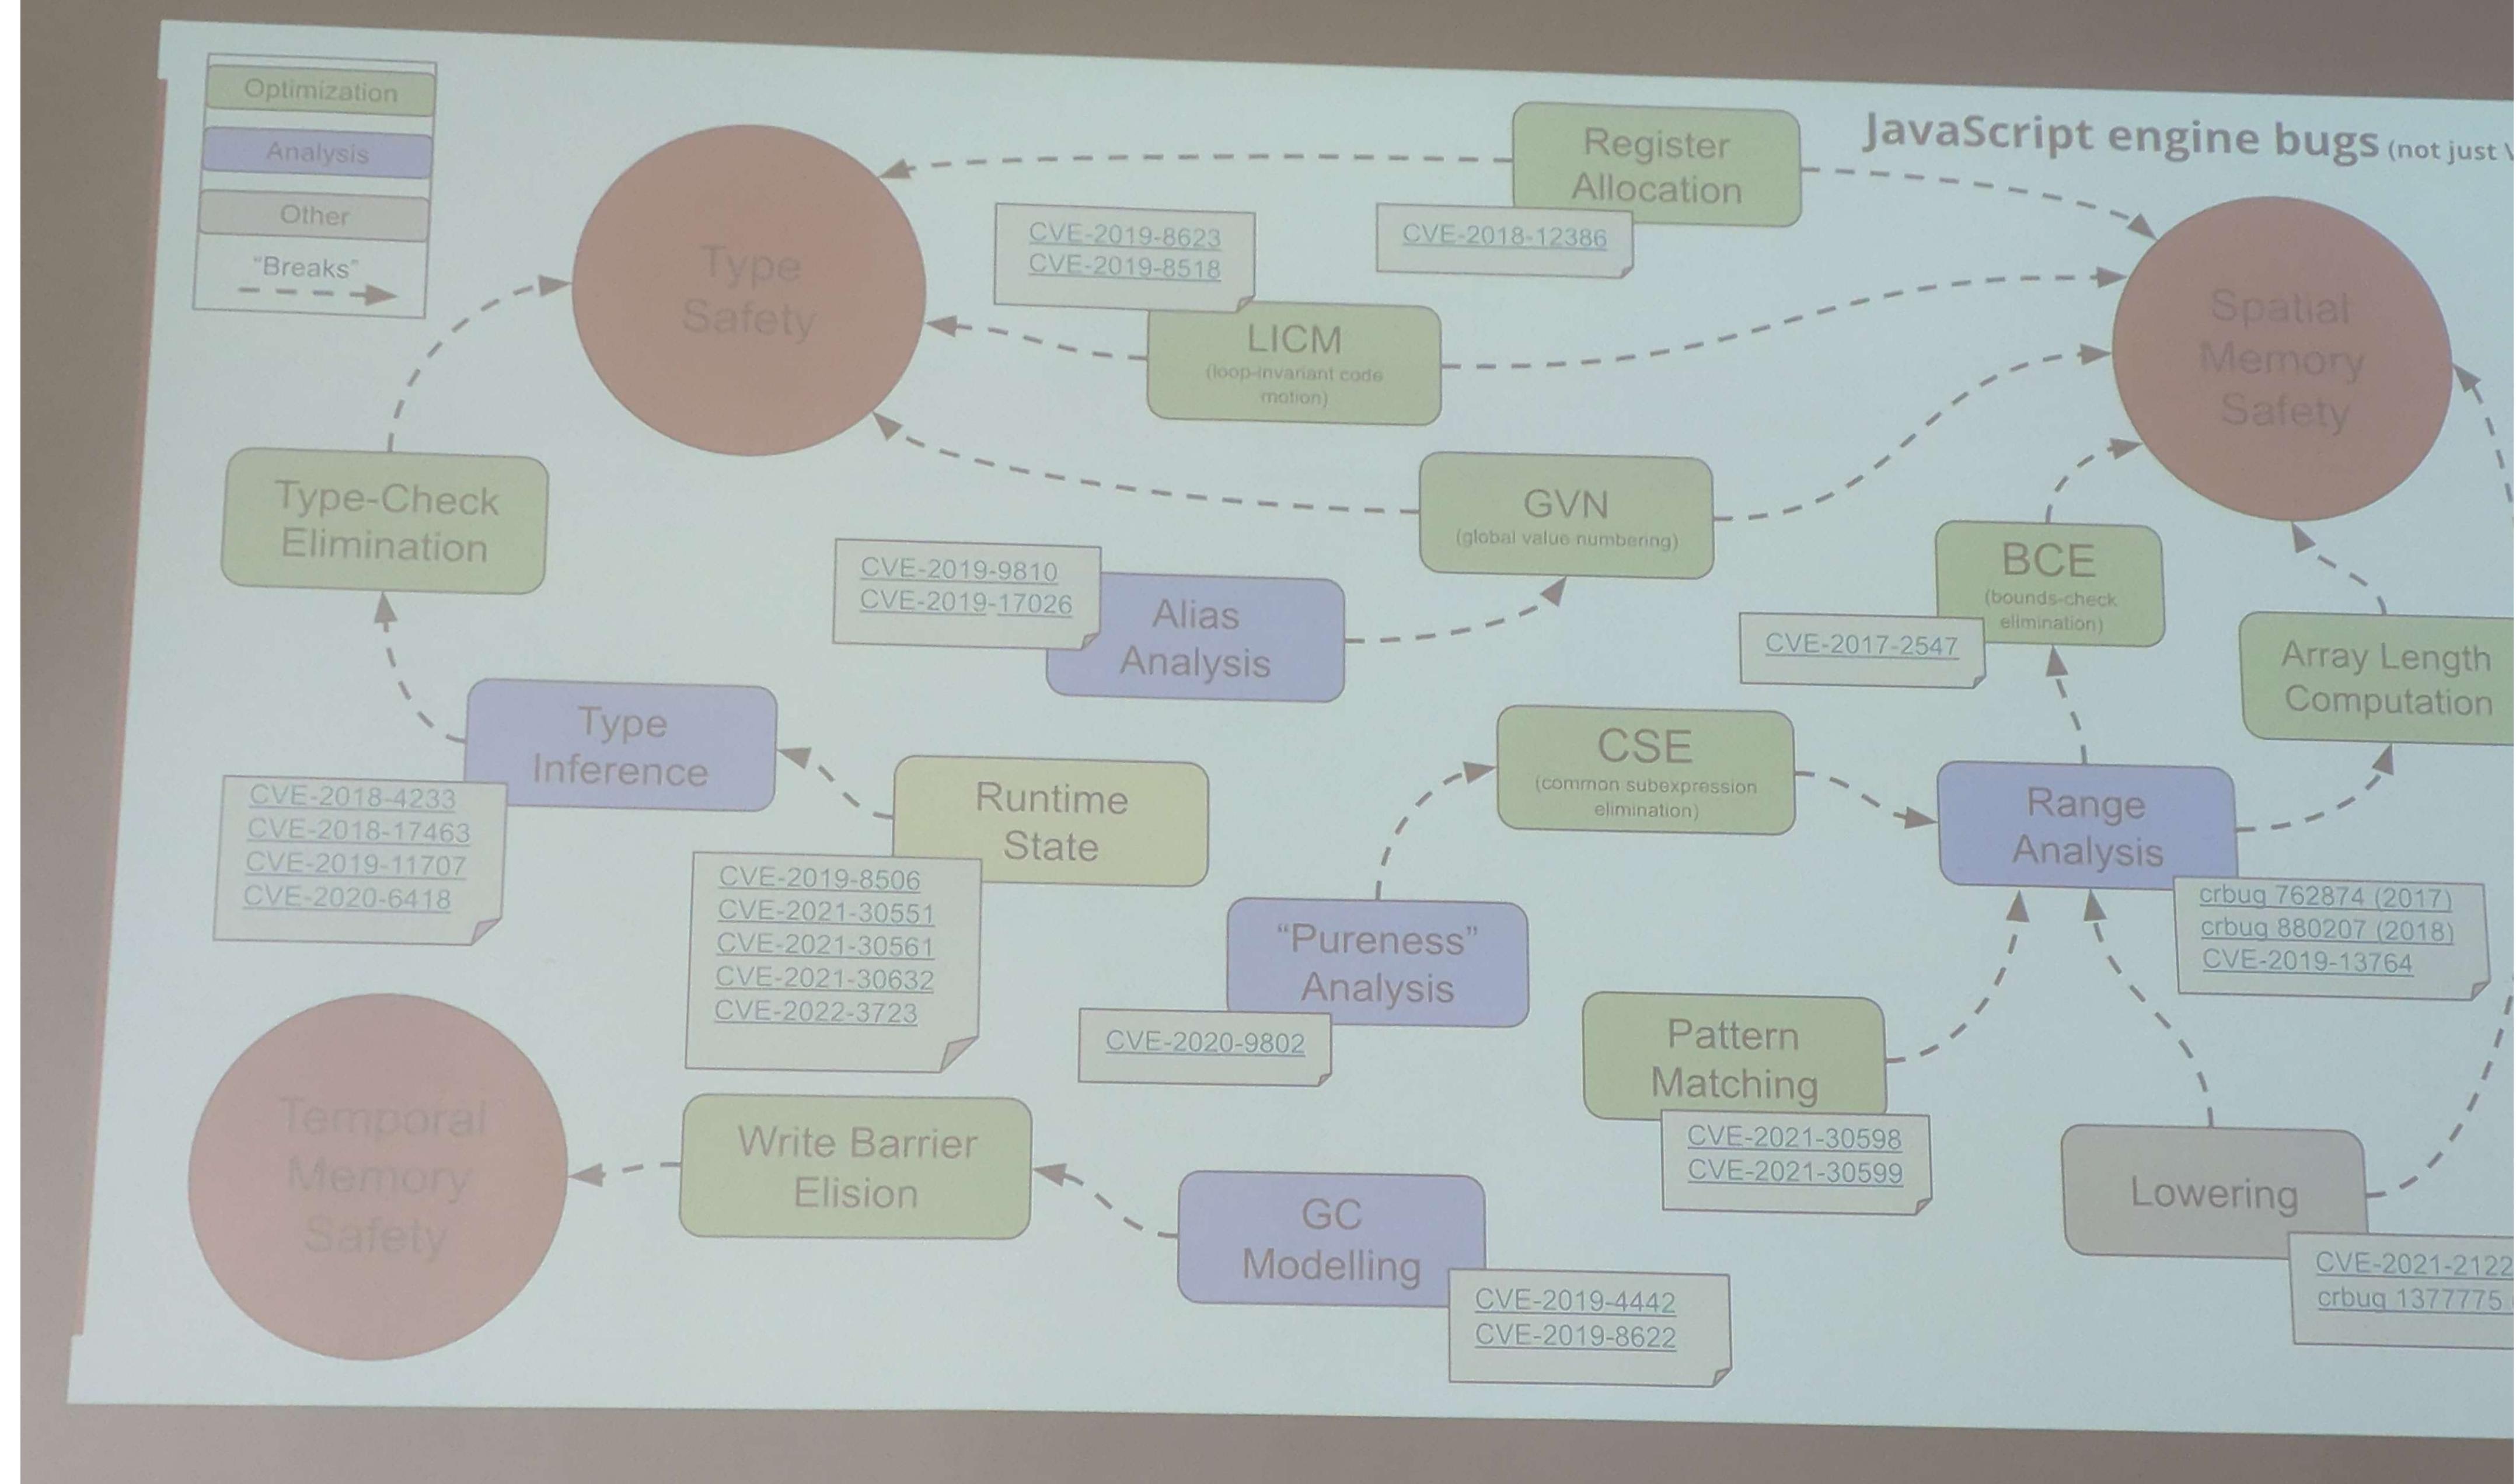


RESULT

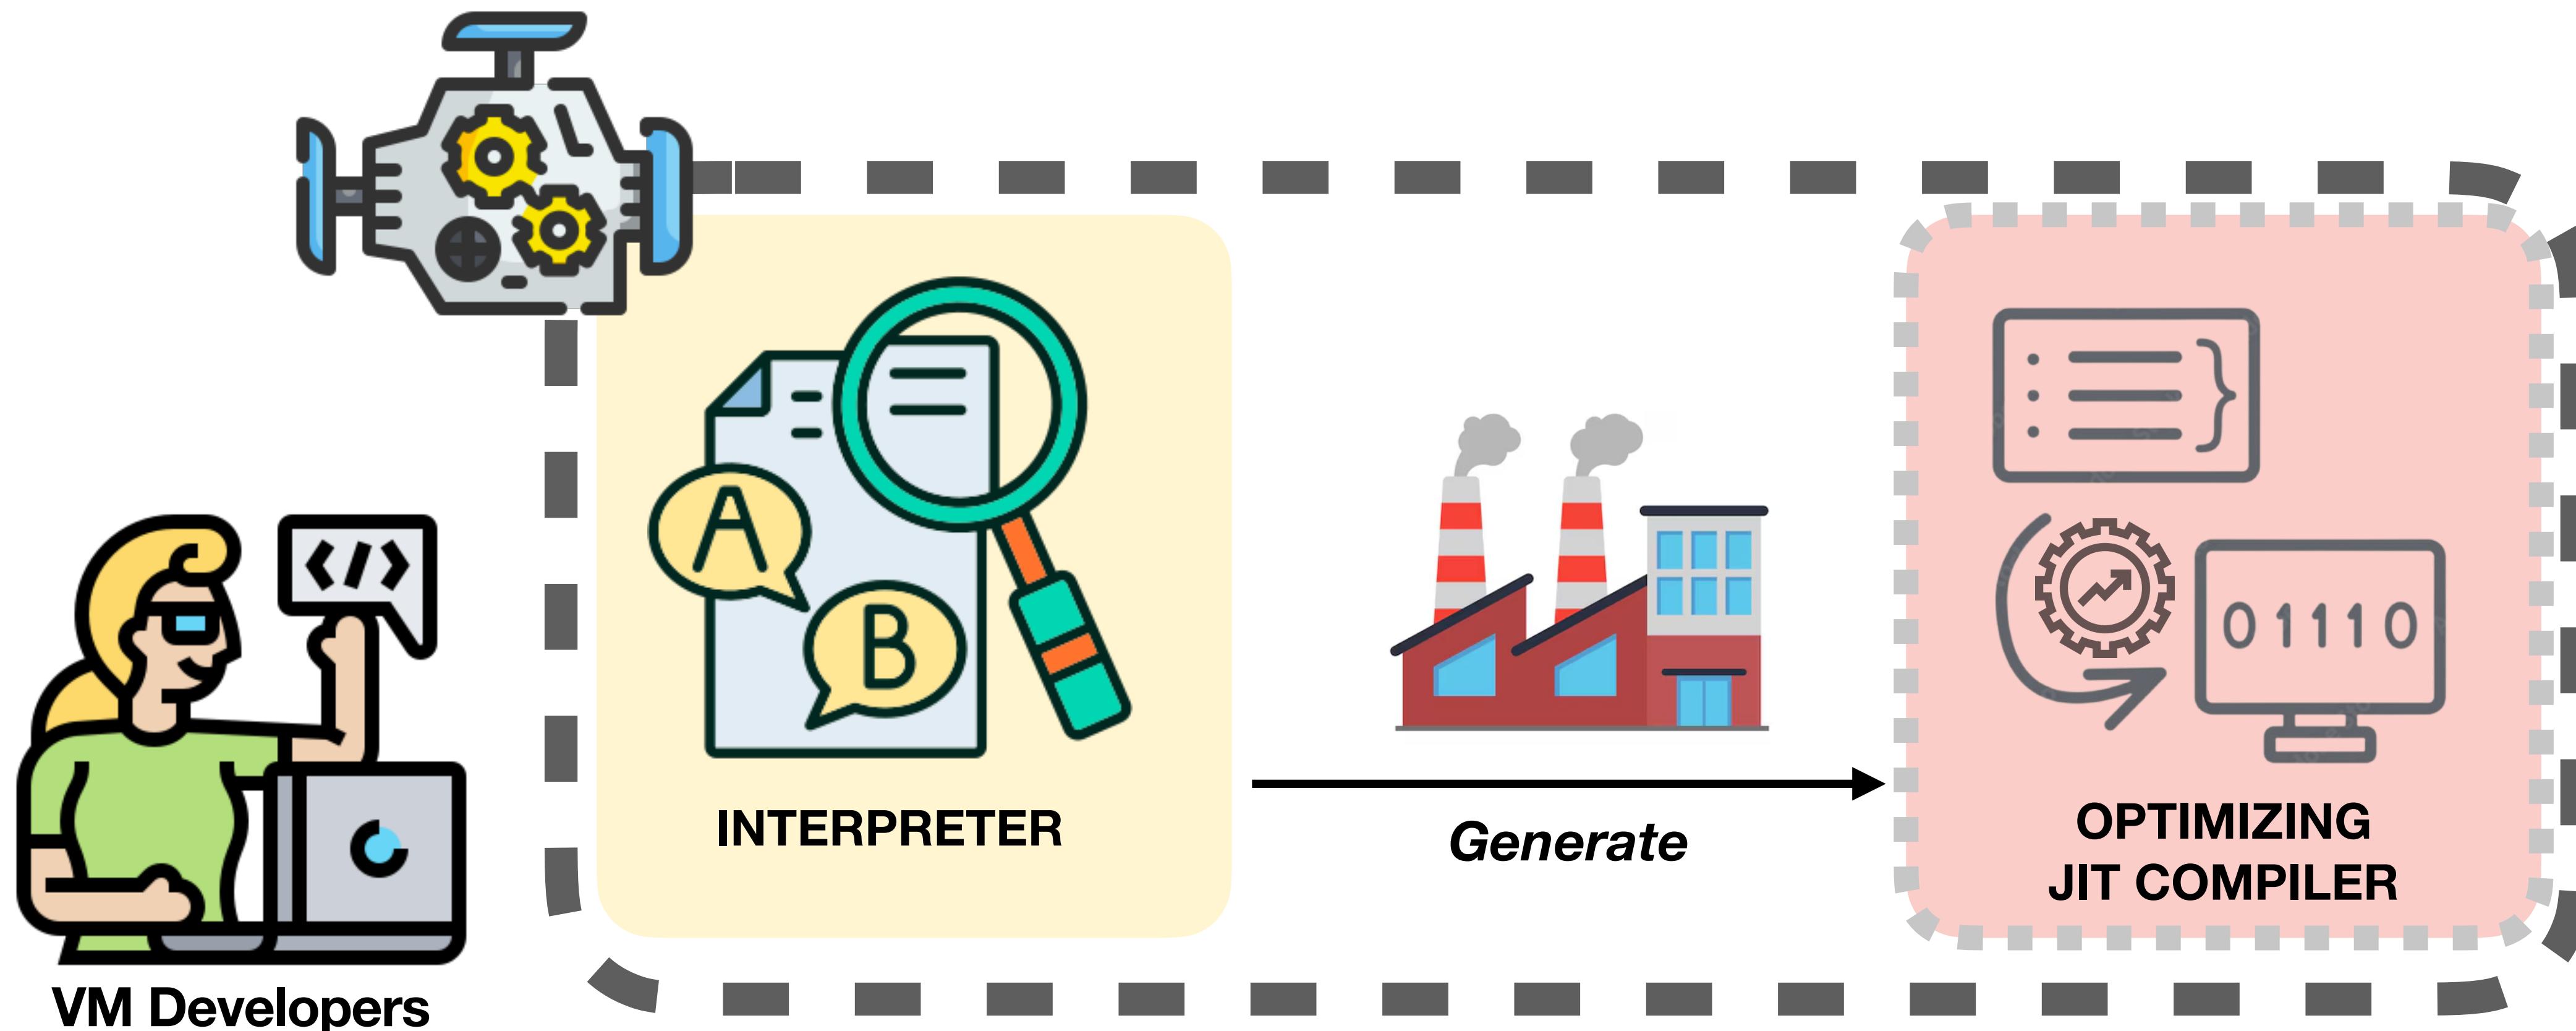


Multi-tier Virtual Machine

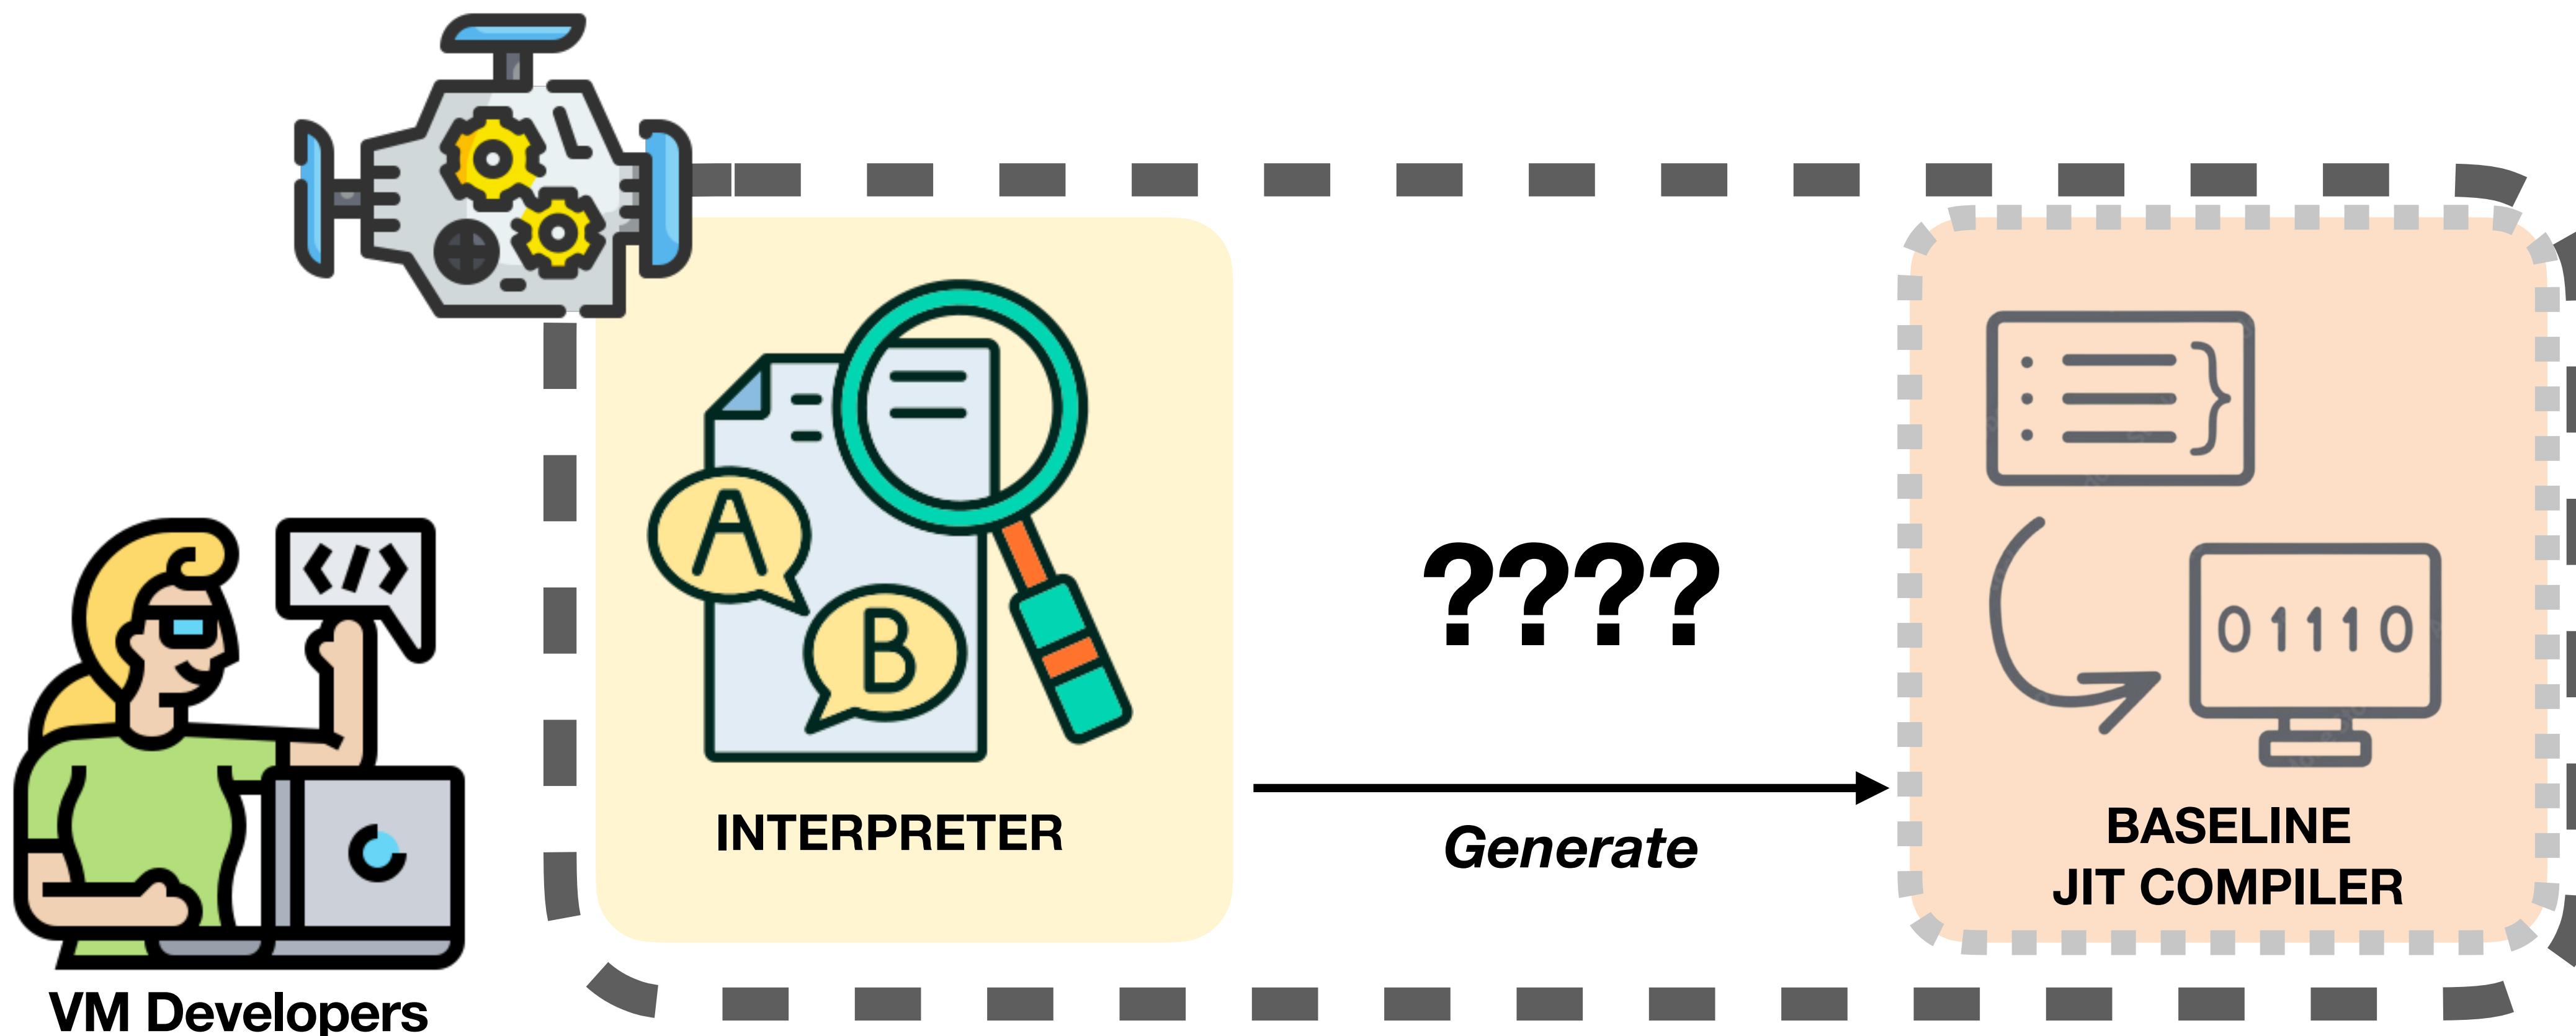




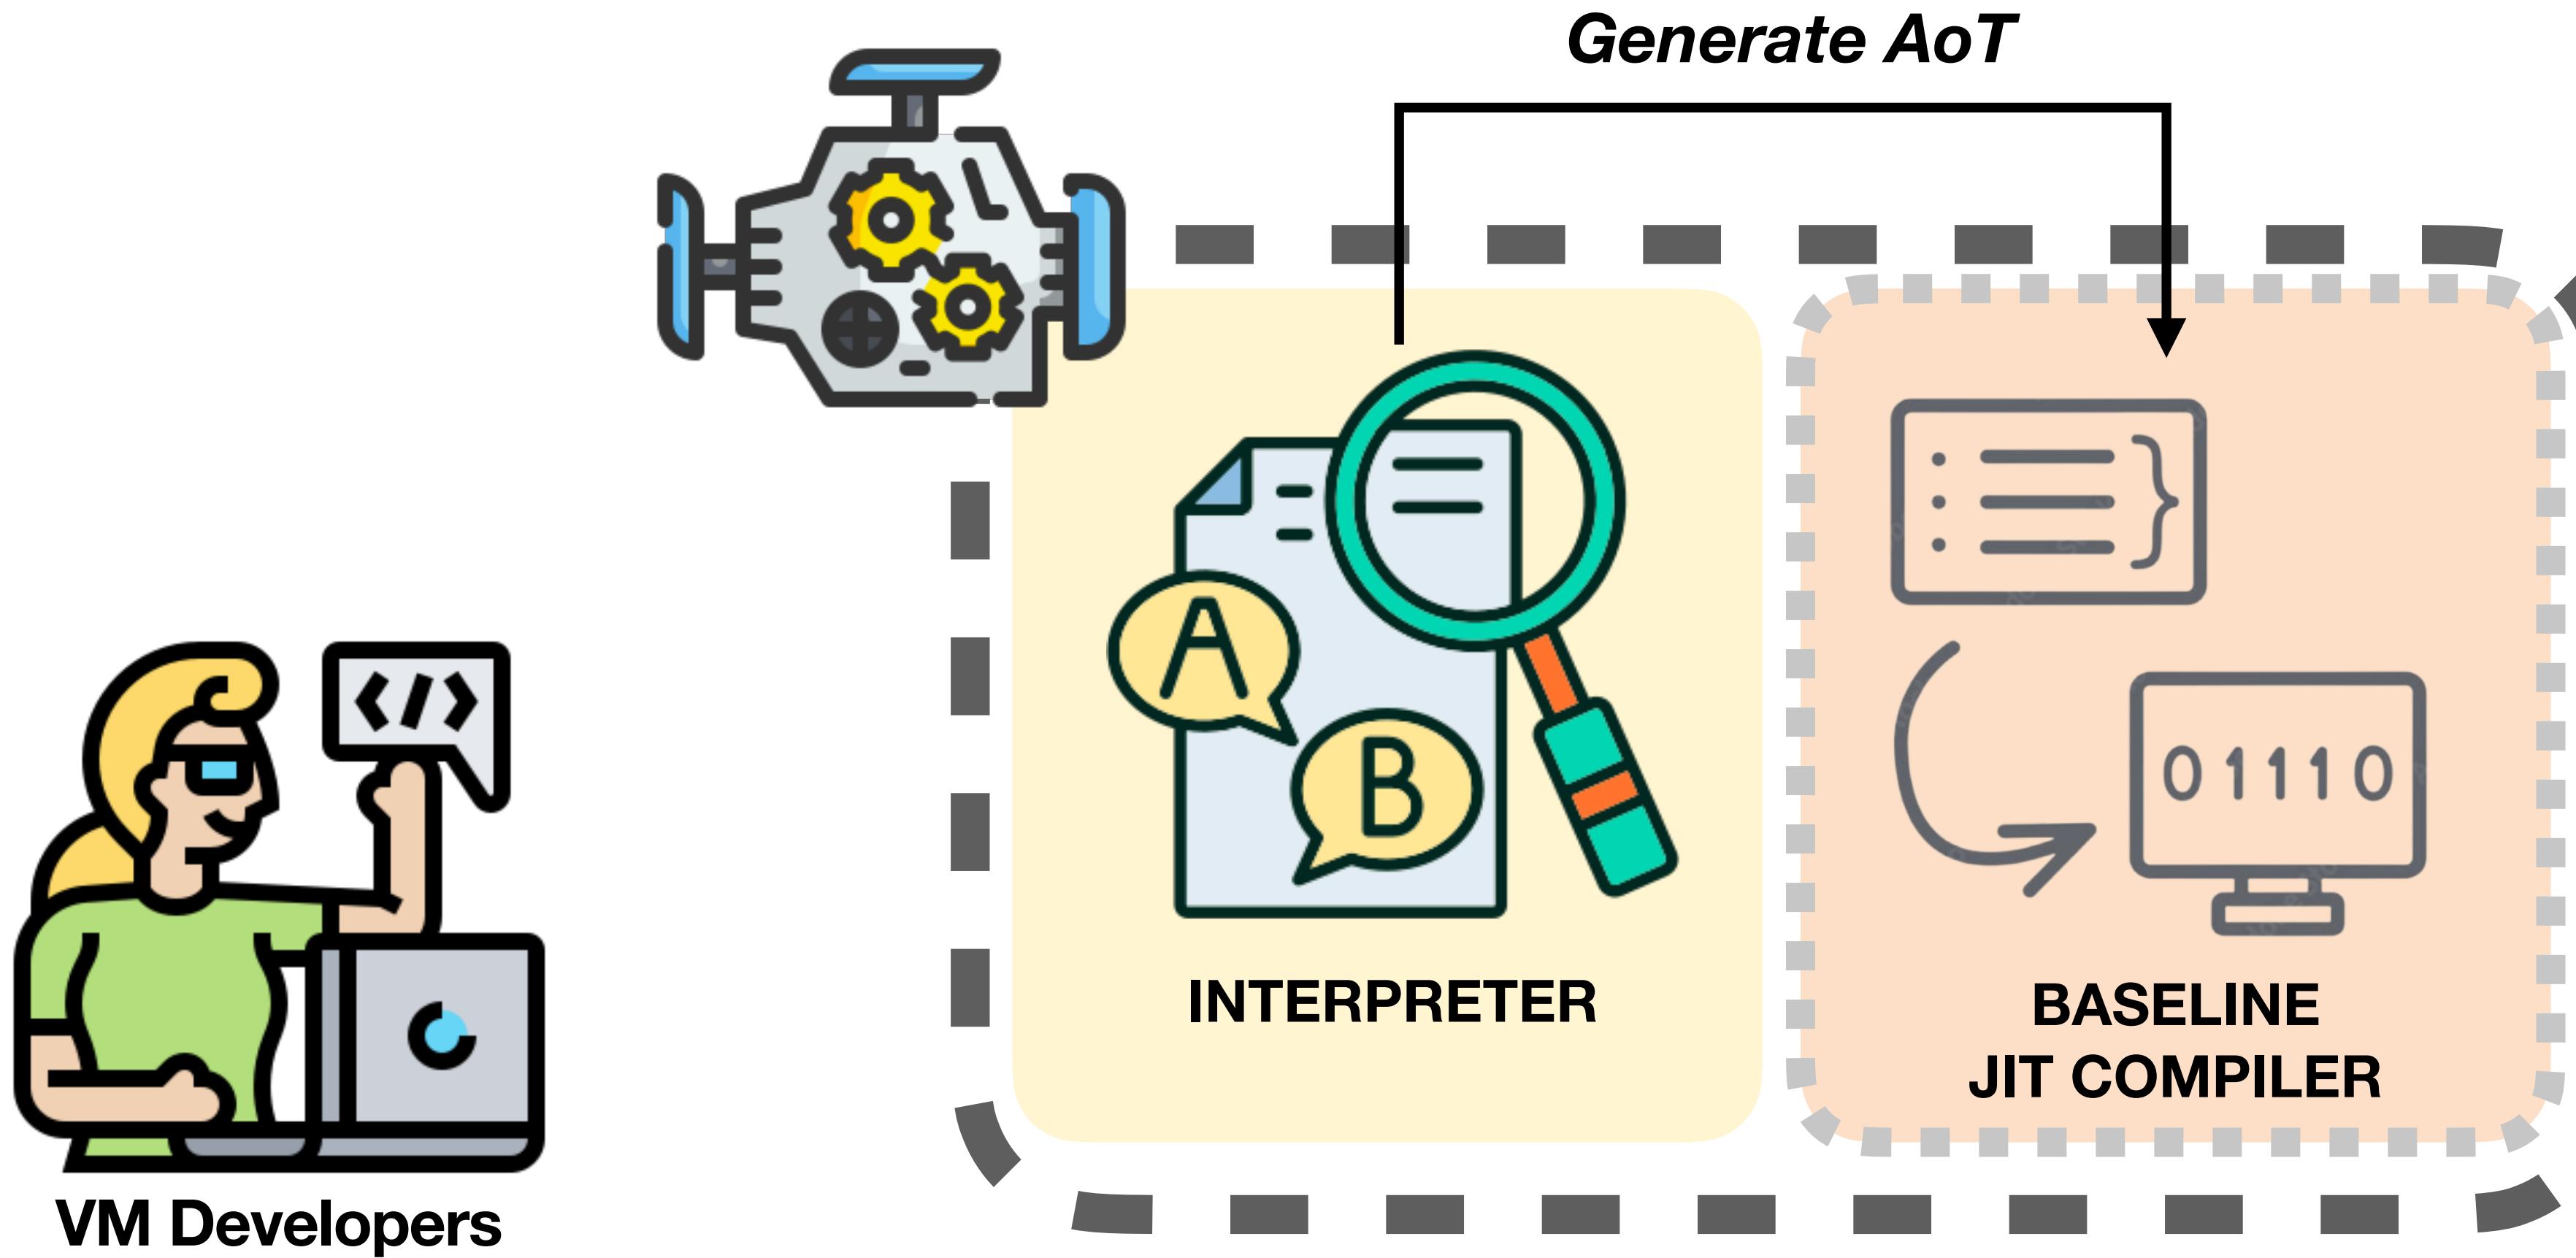
JIT Compiler Generators



JIT Compiler Generators



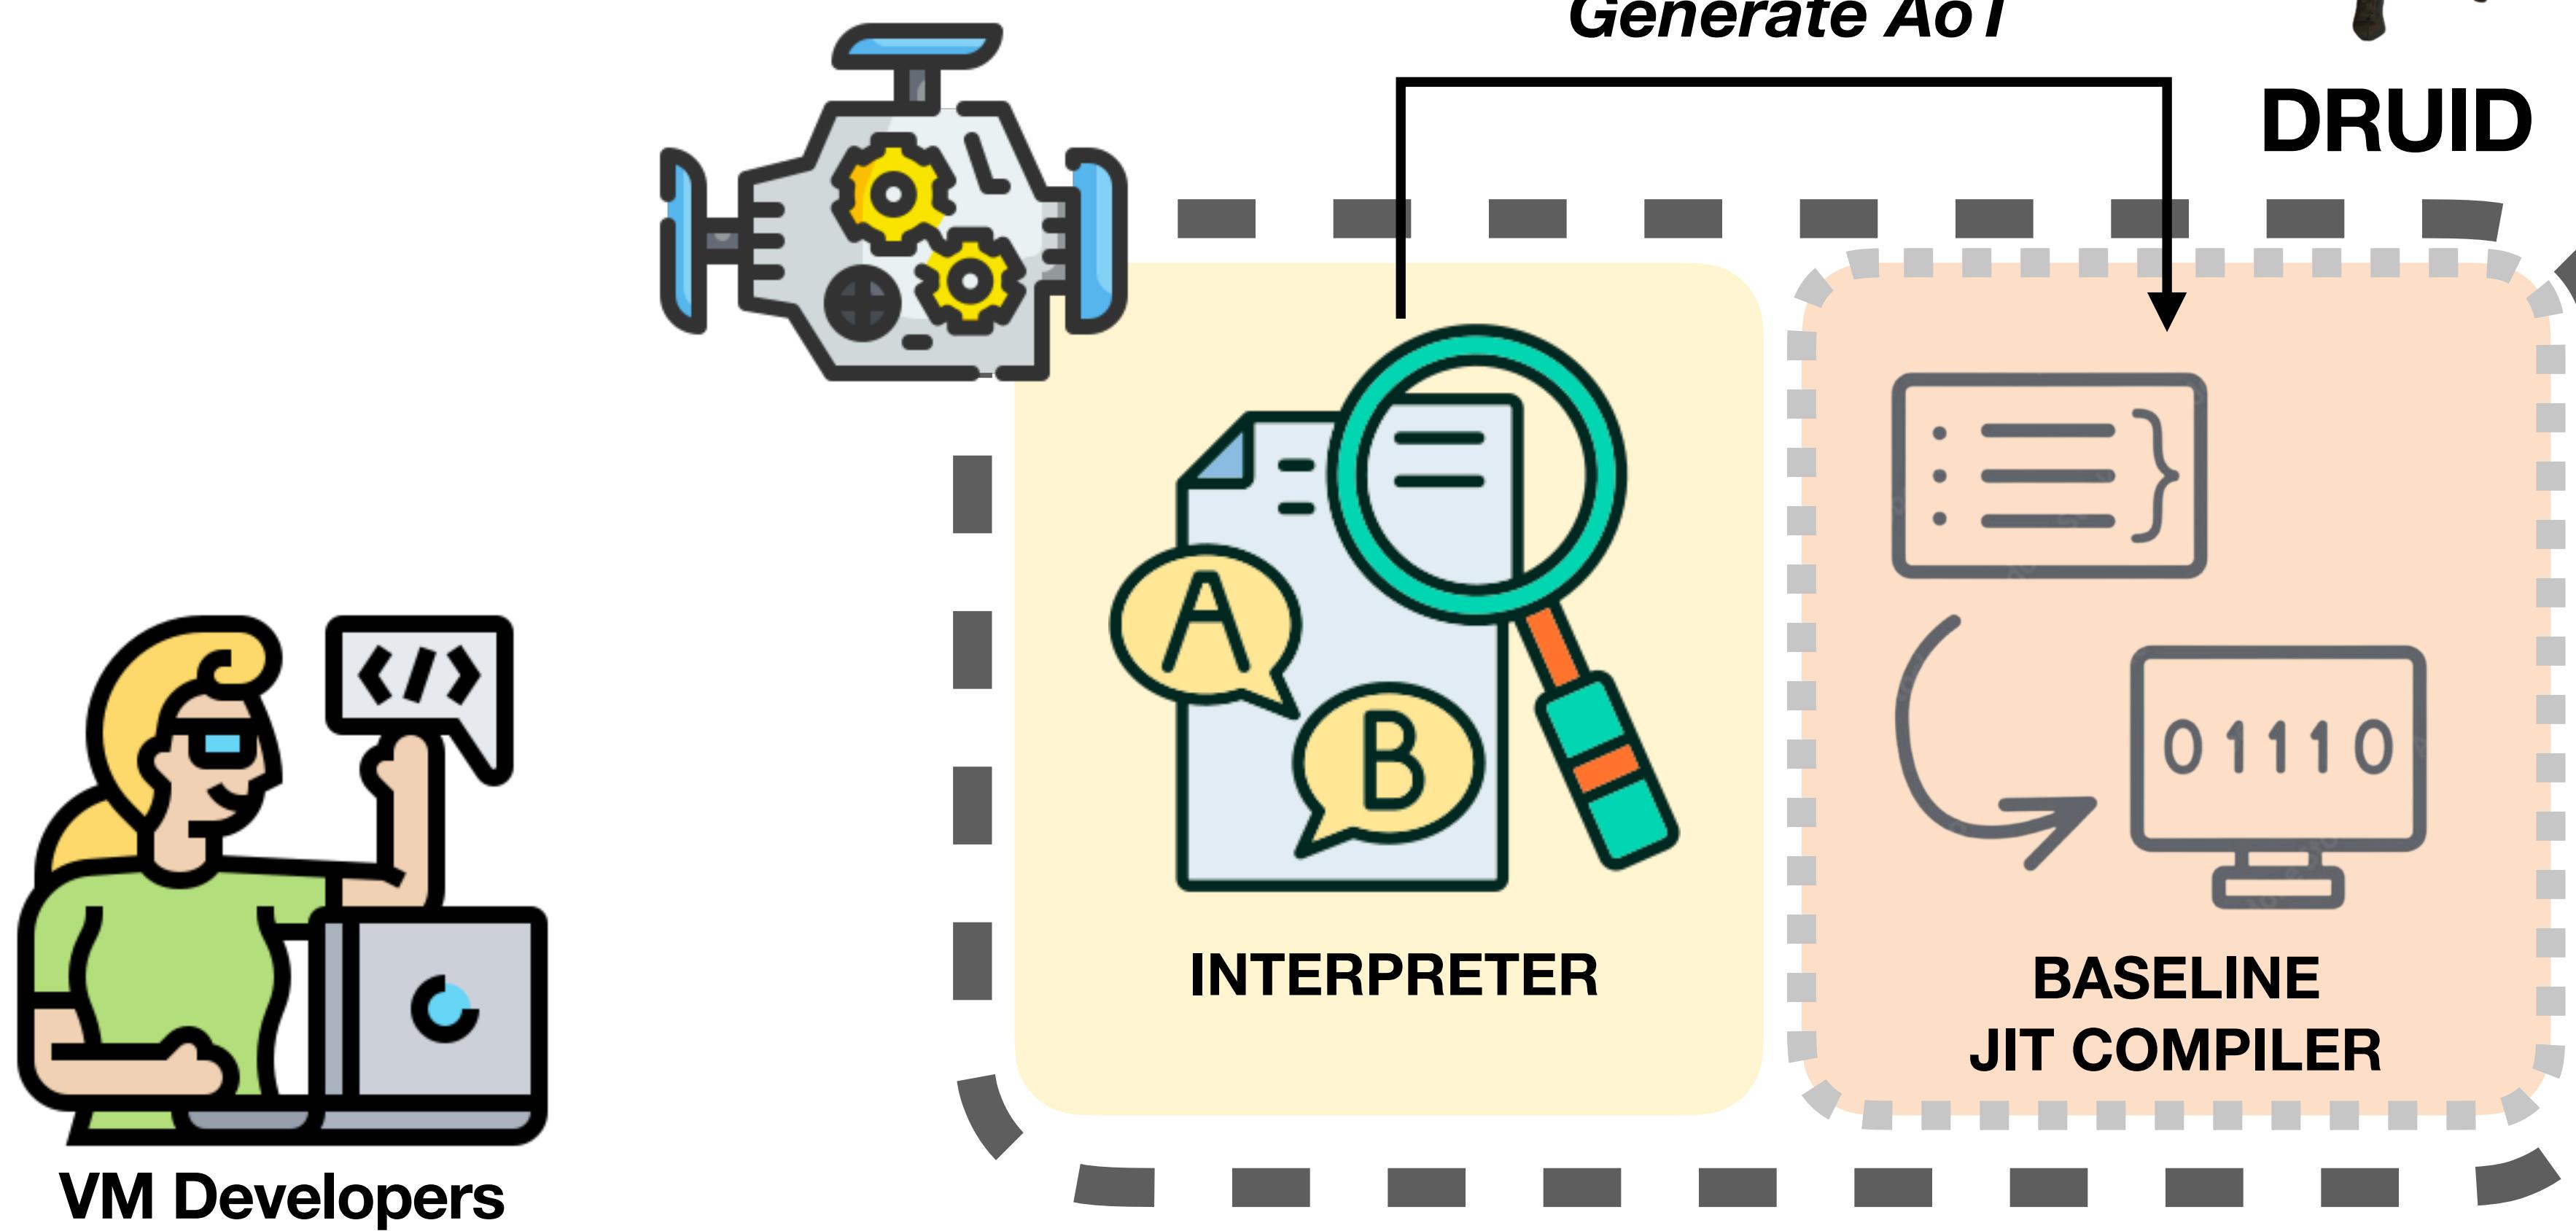
Baseline JIT Compiler Generators



Baseline JIT Compiler Generators

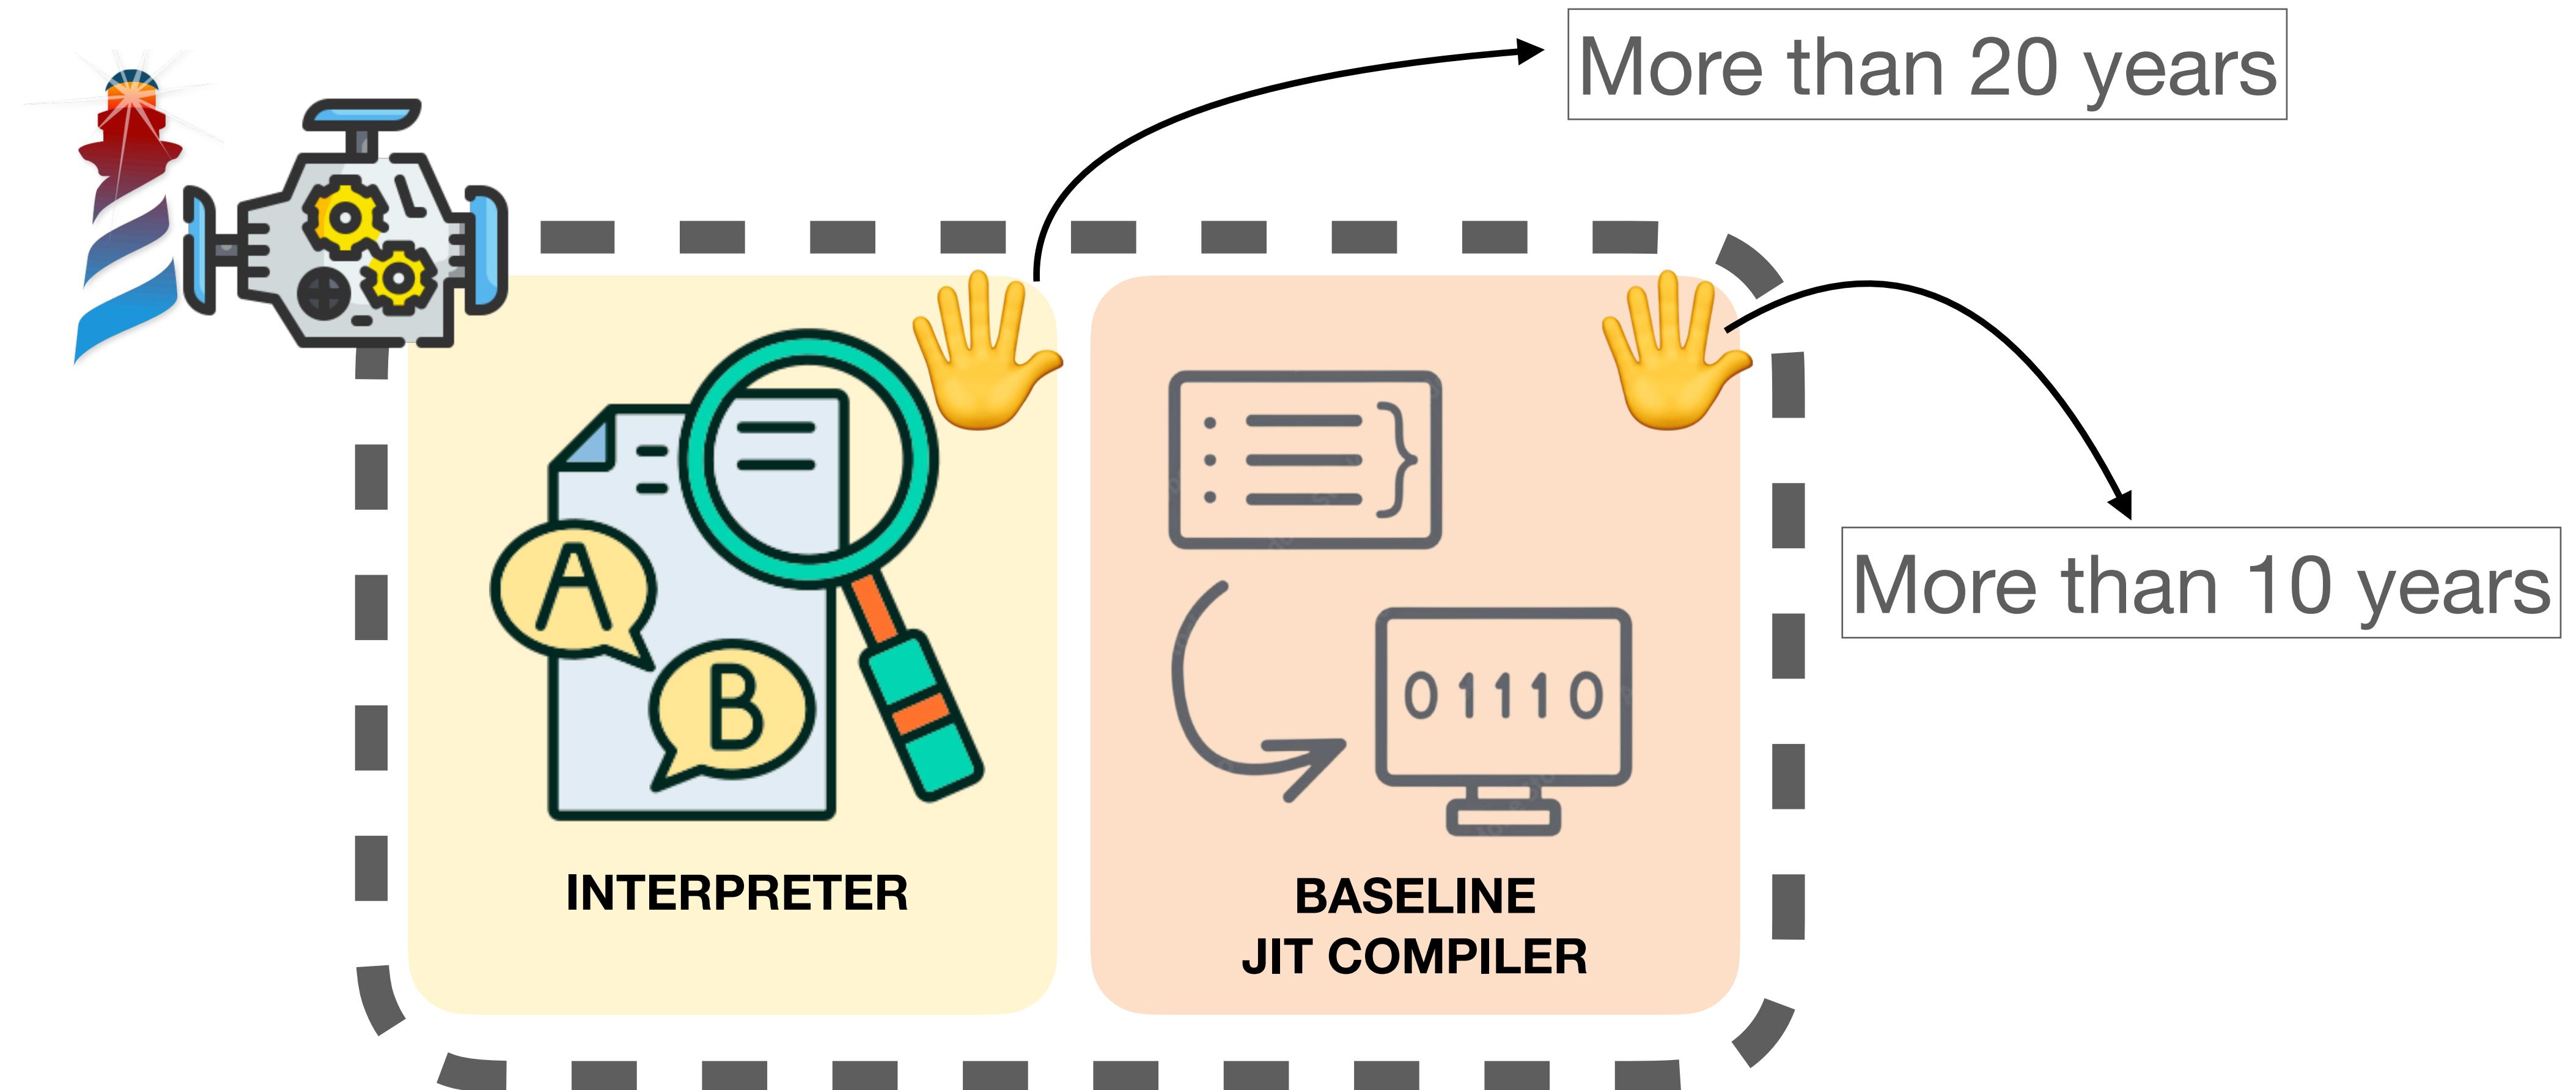


Source-to-source
meta-compiler

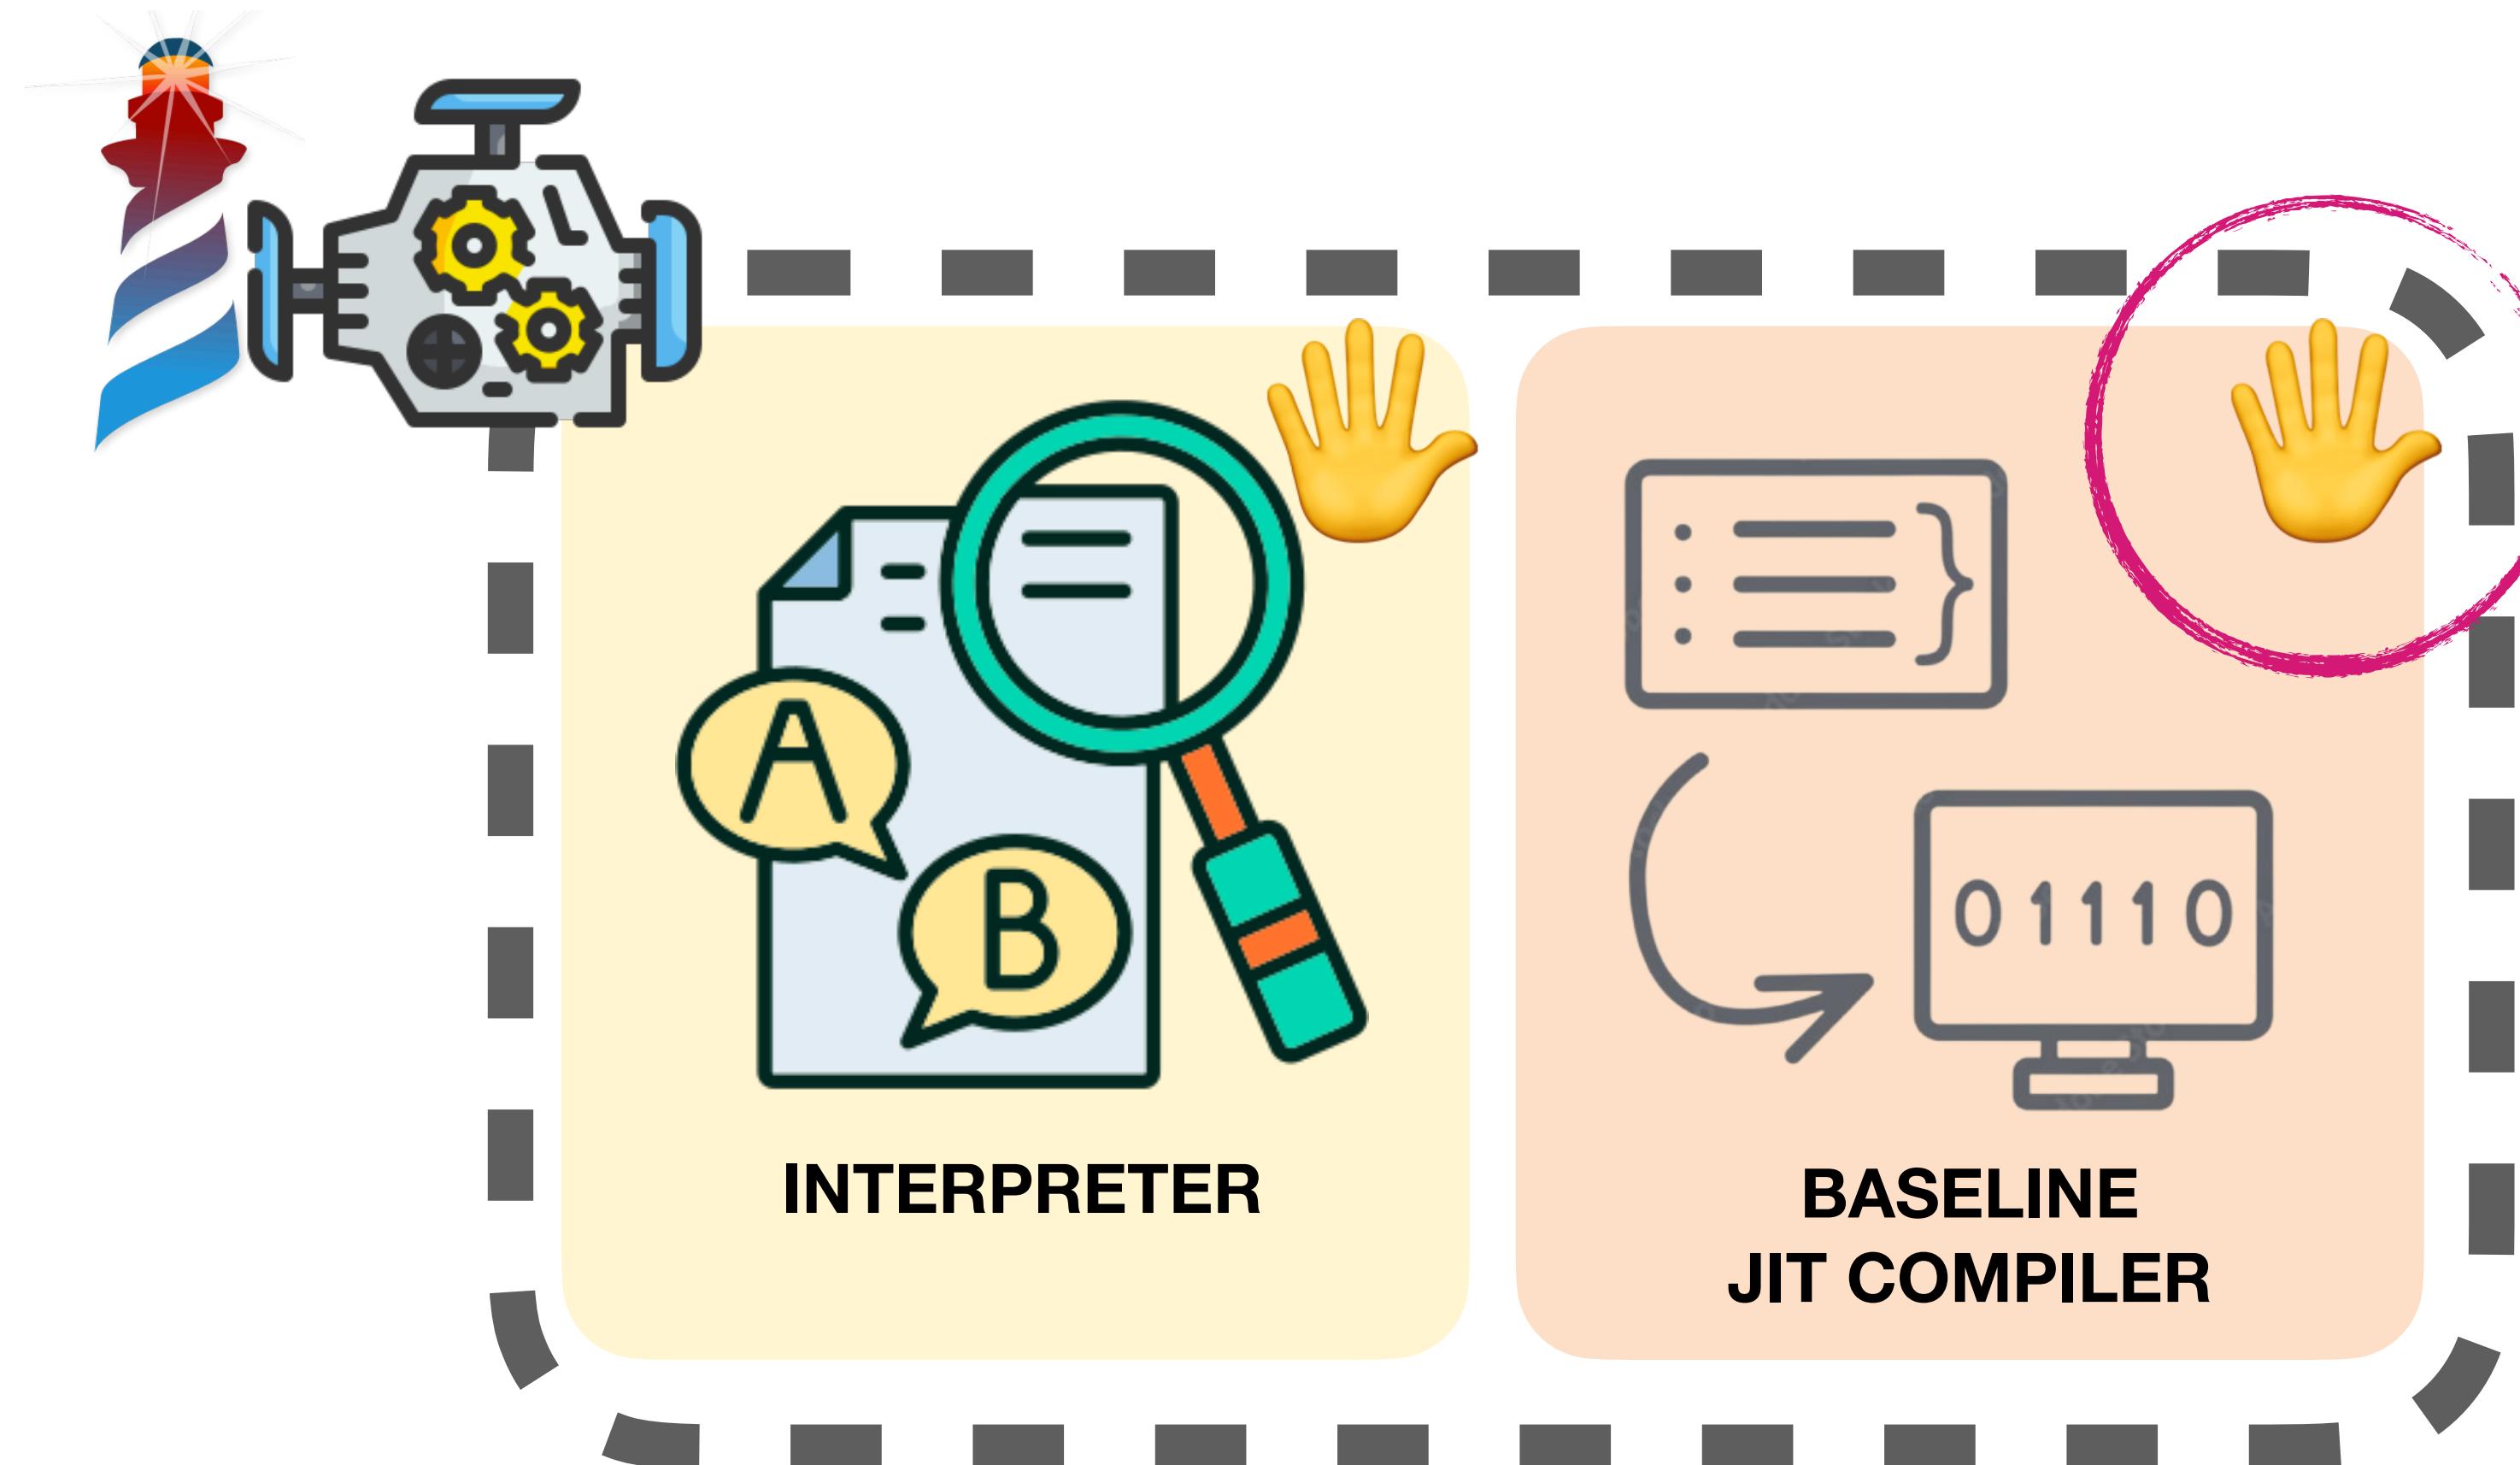


PHARO

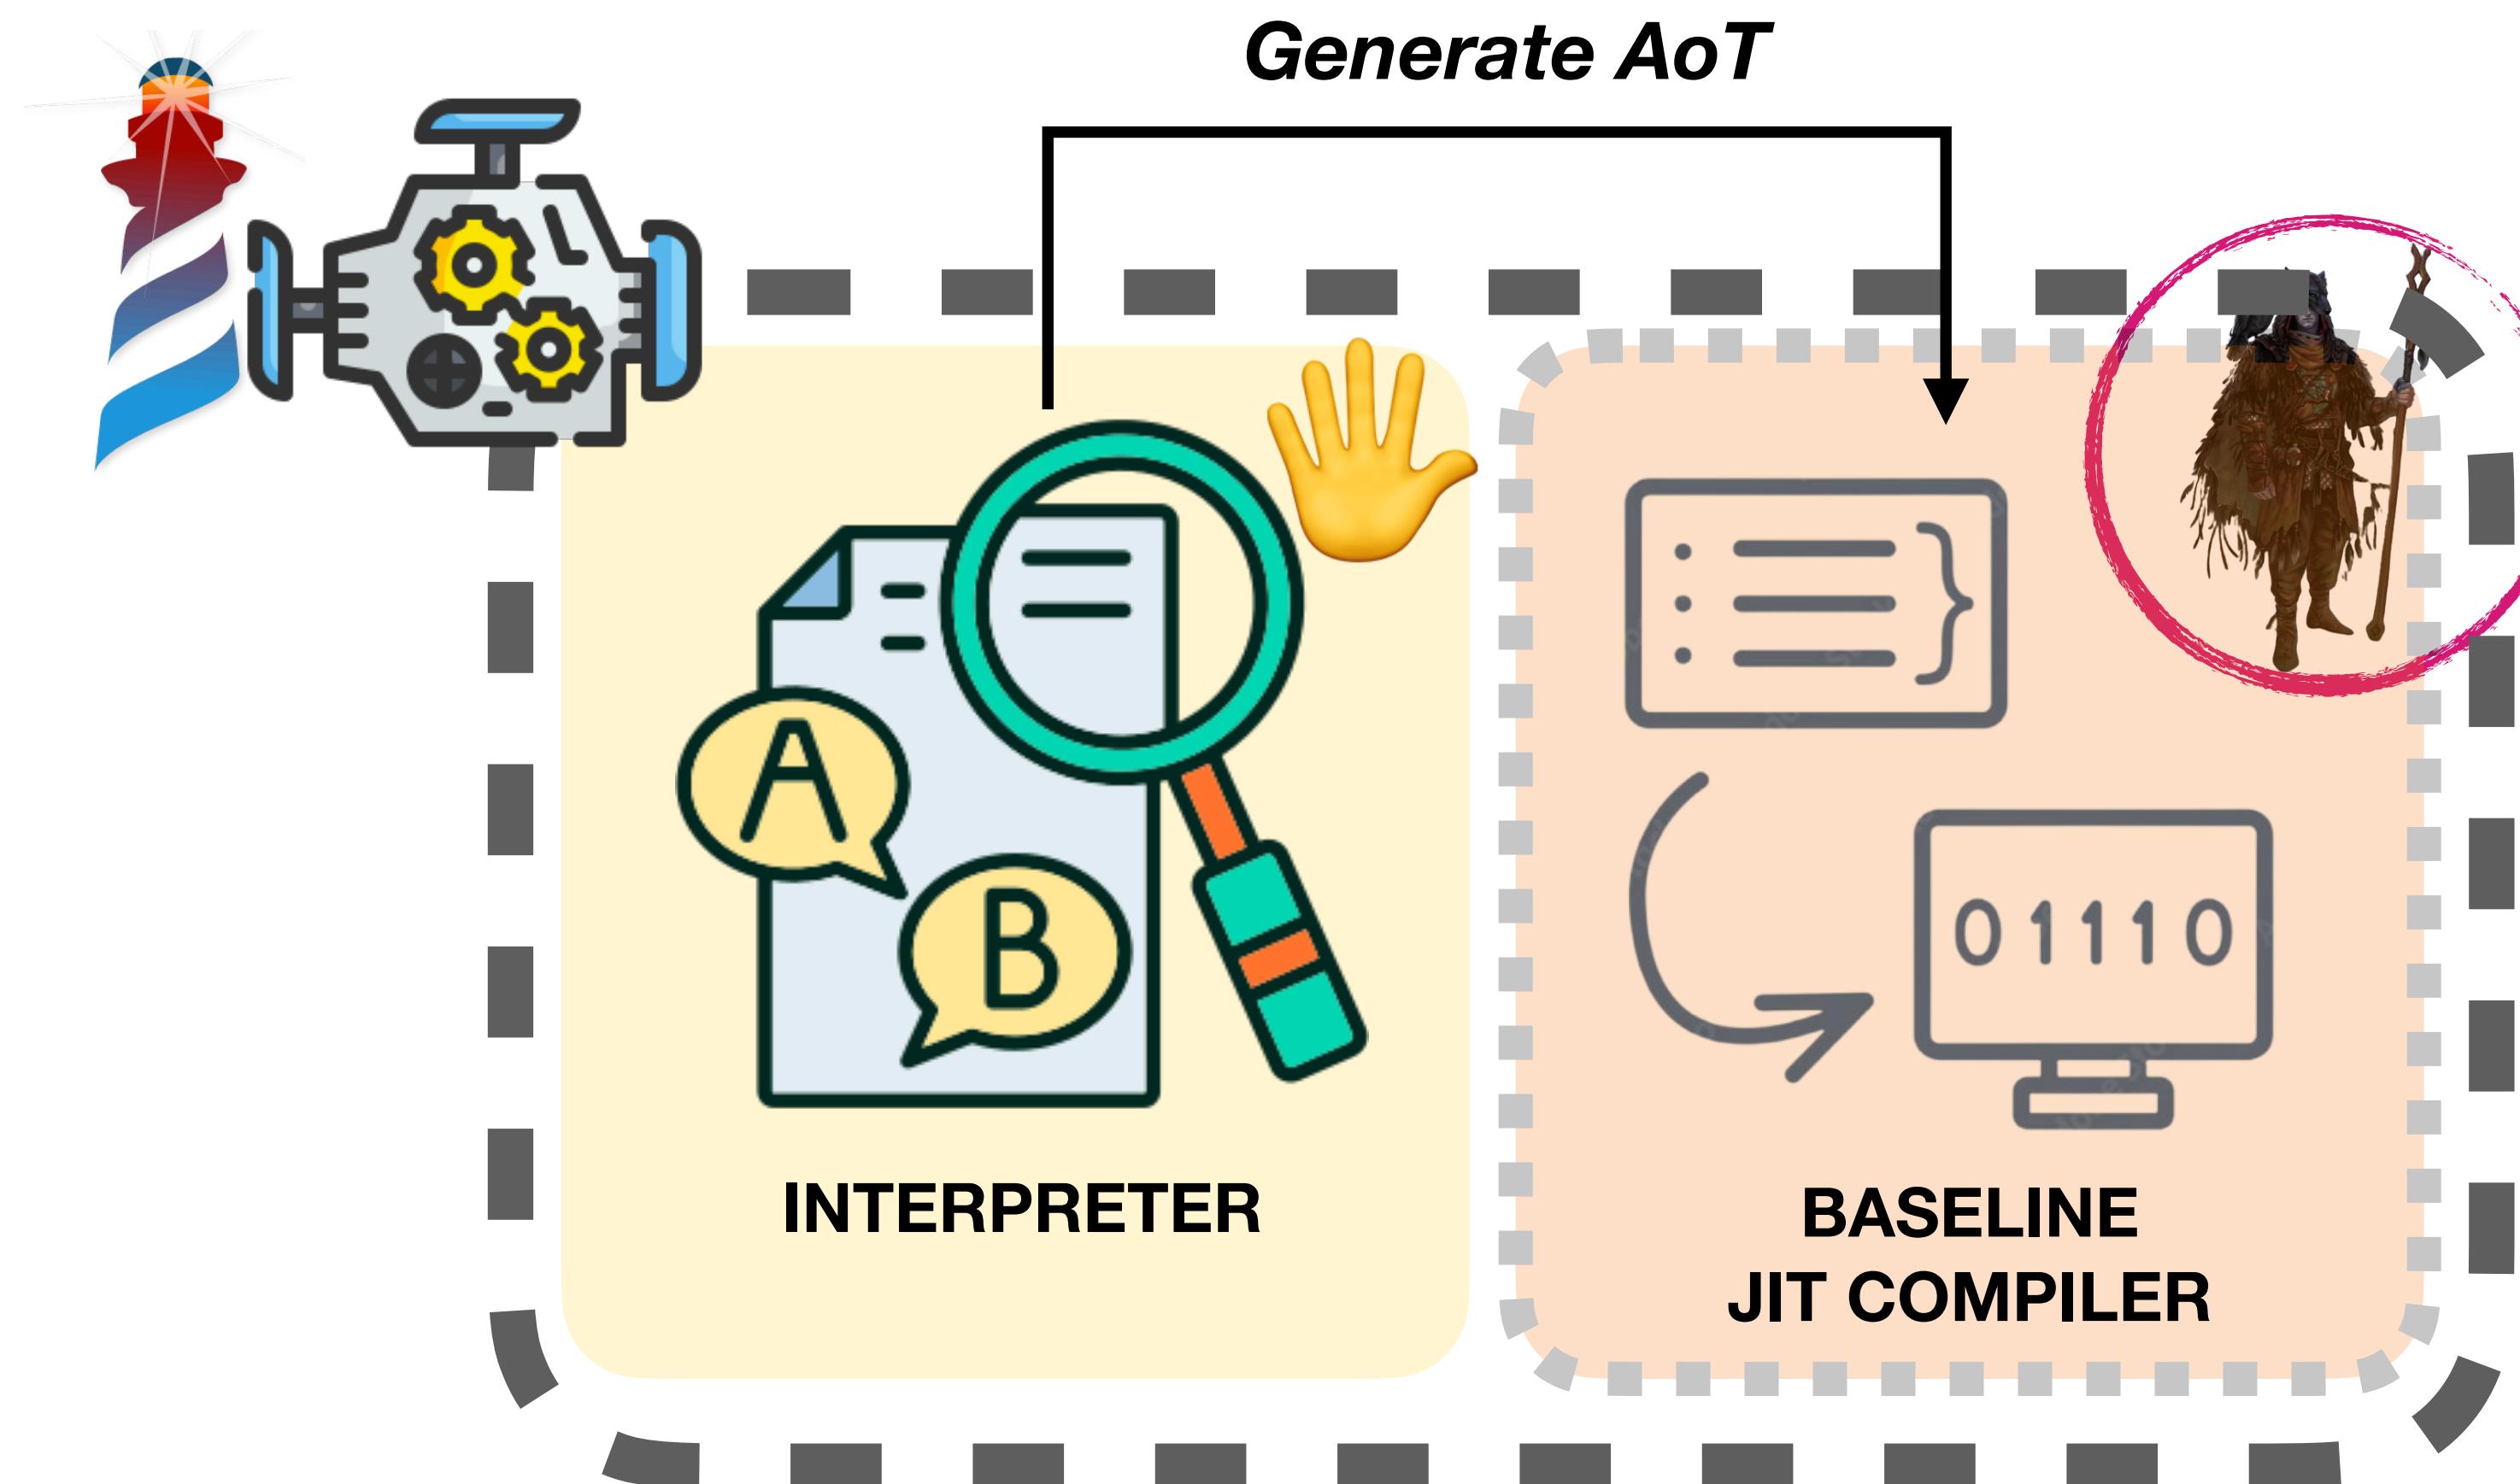
Pharo VM



Pharo VM

