

Ideas for improving meta-tracing warmup



Laurence
Tratt



Lukas
Diekmann



Edd Barrett

KING'S
College
LONDON

Software Development Team
2021-06-09

Virtual Machine Warmup Blows Hot and Cold*

EDD BARRETT, King's College London, UK

CARL FRIEDRICH BOLZ-TEREICK, King's College London, UK

REBECCA KILLIK, Lancaster University, UK

SARAH MOUNT, King's College London, UK

LAURENCE TRATT, King's College London, UK

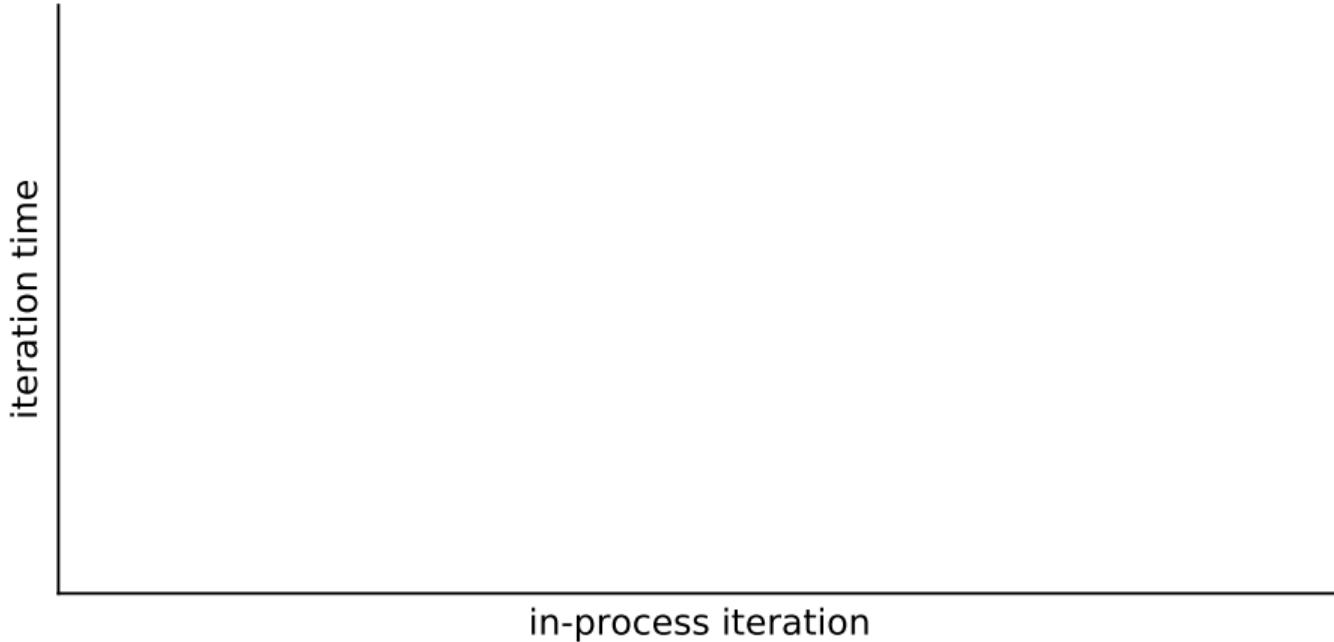
Virtual Machines (VMs) with Just-In-Time (JIT) compilers are traditionally thought to execute programs in two phases: the initial warmup phase determines which parts of a program would most benefit from dynamic compilation, before JIT compiling those parts into machine code; subsequently the program is said to be at a steady state of peak performance. Measurement methodologies almost always discard data collected during the warmup phase such that reported measurements focus entirely on peak performance. We introduce a fully automated statistical approach, based on changepoint analysis, which allows us to determine if a program has reached a steady state and, if so, whether that represents peak performance or not. Using this, we show that even when run in the most controlled of circumstances, small, deterministic, widely studied microbenchmarks often fail to reach a steady state of peak performance on a variety of common VMs. Repeating our experiment on 3 different machines, we found that at most 43.5% of $\langle \text{VM}, \text{benchmark} \rangle$ pairs consistently reach a steady state of peak performance.

CCS Concepts: • Software and its engineering → Software performance; Just-in-time compilers; Interpreters;

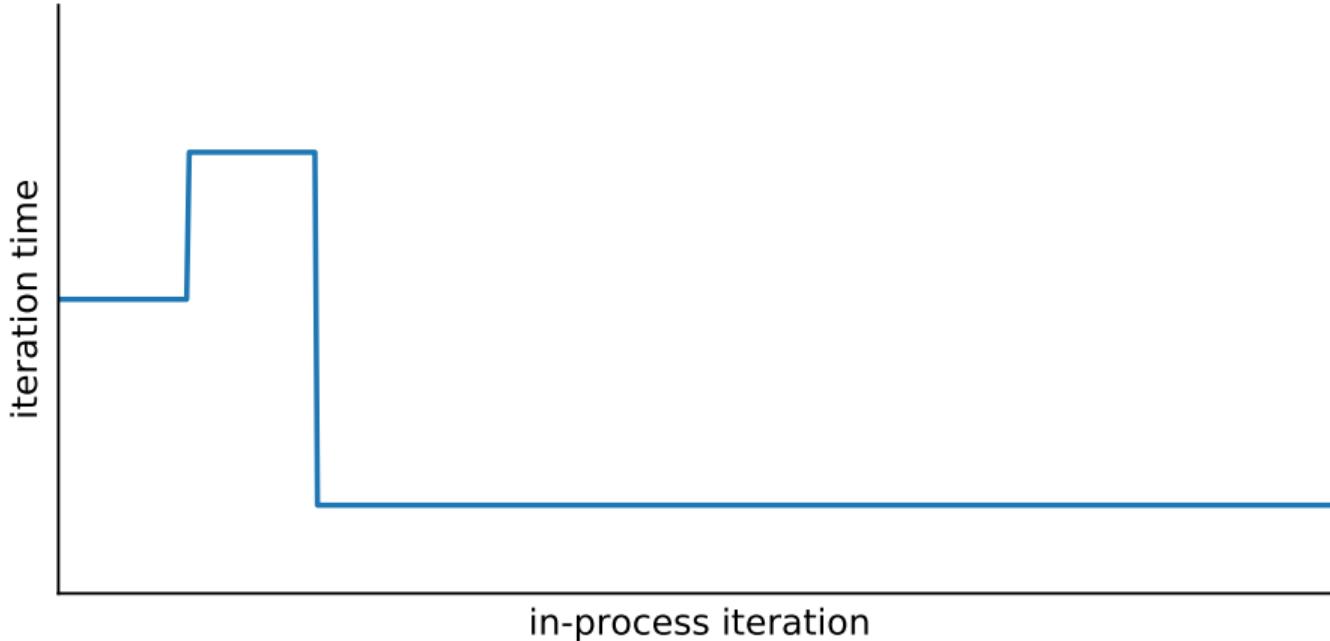
Additional Key Words and Phrases: Virtual machine, JIT, benchmarking, performance

ACM Reference Format:

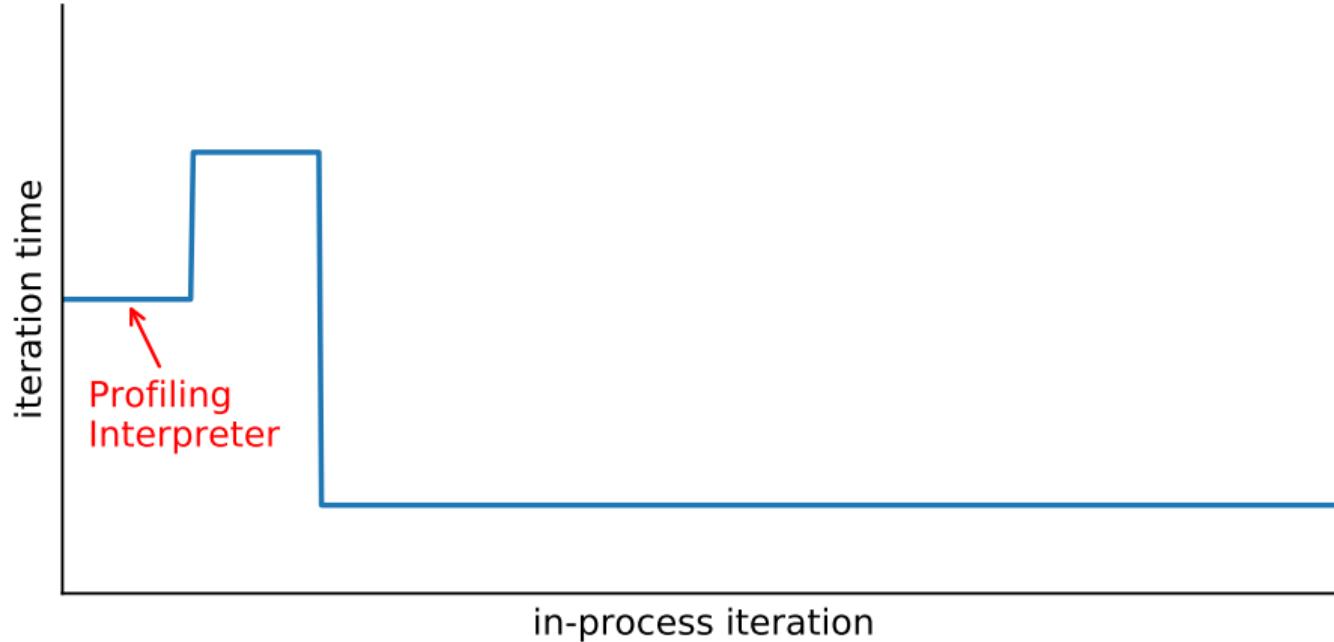
JIT Warmup



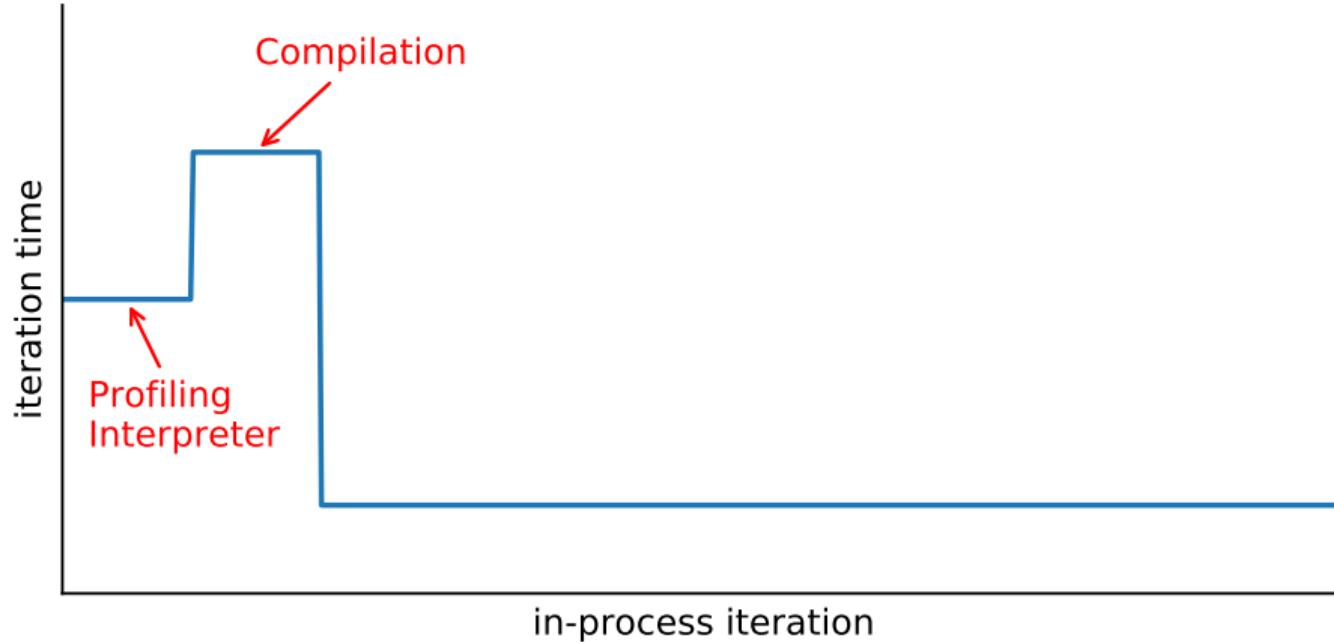
JIT Warmup



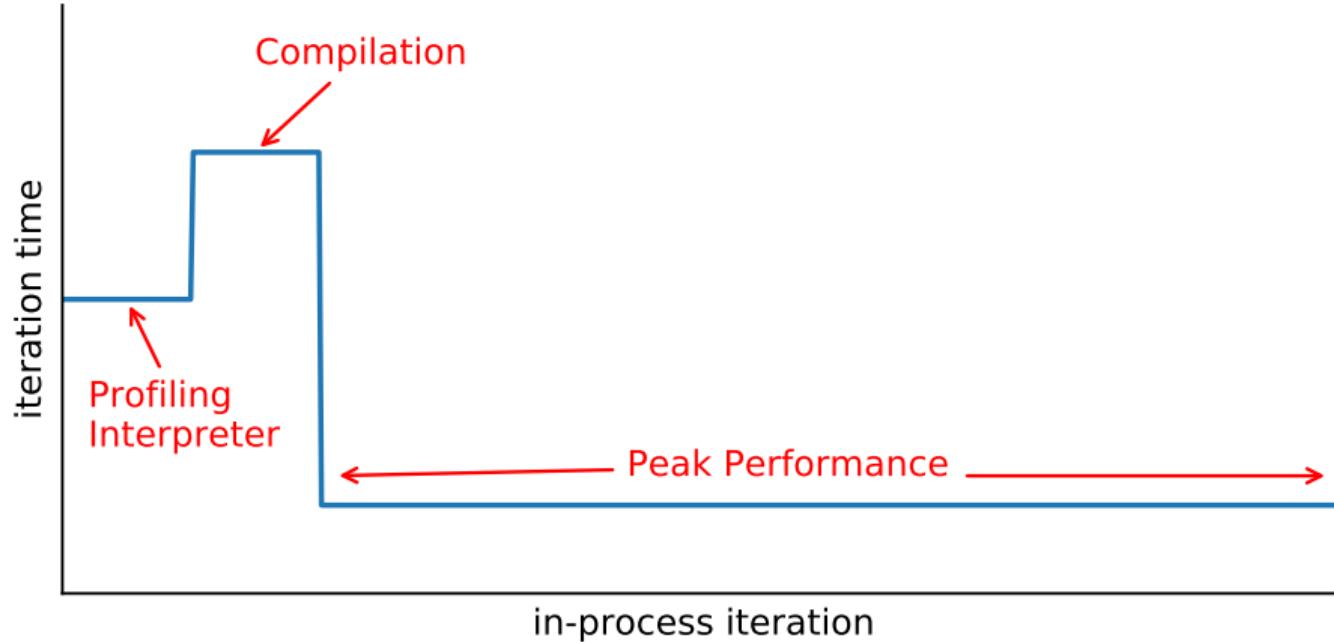
JIT Warmup



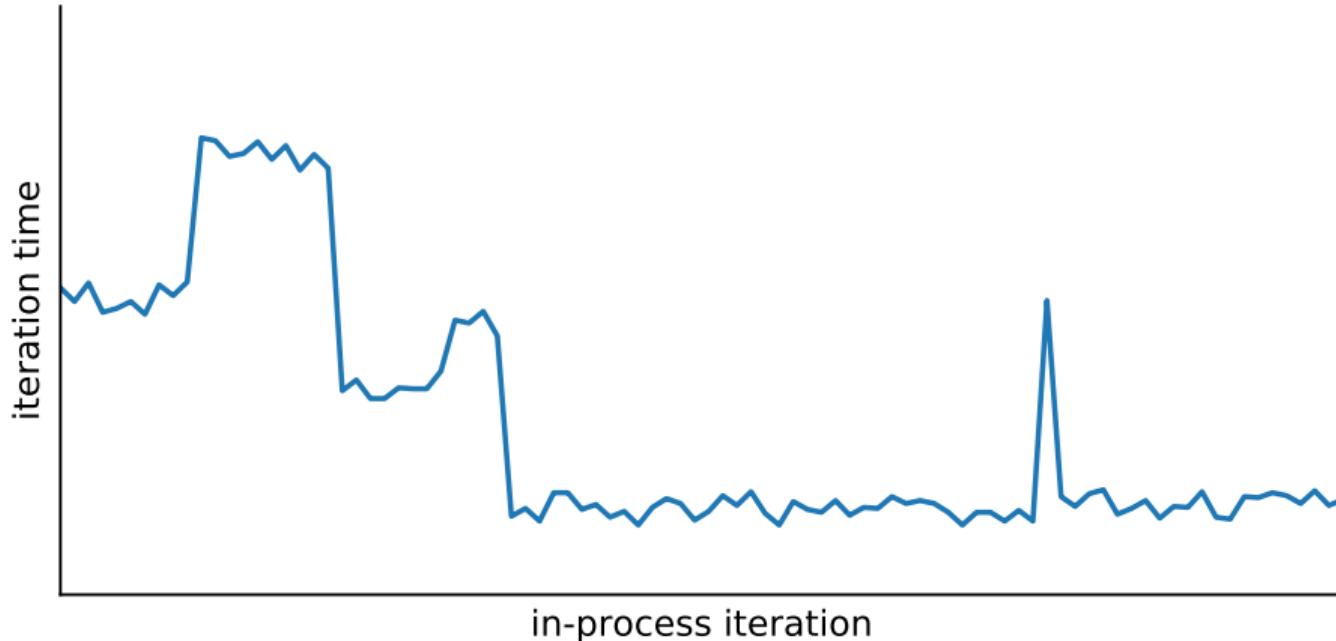
JIT Warmup



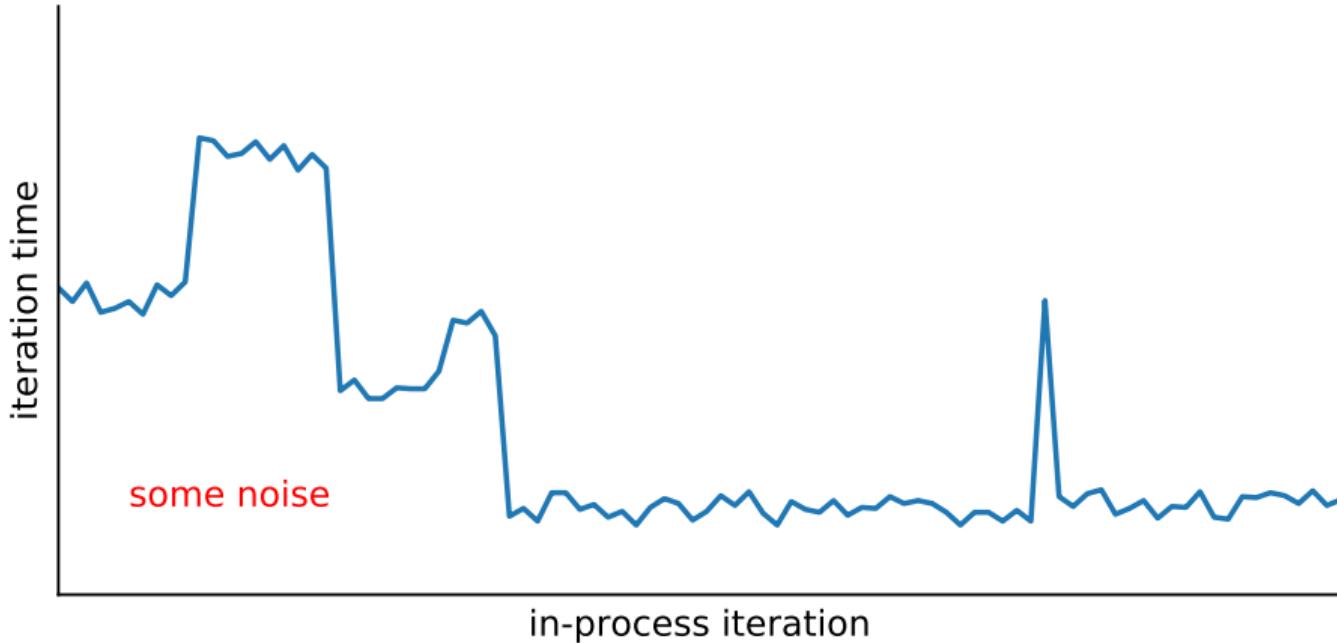
JIT Warmup



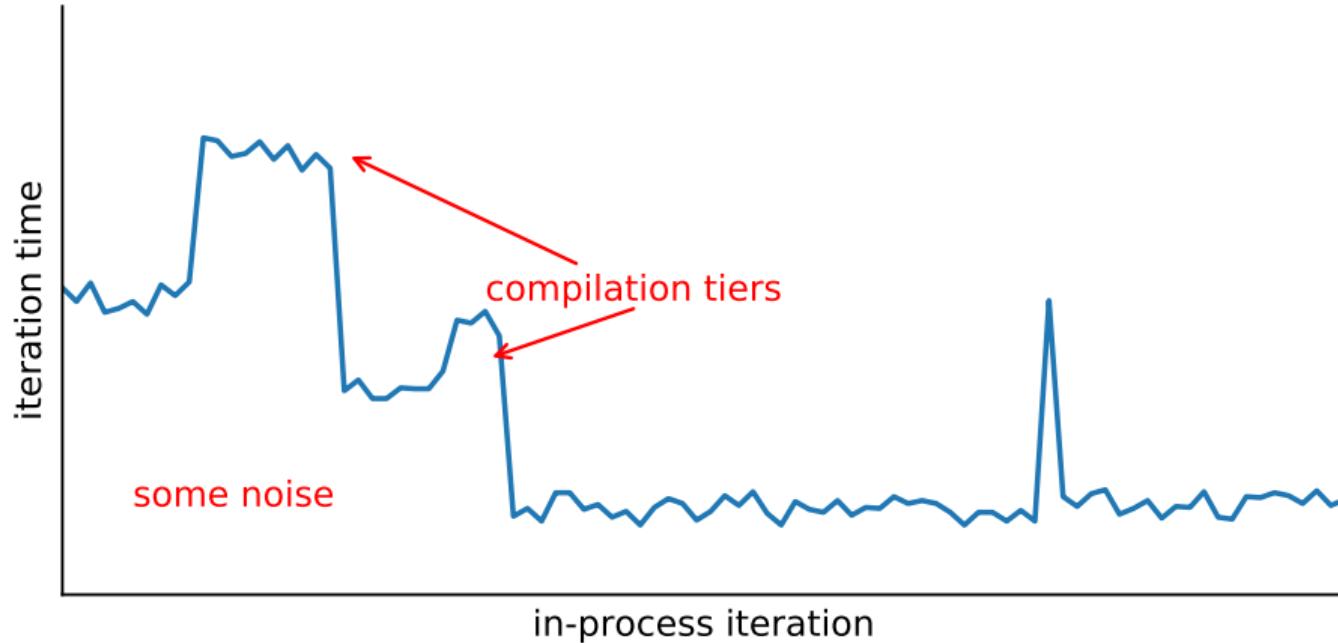
JIT Warmup



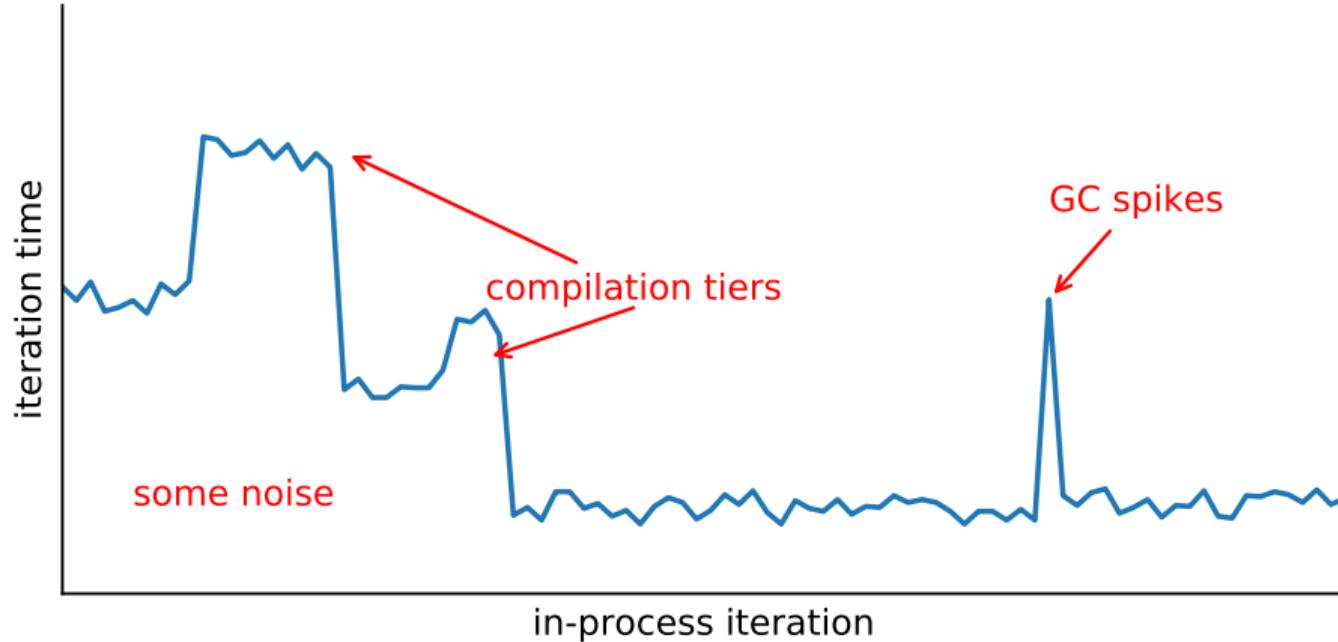
JIT Warmup



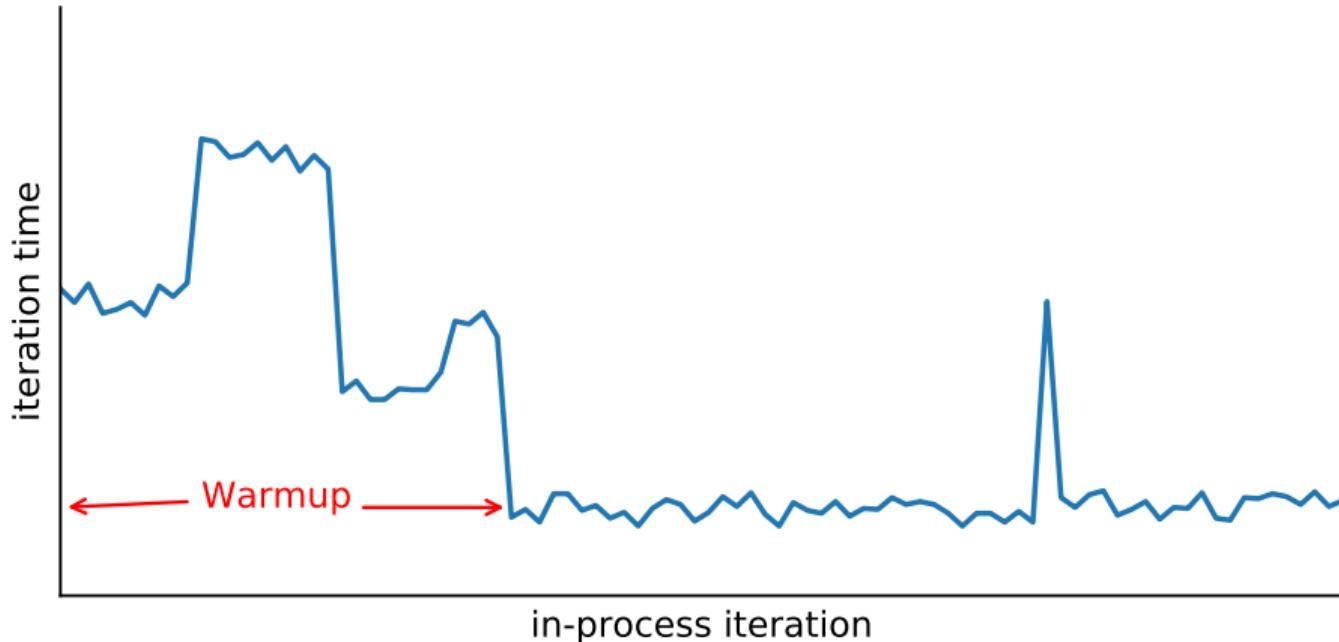
JIT Warmup



JIT Warmup



JIT Warmup



JIT Warmup

JIT Warmup

- ▶ Users dislike poor warmup.