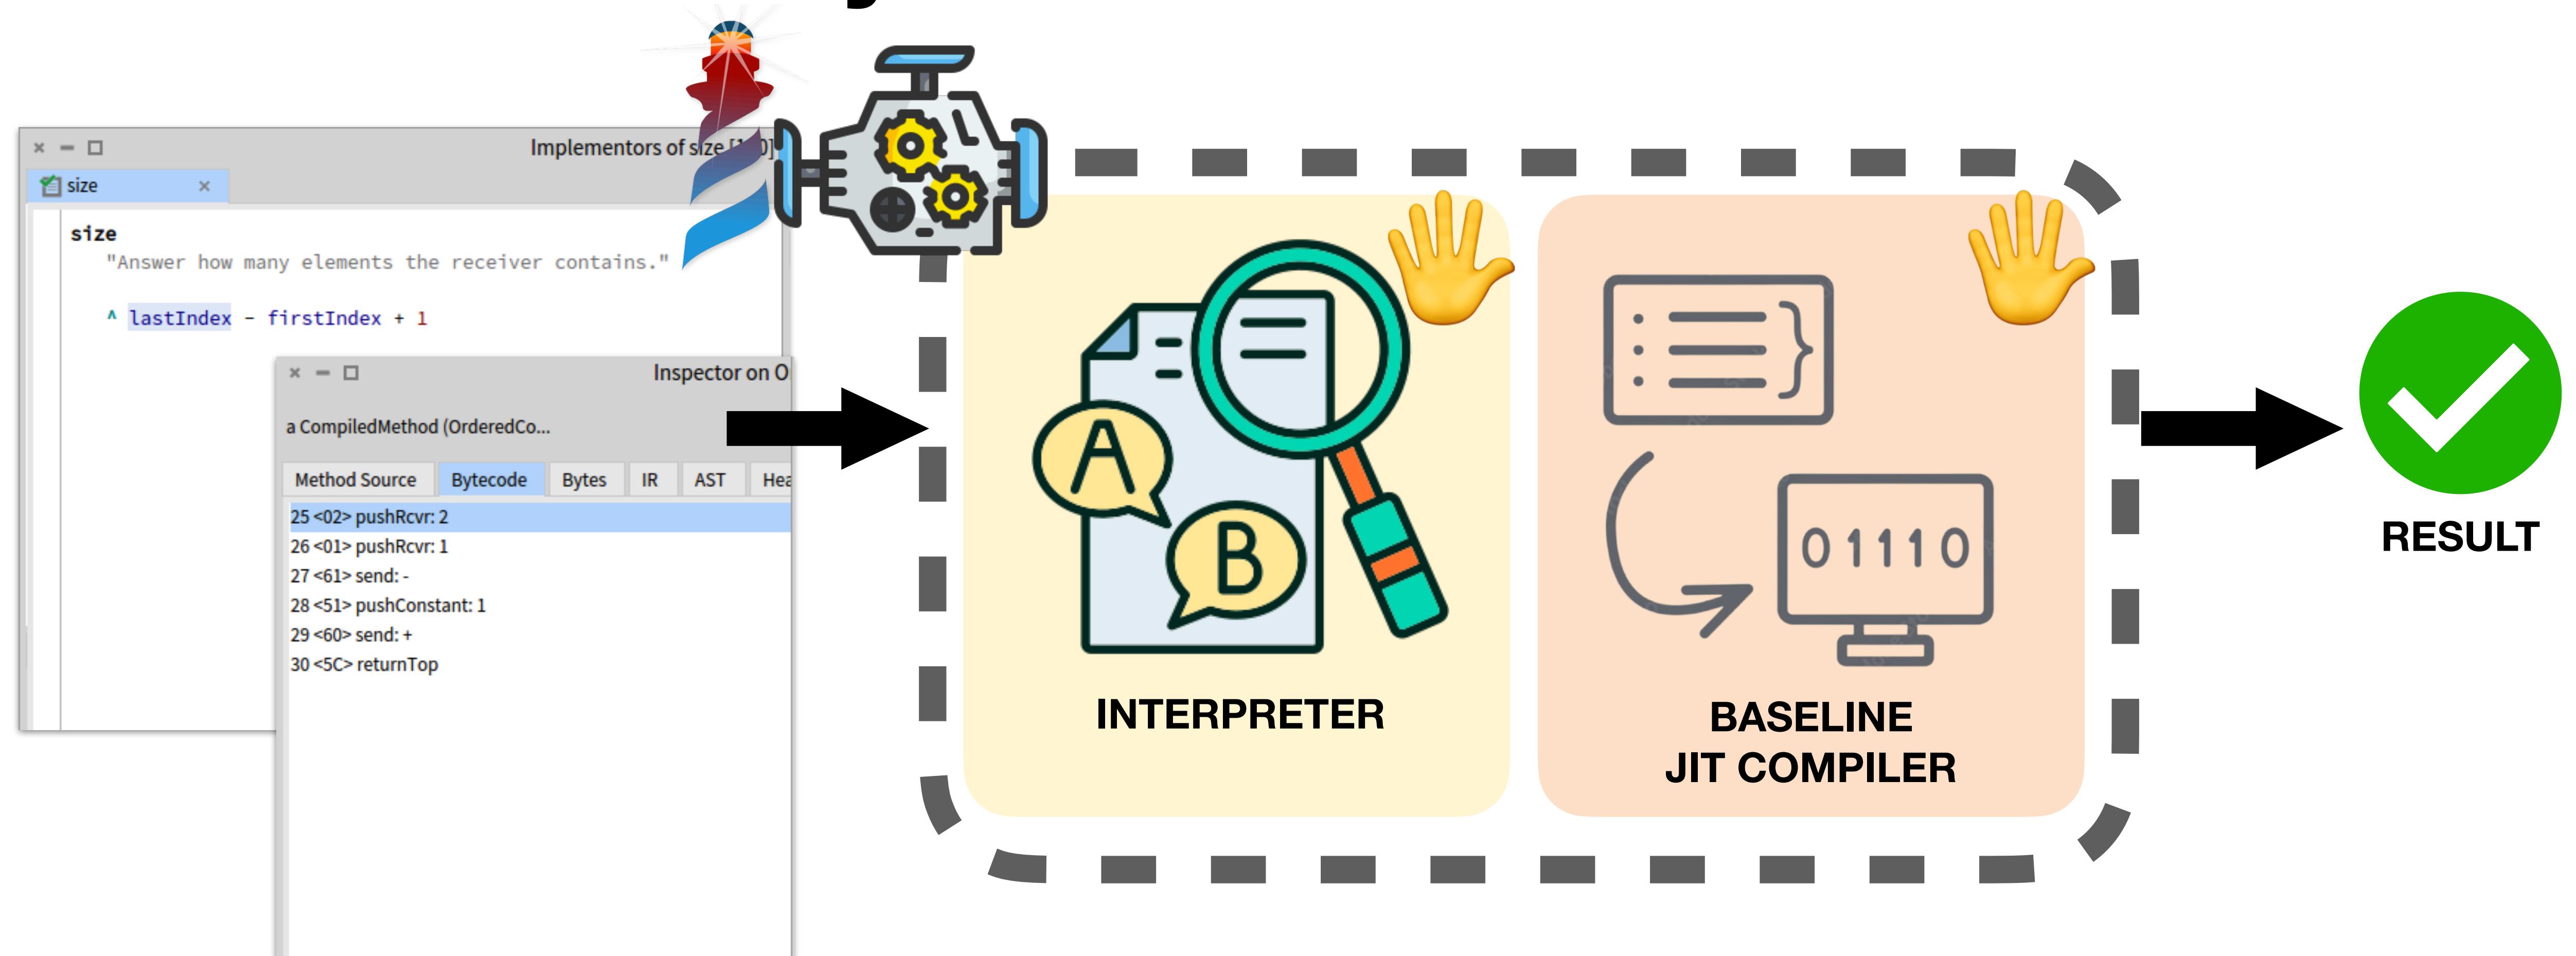


Pharo VM

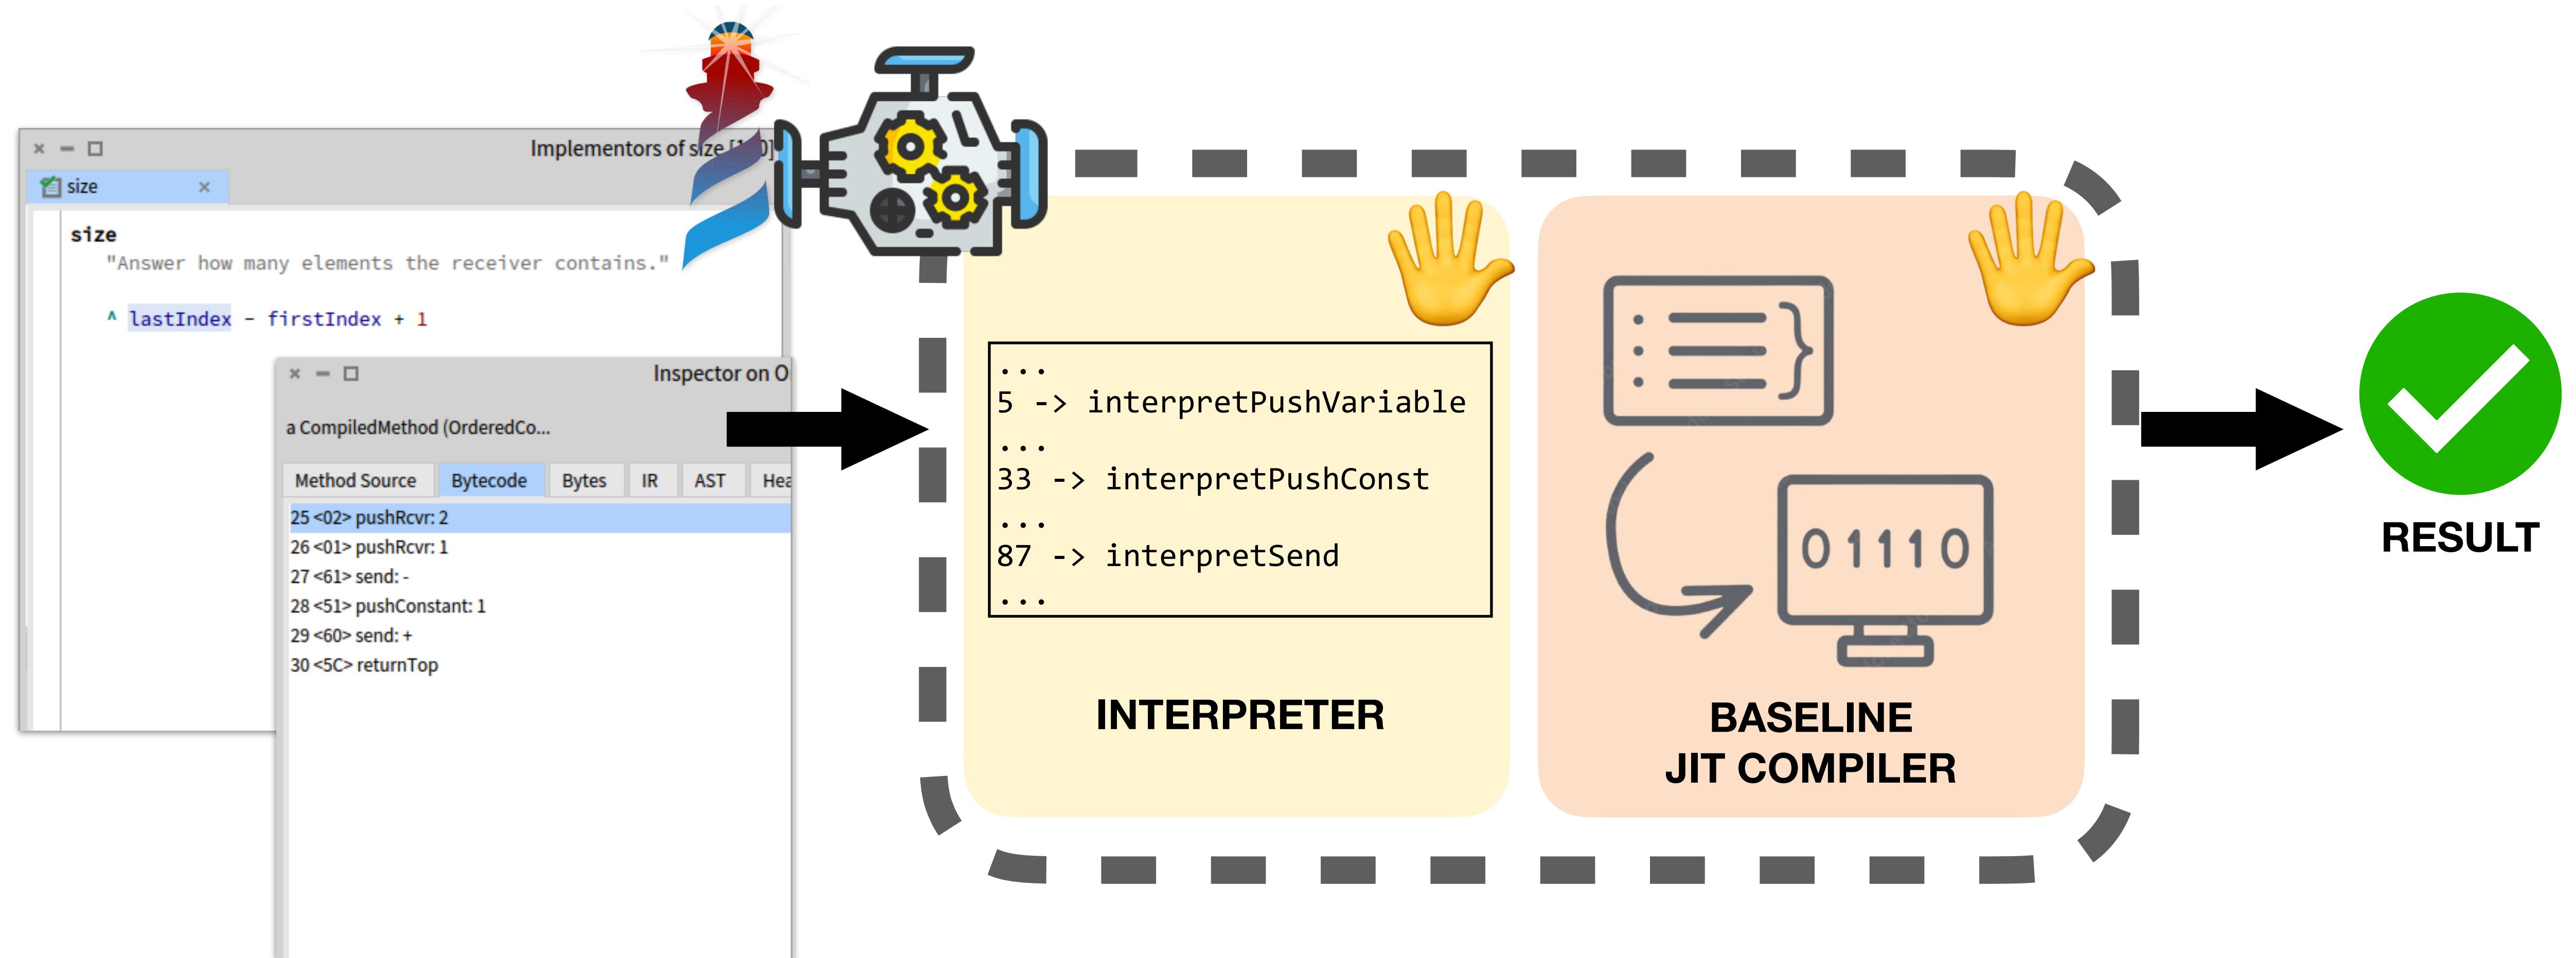


**How do we generate Baseline
JIT Compilers AoT?**

Pharo VM - Bytecodes and Primitives



Pharo VM - Instruction handlers



The addition primitive

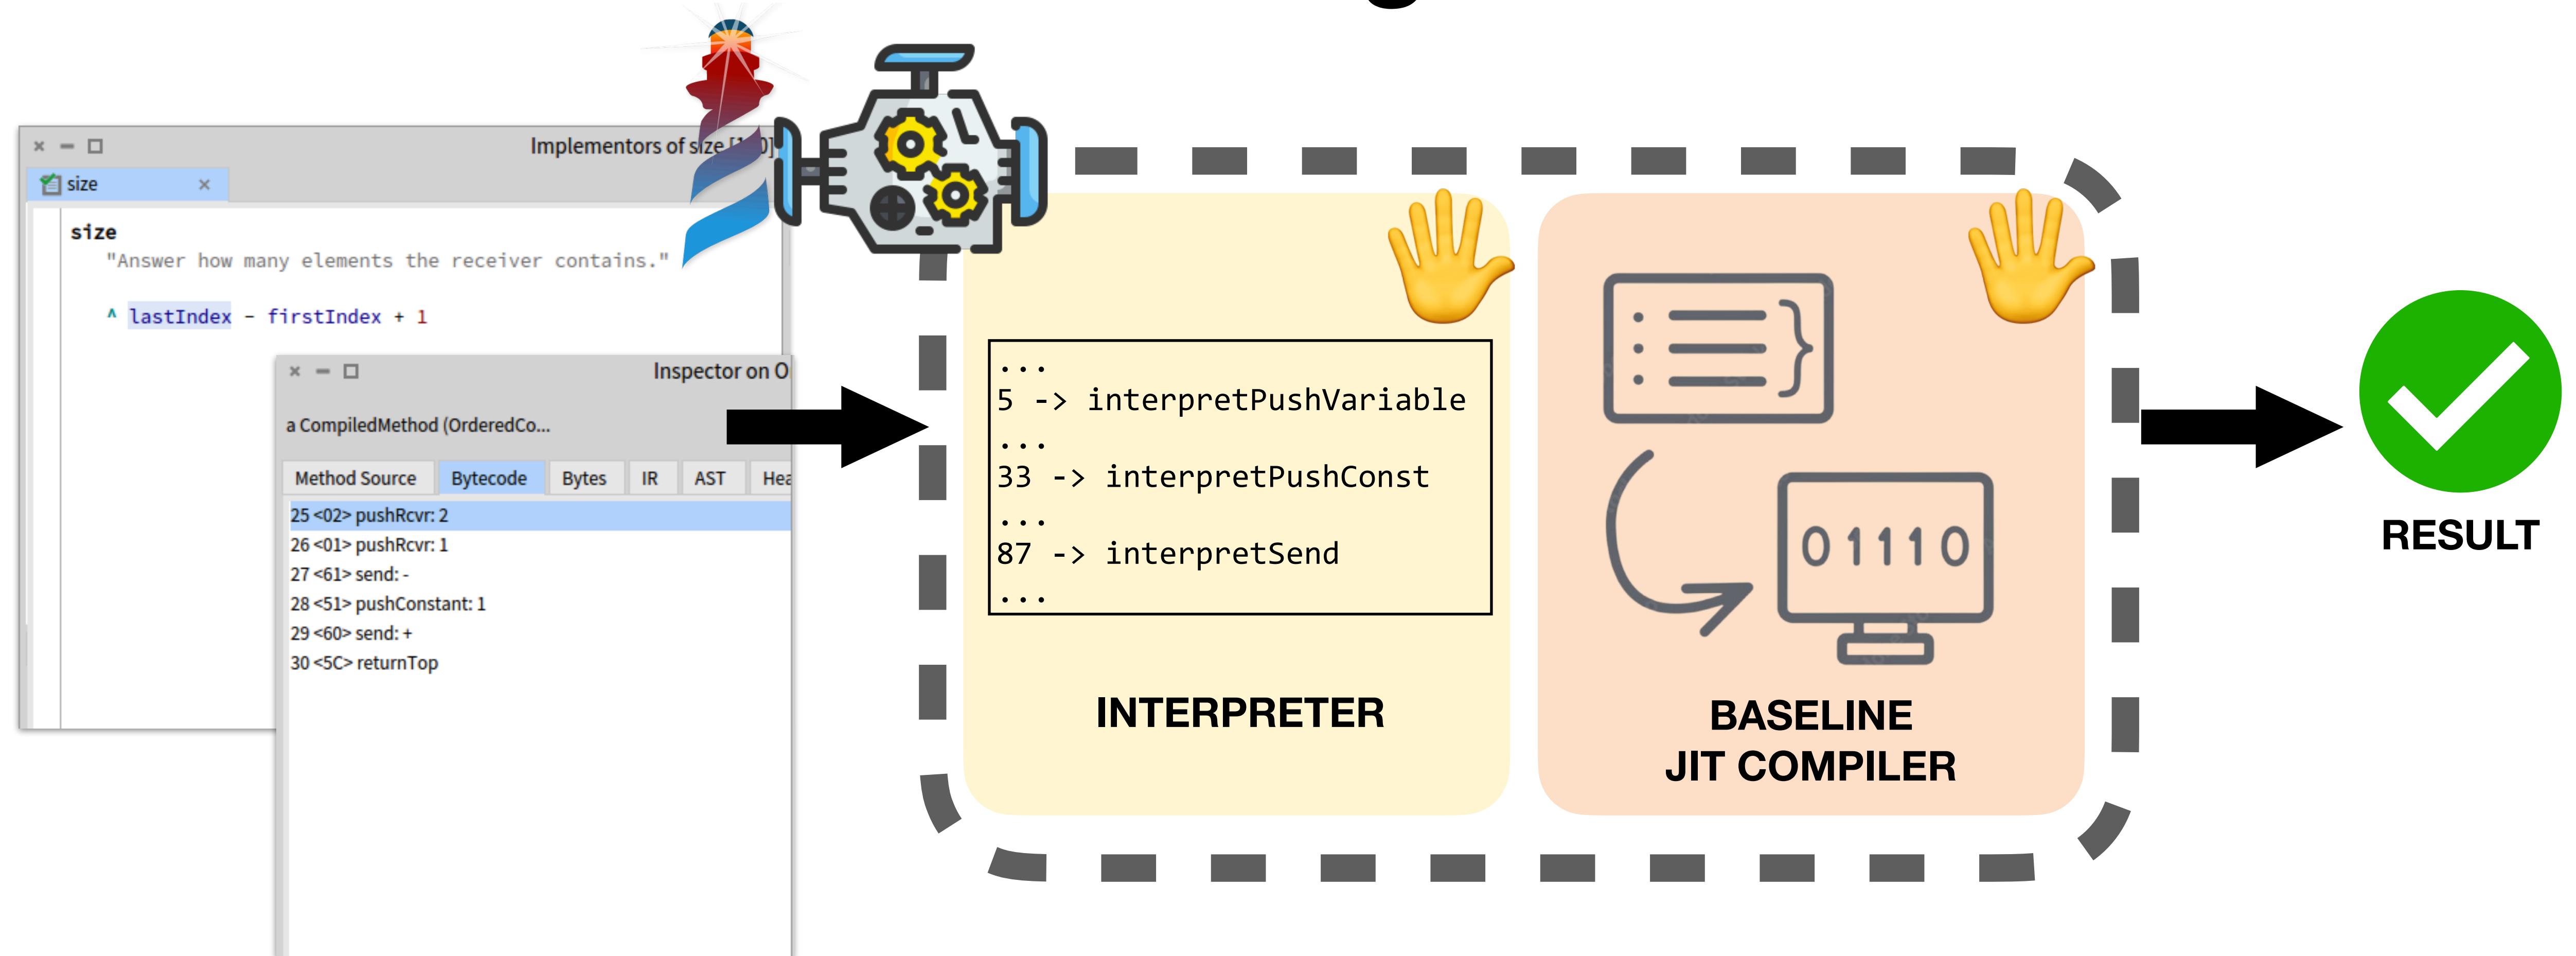
Interpreter

```
1 primitiveAdd
2     <numberOfArguments: 1>
3     <customisedReceiverFor: #smallInteger>
4
5     | maybeSmallInteger maybeSmallInteger2 result |
6
7     maybeSmallInteger := self stackValue: 0.
8     maybeSmallInteger2 := self stackValue: 1.
9
10    (objectMemory isIntegerObject: maybeSmallInteger)
11        ifFalse: [ ^ self primitiveFail ].
12    (objectMemory isIntegerObject: maybeSmallInteger2)
13        ifFalse: [ ^ self primitiveFail ].
14
15    "Check for overflow"
16    result := self
17        sumSmallInteger: maybeSmallInteger
18        withSmallInteger: maybeSmallInteger2
19        ifOverflow: [ ^ self primitiveFail ].
20
21    self pop: 2 thenPush: result
```

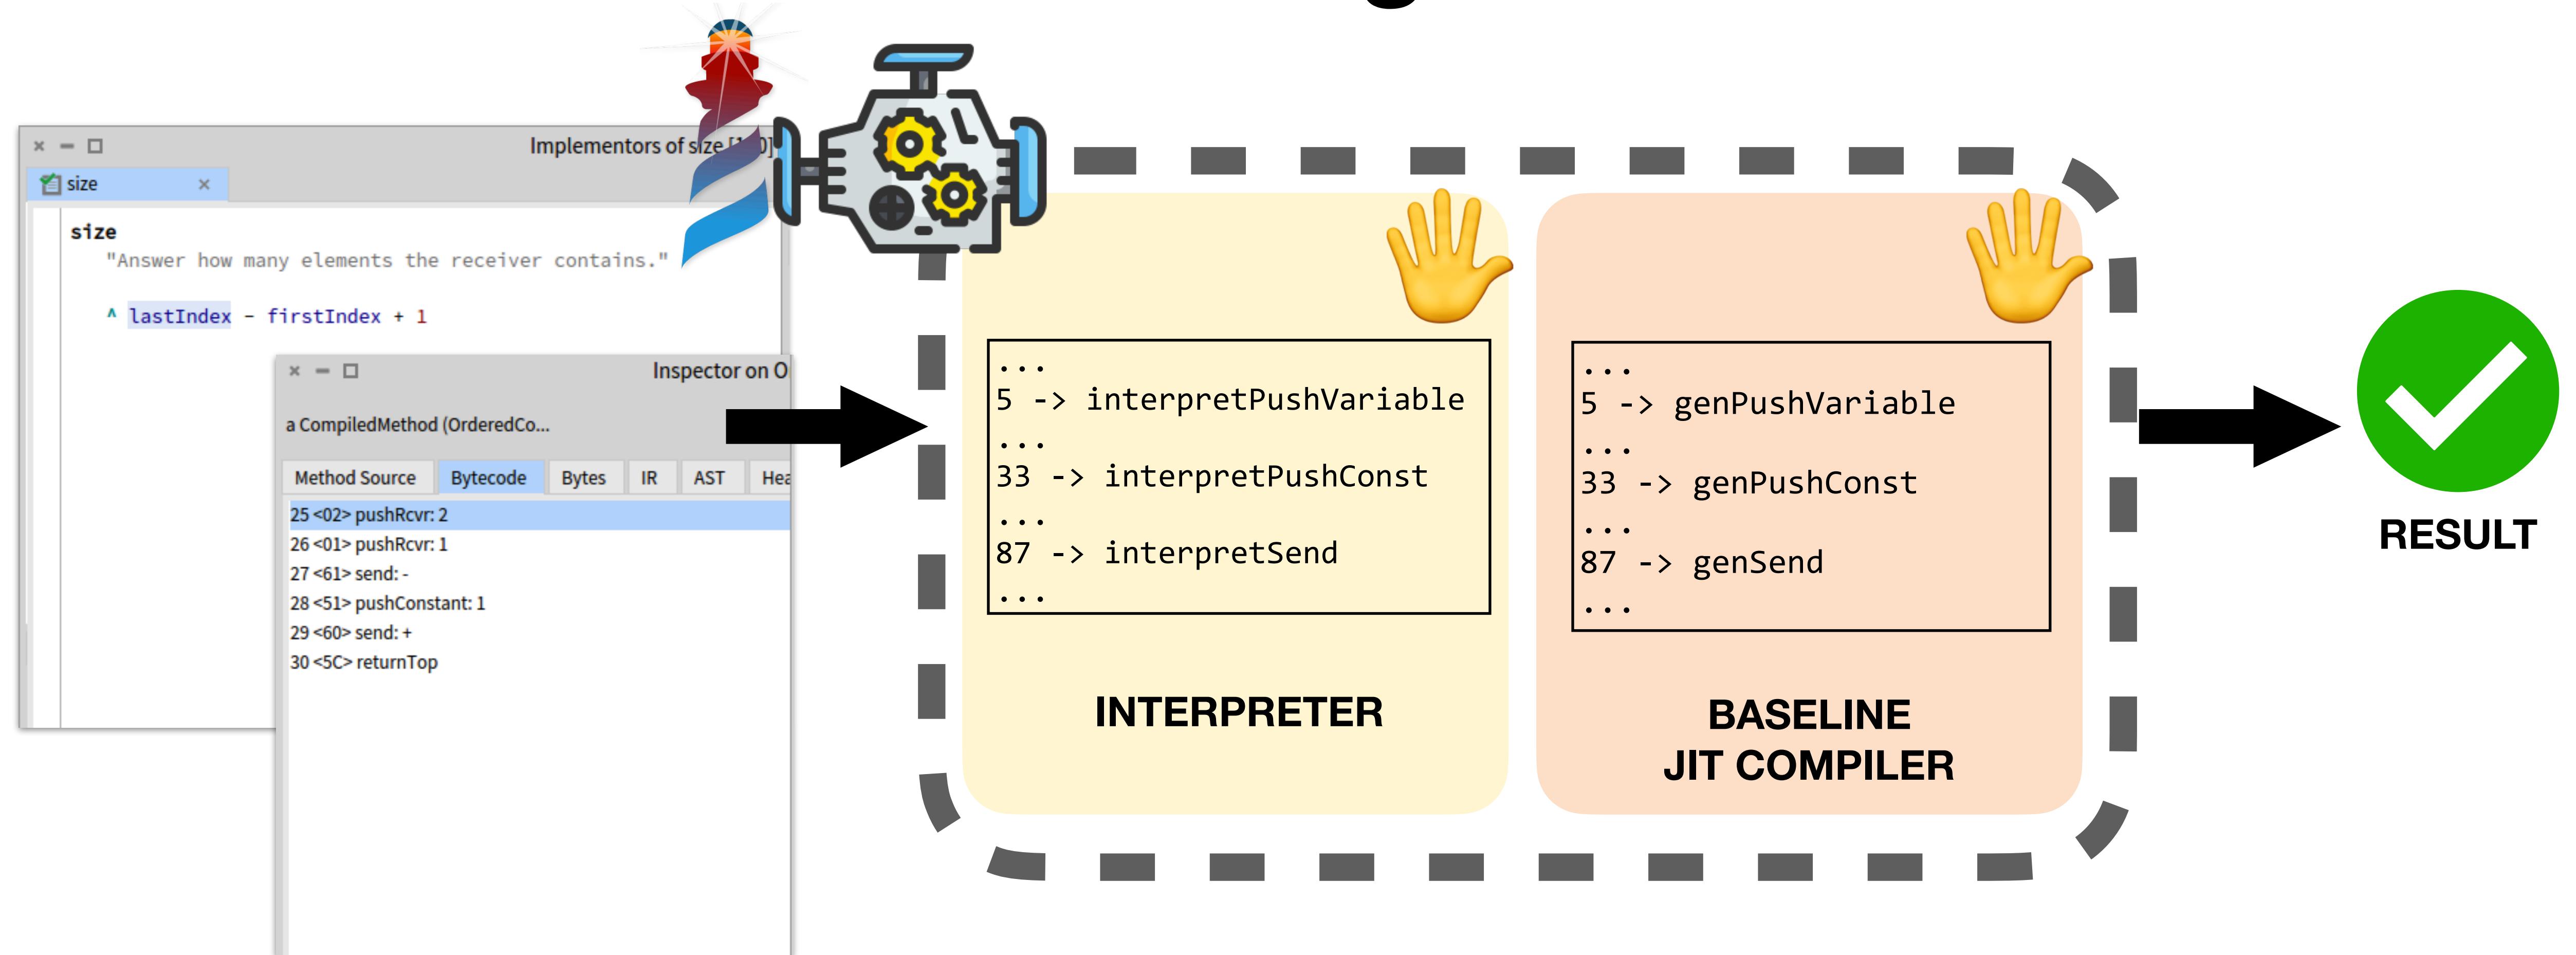


- Pops the two integers from the stack
- Push the addition
- On checks failure, a slow path is taken