- ▶ Users dislike poor warmup.
- ▶ VM authors dislike poor warmup.

The Warmup Experiment

We should measure the warmup of modern language implementations

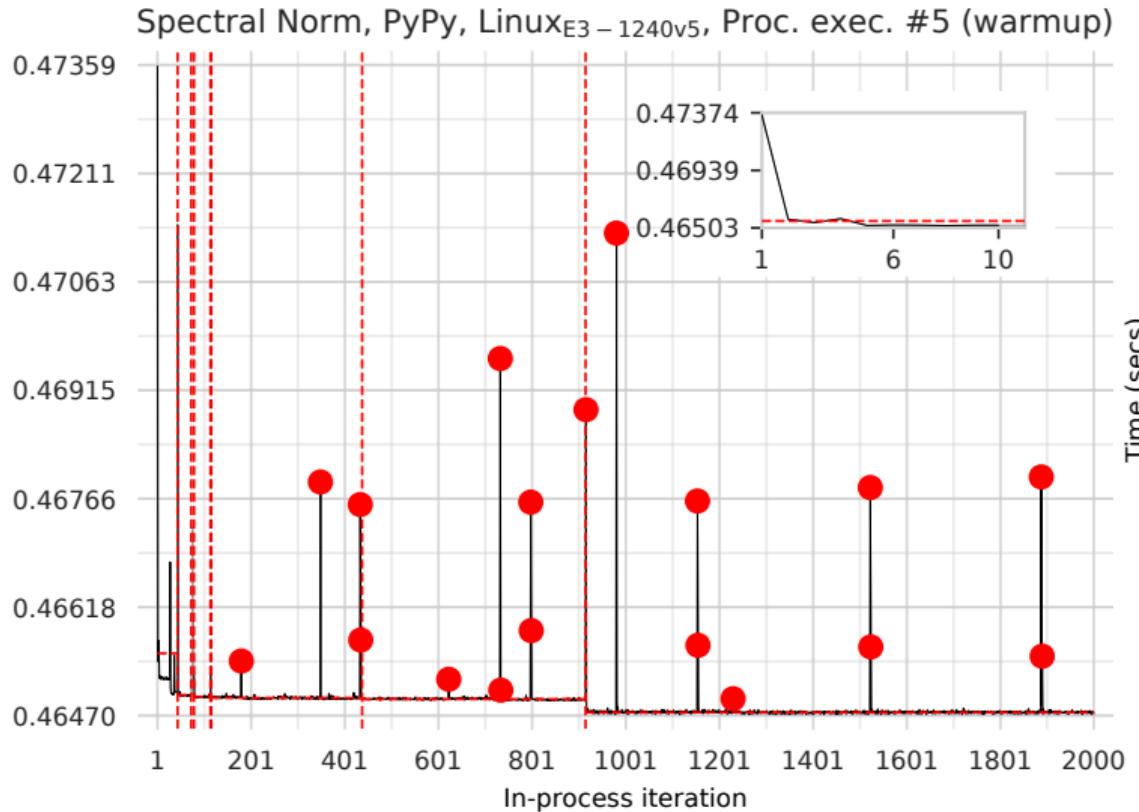
The Warmup Experiment

- ▶ Well-known micro-benchmarks
- ▶ State of the art JIT compilers
 - ▶ Graal, HHVM, TruffleRuby, Hotspot, LuaJIT, PyPy, V8
- ▶ 3 different machines
- ▶ Controlled benchmarking environment
- ▶ Ran the benchmarks for a long time
 - ▶ 2000 in-process iterations.

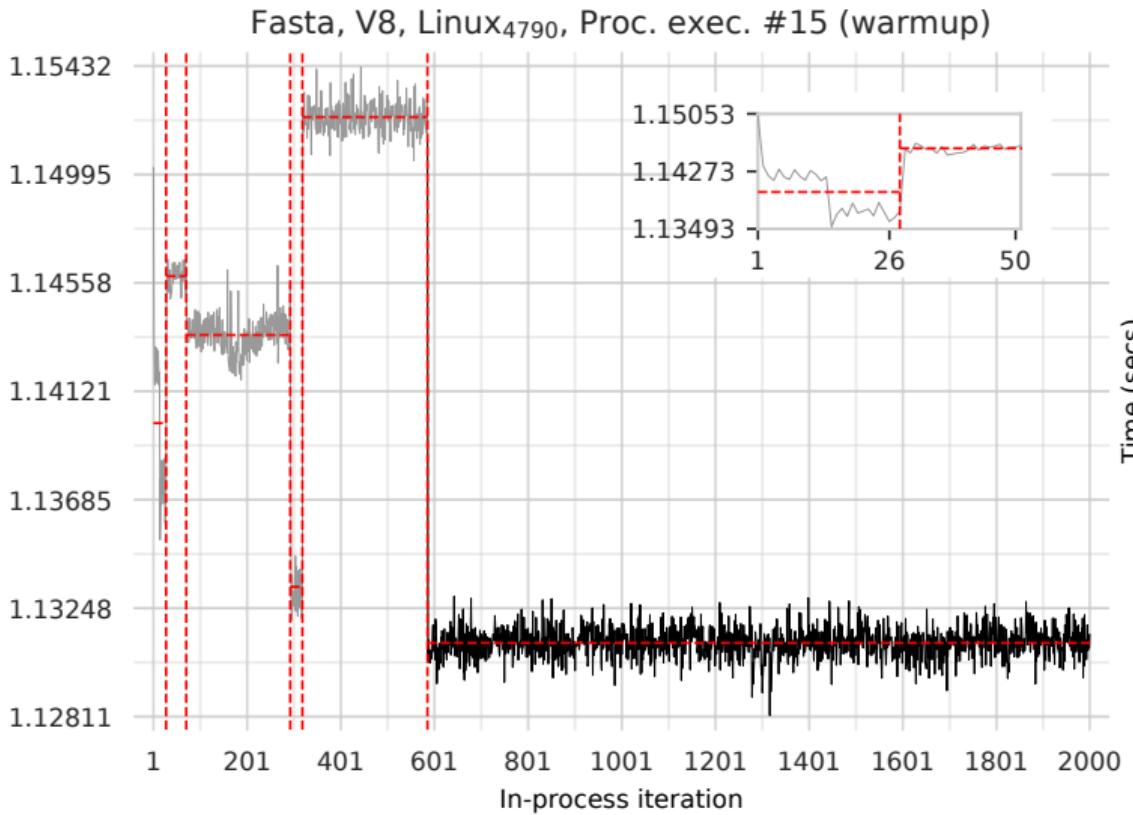
Hypothesis:

Small, deterministic programs reach a steady state of peak performance.