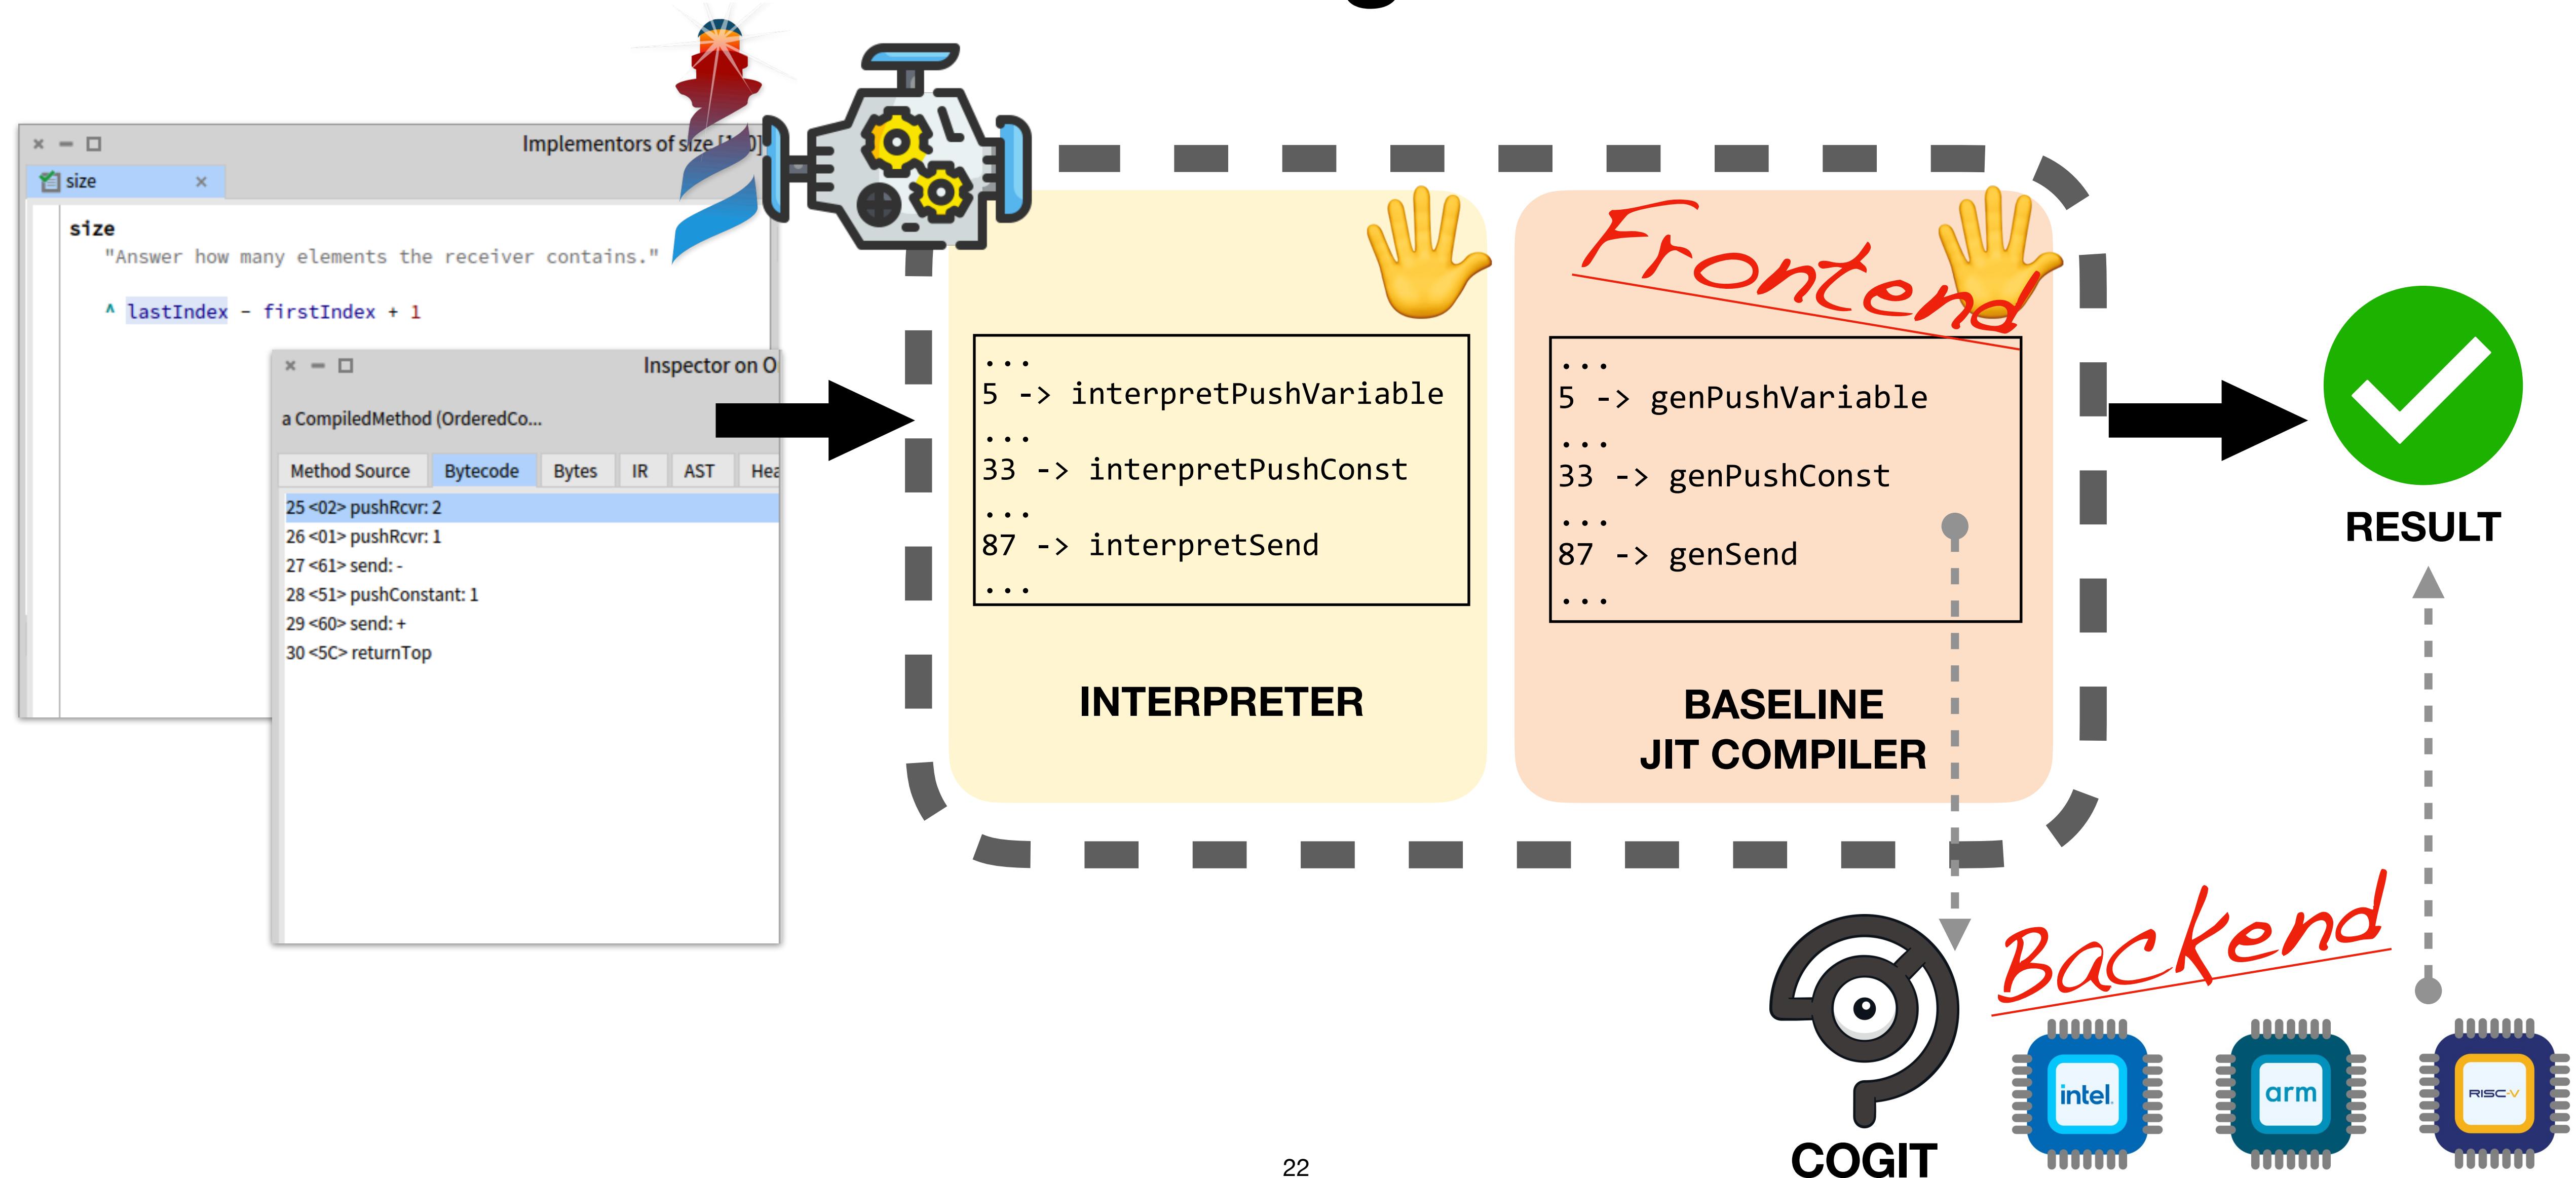
Pharo VM - Instruction generator handlers



Pharo VM - Instruction generator handlers



Pharo VM - Instruction generator handlers



The addition *generator* primitive - Cogit RTL



COGIT

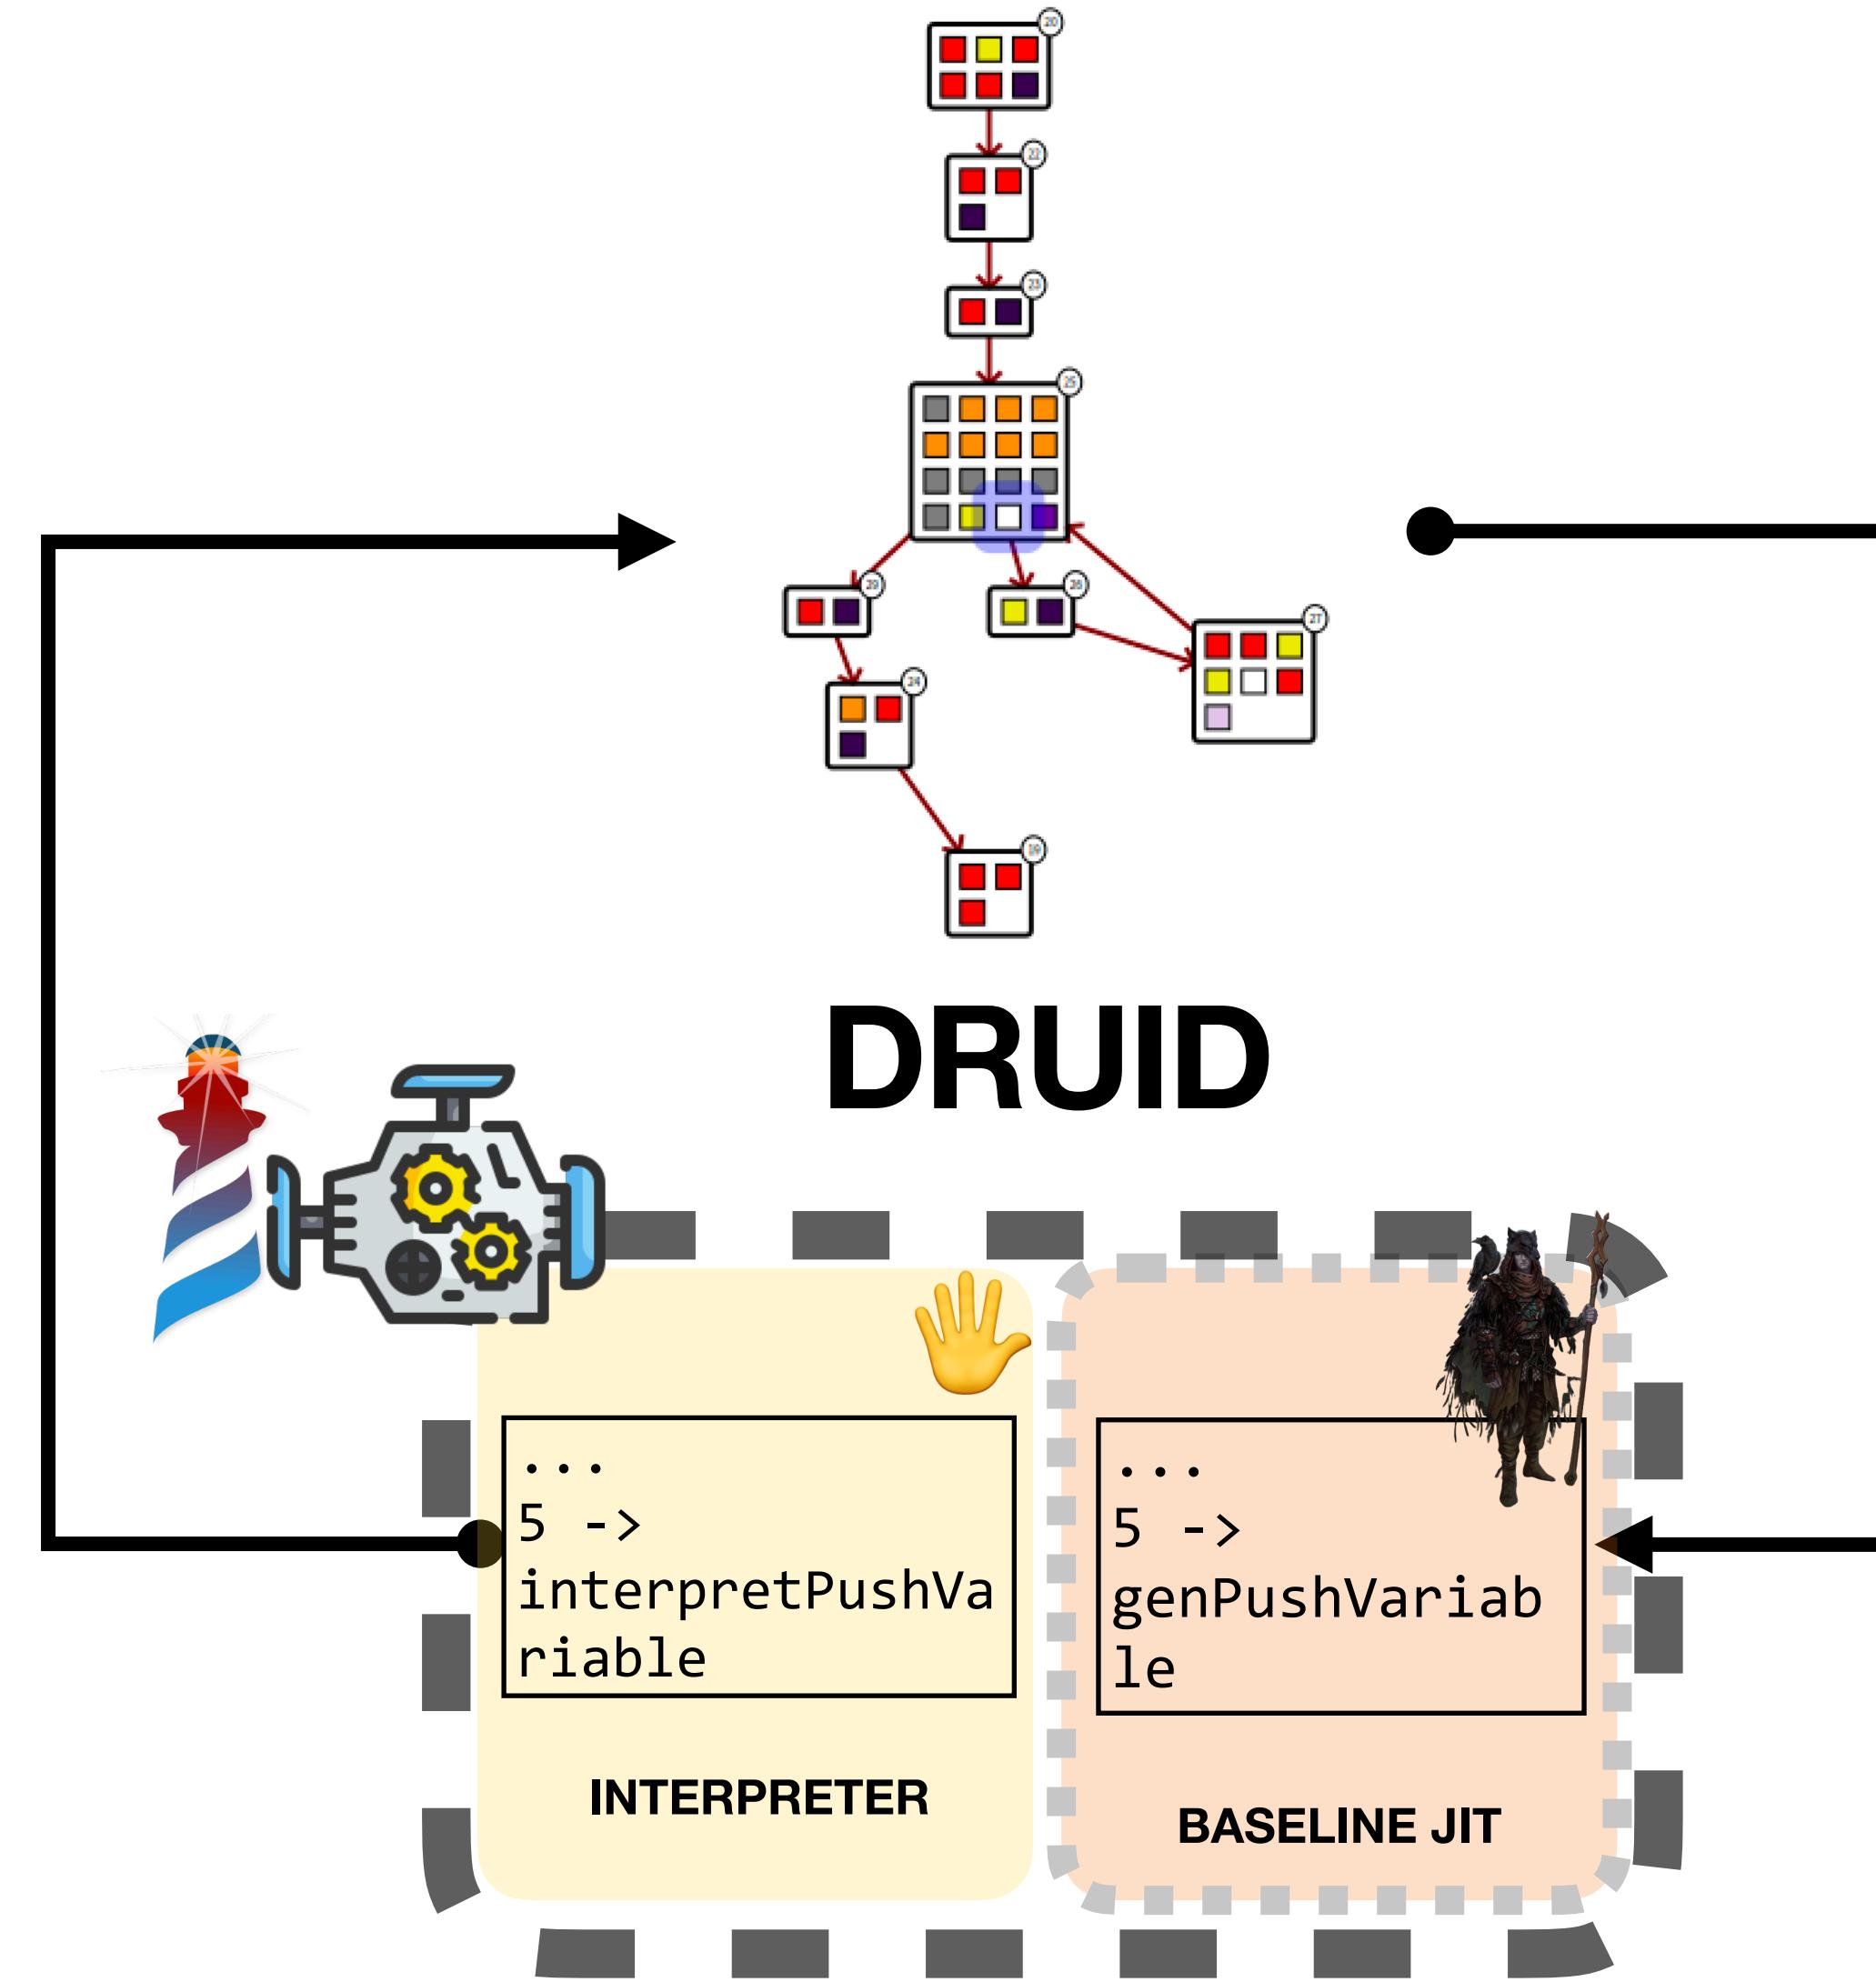
JIT Compiler

```
1 genPrimitiveAdd
2   | jumpNotSI jumpOvfl |
3   <var: #jumpNotSI type: #'AbstractInstruction *>
4   <var: #jumpOvfl type: #'AbstractInstruction *>
5   cogit mclassIsSmallInteger ifFalse:
6     [^UnimplementedPrimitive].
7
8   cogit genLoadArgAtDepth: 0 into: Arg0Reg.
9   cogit MoveR: Arg0Reg R: ClassReg.
10  jumpNotSI := self
11    genJumpNotSmallInteger: Arg0Reg scratchReg: TempReg.
12
13  self genRemoveSmallIntegerTagsInScratchReg: ClassReg.
14  cogit AddR: ReceiverResultReg R: ClassReg.
15  jumpOvfl := cogit JumpOverflow: 0.
16
17  cogit MoveR: ClassReg R: ReceiverResultReg.
18  cogit genPrimReturn.
19
20  jumpOvfl jmpTarget: (jumpNotSI jmpTarget: cogit Label).
21  ^CompletePrimitive
```



- Generate code
- Virtual register-based (2AC)
- Low level machine code

Meta-compile a baseline JIT compiler for Pharo



Code generated by Druid

Interpreter

```
1 primitiveAdd
2   <numberOfArguments: 1>
3   <customisedReceiverFor: #smallInteger>
4
5   | maybeSmallInteger maybeSmallInteger2 result |
6
7   maybeSmallInteger := self stackValue: 0.
8   maybeSmallInteger2 := self stackValue: 1.
9
10  (objectMemory isIntegerObject: maybeSmallInteger)
11    ifFalse: [ ^ self primitiveFail ].
12  (objectMemory isIntegerObject: maybeSmallInteger2)
13    ifFalse: [ ^ self primitiveFail ].
14
15  "Check for overflow"
16  result := self
17    sumSmallInteger: maybeSmallInteger
18    withSmallInteger: maybeSmallInteger2
19    ifOverflow: [ ^ self primitiveFail ].
20
21  self pop: 2 thenPush: result
```