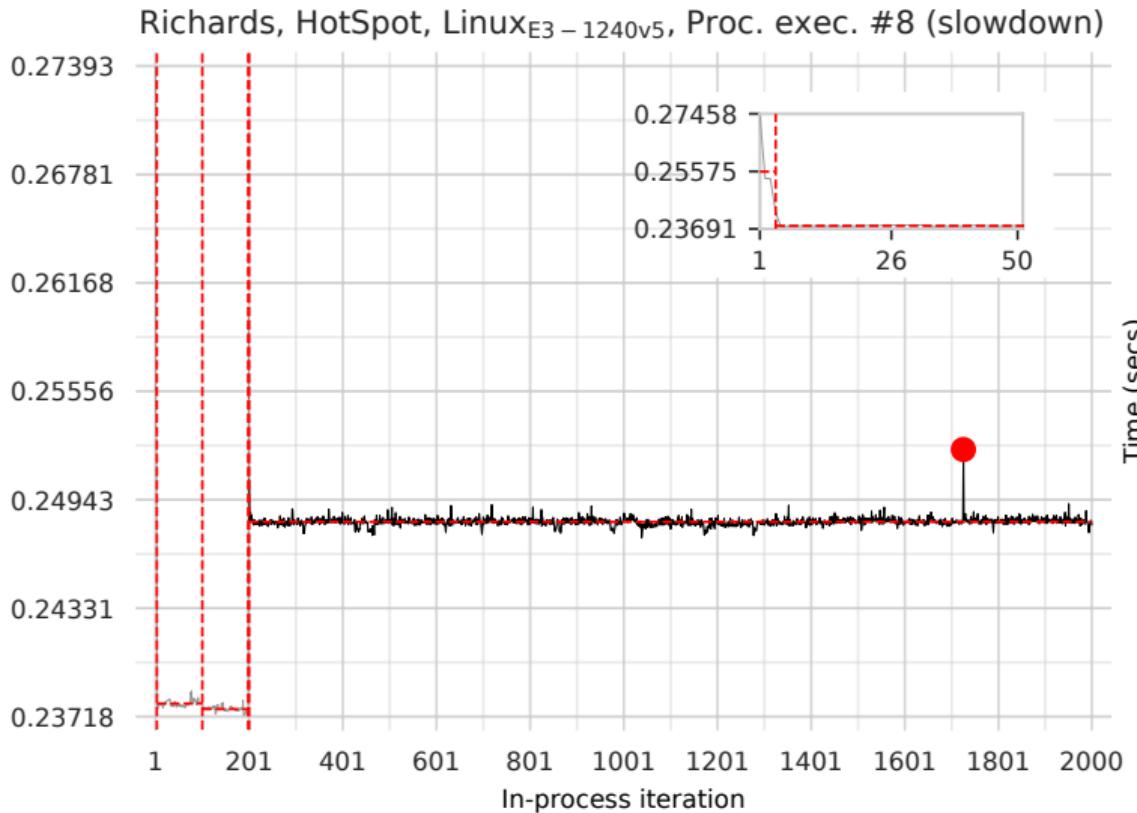
Results: Warmup



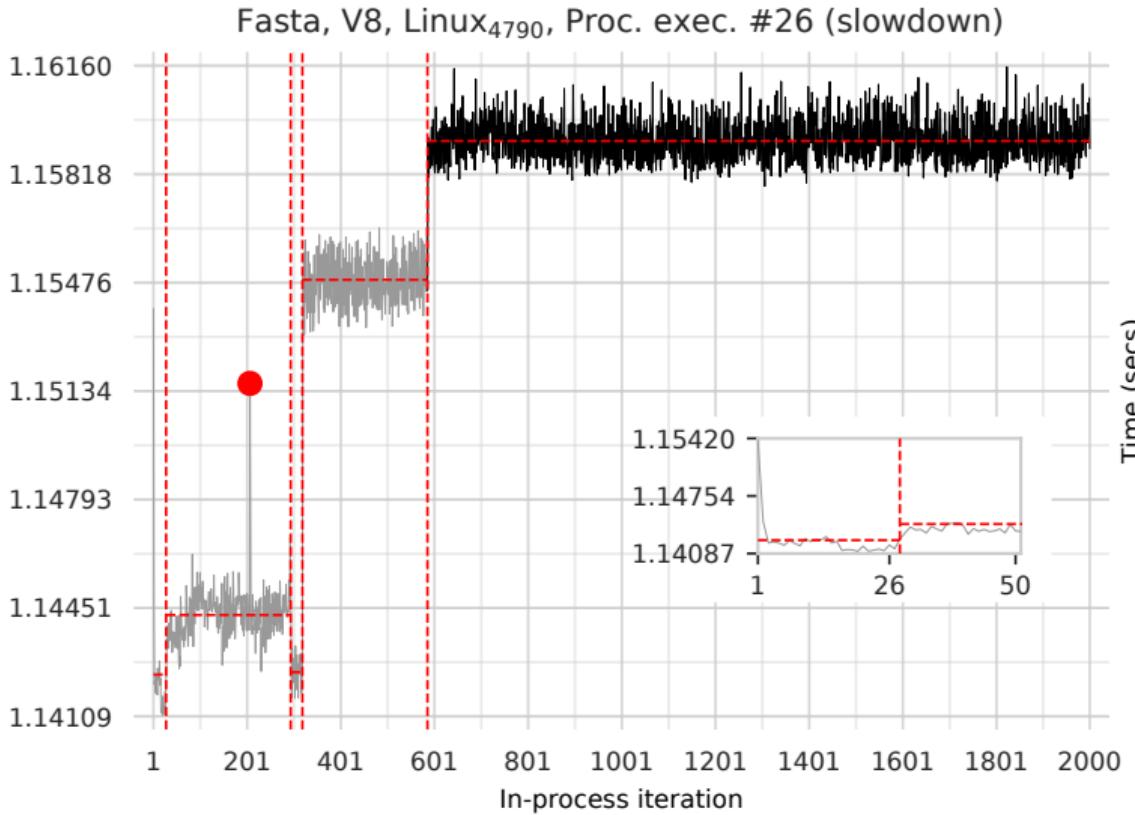
Results: Warmup



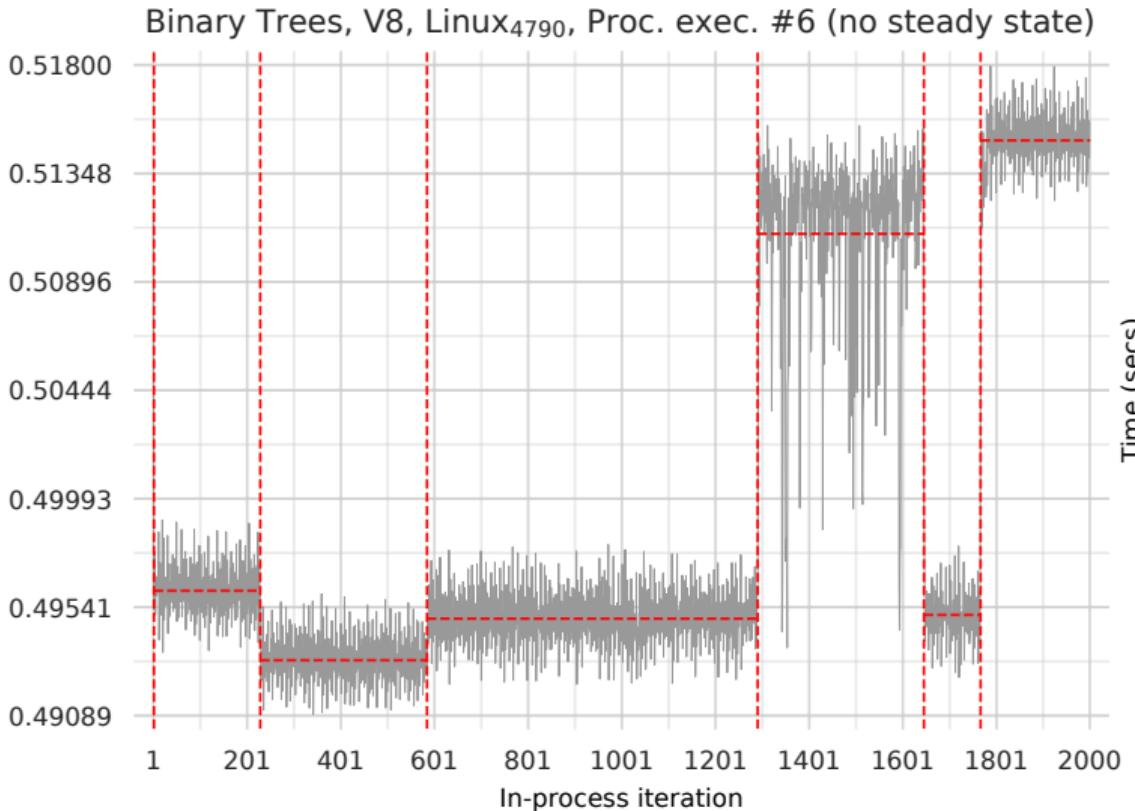
Results: Slowdown



Results: Slowdown

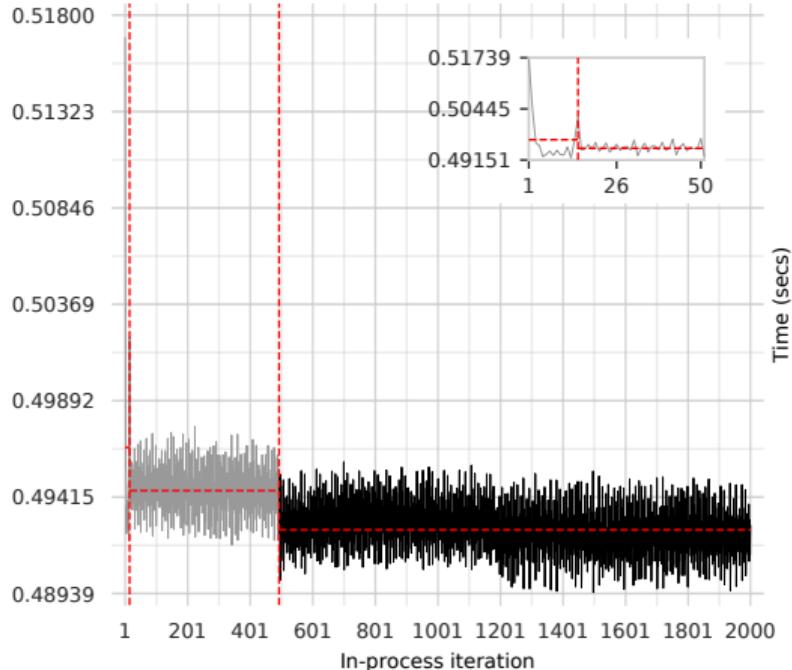


Results: No Steady State

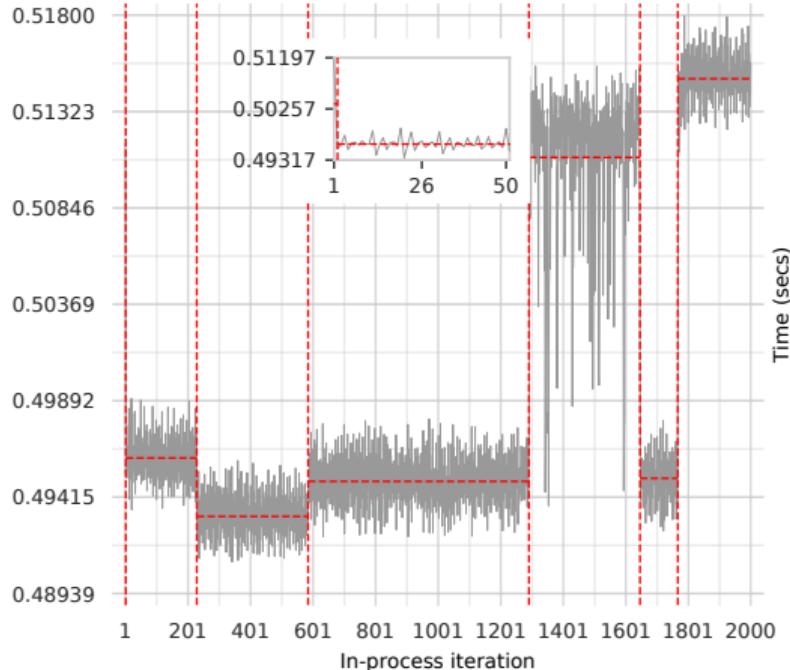


Results: Inconsistent Process-executions

Binary Trees, V8, Linux₄₇₉₀, Proc. exec. #15 (warmup)

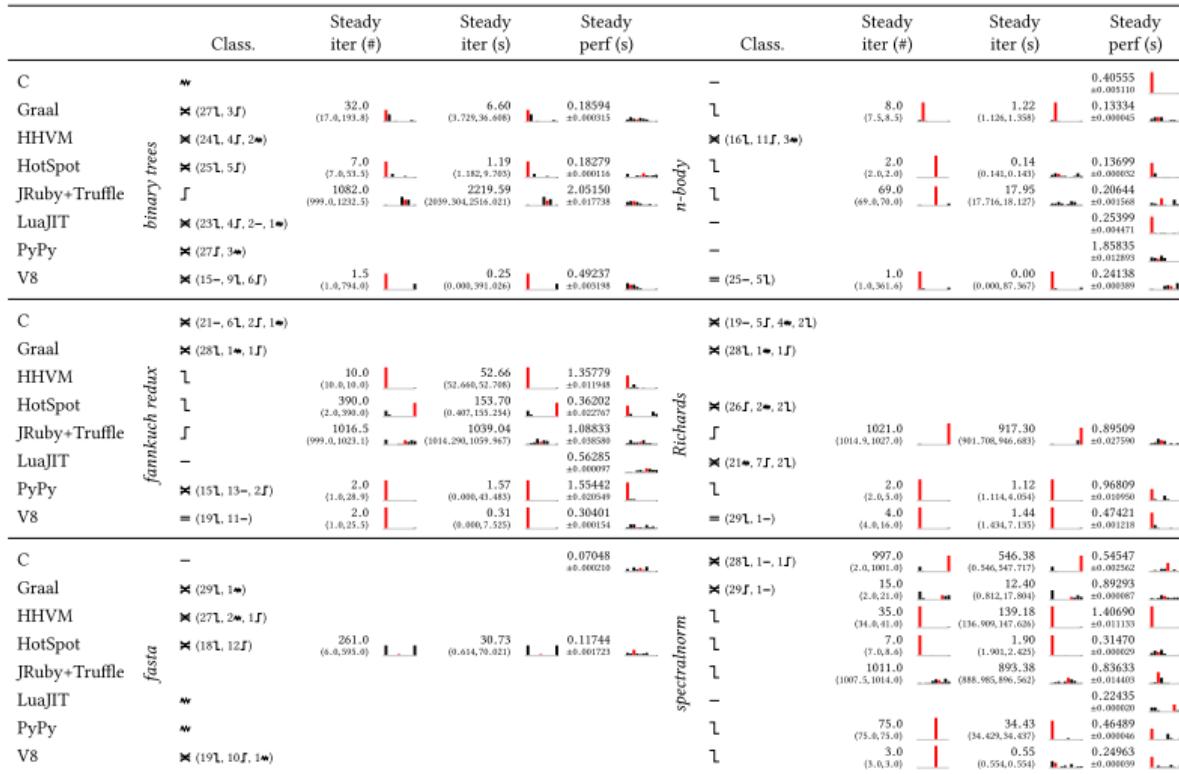


Binary Trees, V8, Linux₄₇₉₀, Proc. exec. #6 (no steady state)



(Same machine)

Results



Results

		Class.	Steady iter (#)	Steady iter (s)	Steady perf (s)
C		\times (27L, 2-, 1J)	775.0 (1.5, 780.0)	425.16 (0.246, 426.809)	0.54581 ± 0.033116
Graal		J	14.0 (2.0, 94.6)	13.60 (0.830, 98.737)	1.05685 ± 0.000126
HHVM	<i>spectralnorm</i>	\times (29L, 1w)			
HotSpot		L	7.0 (7.0, 7.5)	1.91 (1.902, 3.645)	0.31472 ± 0.169143
LuaJIT		-			0.22181 ± 0.000039
PyPy		= (27-, 3L)	1.0 (1.0, 45.2)	0.00 (0.000, 20.597)	0.46480 ± 0.000085
TruffleRuby		\times (25J, 5w)			
V8		L	3.0 (3.0, 3.0)	0.52 (0.523, 0.526)	0.25362 ± 0.000034

Summary of Results

“Good warmup” (flat/warmup) occurs for only:

Summary of Results

“Good warmup” (flat/warmup) occurs for only:

67.2%–70.3% of process executions

Summary of Results

“Good warmup” (flat/warmup) occurs for only:

67.2%–70.3% of process executions

37.8%–40.0% of $\langle \text{VM}, \text{benchmark} \rangle$ pairs

Hypothesis:

Small, deterministic programs reach a steady state of peak performance

In the short term:

- ▶ Run benchmarks for longer to uncover issues.
- ▶ Be prepared to discard some of your results.
- ▶ Report the warmup time in results.

In the longer term:

- ▶ Find ways to improve warmup.

Improving Warmup

CHAPTER 35

INTEL® PROCESSOR TRACE

35.1 OVERVIEW

Intel® Processor Trace (**Intel PT**) is an extension of Intel® Architecture that captures information about software execution using dedicated hardware facilities that cause only minimal performance perturbation to the software being traced. This information is collected in **data packets**. The initial implementations of Intel PT offer **control flow tracing**, which generates a variety of packets to be processed by a software decoder. The packets include timing, program flow information (e.g. branch targets, branch taken/not taken indications) and program-induced mode related information (e.g. Intel TSX state transitions, CR3 changes). These packets may be buffered internally before being sent to the memory subsystem or other output mechanism available in the platform. Debug software can process the trace data and reconstruct the program flow.

Later generations include additional trace sources, including software trace instrumentation using PTWRITE, and Power Event tracing.

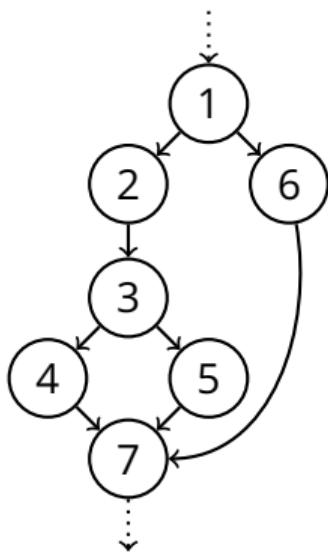
Last Branch Record (LBR) Low overhead, but can only record short traces.

Branch Trace Store (BTS) Can record long traces, but high overhead.