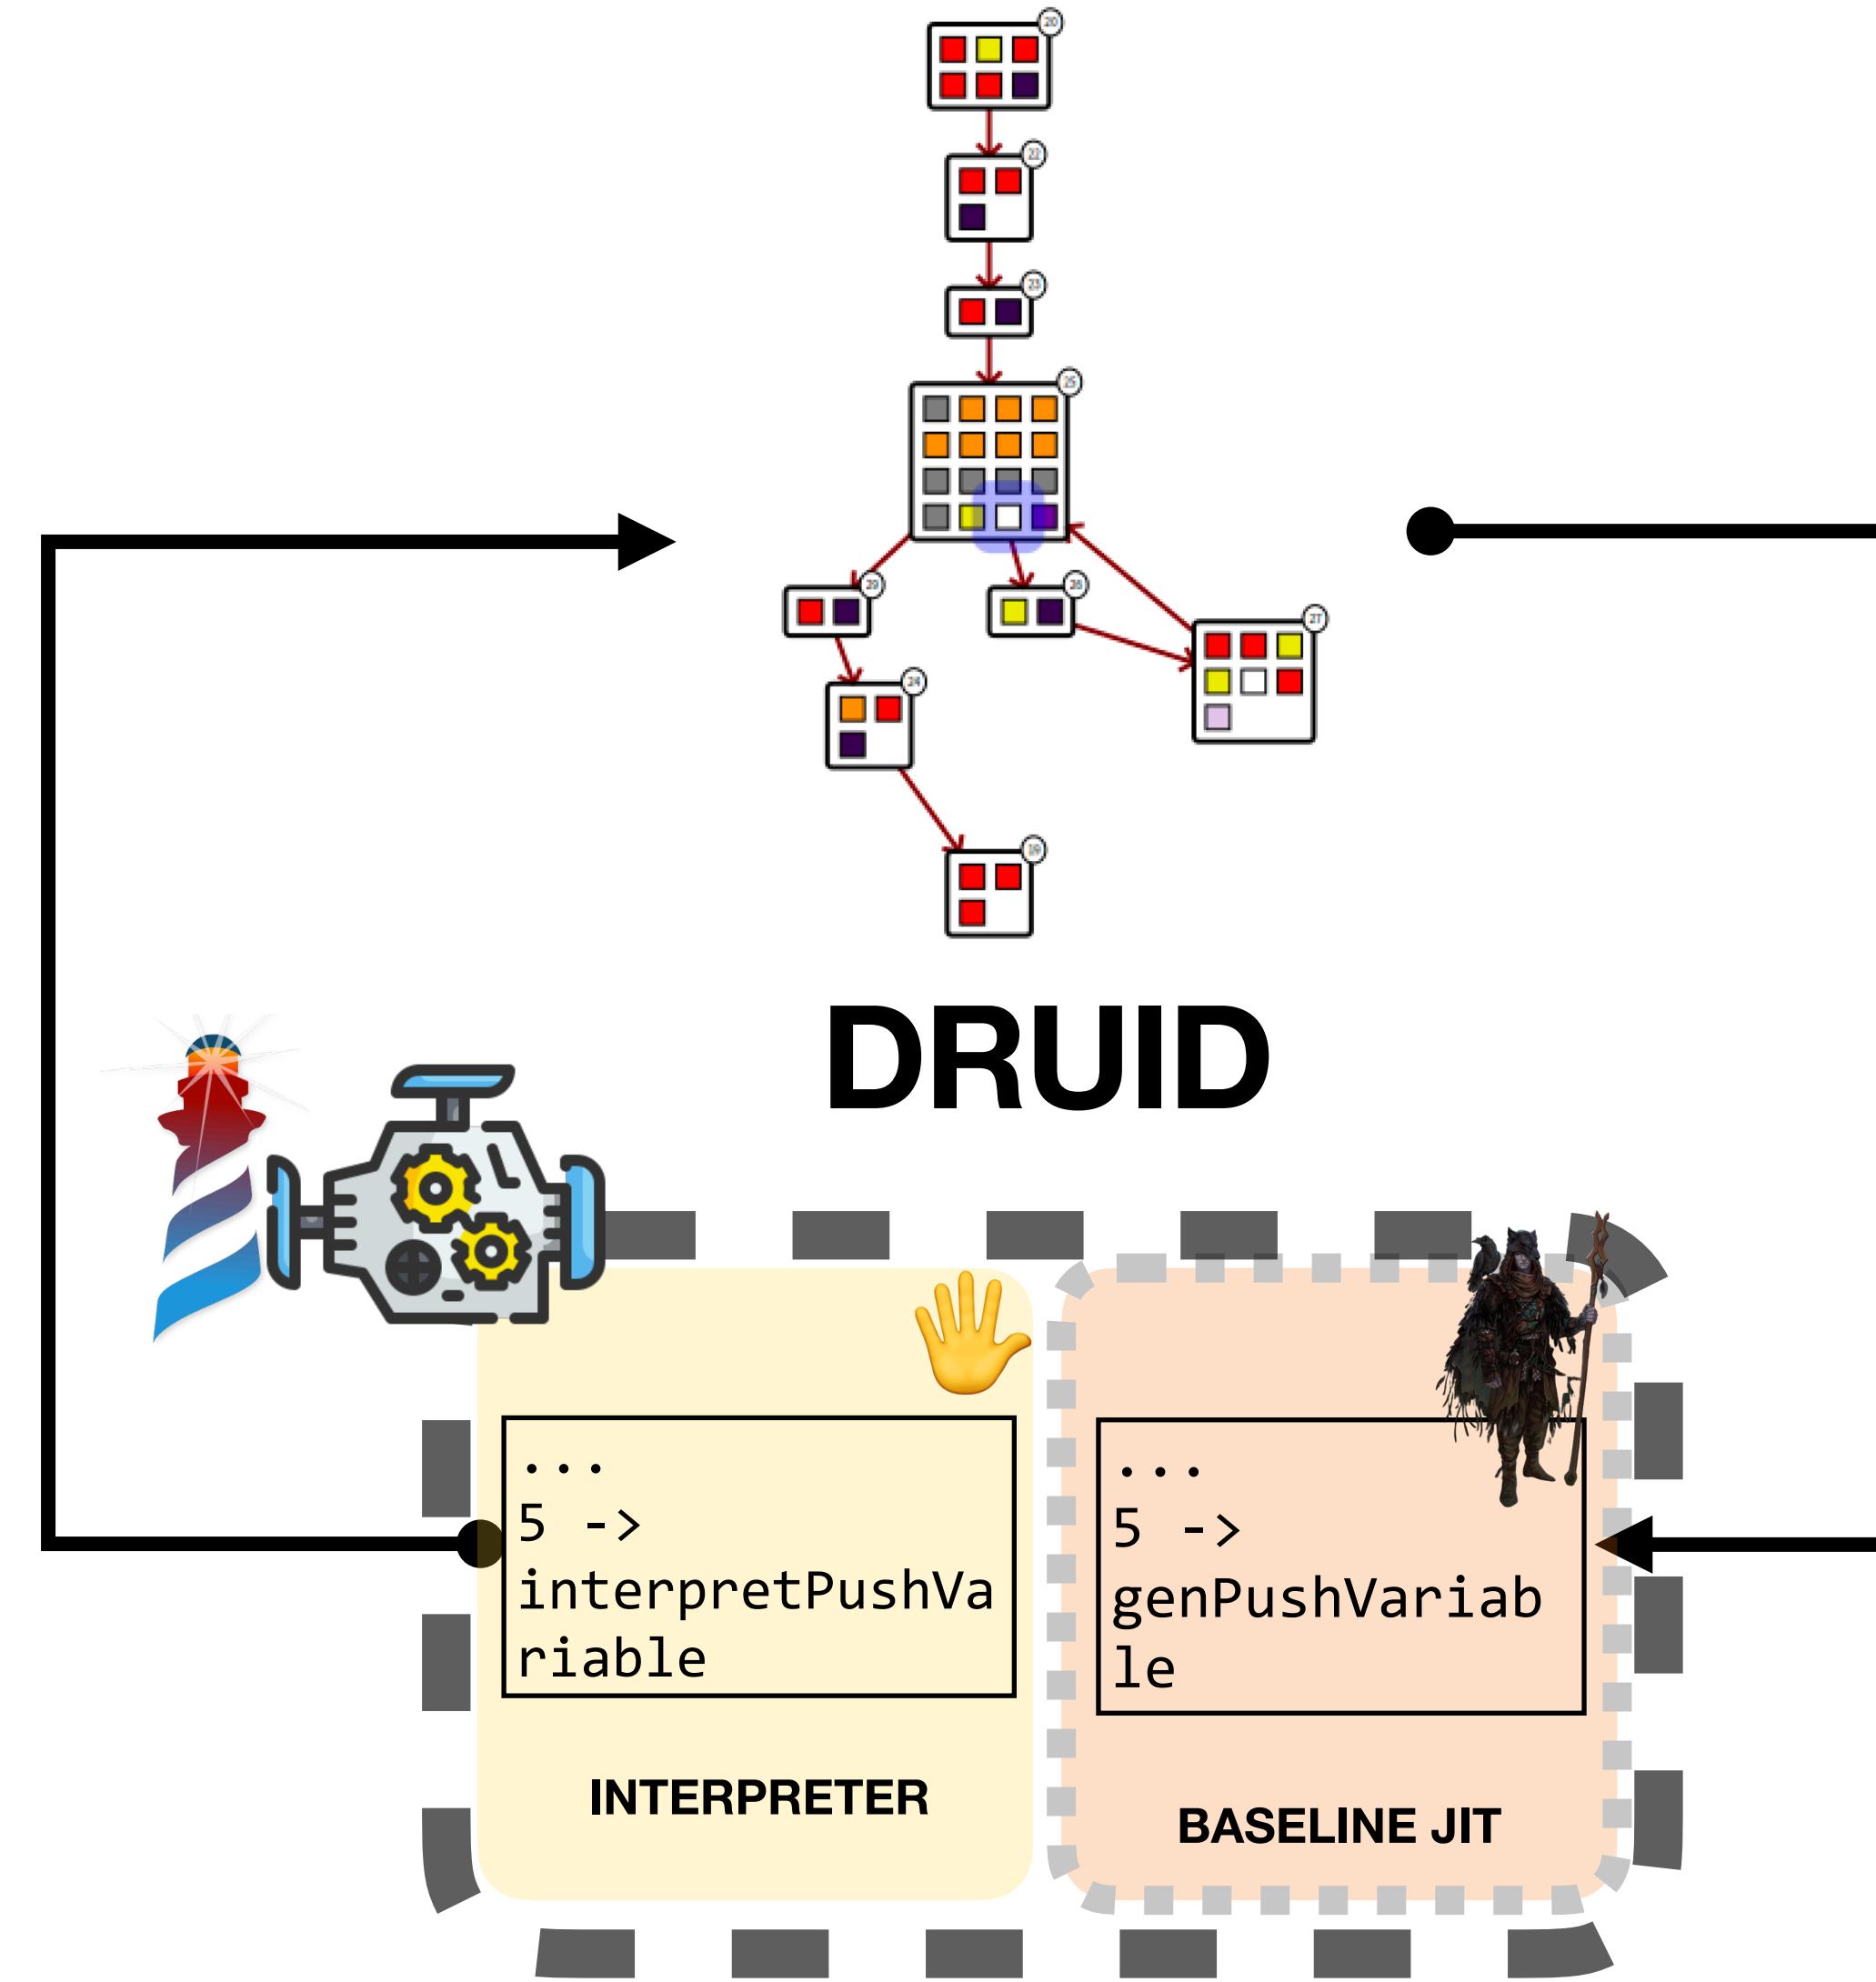
Baseline JIT Compiler



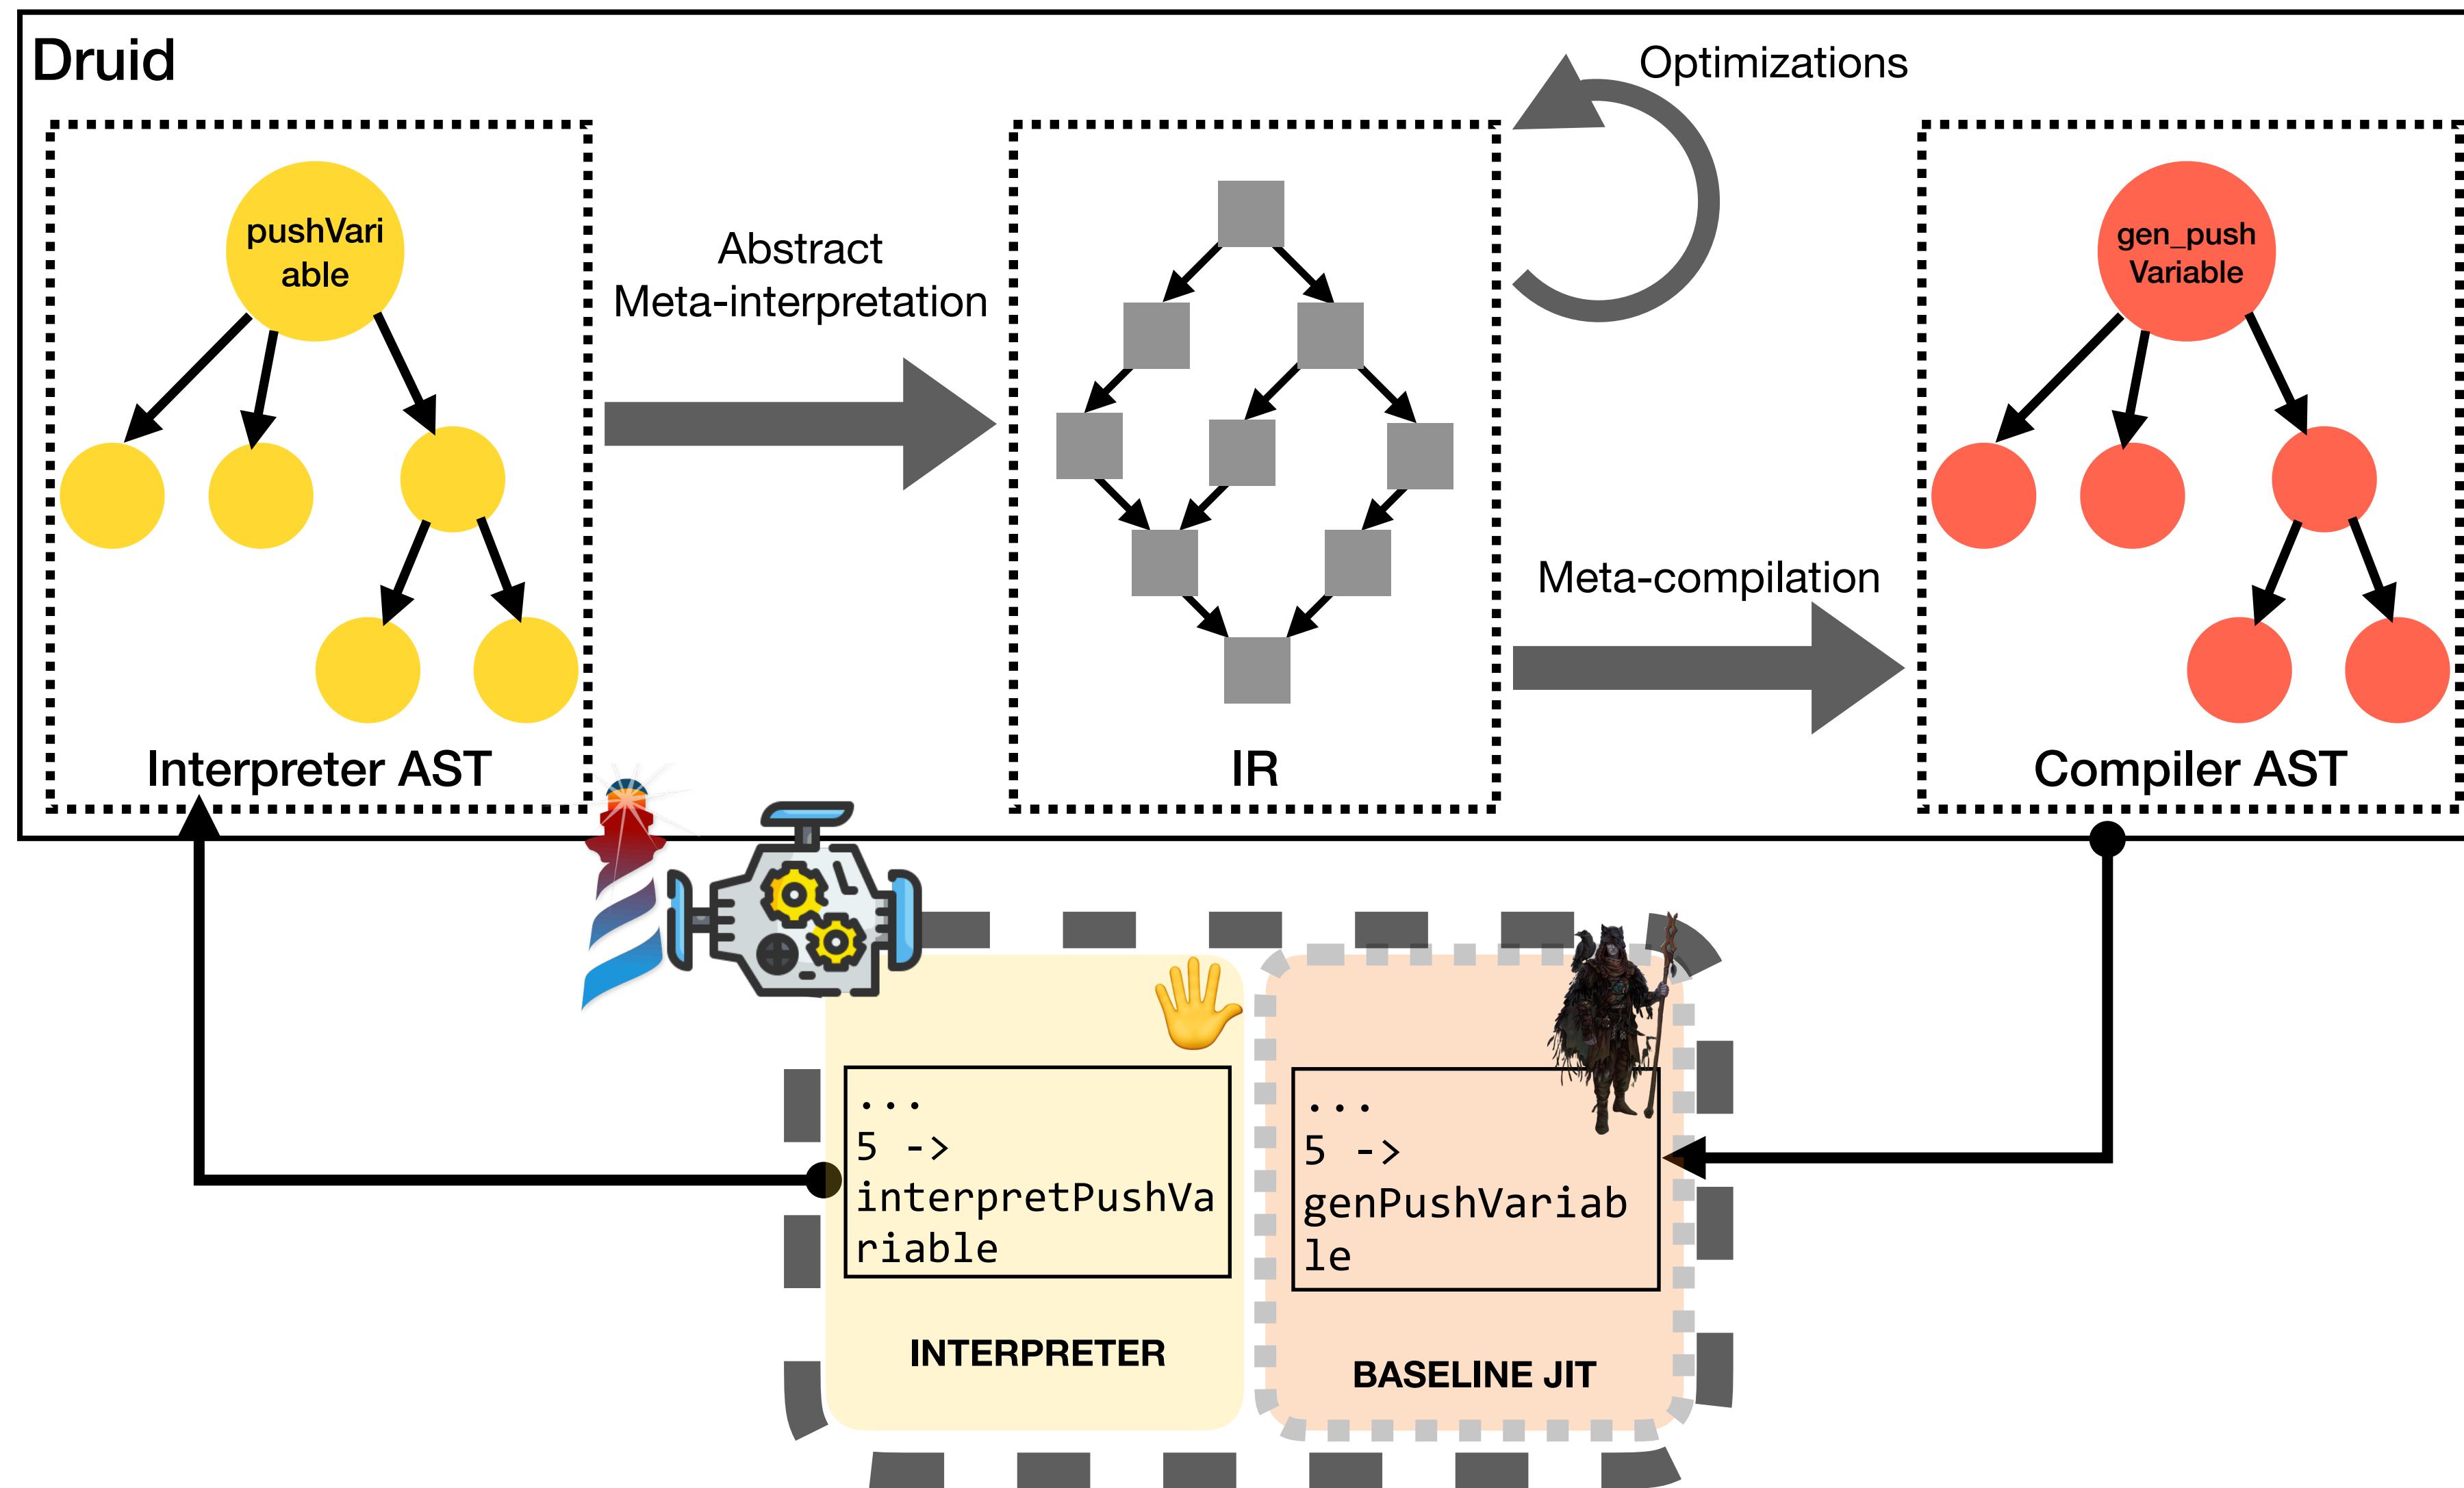
```
1 gen_PrimitiveAdd
2   "AutoGenerated by Druid"
3
4   | jump1 jump2 currentBlock |
5   self mclassIsSmallInteger ifFalse: [ ^ UnimplementedPrimitive ].
6   self MoveR: ReceiverResultReg R: ClassReg.
7   self MoveR: Arg0Reg R: SendNumArgsReg.
8   self TstCq: 1 R: SendNumArgsReg.
9   jump1 := self JumpZero: 0.
10  self AddCq: -1 R: SendNumArgsReg.
11  self AddR: ClassReg R: SendNumArgsReg.
12  jump2 := self JumpOverflow: 0.
13  self MoveR: SendNumArgsReg R: ReceiverResultReg.
14  self genPrimReturn.
15  currentBlock := self Label.
16  jump1 jmpTarget: currentBlock.
17  jump2 jmpTarget: currentBlock.
18  ^ CompletePrimitive
```



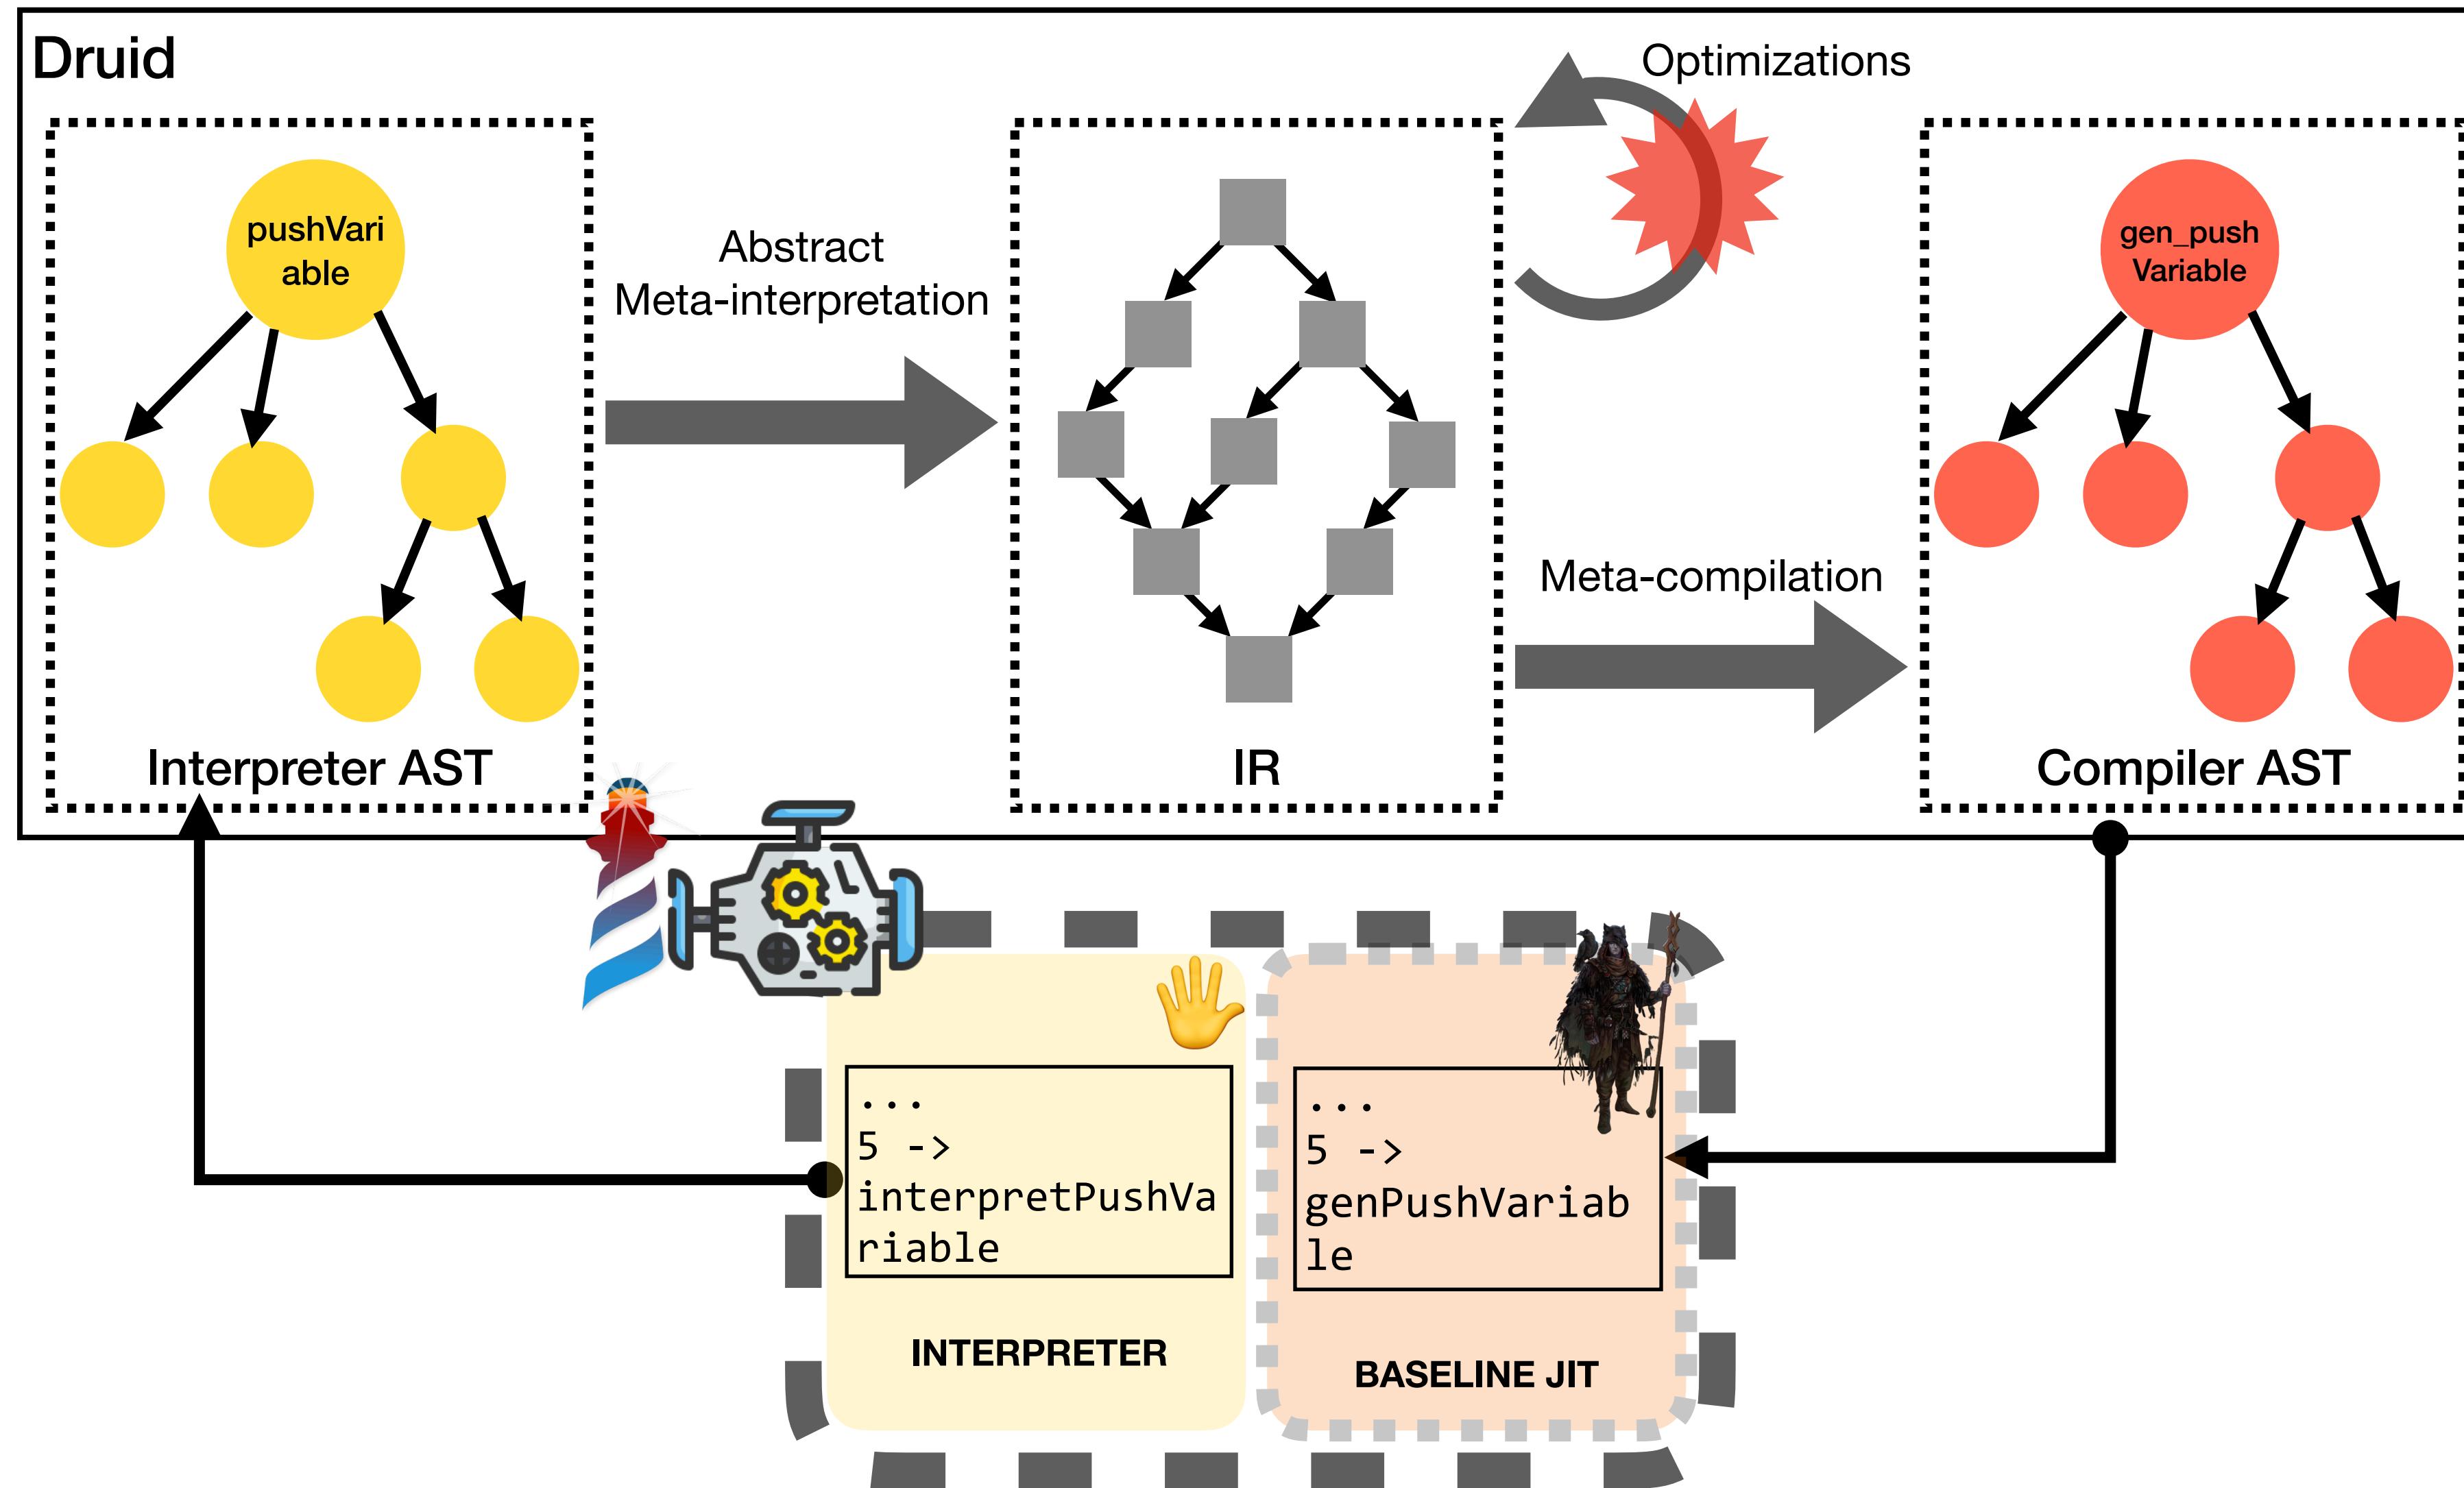
Meta-compile a baseline JIT compiler for Pharo



Meta-compile a baseline JIT compiler for Pharo



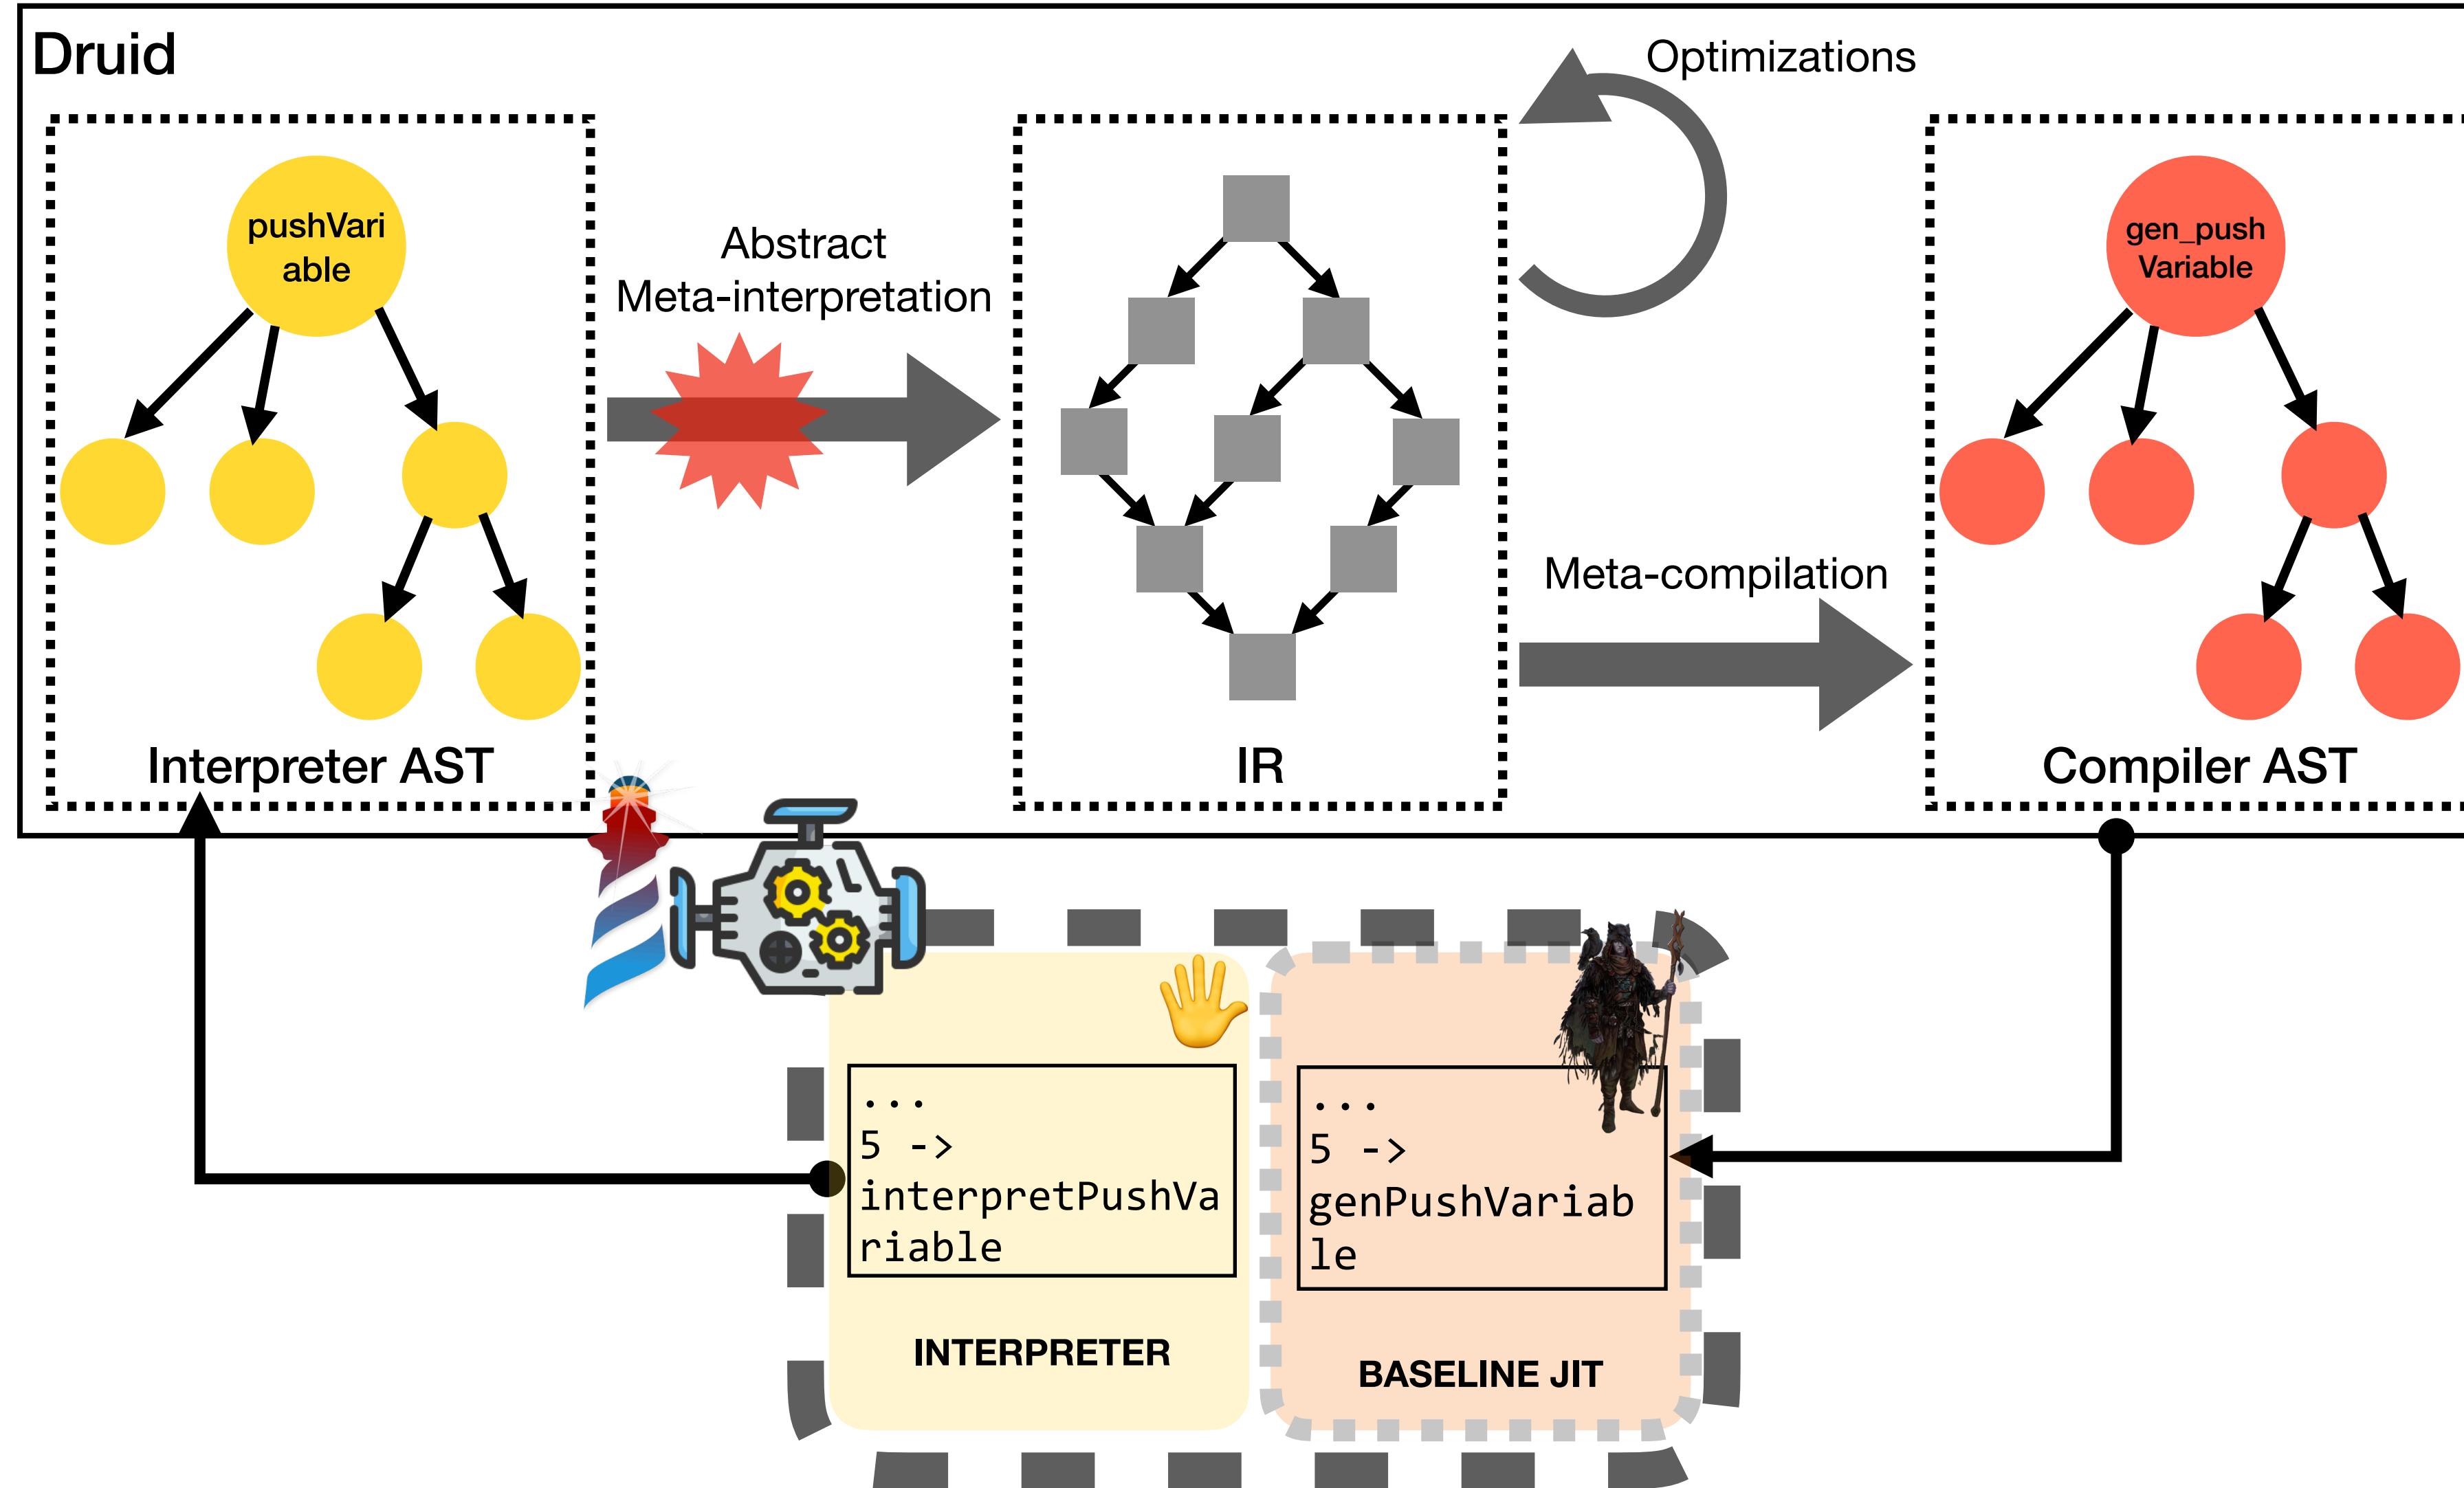
Meta-compile a baseline JIT compiler for Pharo



Druid Optimizations

Name	Description	Open opportunities
BRANCH COLLAPSE	Inline conditions in conditional jumps	Dead path analysis
CLEAN CONTROL FLOW	Merge instructions in consecutive blocks to avoid unnecessary jumps	Better local analysis
COPY PROPAGATION	Replace copy instructions in operands by real value	Better dependency analysis and possible unused code
DEAD BLOCK ELIMINATION	Remove inaccessible blocks	-
DEAD BRANCH ELIMINATION	Remove branches with only dead paths	Better propagations
DEAD CODE ELIMINATION	Remove unused instructions	-
DEAD EDGE SPLITTING	Duplicate blocks with dead and not-dead paths	Create branches with only dead paths
FAILURE CODE TAIL DUPLICATION	Tail duplicate block with resulted code	Divide fail and success paths
PHI SIMPLIFICATION	Replace one operand phis with copy instruction	Better propagations
REDUNDANT COPY ELIMINATION	Remove copies of form $x := x$	-
SCCP	Performs constants folding and propagation	Better dependency analysis and possible unused code