Architecture Event Trace (AET) Requires specialised hardware probe.

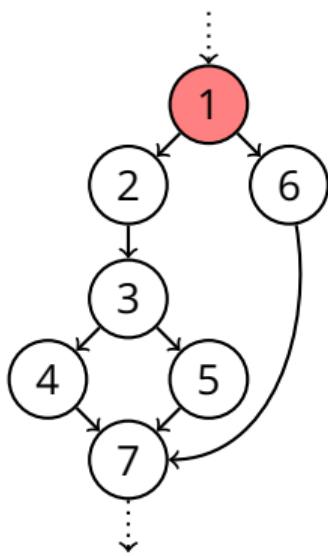
Processor Trace (PT) Long traces, low *collection* overhead.

Tracing JITs



Profiling

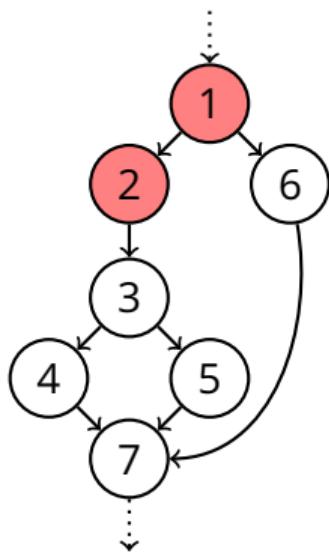
Tracing JITs



- interp(1)
- profile(1)

Profiling

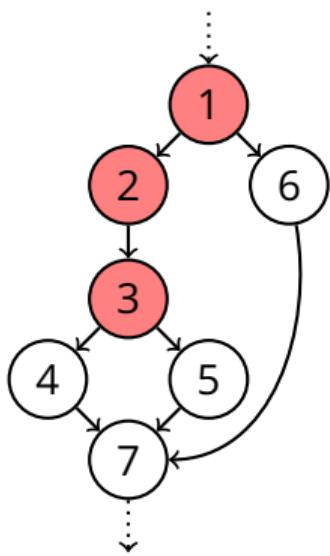
Tracing JITs



- interp(1)
- profile(1)
- interp(2)
- profile(2)

Profiling

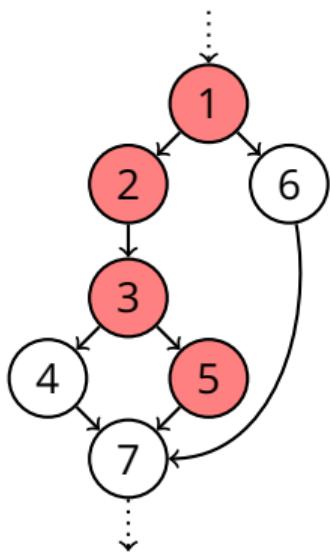
Tracing JITs



- interp(1)
- profile(1)
- interp(2)
- profile(2)
- interp(3)
- profile(3)

Profiling

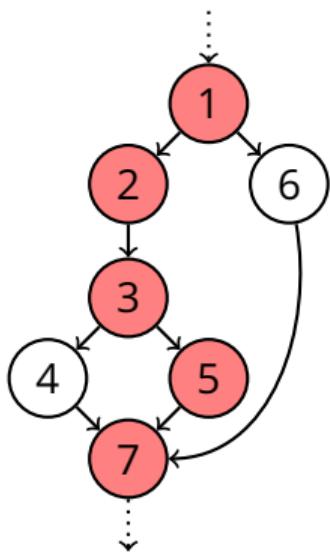
Tracing JITs



- interp(1)
- profile(1)
- interp(2)
- profile(2)
- interp(3)
- profile(3)
- interp(5)
- profile(5)

Profiling

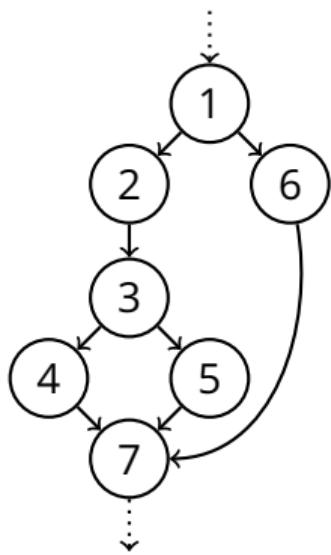
Tracing JITs



- interp(1)
- profile(1)
- interp(2)
- profile(2)
- interp(3)
- profile(3)
- interp(5)
- profile(5)
- interp(7)
- profile(7)

Profiling

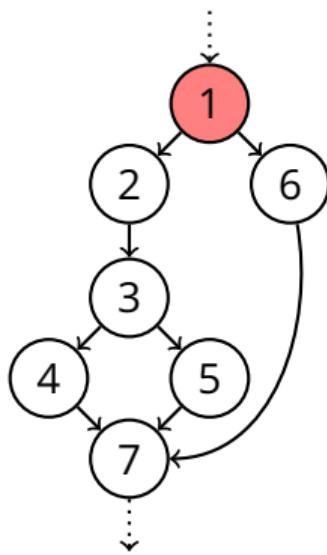
Tracing JITs



Tracing

Trace buffer = < >

Tracing JITs

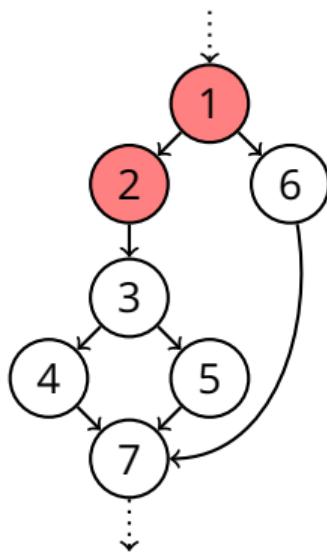


- interp(1)
- record(1)

Tracing

Trace buffer = ⟨ 1 ⟩

Tracing JITs

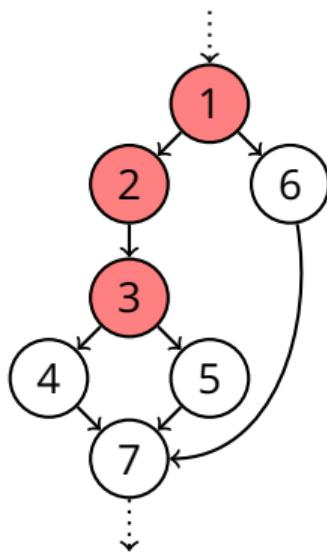


- interp(1)
- record(1)
- interp(2)
- record(2)

Tracing

Trace buffer = ⟨ 1 , 2 ⟩

Tracing JITs

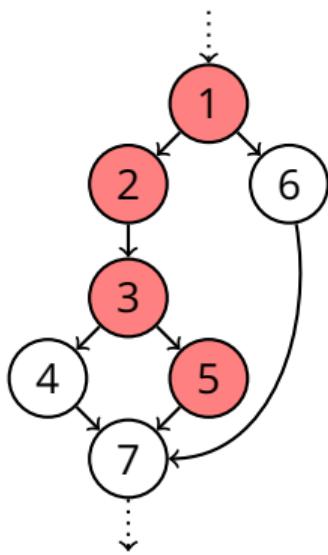


- interp(1)
- record(1)
- interp(2)
- record(2)
- interp(3)
- record(3)

Tracing

Trace buffer = ⟨ 1 , 2 , 3 ⟩

Tracing JITs

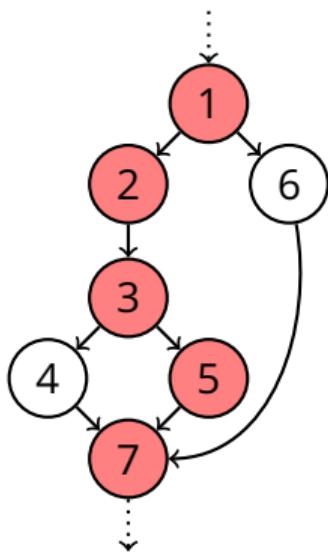


- interp(1)
- record(1)
- interp(2)
- record(2)
- interp(3)
- record(3)
- interp(5)
- record(5)

Tracing

Trace buffer = ⟨ 1 , 2 , 3 , 5 ⟩

Tracing JITs

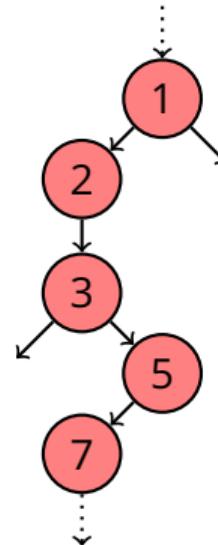
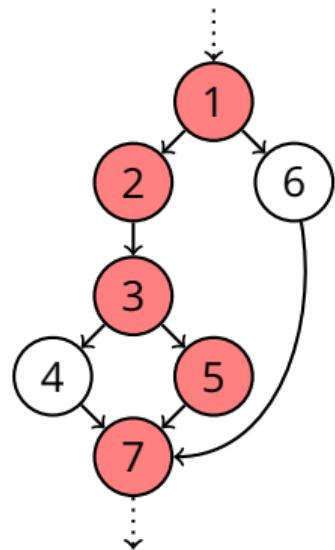


- interp(1)
- record(1)
- interp(2)
- record(2)
- interp(3)
- record(3)
- interp(5)
- record(5)
- interp(7)
- record(7)

Tracing

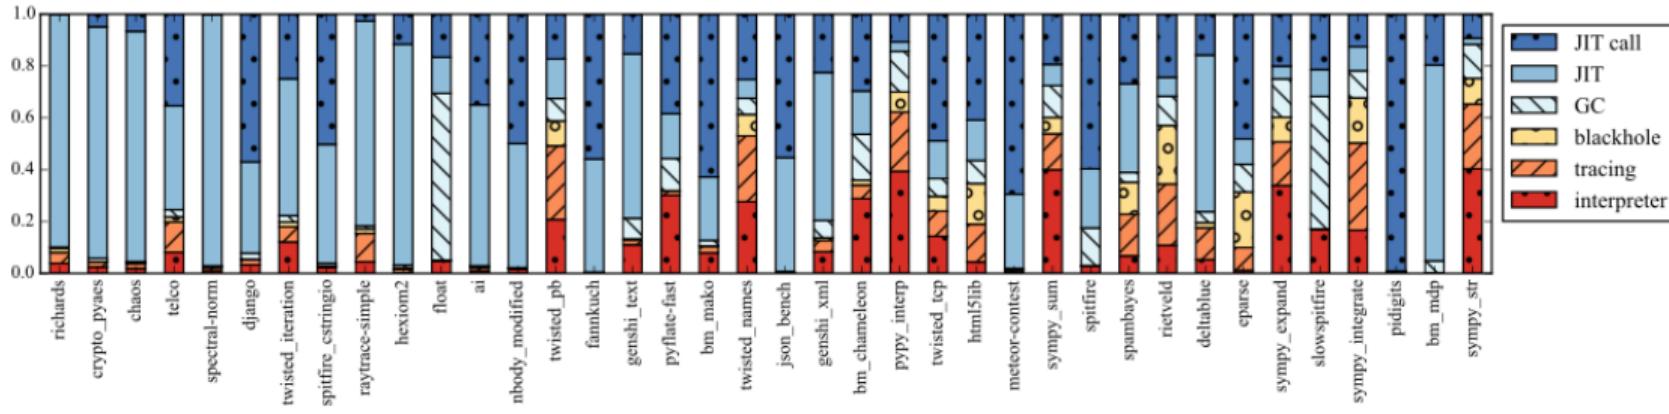
Trace buffer = ⟨ 1 , 2 , 3 , 5 , 7 ⟩

Tracing JITs



Compiling

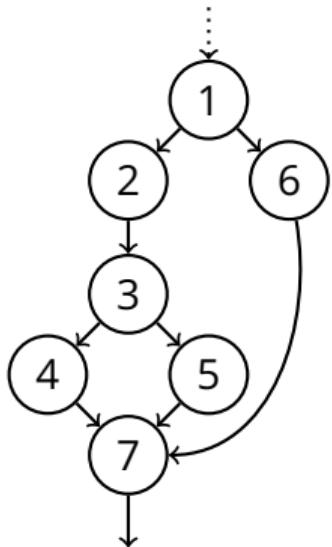
Tracing JITs



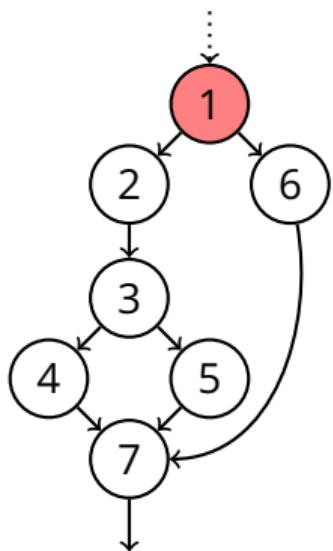
Cross-Layer Workload Characterization of Meta-Tracing JIT VMs. Ilbeyi et al.

PT-enabled Trace Collection

- Start Intel PT

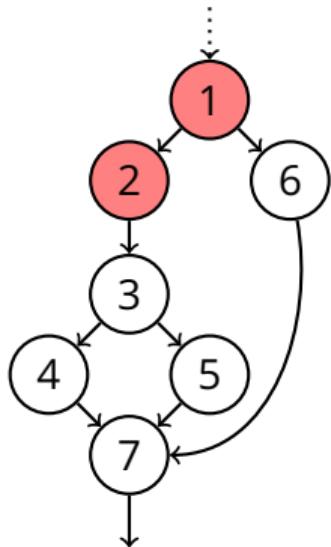


PT-enabled Trace Collection



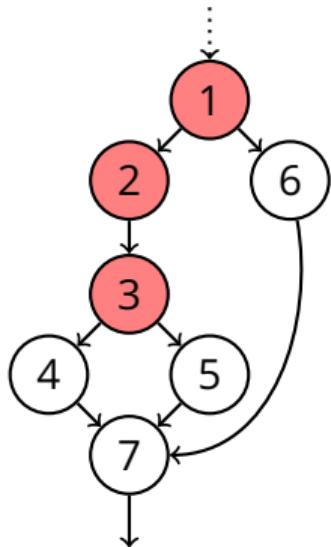
- Start Intel PT
- interp(1)

PT-enabled Trace Collection



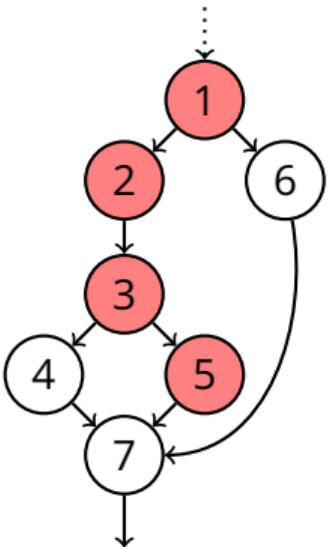
- Start Intel PT
- interp(1)
- interp(2)

PT-enabled Trace Collection



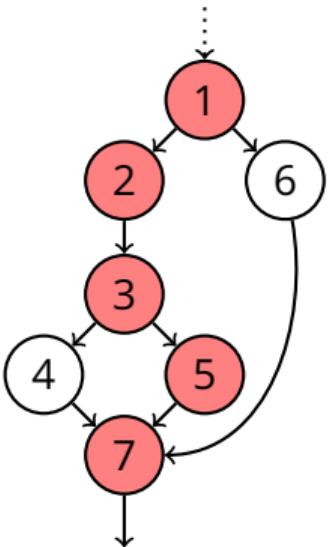
- Start Intel PT
- interp(1)
- interp(2)
- interp(3)

PT-enabled Trace Collection



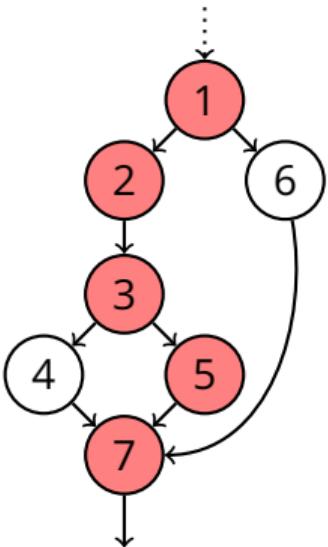
- Start Intel PT
- interp(1)
- interp(2)
- interp(3)
- interp(5)

PT-enabled Trace Collection



- Start Intel PT
- interp(1)
- interp(2)
- interp(3)
- interp(5)
- interp(7)

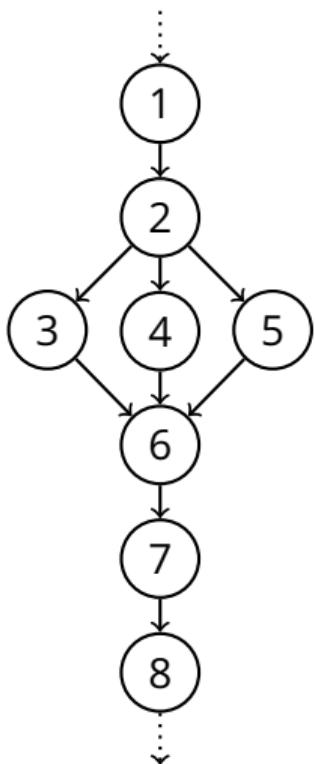
PT-enabled Trace Collection



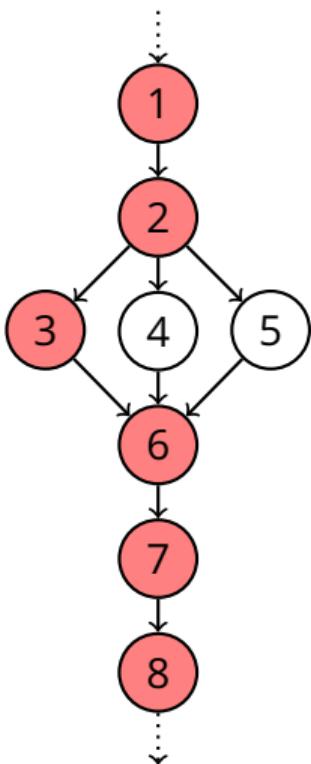
- Start Intel PT
- interp(1)
- interp(2)
- interp(3)
- interp(5)
- interp(7)
- Stop Intel PT

If tracing can be cheap, new opportunities arise

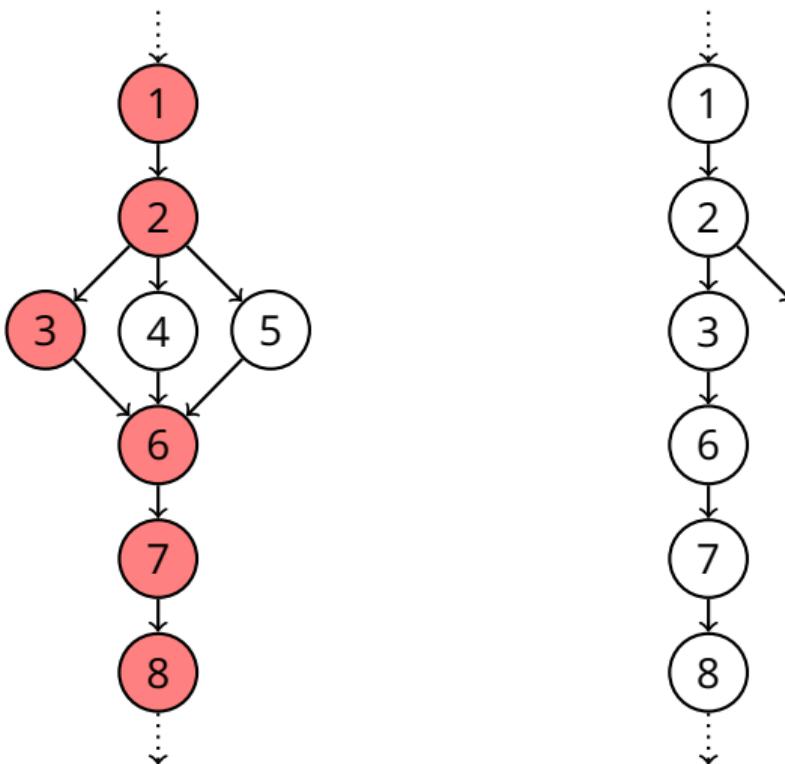
Avoiding trace tree explosions with trace sampling



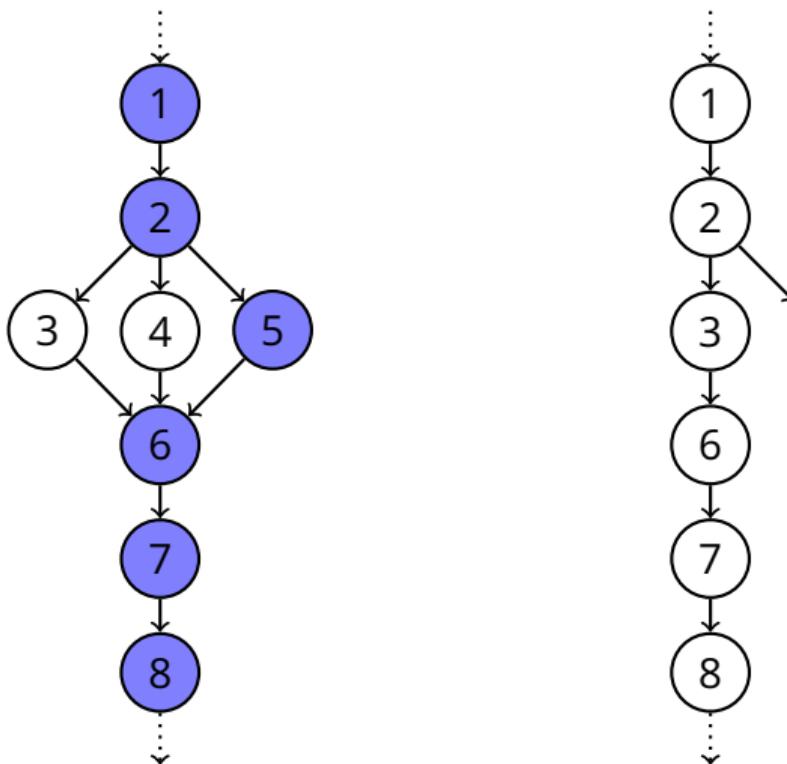
Avoiding trace tree explosions with trace sampling



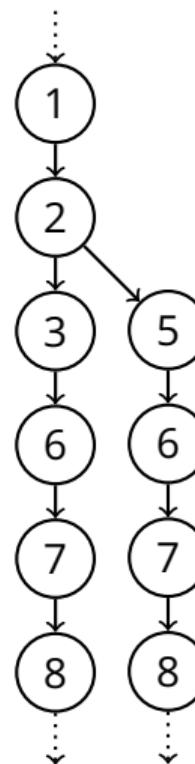
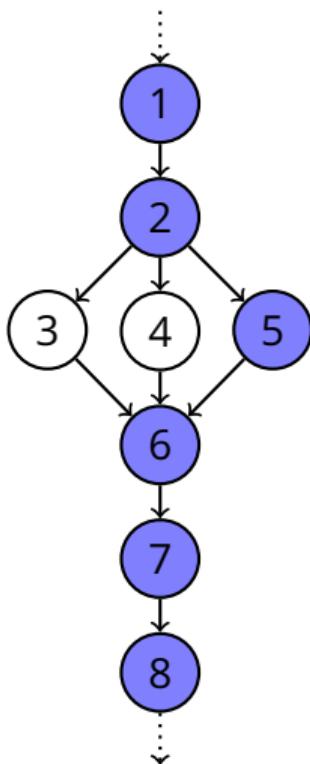
Avoiding trace tree explosions with trace sampling



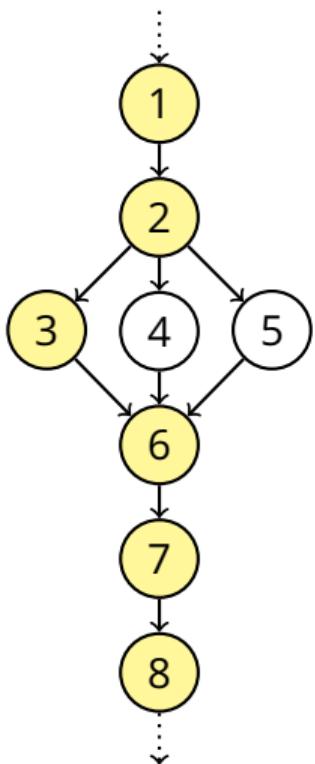
Avoiding trace tree explosions with trace sampling



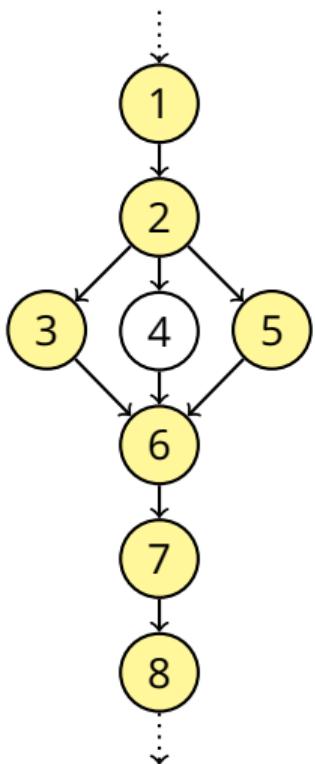
Avoiding trace tree explosions with trace sampling