Meta-interpretation guided by intrinsics



Meta-compilation guided by intrinsics

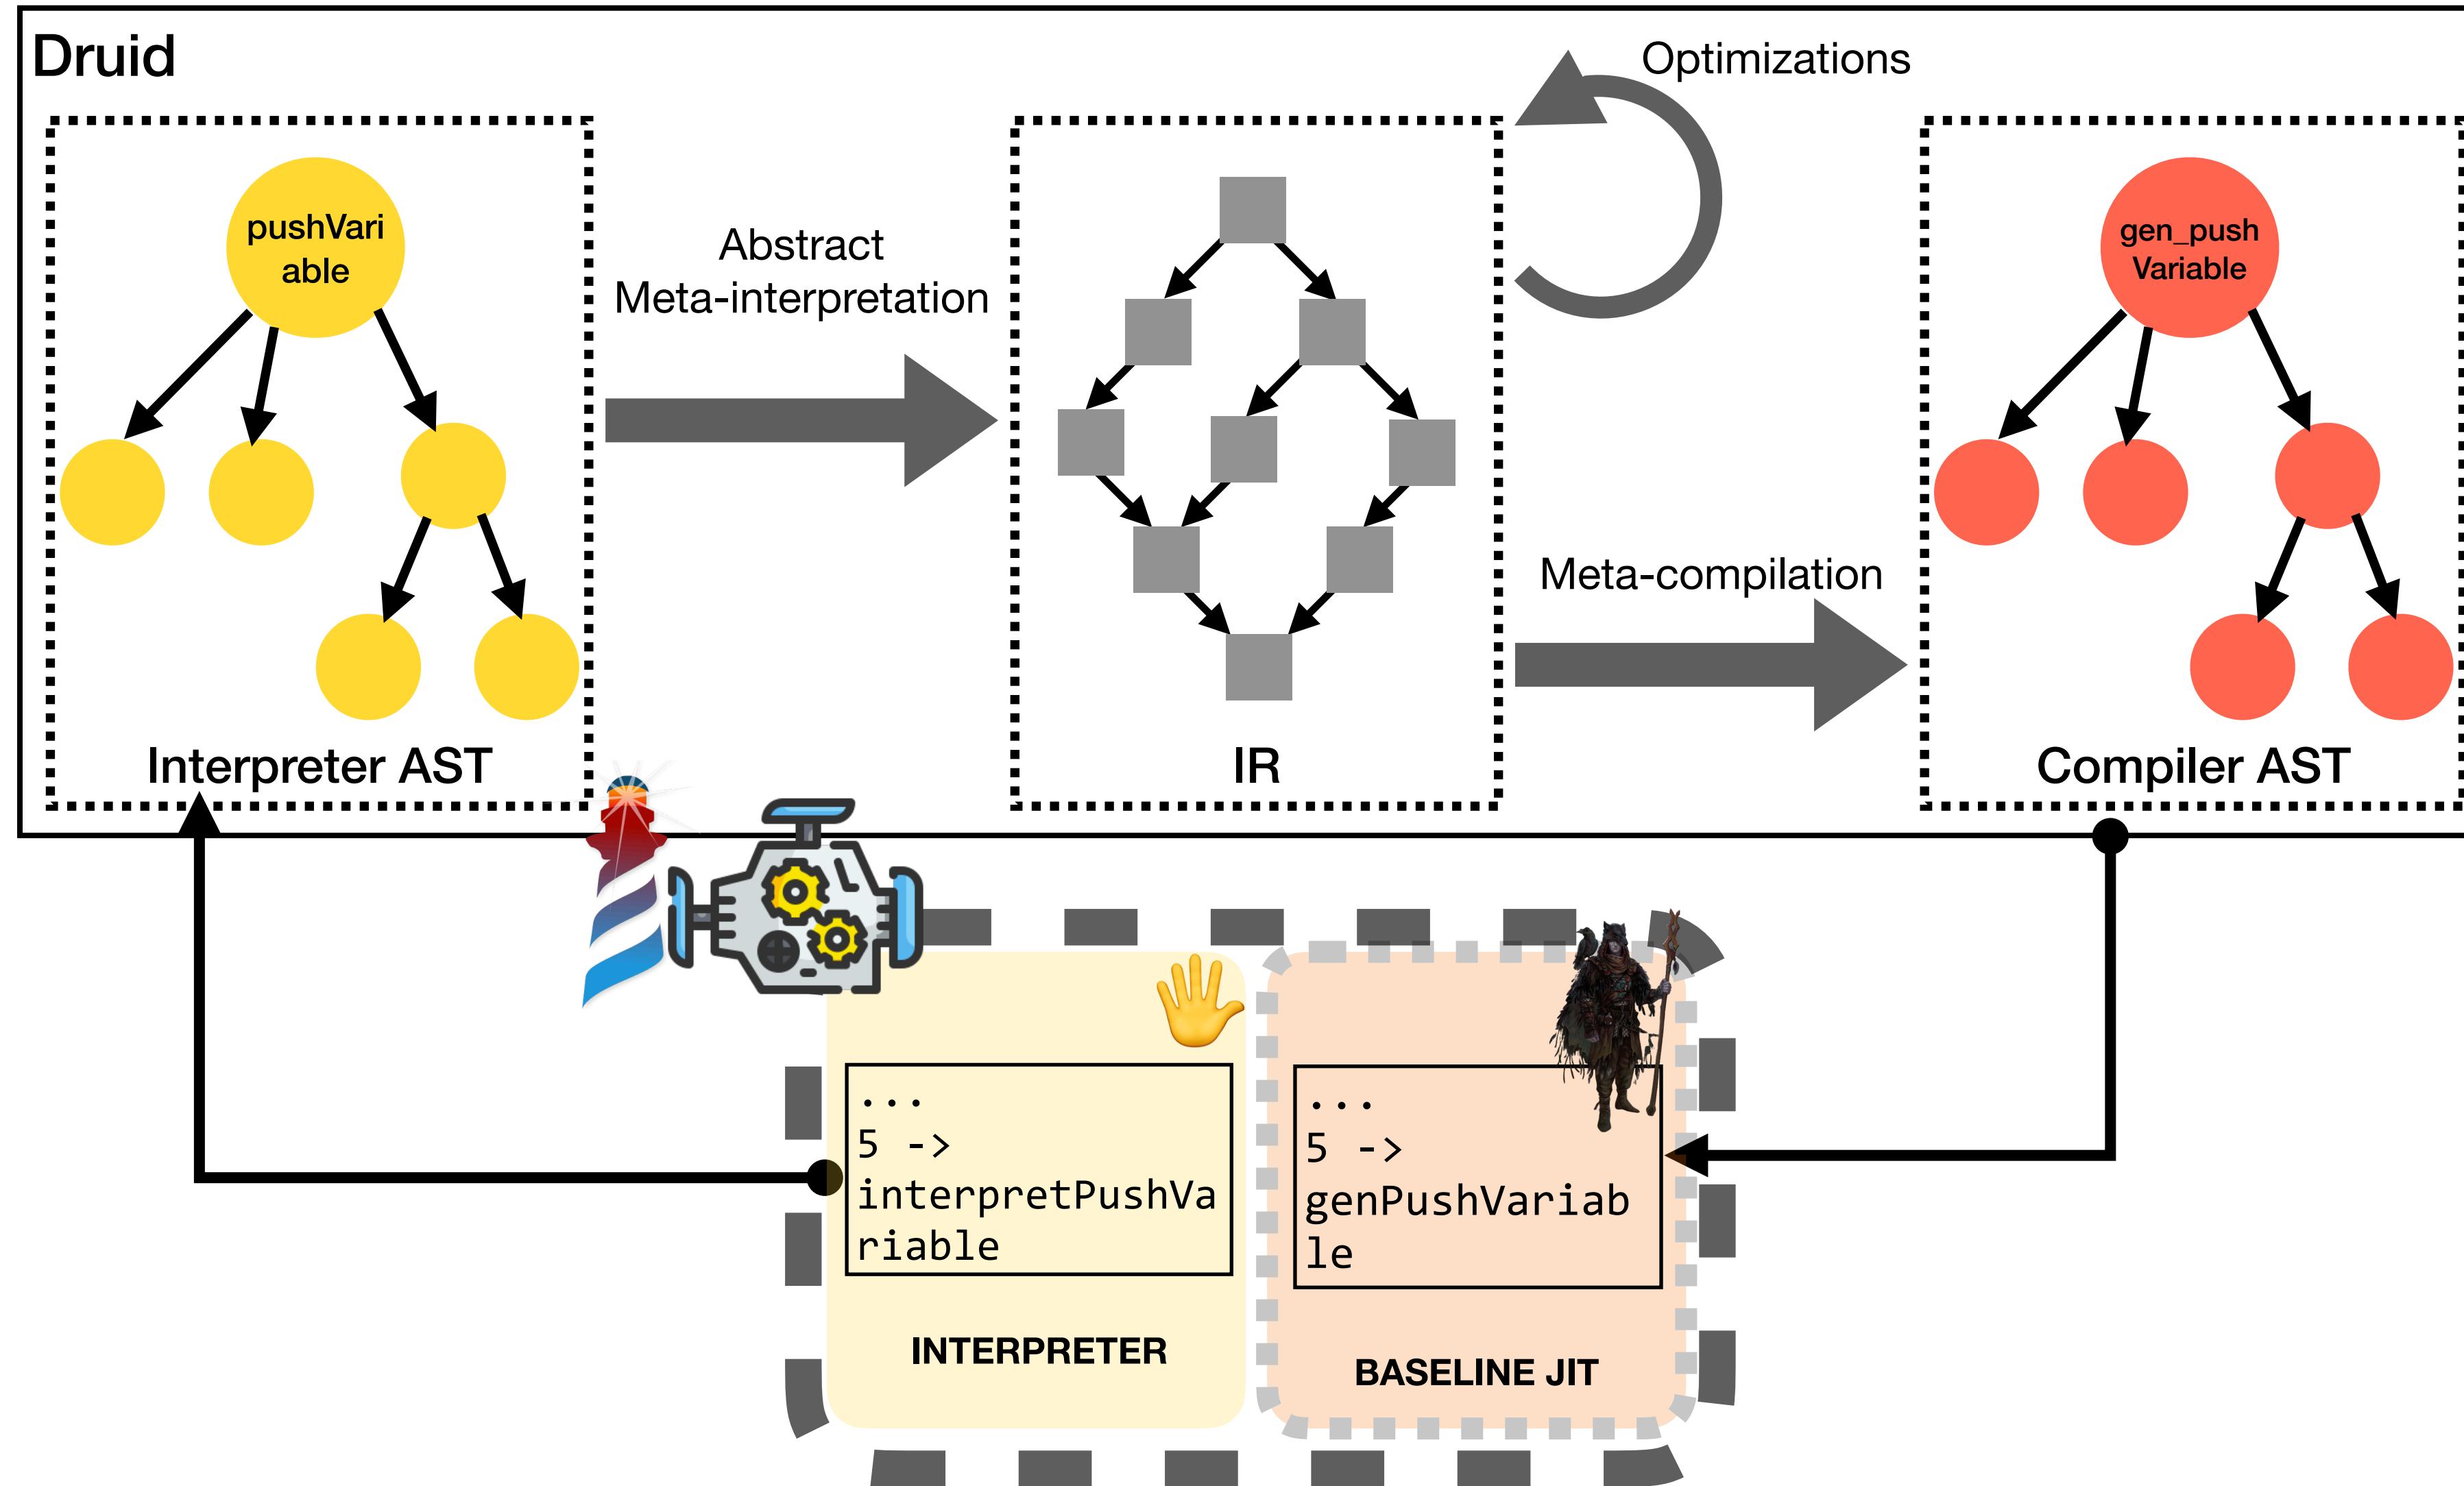
Interpreter

```
pushMaybeCon x

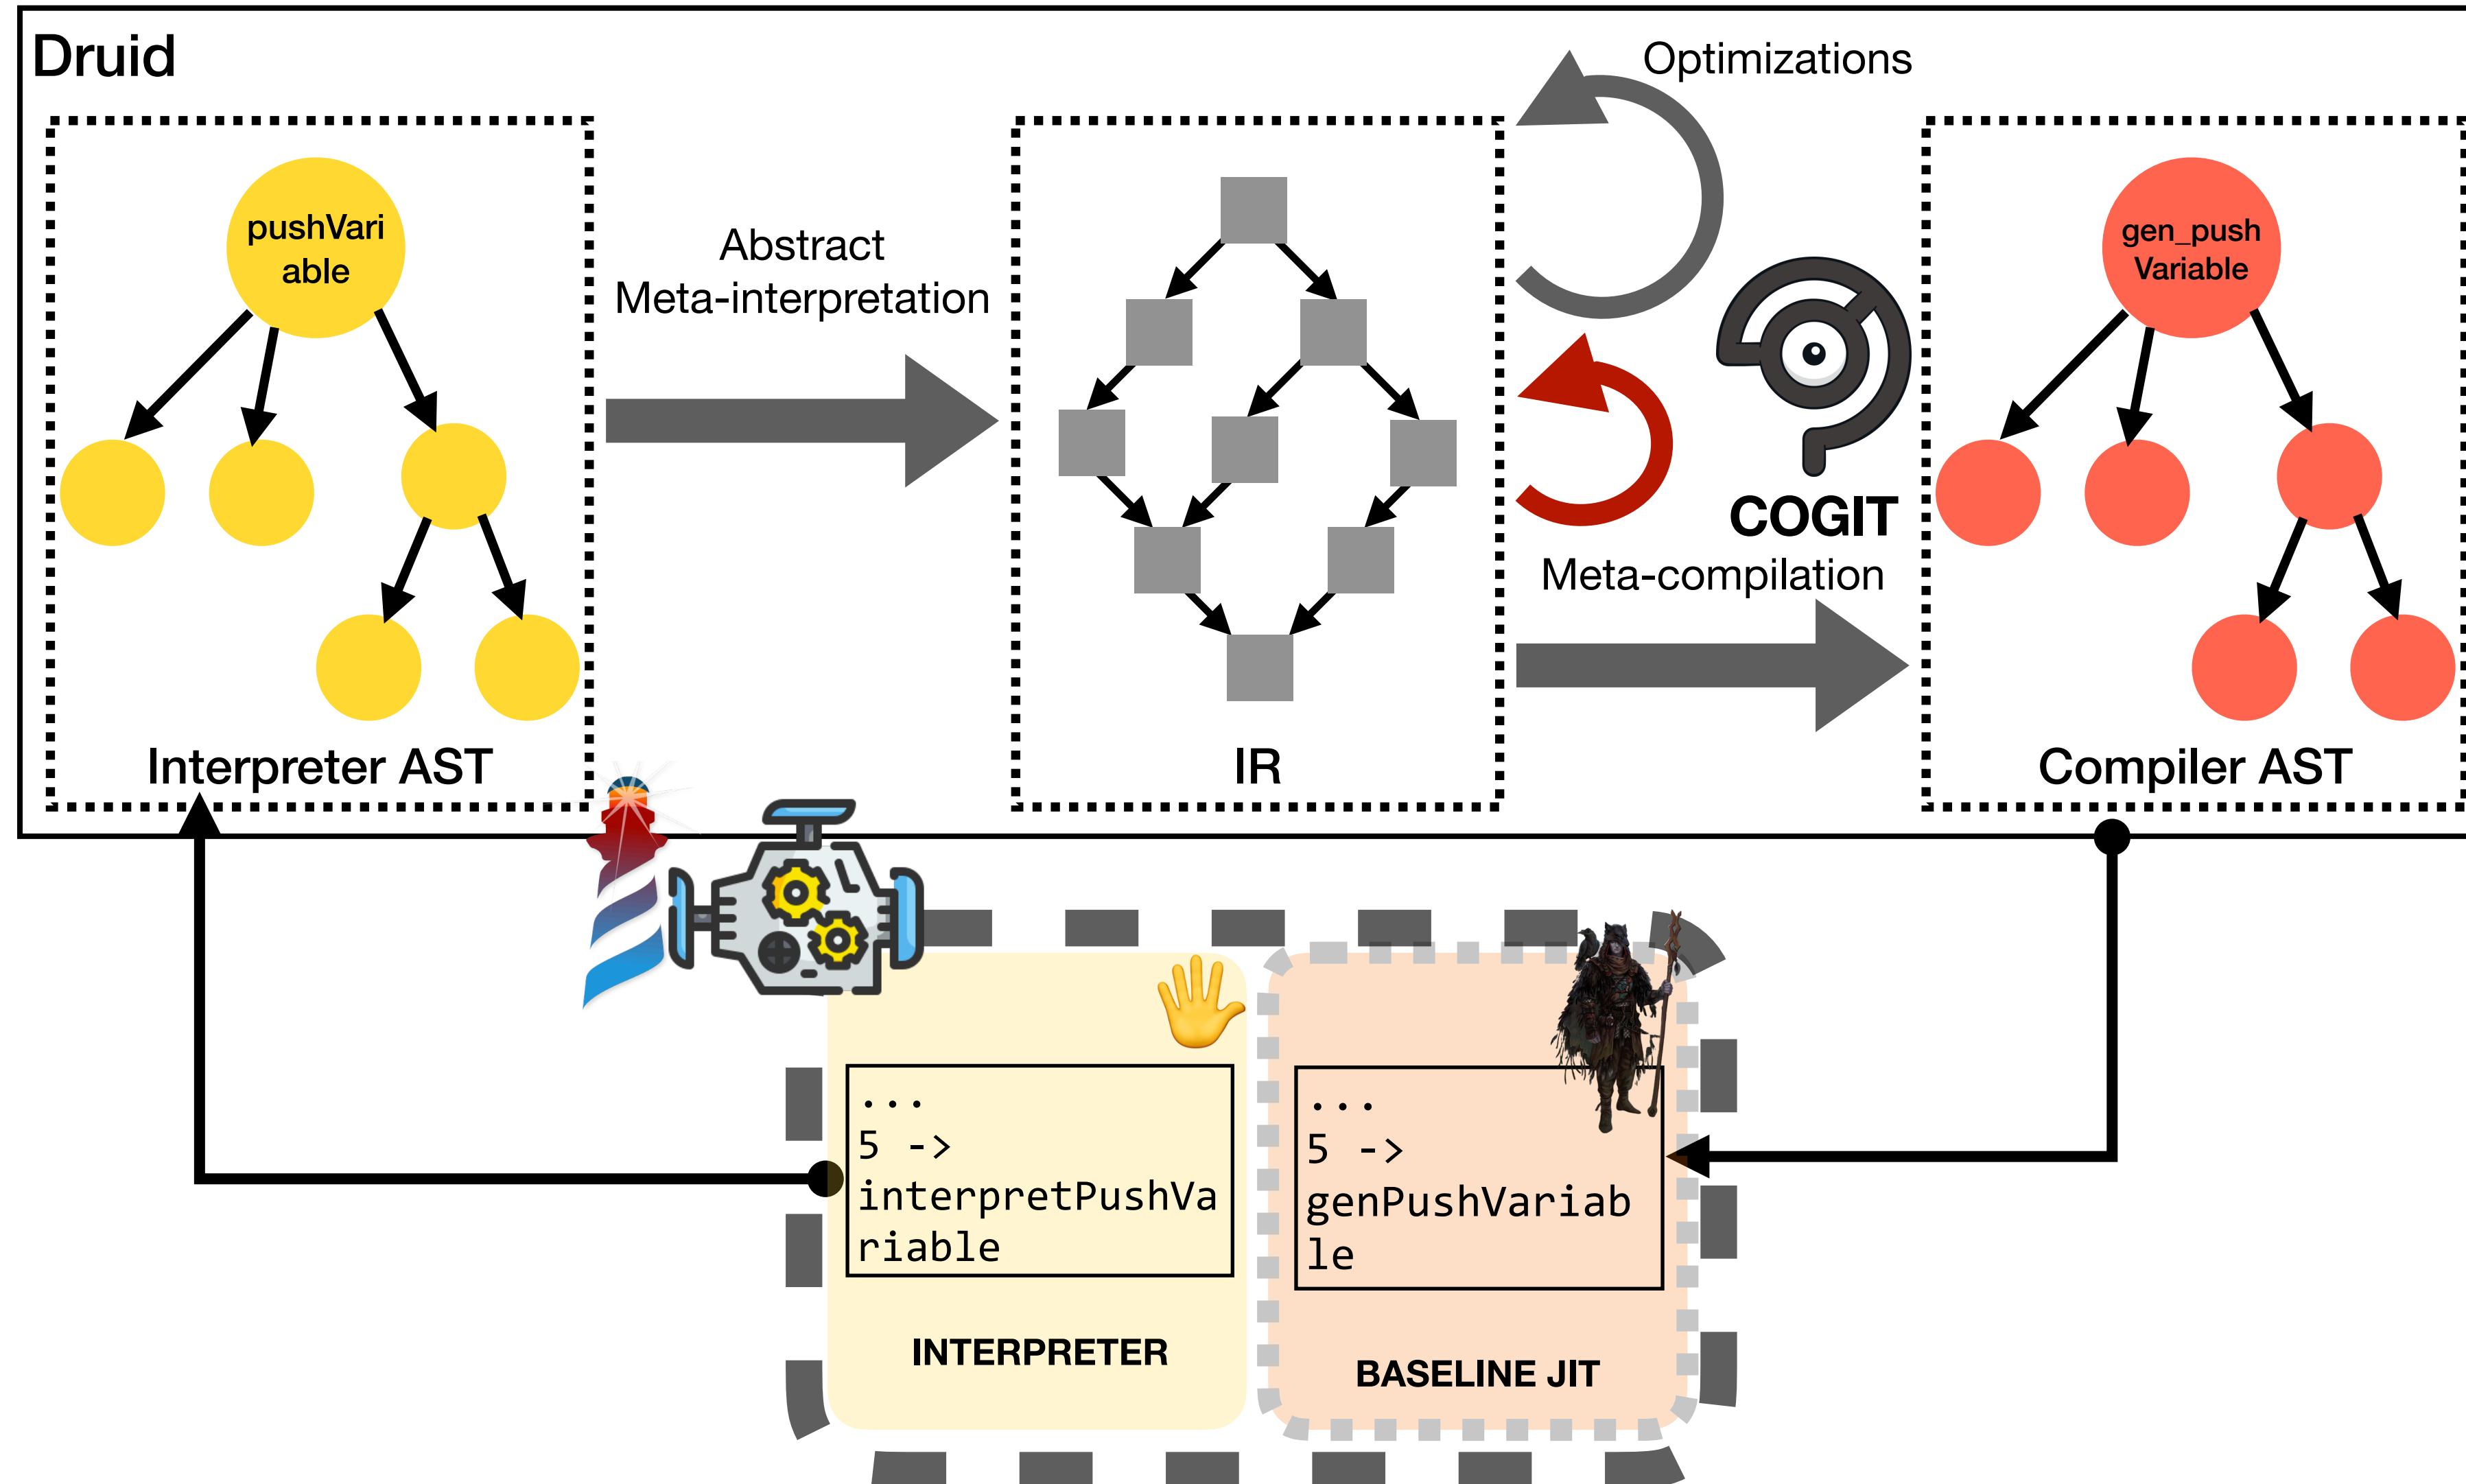
1 pushMaybeContext: obj receiverVariable: fieldIndex
2   "Must trap accesses to married and widowed contexts.
3   But don't want to check on all inst var accesses. This
4   method is only used by the long-form bytecodes, evading
5   the cost. Note that the method, closure and receiver fields
6   of married contexts are correctly initialized so they don't
7   need special treatment on read. Only sender, instruction
8   pointer and stack pointer need to be intercepted on reads."
9
10 <inline: true>
11 ((self isReadMediatedContextInstVarIndex: fieldIndex) and: [
12   objectMemory isContextNonImm: obj ])
13 ifTrue: [
14   self druidForceInterpretation.
15   self druidIgnore: [
16     self push: (self instVar: fieldIndex ofContext: obj) ] ]
17 ifFalse: [
18   self push: (objectMemory fetchPointer: fieldIndex ofObject: obj) ]
```



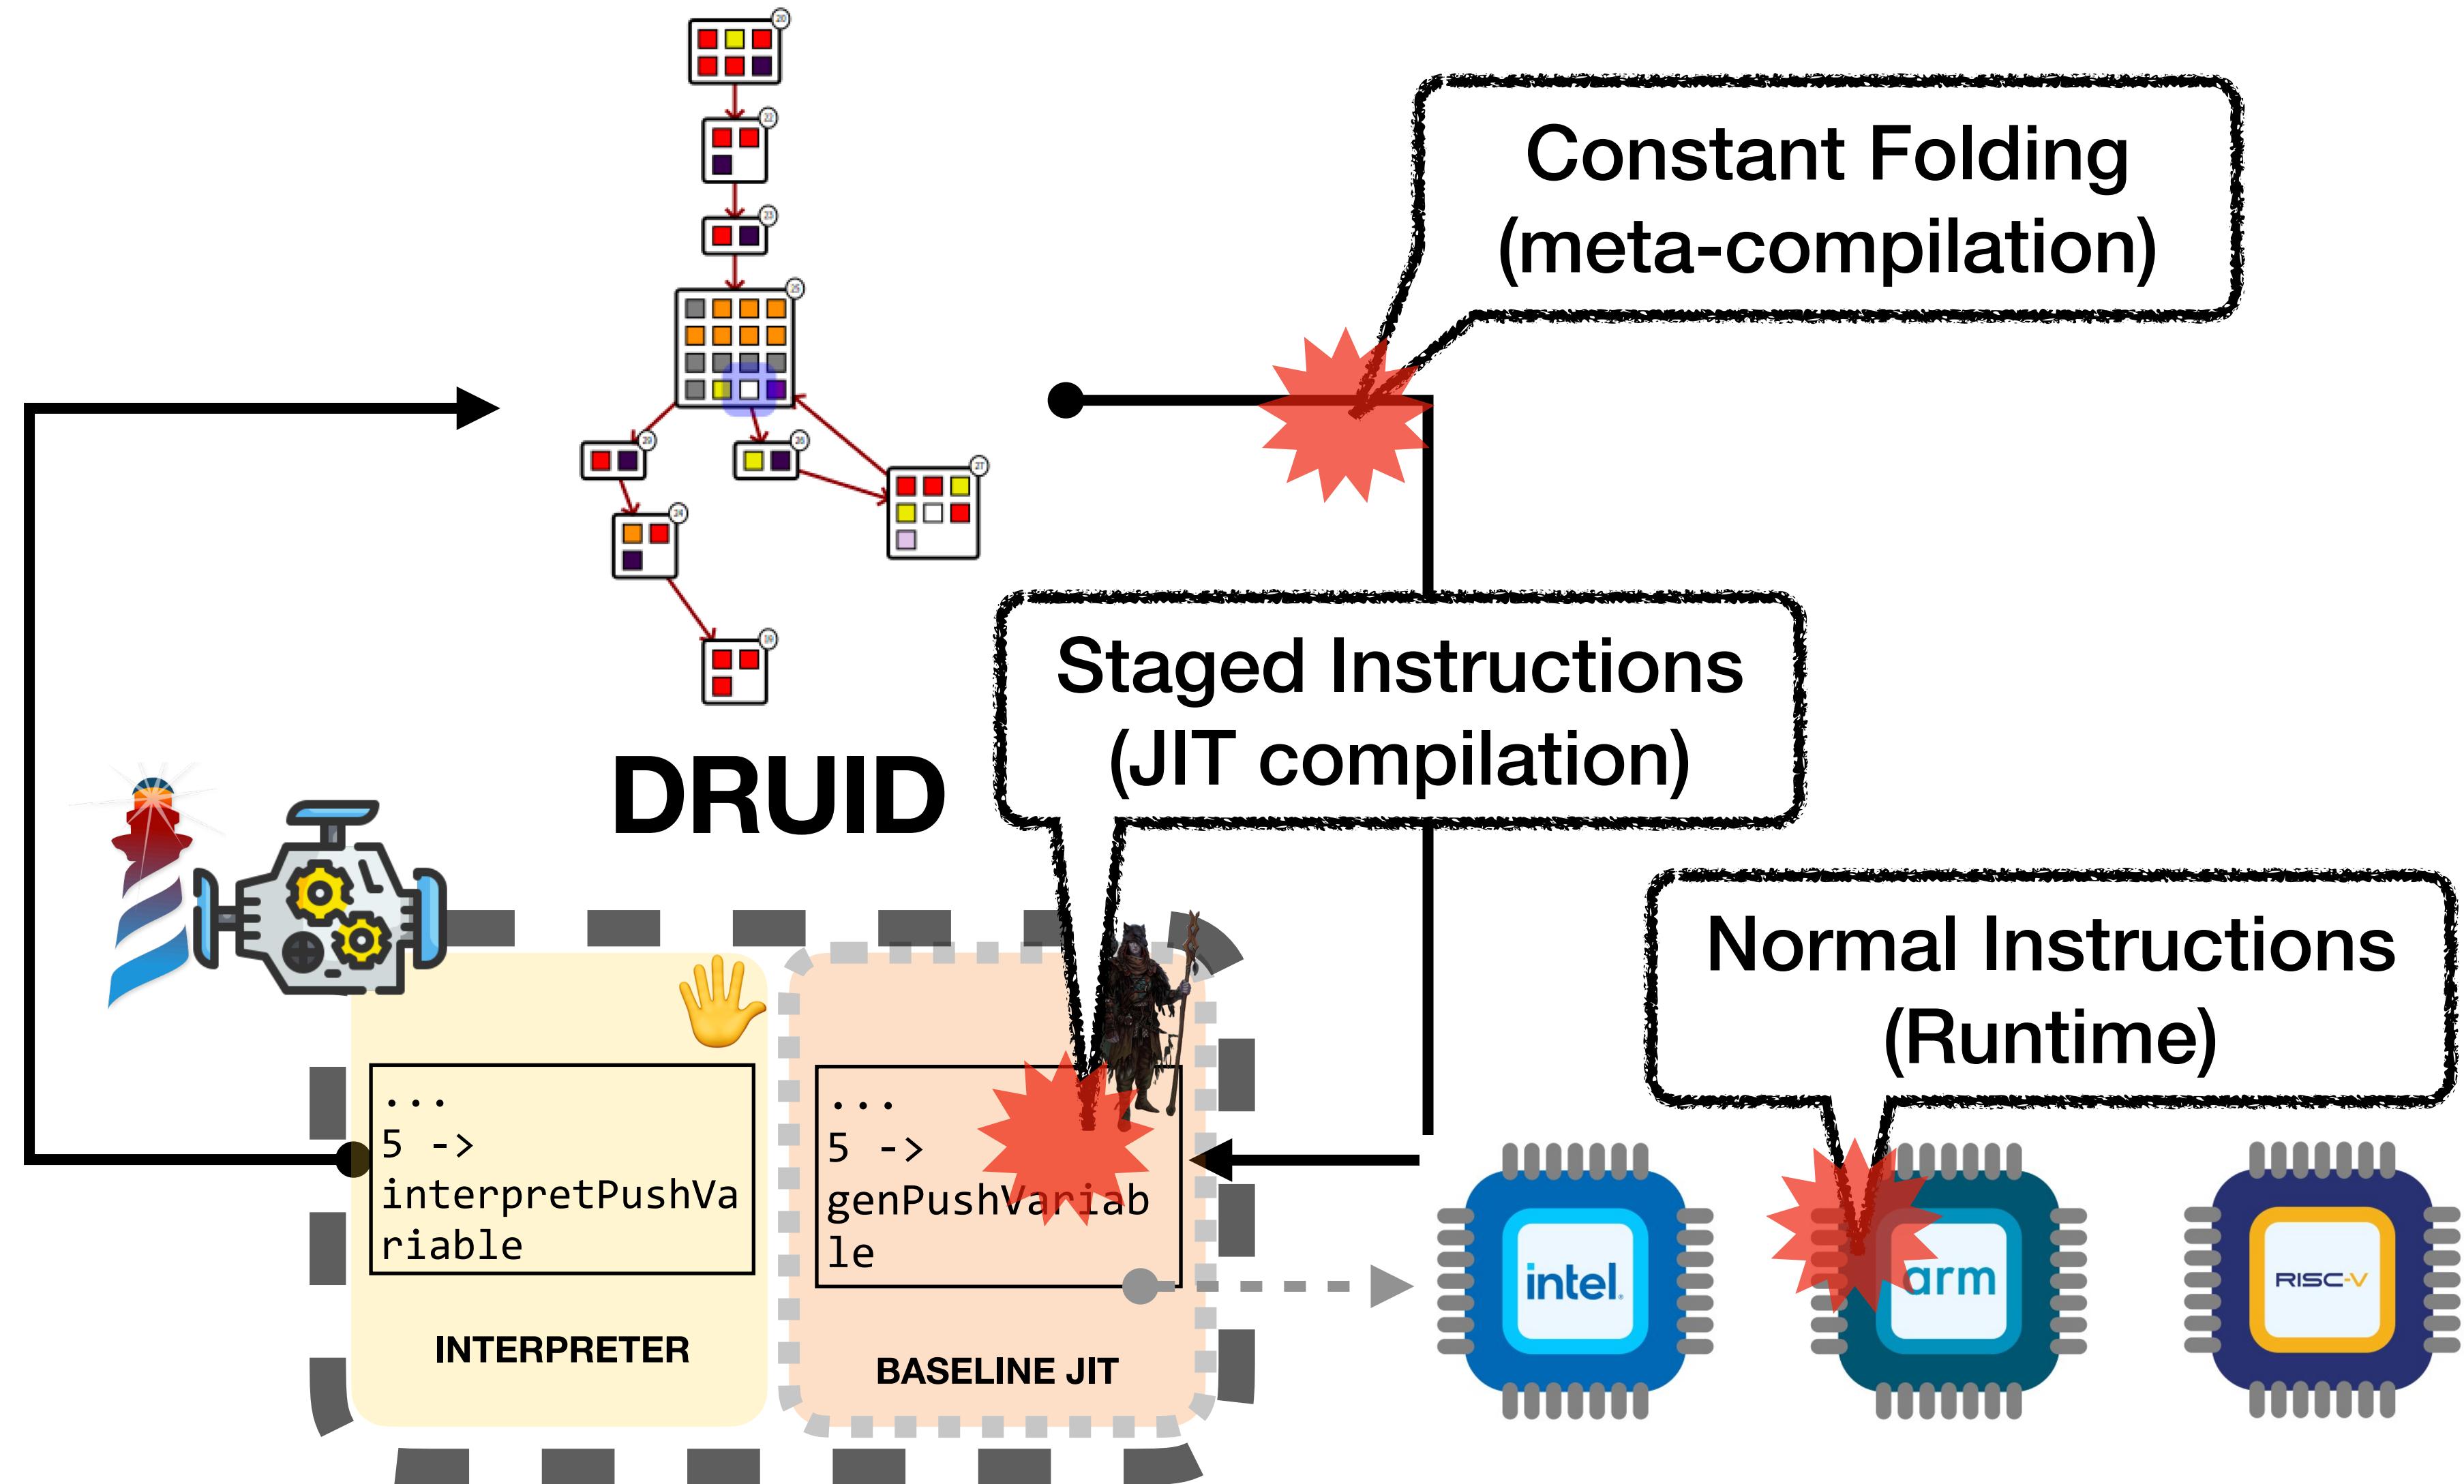
Target Cogit framework



Target Cogit framework



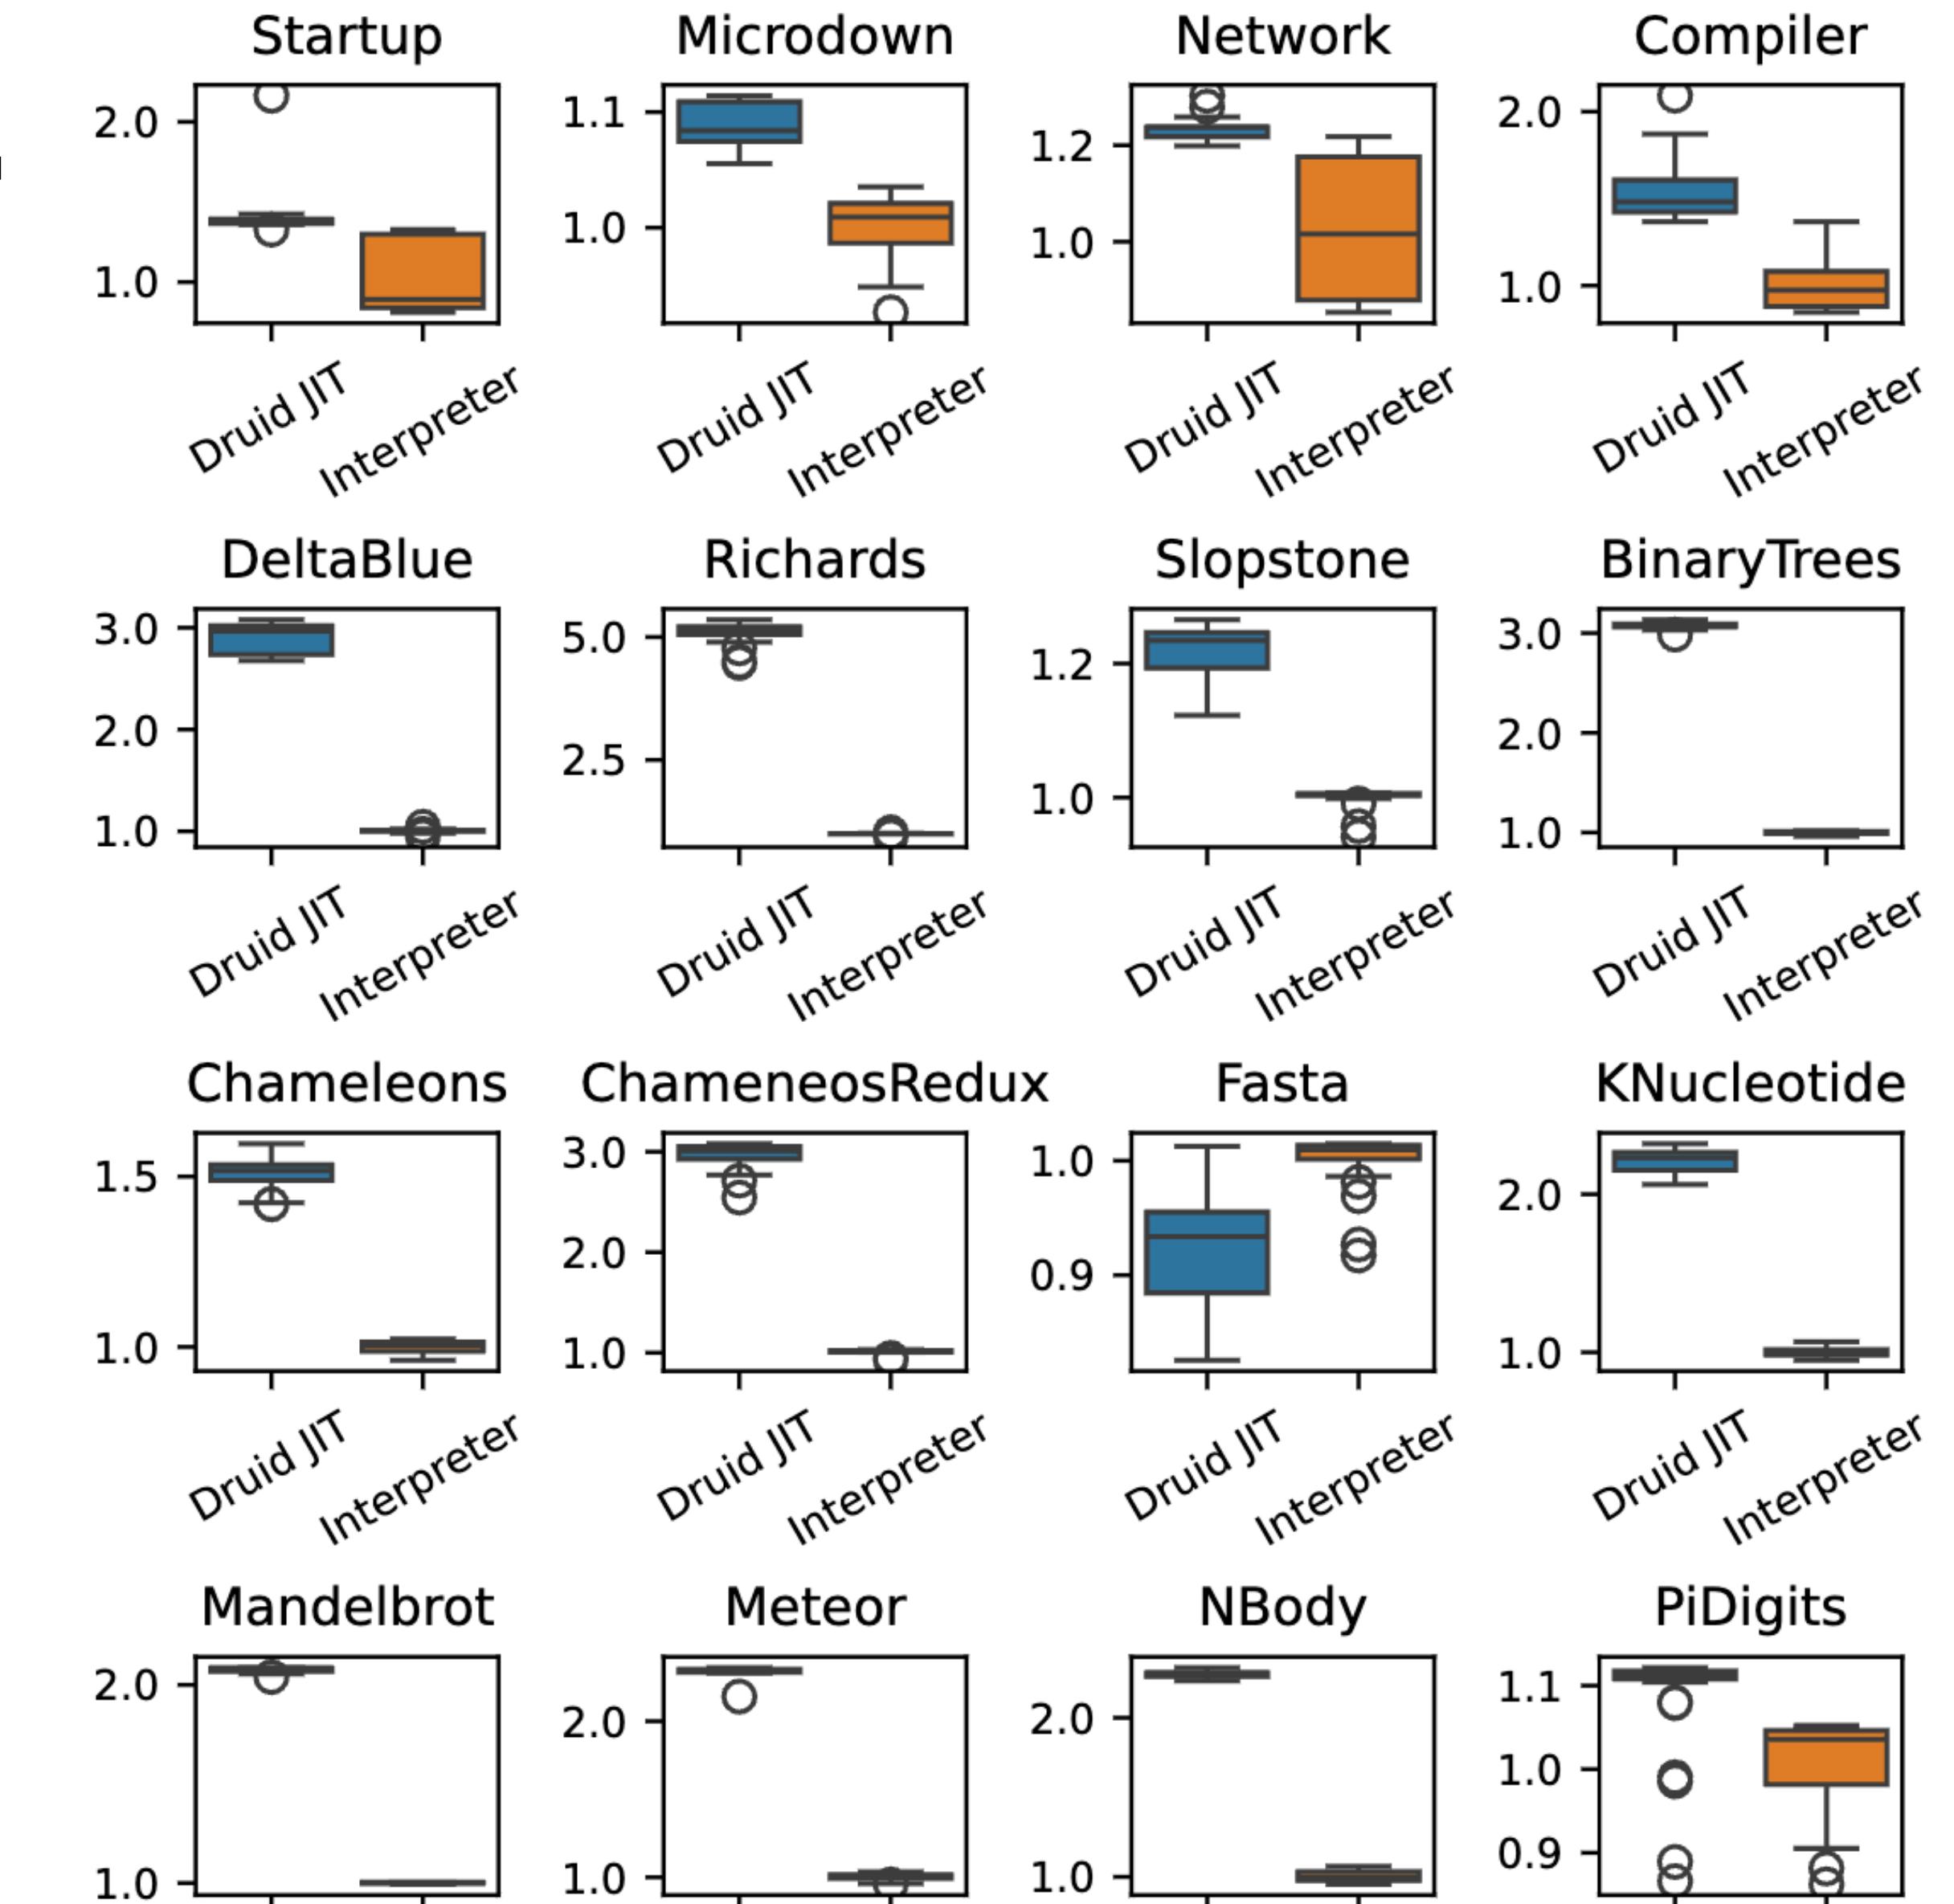
Staging



Evaluation



Druid JIT vs Interpreter

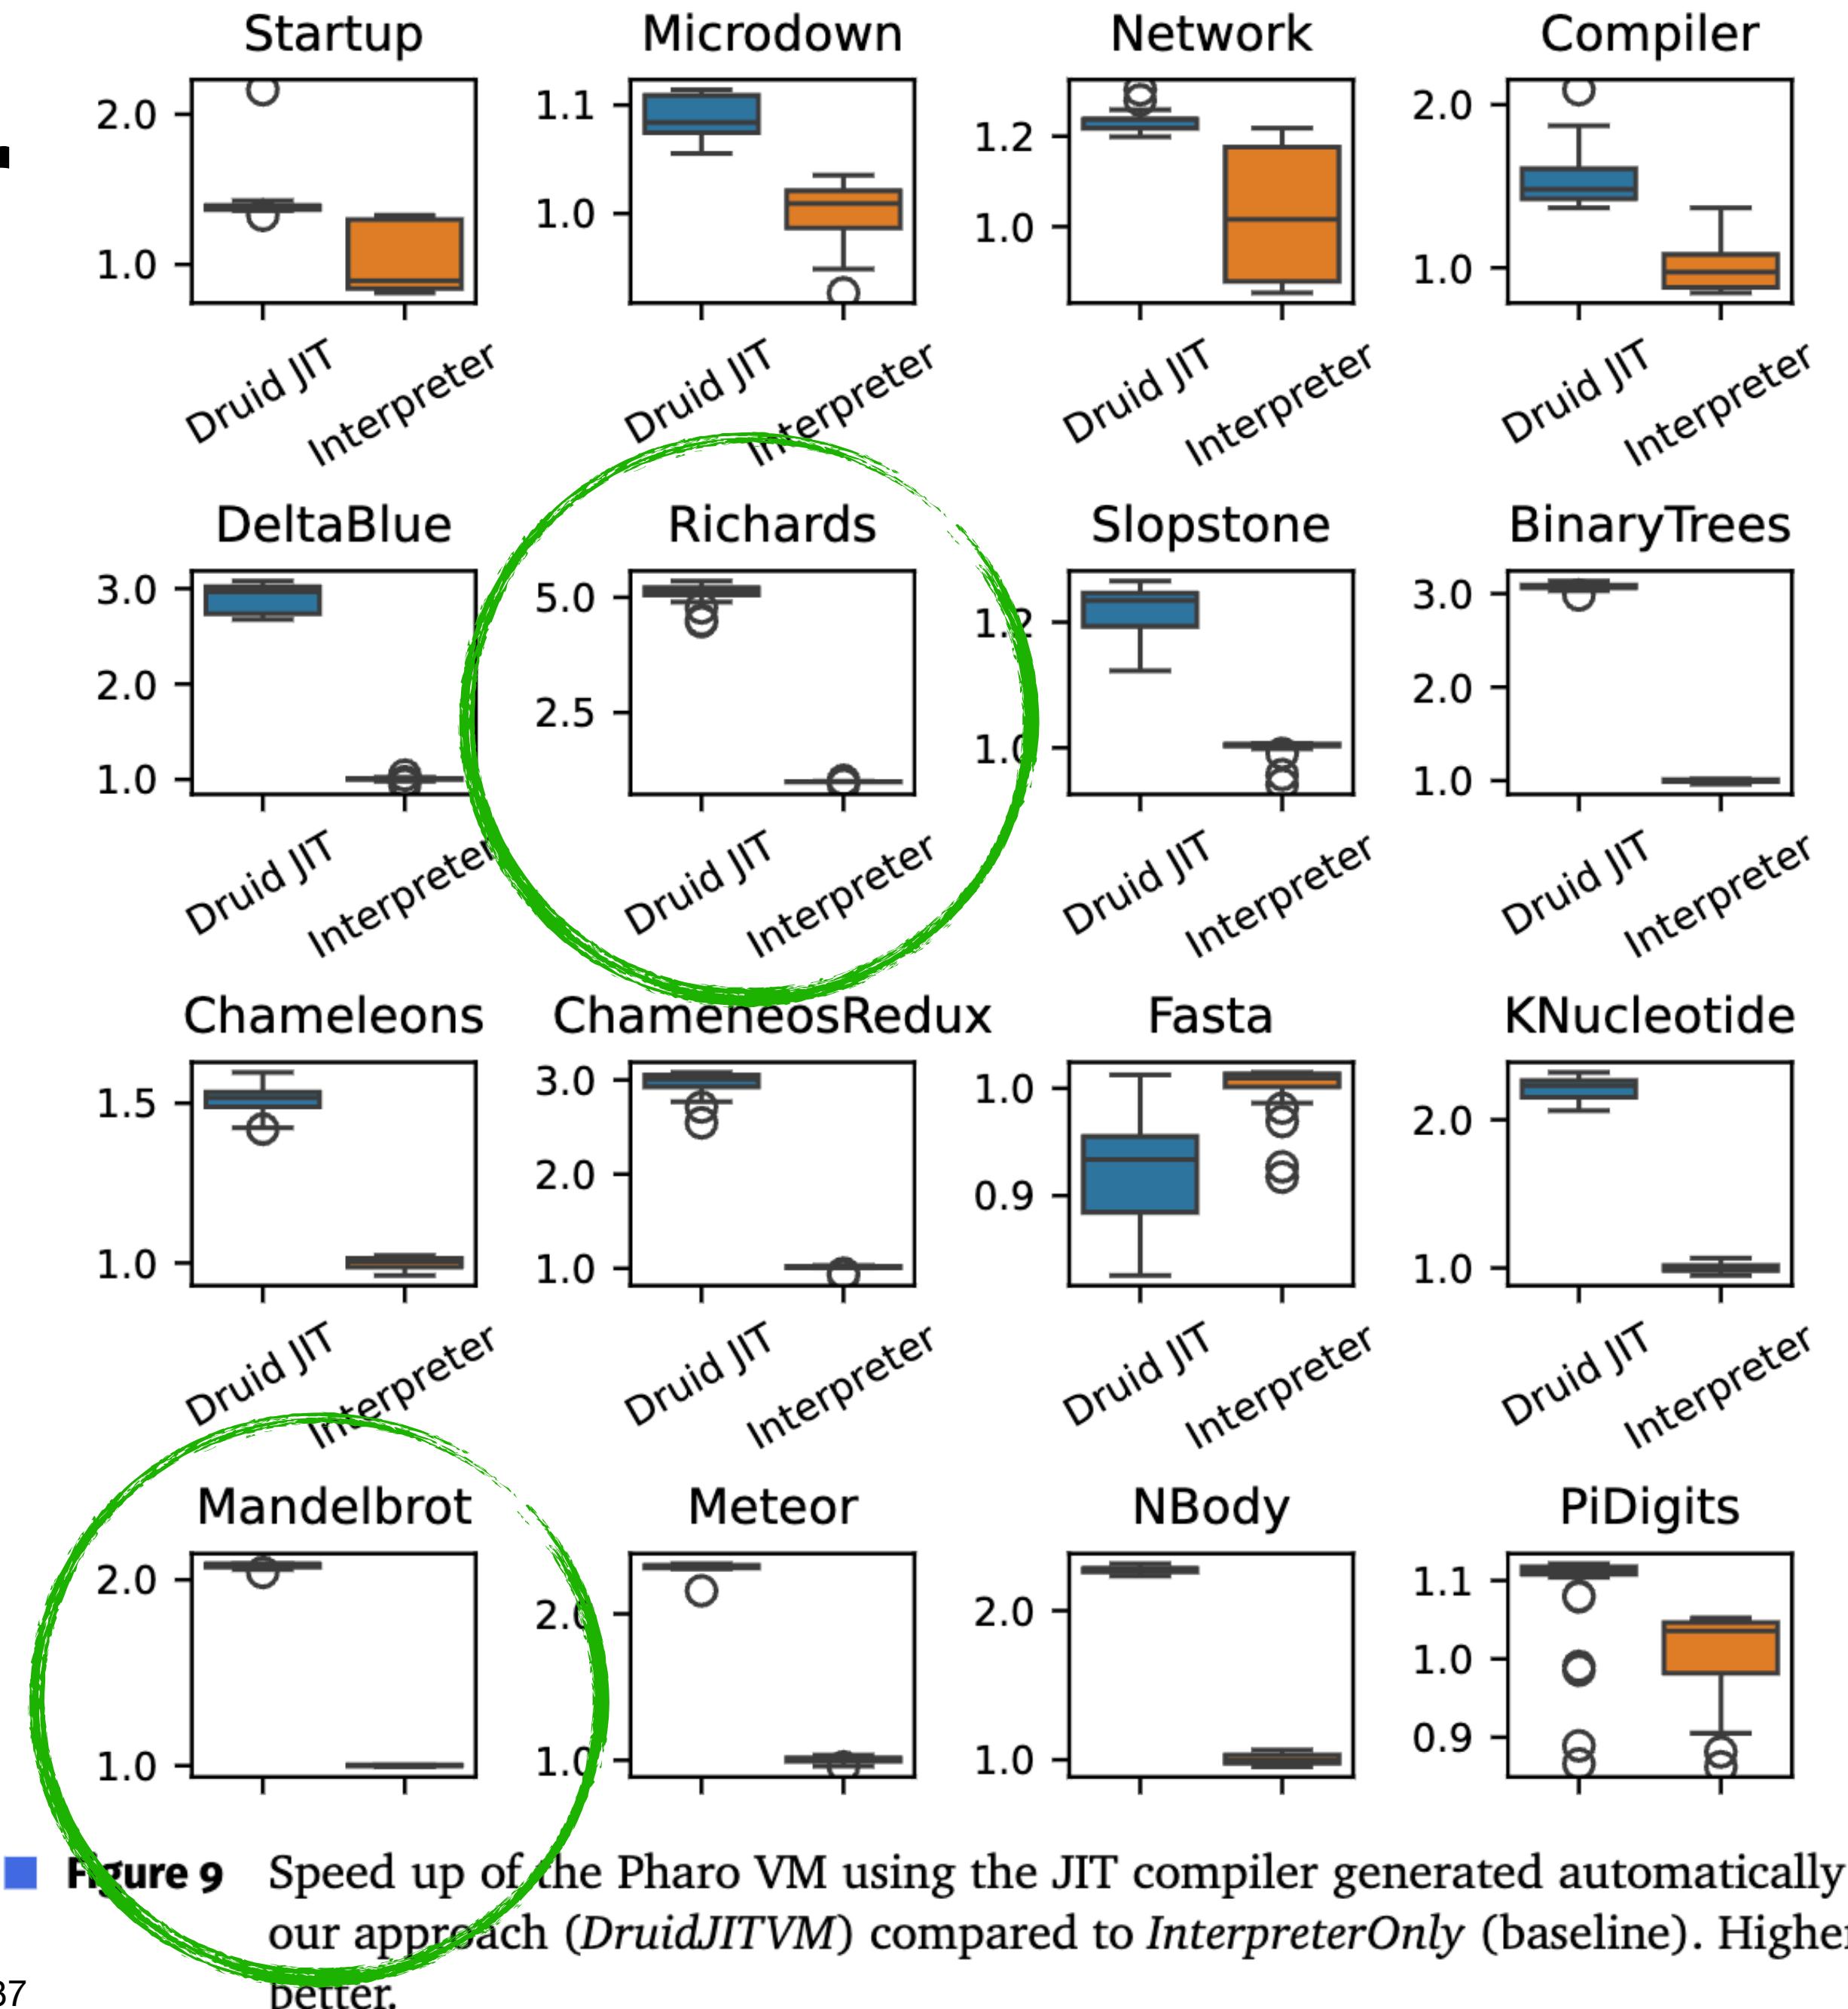


■ **Figure 9** Speed up of the Pharo VM using the JIT compiler generated automatically by our approach (*DruidJITVM*) compared to *InterpreterOnly* (baseline). Higher is better.

Druid JIT vs Interpreter

Faster than the Interpreter

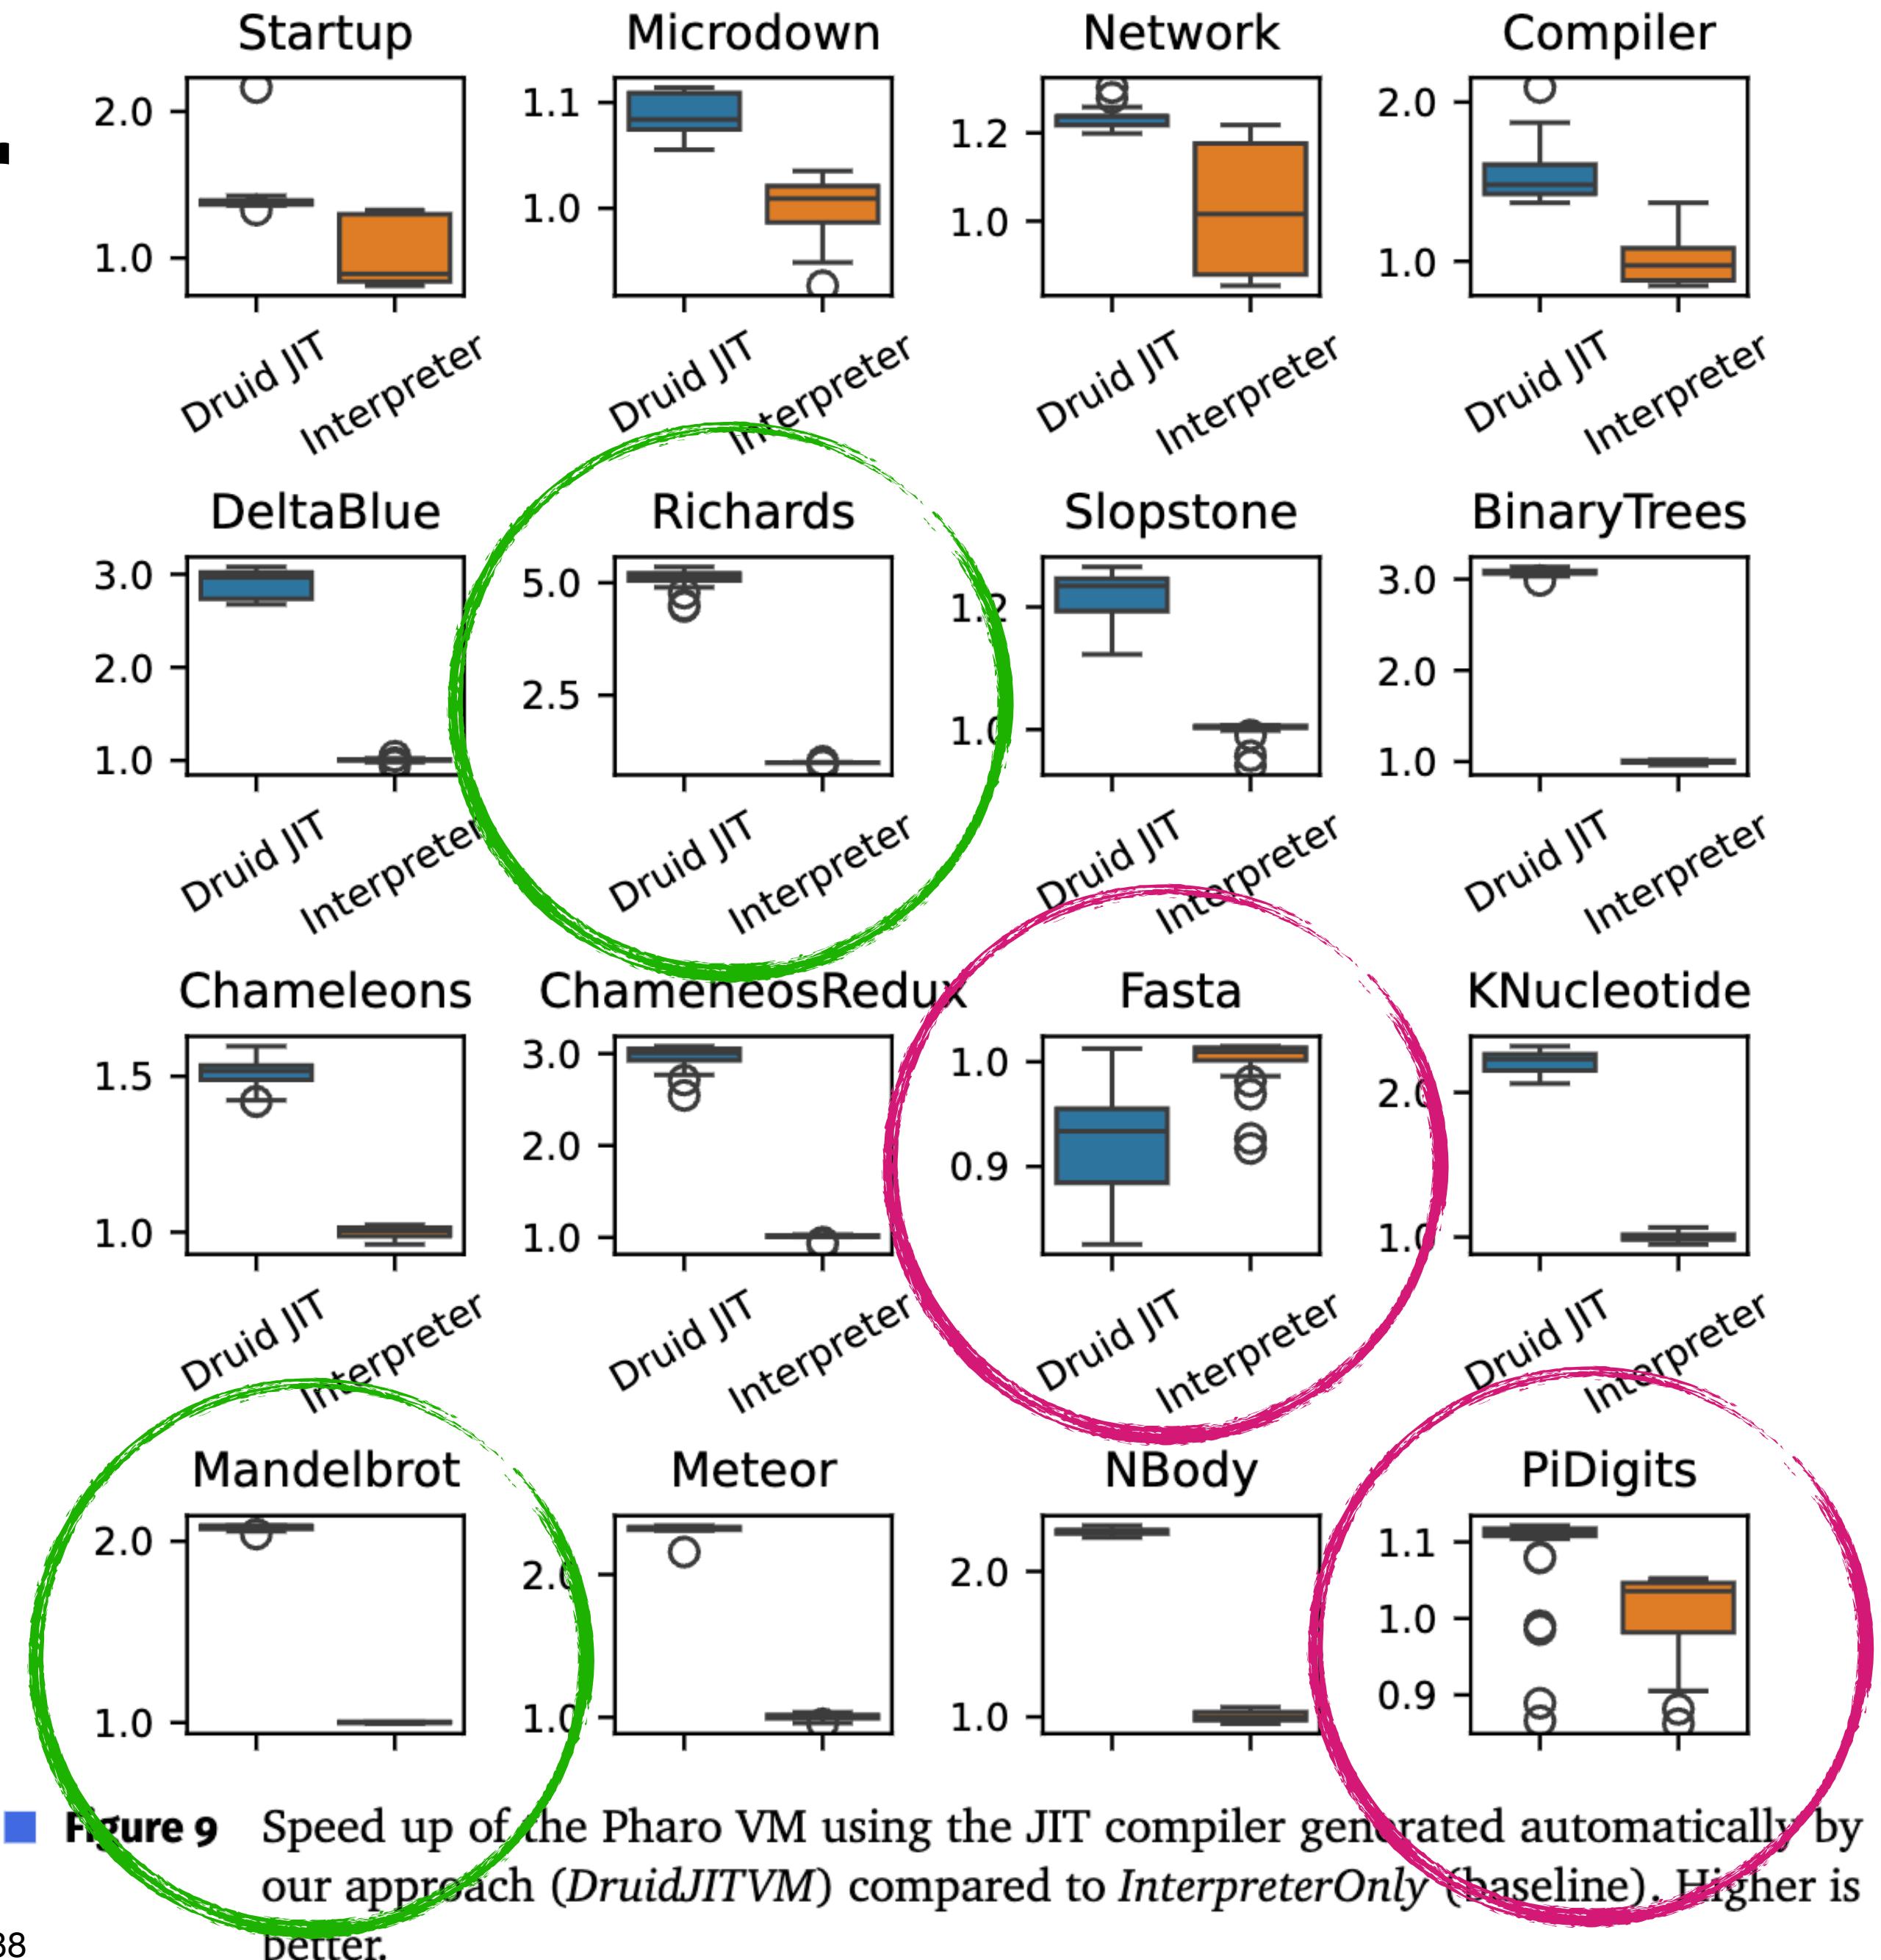
- 2x on average
- Up to 5x



Druid JIT vs Interpreter

Faster than the Interpreter

- 2x on average
- Up to 5x



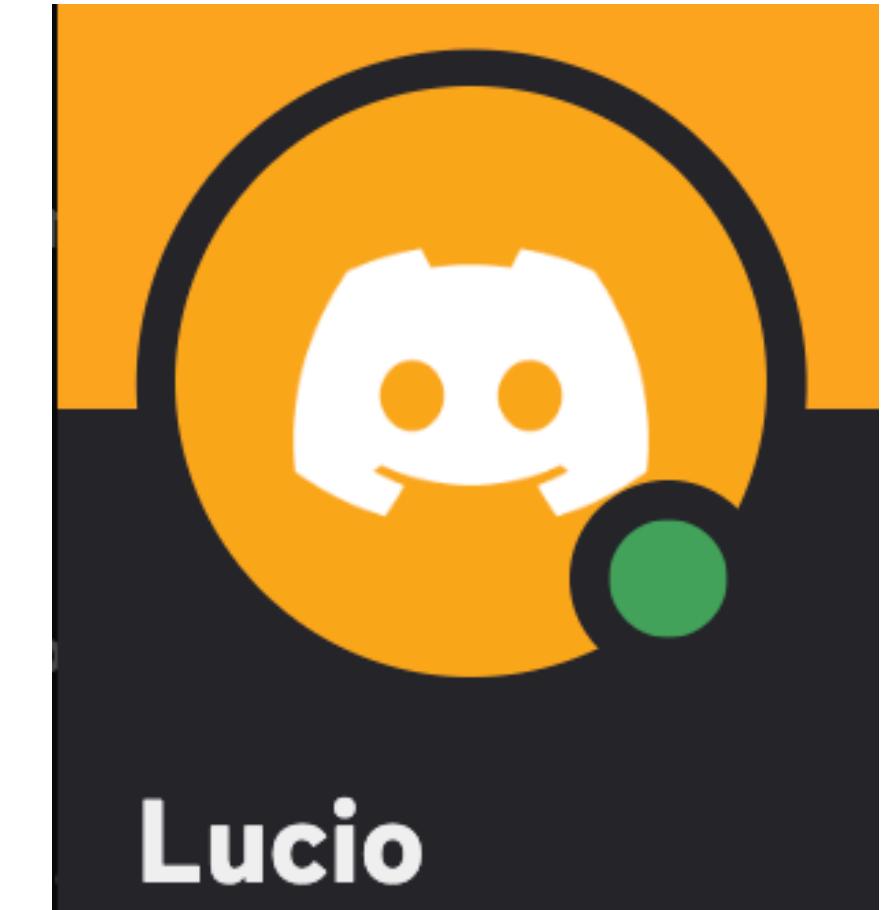
What else?



Extra!

We are also using Druid as an

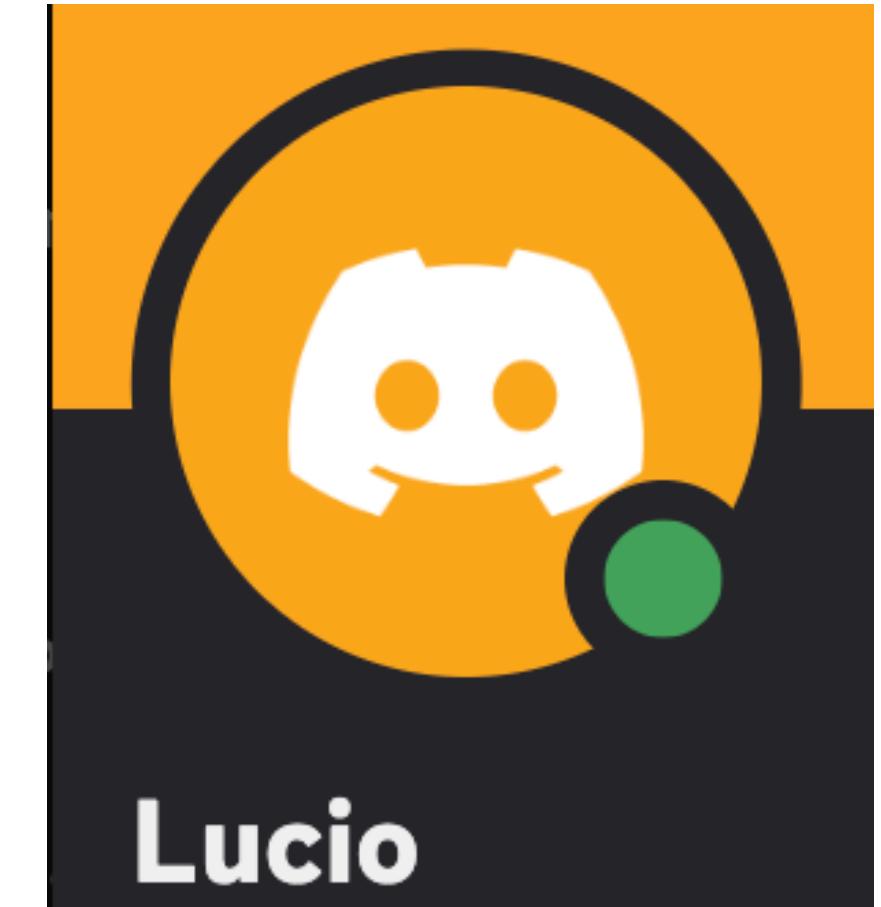
Optimizing Bytecode Compiler (WIP)



Method Source	Bytes	Bytecode	IR	AST	Header	Method Source	Bytes	Bytecode	IR	AST	Header
1 slowCount						1 slowCount					
2						2					
3 r						3 tmp2 tmp3 tmp4 tmp5 tmp6 tmp7					
4 r := 0.						4 tmp3 := 1.					
5 self do: [:each r := r + 1].						5 [
6 ^ r						6 tmp4 := tmp4.					
						7 tmp7 := self size.					
						8 tmp3 <= tmp7] whileTrue: [
						9 self at: tmp3.					
						10 tmp2 := tmp2 + 1.					
						11 tmp3 := tmp3 + 1].					
						12 ^ tmp2					

Optimizing Bytecode Compiler

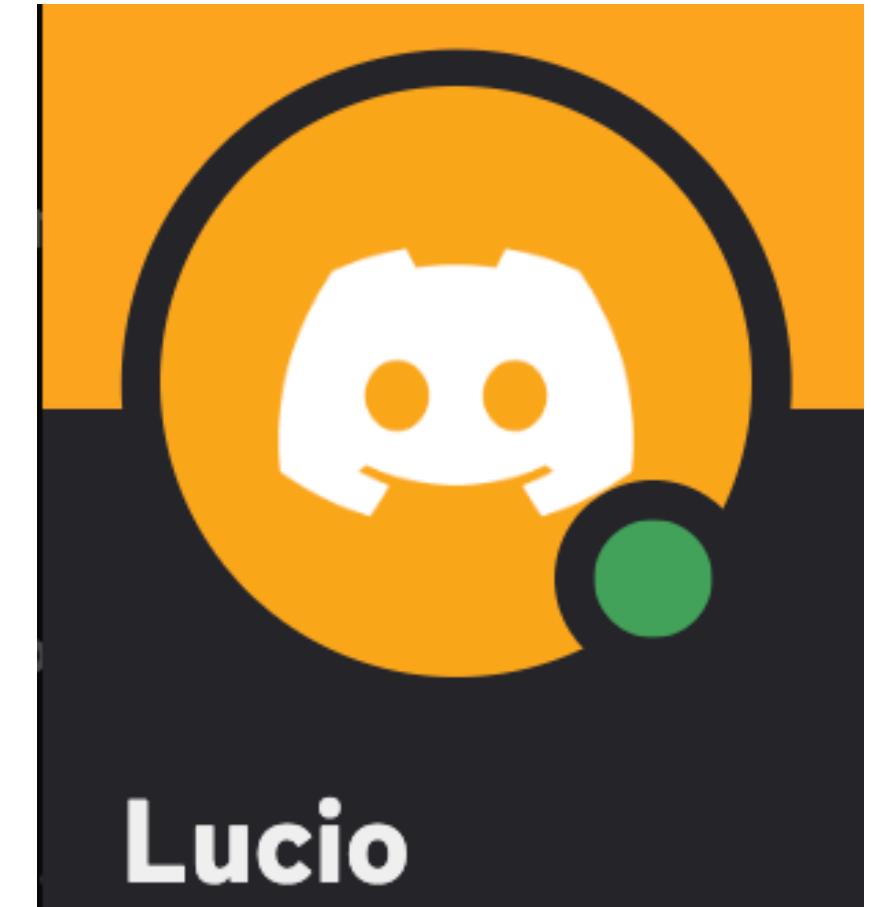
With Inlinings



```
Method Source Bytes Bytecode IR AST Header <> Method Source Bytes Bytecode IR AST <>
1 slowCount
2
3 | r |
4 r := 0.
5 self do: [ :each | r := r + 1 ].  
6 ^ r
1 slowCount
2
3 | tmp2 tmp3 tmp4 tmp5 tmp6 tmp7 |
4 tmp3 := 1.
5 [
6 tmp4 := tmp4.
7 tmp7 := self size.
8 tmp3 <= tmp7 ] whileTrue: [
9   self at: tmp3.
10  tmp2 := tmp2 + 1.
11  tmp3 := tmp3 + 1 ].  
12 ^ tmp2
```

Optimizing Bytecode Compiler

With Inlinings



Why is this important?