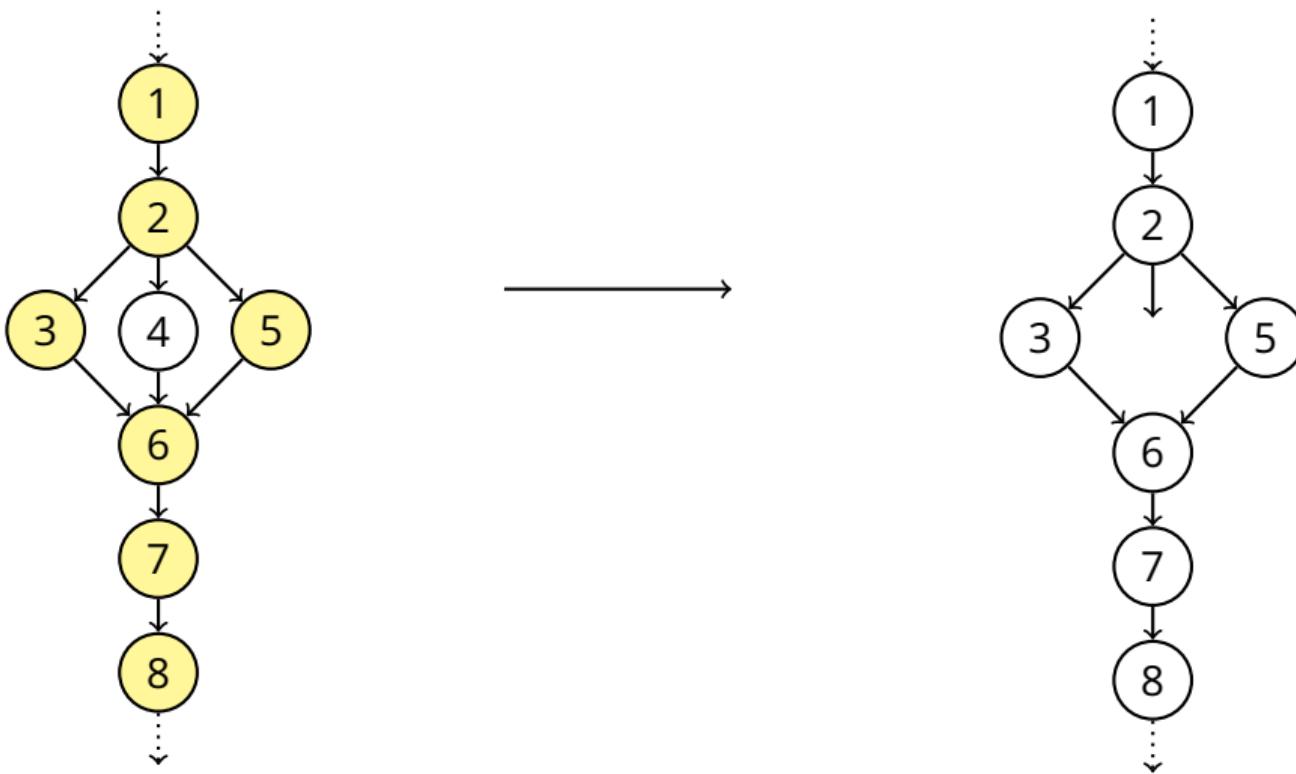
Avoiding trace tree explosions with trace sampling



Avoiding trace tree explosions with trace sampling



Avoiding trace tree explosions with trace sampling



Language Implementation

Language Implementation

Lang1 Interpreter

Language Implementation

Lang1 Interpreter

Lang2 Interpreter

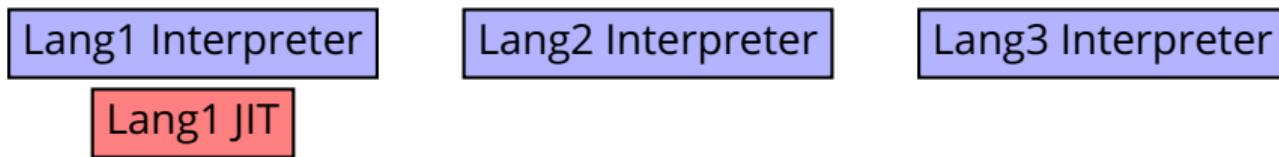
Language Implementation

Lang1 Interpreter

Lang2 Interpreter

Lang3 Interpreter

Language Implementation



Language Implementation



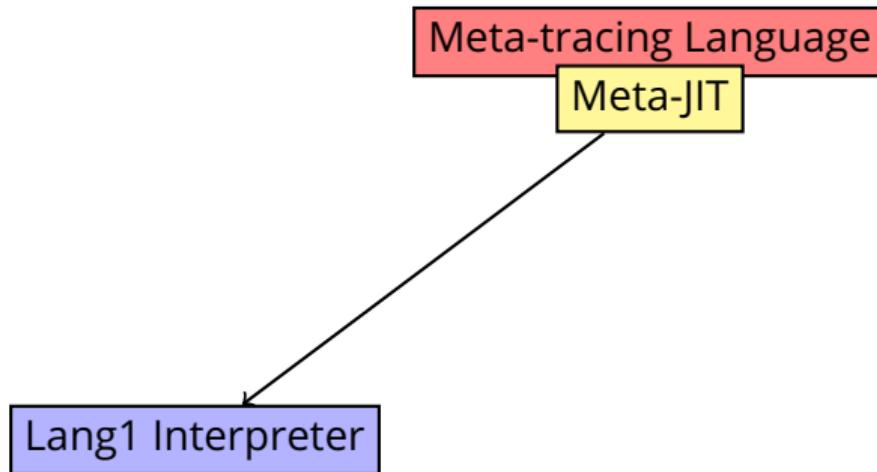
Language Implementation



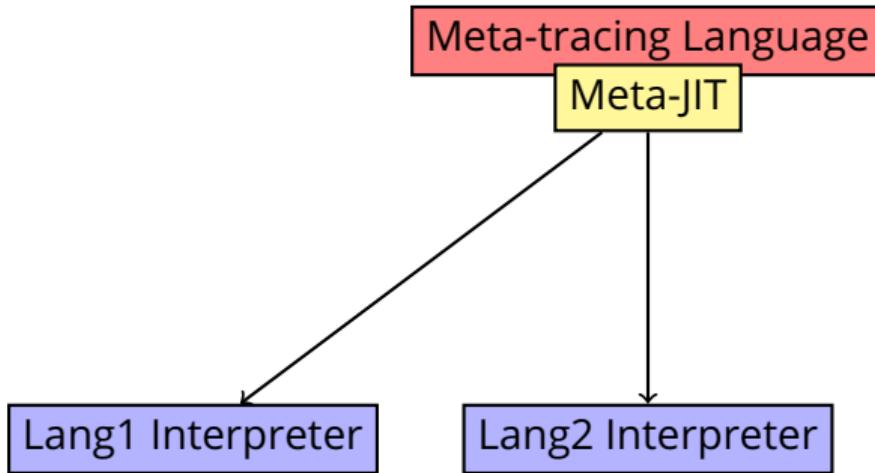
Language Implementation with Meta-Tracing



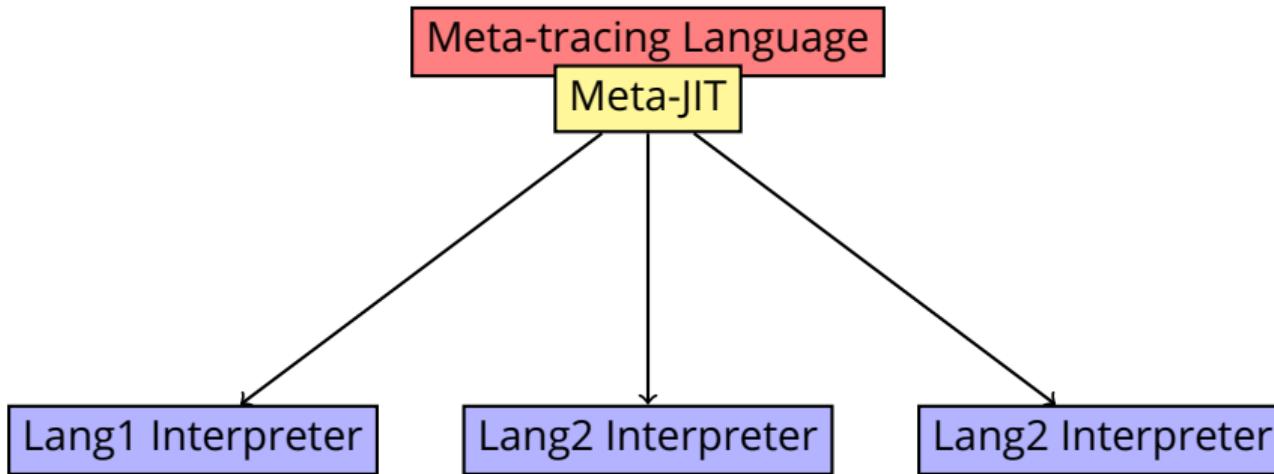
Language Implementation with Meta-Tracing



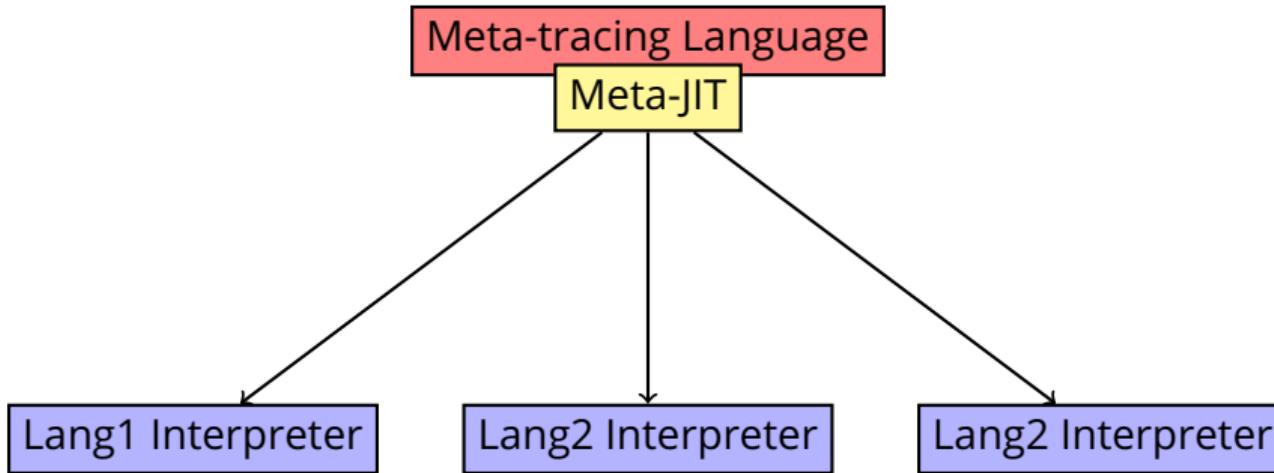
Language Implementation with Meta-Tracing



Language Implementation with Meta-Tracing



Language Implementation with Meta-Tracing



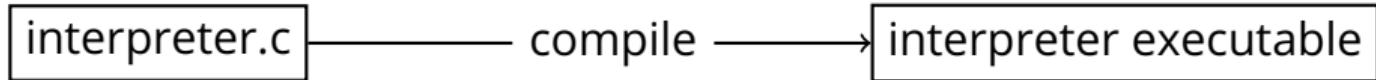
Free JITs!