<https://github.com/ESUG/esug.github.io/blob/source/2024-Conference/slides/show-us-your-project/show-us-your-project-esug-2024-dropal.pdf>

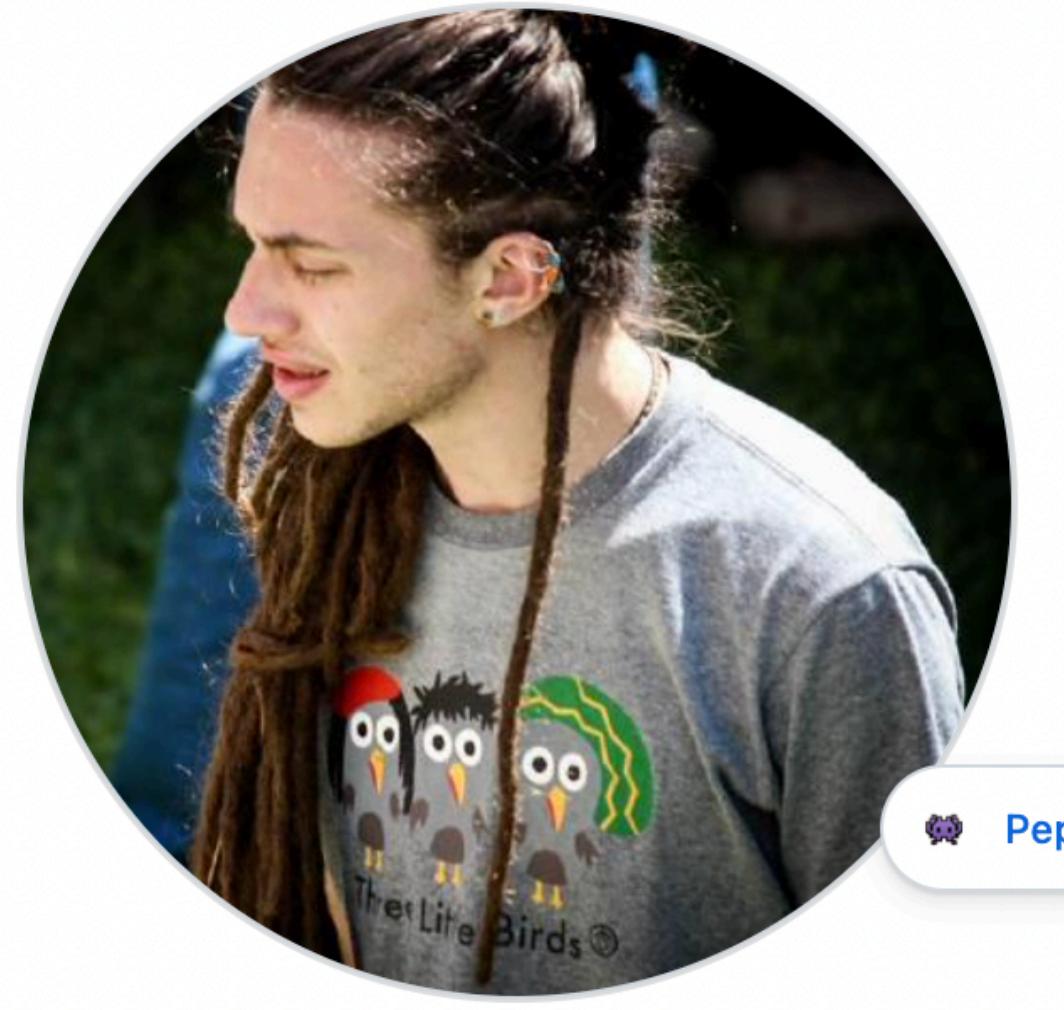
```
4      r := 0.
5      self do: [ :each | r := r + 1 ].
6      ^ r
```

The screenshot shows a debugger interface with tabs for Method, Source, Bytes, Bytecode, IR, and AST. The 'Source' tab is selected. A callout arrow points from the text "Why is this important?" to the 'Source' tab. The code in the source tab is:

```
1 slowCount
2
3 | tmp2 tmp3 tmp4 tmp5 tmp6 tmp7 |
4 tmp3 := 1.
5 [
6 tmp4 := tmp4.
7 tmp7 := self size.
8 tmp3 <= tmp7 ] whileTrue: [
9   self at: tmp3.
10  tmp2 := tmp2 + 1.
11  tmp3 := tmp3 + 1 ].
12 ^ tmp2
```

Conclusions

- AoT meta-compilation of Baseline JITs is feasible
- It achieves decent performance
 - Traditional compiler optimizations
 - Intrinsics + Staging
- Extra!
 - Optimizing Bytecode Compiler



Pinne



Woll



W



A m



Pepita ❤️ T

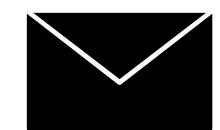


Cou



Nahuel Palumbo

PalumboN



nahuel.palumbo@inria.fr



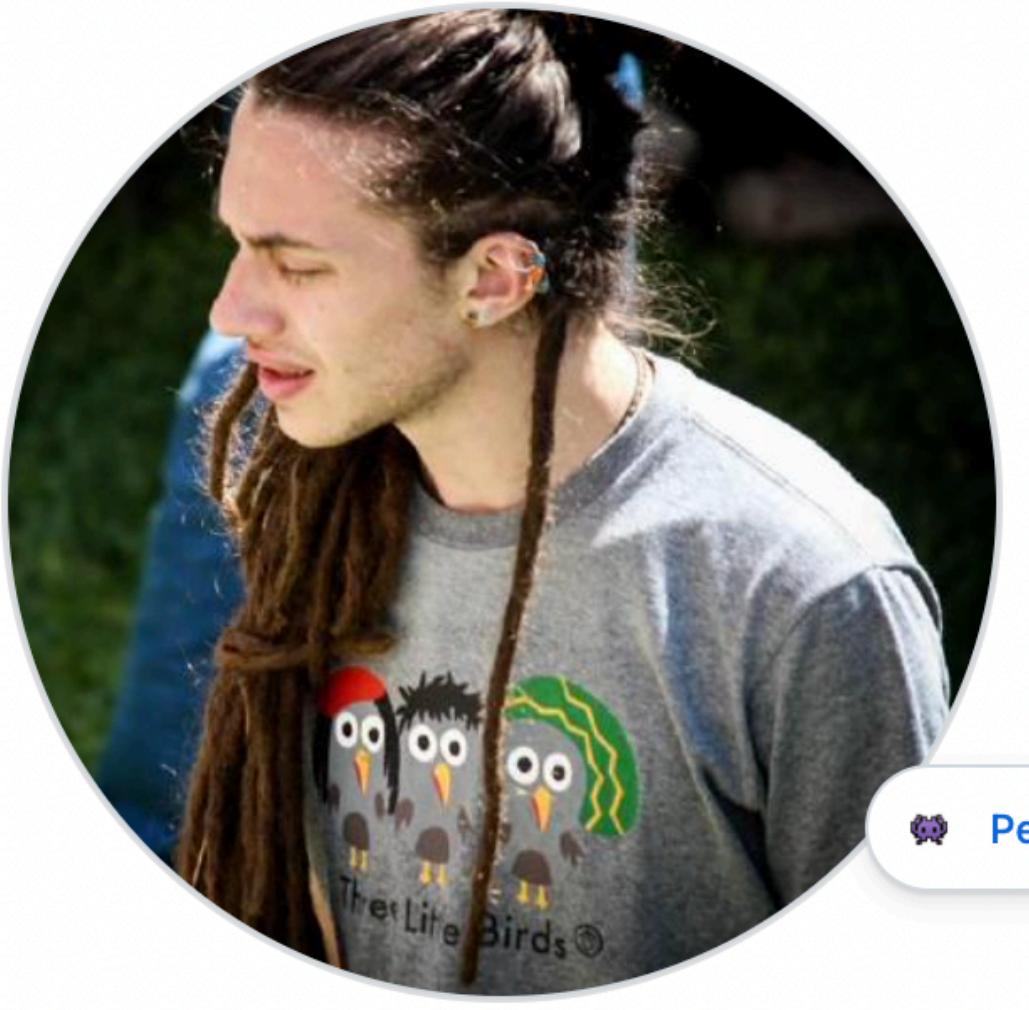
/Alamvic/druid

/pharo-project/pharo-vm

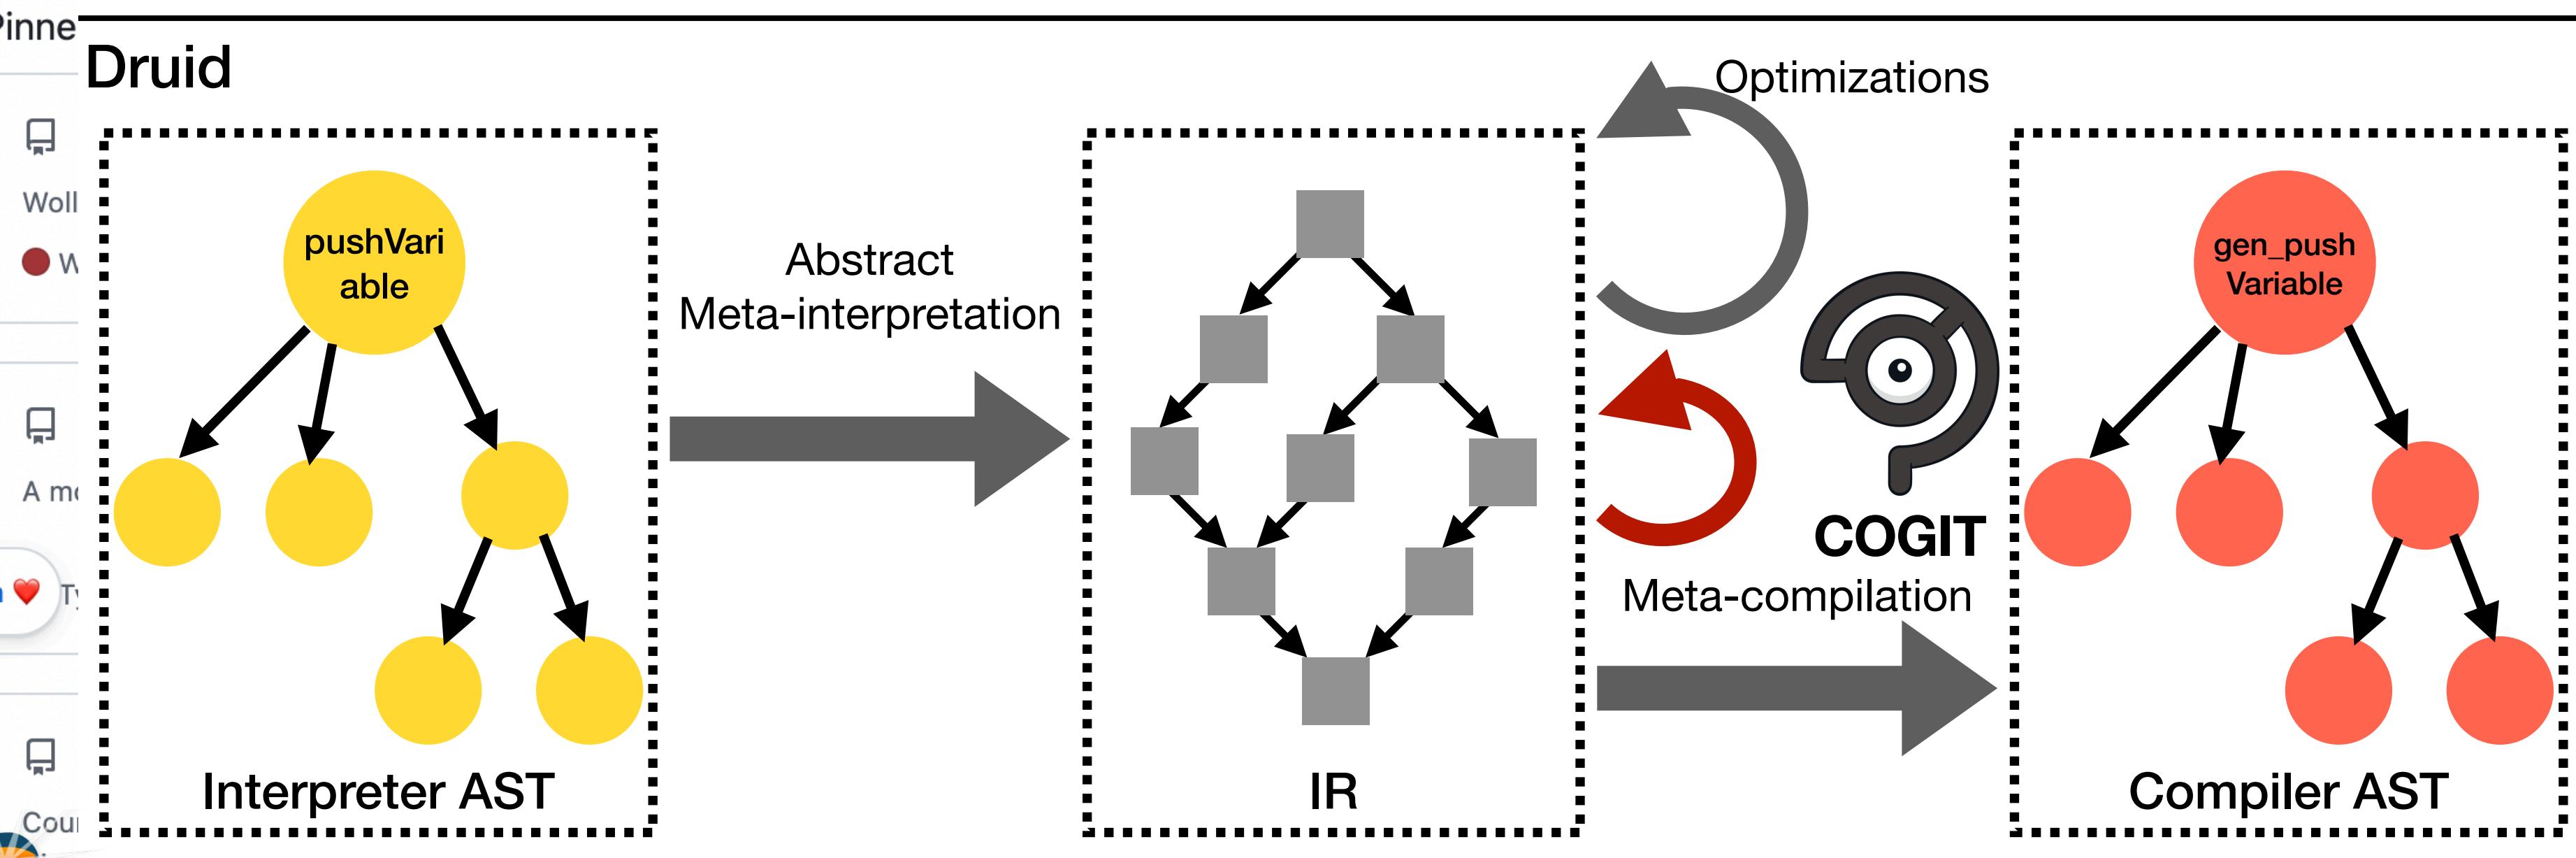


*Let's keep
in contact!*

<https://discord.gg/QewZMza>



Nahuel Palumbo
PalumboN



nahuel.palumbo@inria.fr

/Alamvic/druid

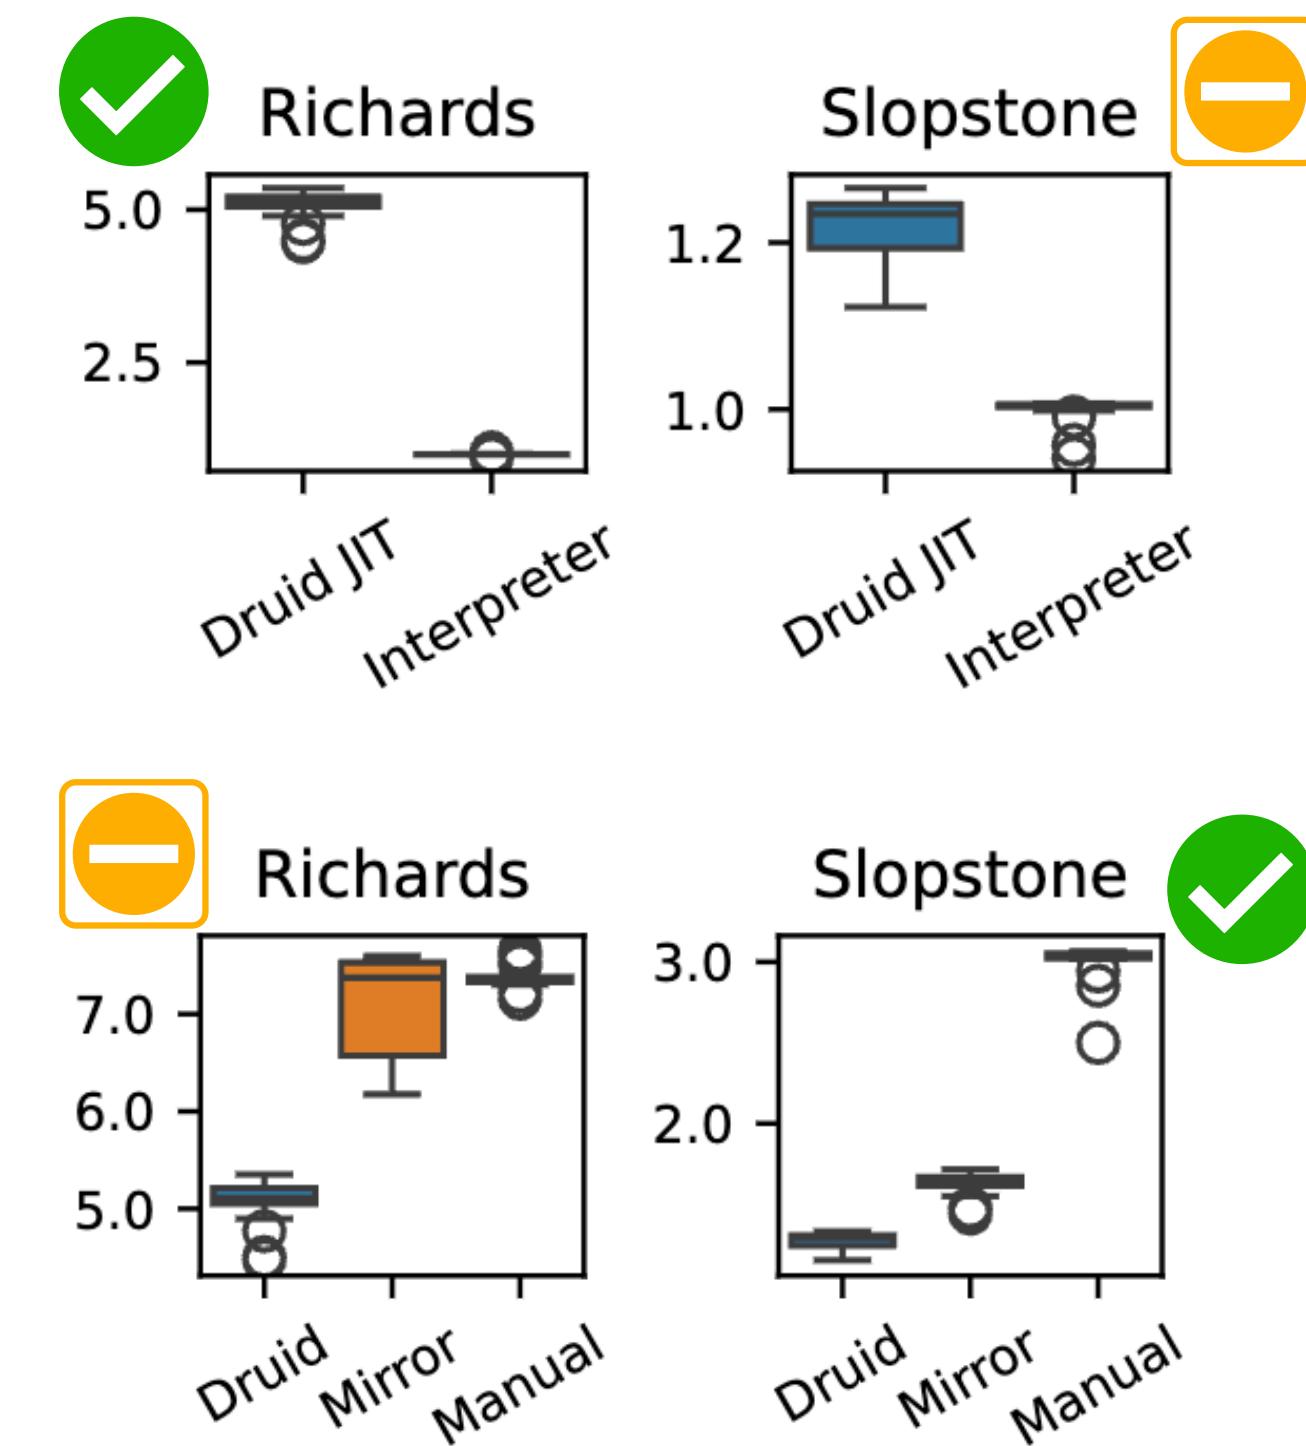
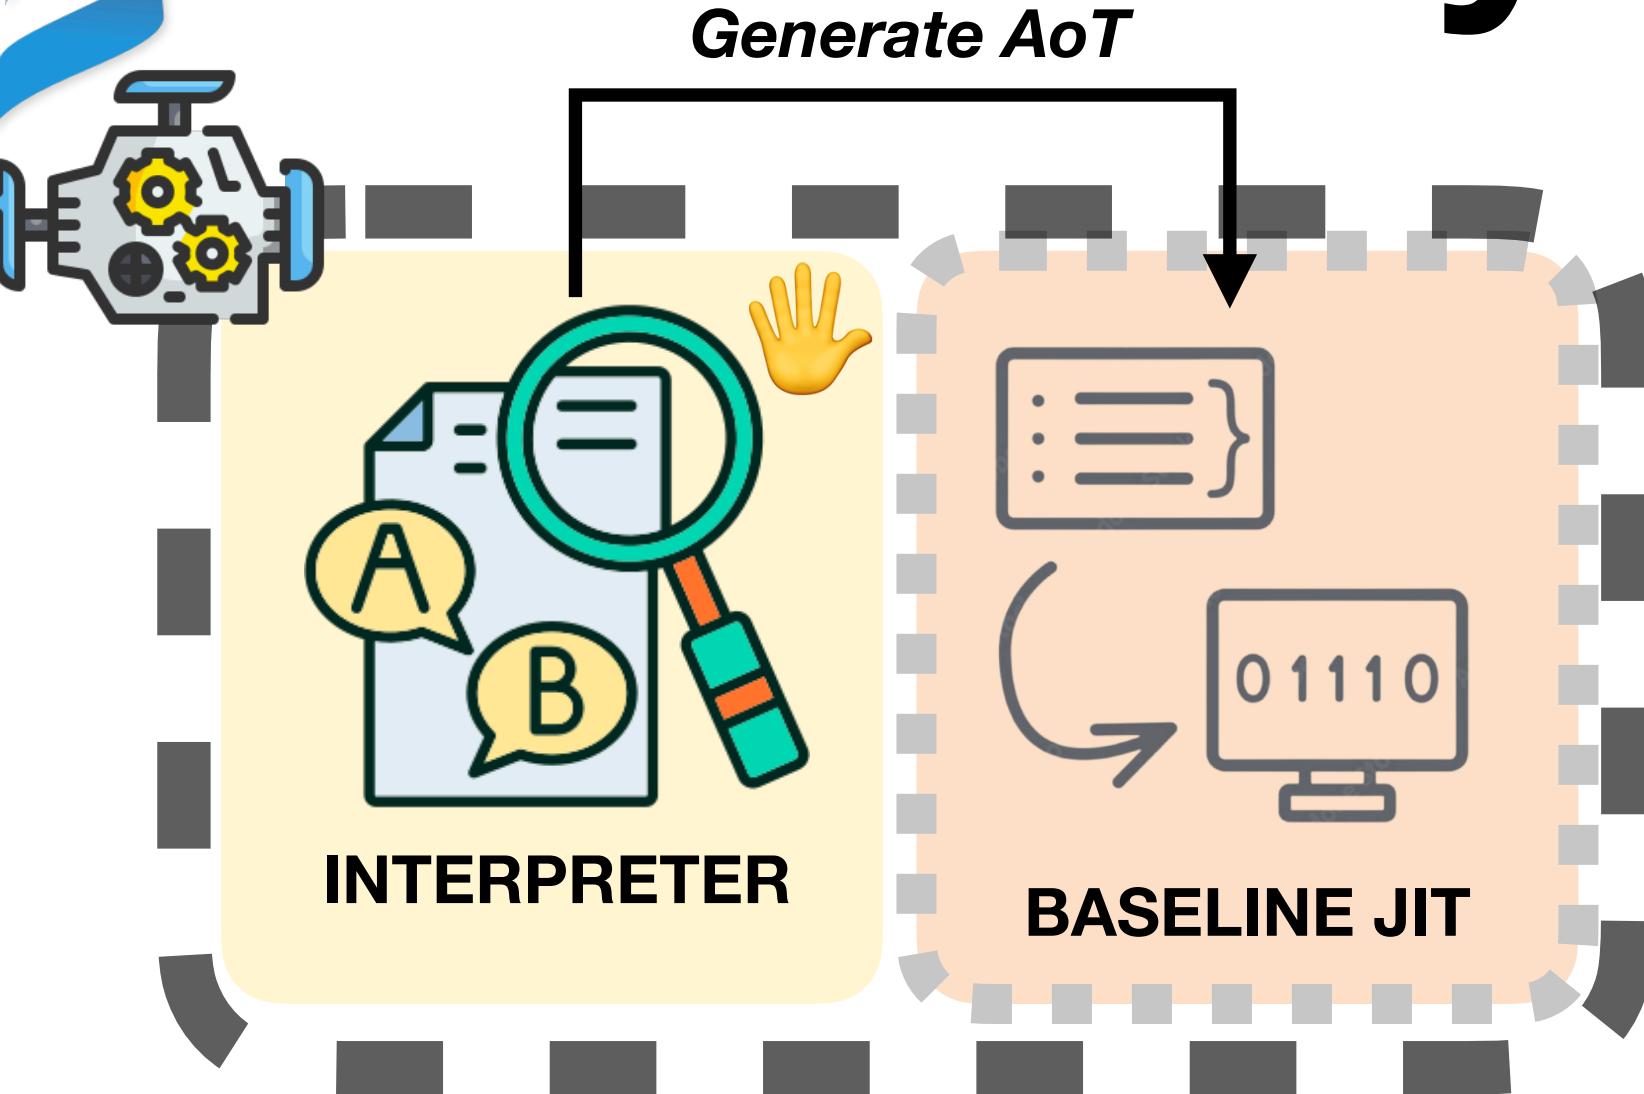
/pharo-project/pharo-vm

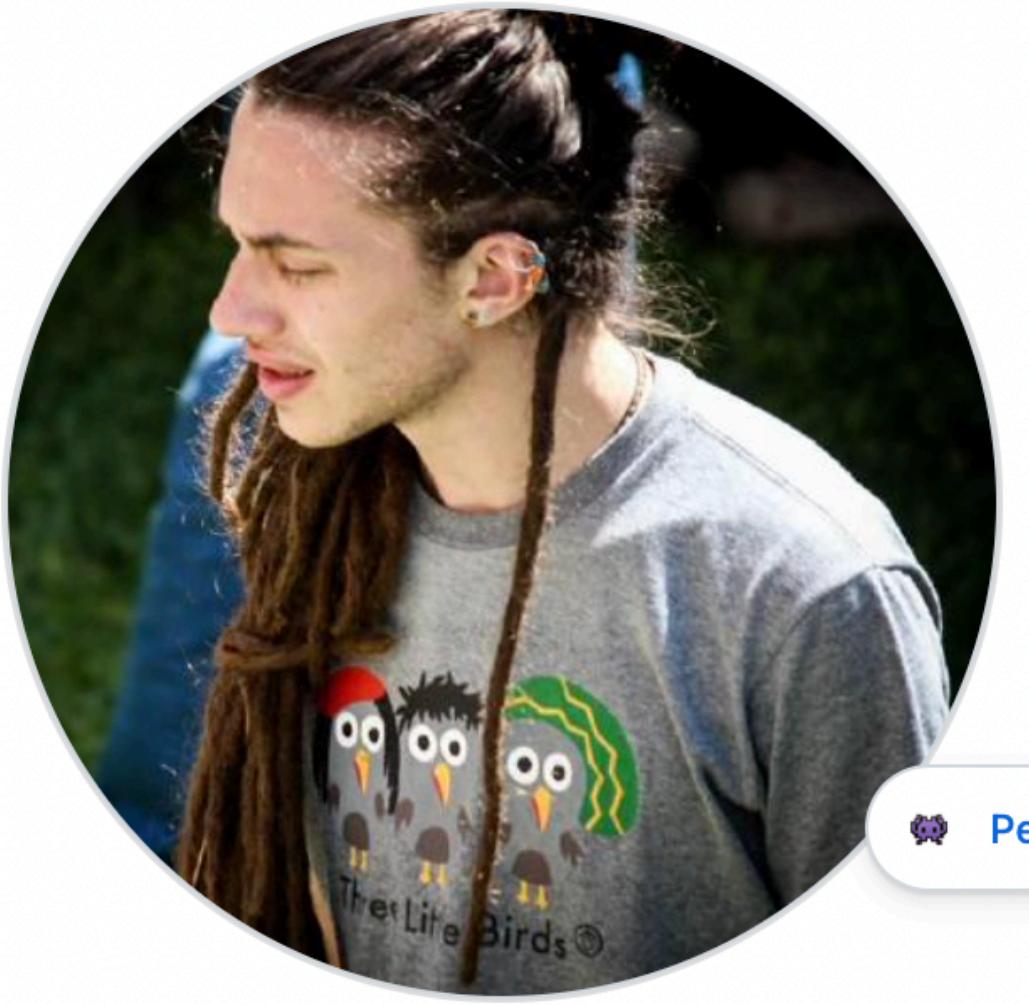


*Let's keep
in contact!*

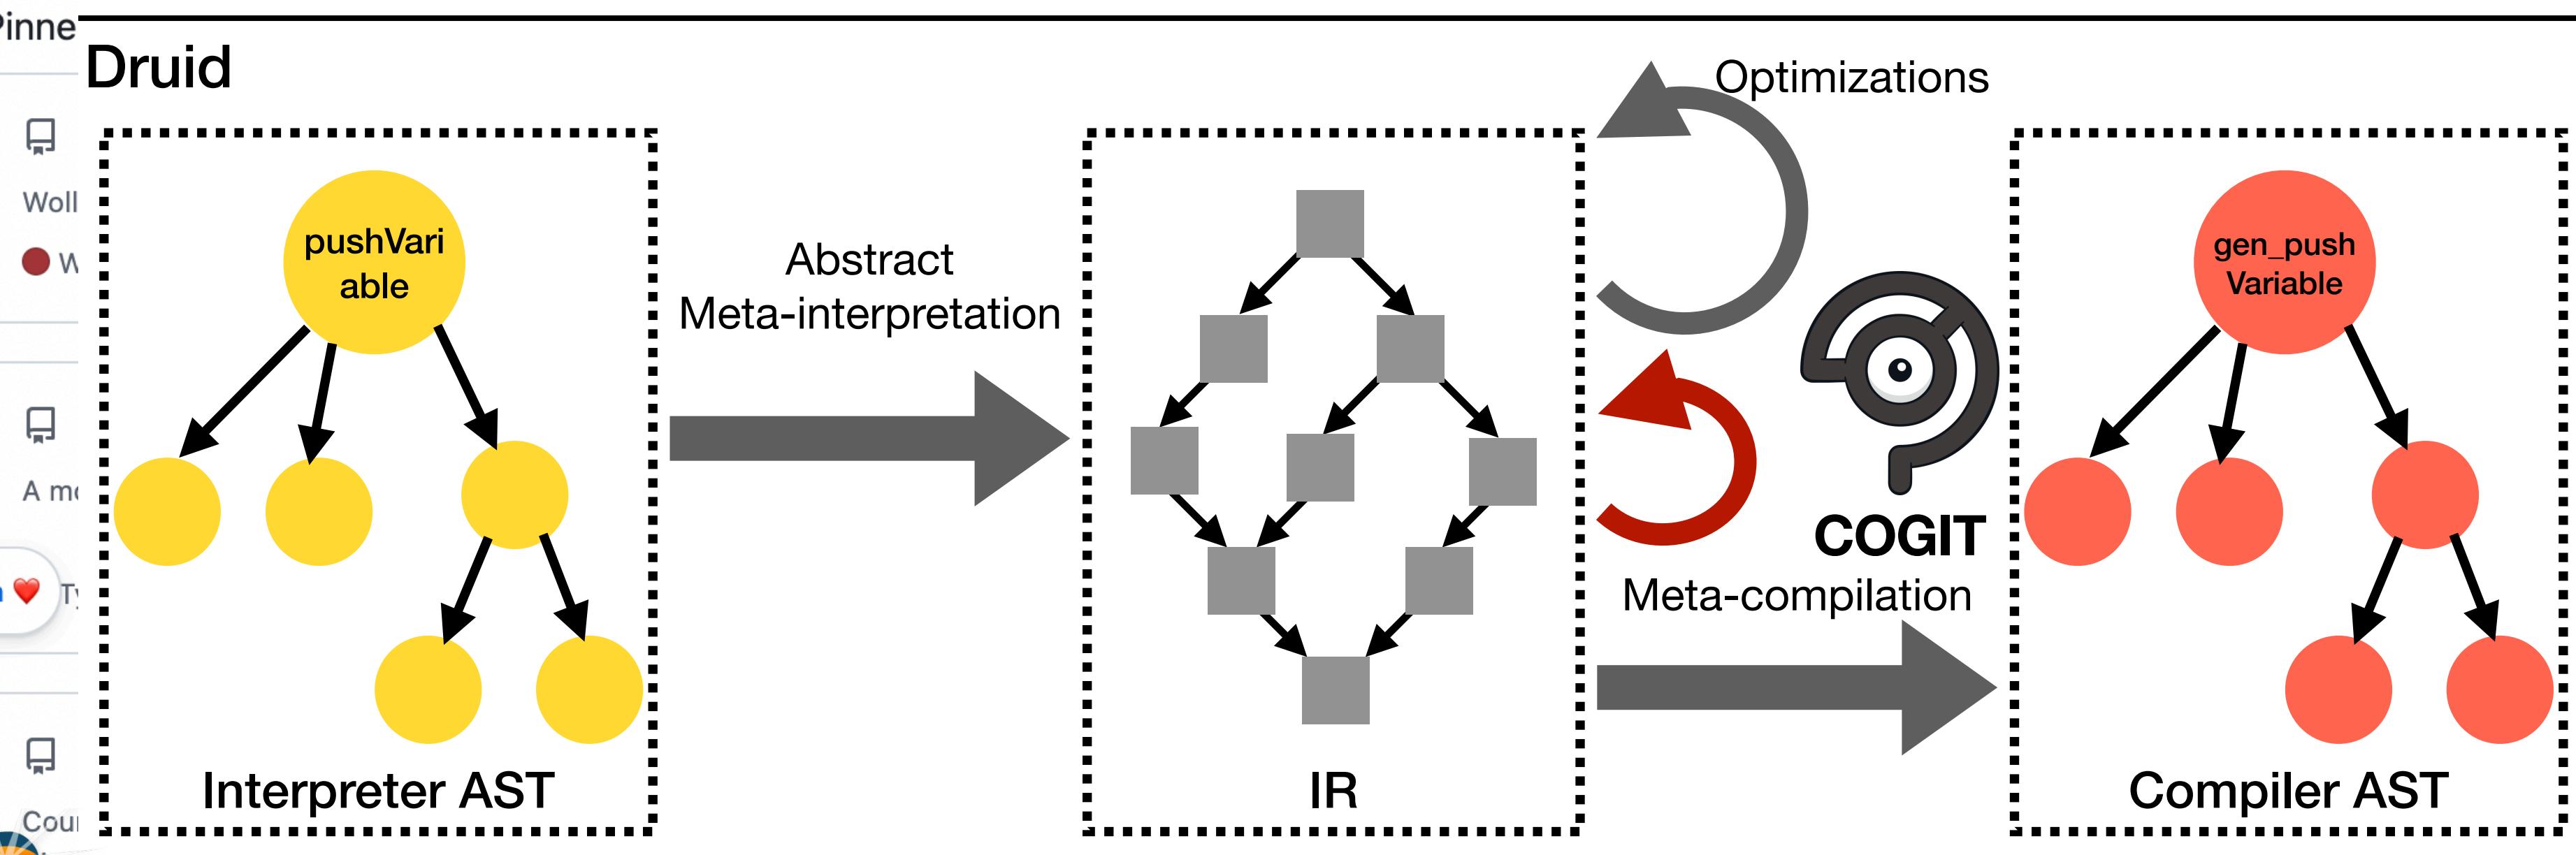
<https://discord.gg/QewZMza>

Thank you!





Nahuel Palumbo
PalumboN



nahuel.palumbo@inria.fr

[/Alamvic/druid](https://github.com/Alamvic/druid)

[/pharo-project/pharo-vm](https://github.com/pharo-project/pharo-vm)



*Let's keep
in contact!*

<https://discord.gg/QewZMza>