yk

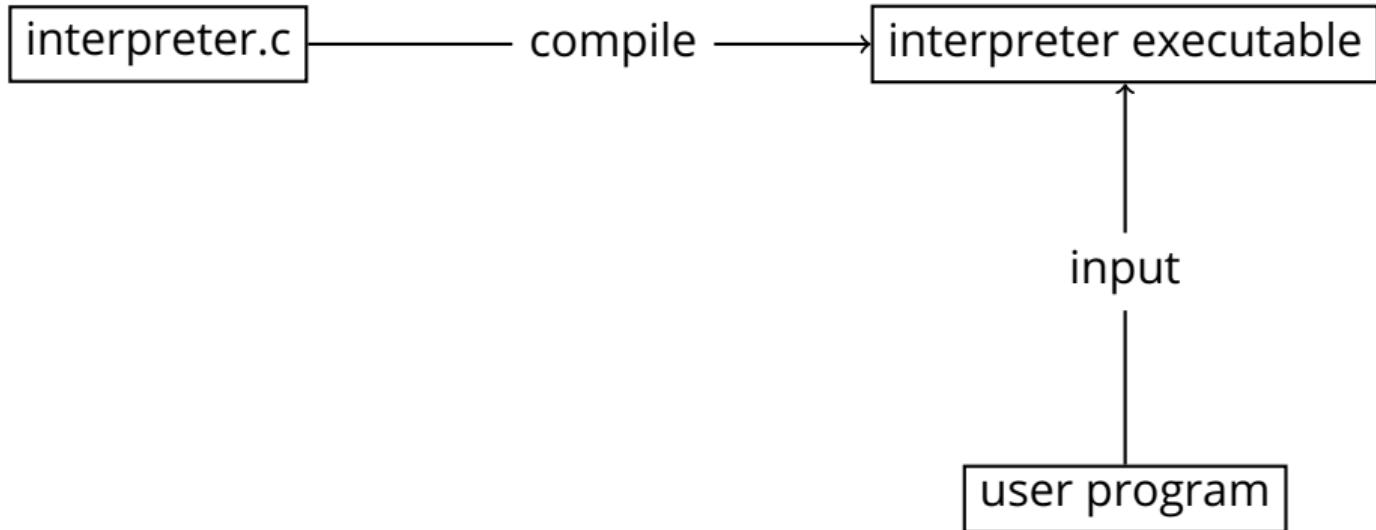
yk: high-level view

interpreter.c

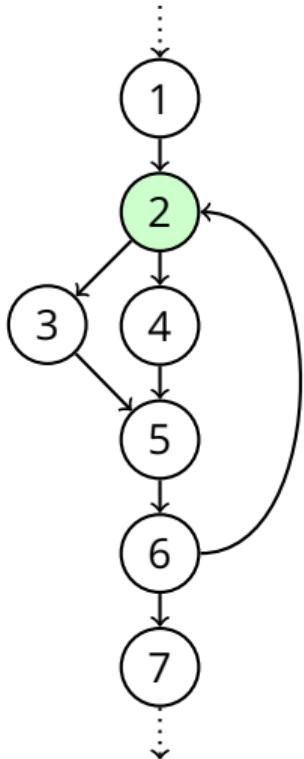
yk: high-level view



yk: high-level view

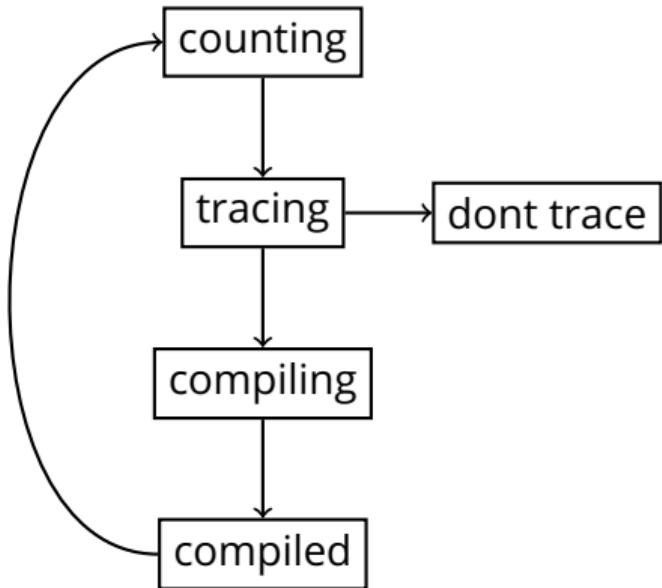


Feeding-back to the JIT

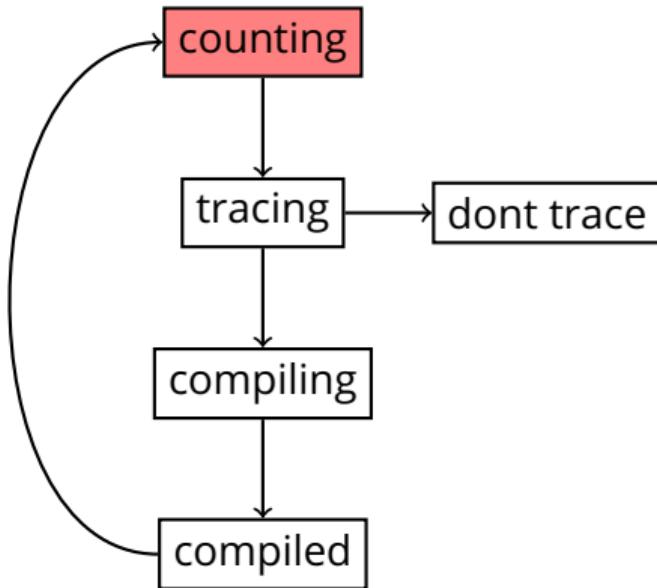


```
YkLocation *locs = calloc(  
    len(prog), sizeof(YkLocation *)) ;  
locs[2] = yk_location_new();  
...  
yk_control_point(locs[pc])
```

Location: state machine



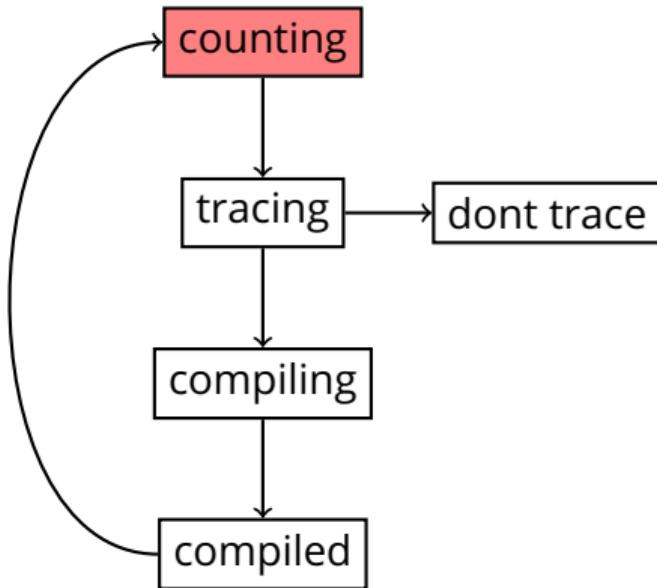
Location: state machine



Counting

- ▶ Atomic increment of hot count.

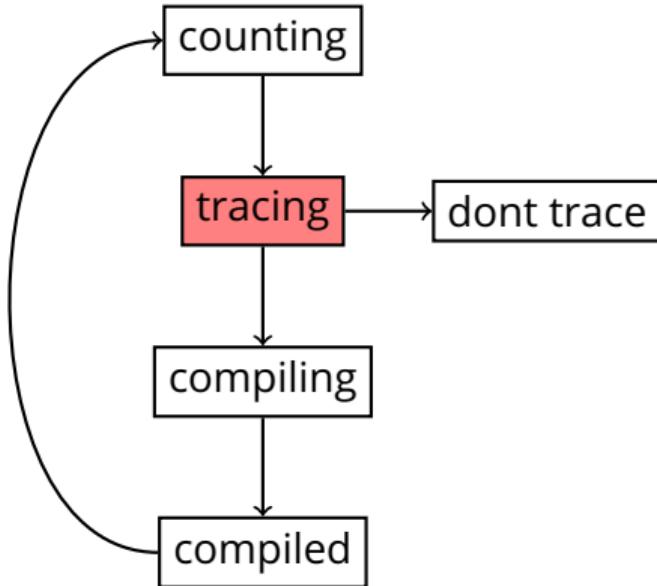
Location: state machine



Counting

- ▶ Atomic increment of hot count.
- ▶ Transition once hot threshold met.

Location: state machine

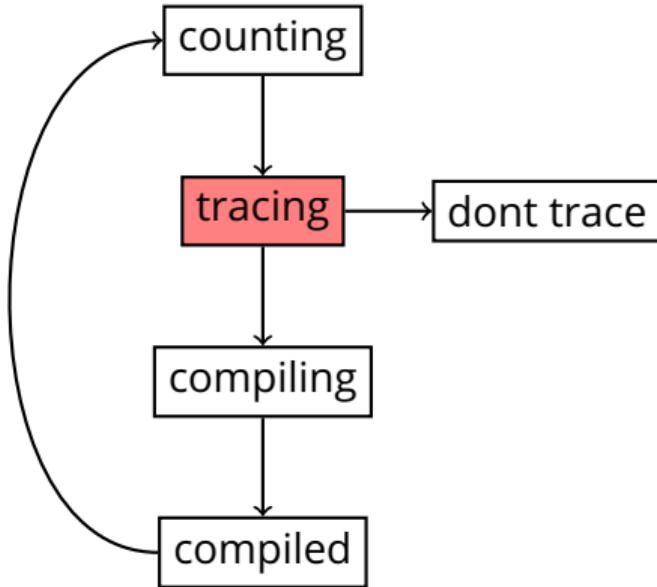


Tracing

As the location enters the tracing state:

- ▶ Turn on PT.

Location: state machine

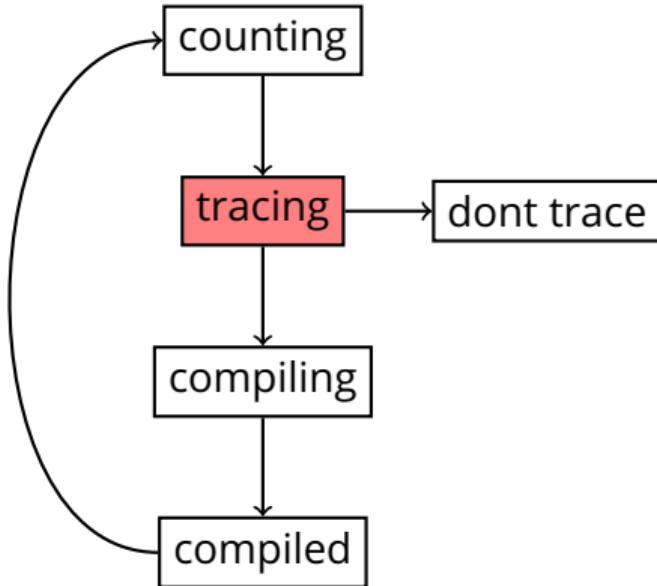


Tracing

As the location enters the tracing state:

- ▶ Turn on PT.
- ▶ Carry on interpreting.

Location: state machine



Tracing

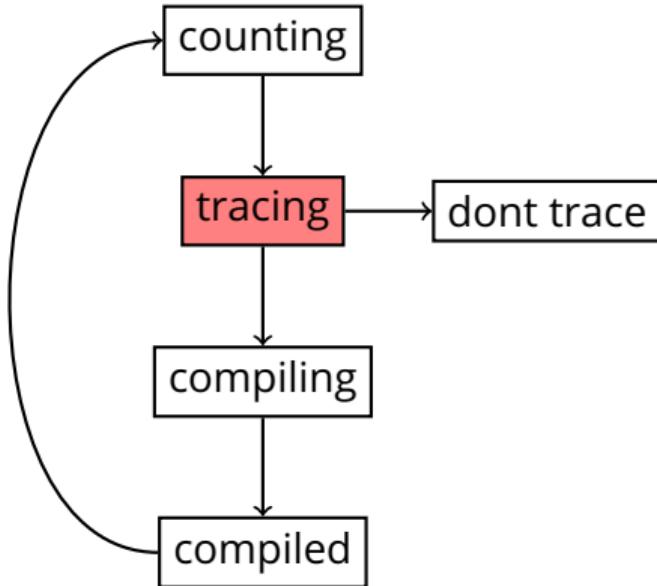
As the location enters the tracing state:

- ▶ Turn on PT.
- ▶ Carry on interpreting.

When we see this location again:

- ▶ Stop PT.

Location: state machine



Tracing

As the location enters the tracing state:

- ▶ Turn on PT.
- ▶ Carry on interpreting.

When we see this location again:

- ▶ Stop PT.
- ▶ PT payload is decoded.
- ▶ Binary trace is “mapped”.

yk

\langle  ,  ,  ,  ,  ,  ,  , \dots \rangle

 IPT block decoder

\langle 0x7ffff7fd6090, 0x7ffff7fd609c, 0x7ffff7fd60ba, ... \rangle

```
clang -fuse-ld=lld \
-flto \
-Wl,--plugin-opt=-lto-embed-bitcode=optimized \
-Wl,--lto-basic-block-sections=labels \
...
interpreter.c
```

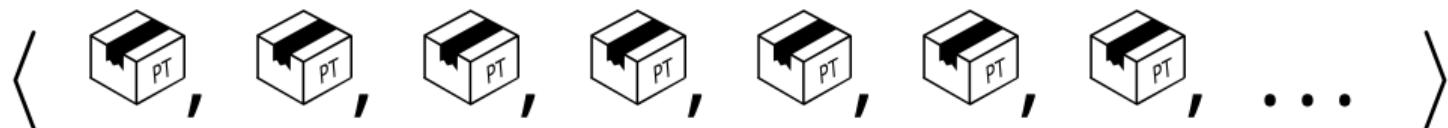
SHT_LLVM_BB_ADDR_MAP Section (basic block address map)

This section stores the binary address of basic blocks along with other related metadata. This information can be used to map binary profiles (like perf profiles) directly to machine basic blocks. This section is emitted with `-basic-block-sections=labels` and will contain a BB address map table for every function which may be constructed as follows:

```
.section ".llvm_bb_addr_map","",@llvm_bb_addr_map
.quad .Lfunc_begin0           # address of the function
.byte 2                      # number of basic blocks
# BB record for BB_0
.uleb128 .Lfunc_beign0-.Lfunc_begin0 # BB_0 offset relative to function entry (always zero)
.uleb128 .LBB_END0_0-.Lfunc_begin0 # BB_0 size
.byte x                      # BB_0 metadata
# BB record for BB_1
.uleb128 .LBB0_1-.Lfunc_begin0      # BB_1 offset relative to function entry
.uleb128 .LBB_END0_1-.Lfunc_begin0 # BB_1 size
.byte y                      # BB_1 metadata
```

This creates a BB address map table for a function with two basic blocks.

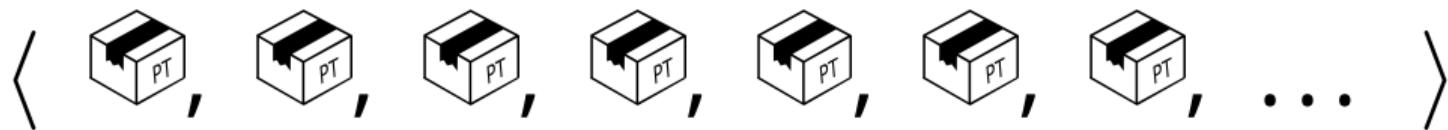
yk



↓ IPT block decoder

⟨ 0x7fffff7fd6090, 0x7fffff7fd609c, 0x7fffff7fd60ba, ... ⟩

yk



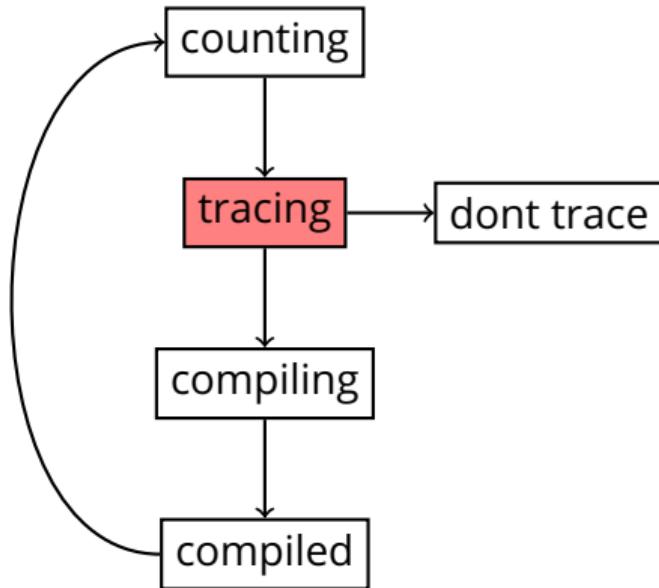
↓ IPT block decoder

`(0x7ffff7fd6090, 0x7ffff7fd609c, 0x7ffff7fd60ba, ...)`

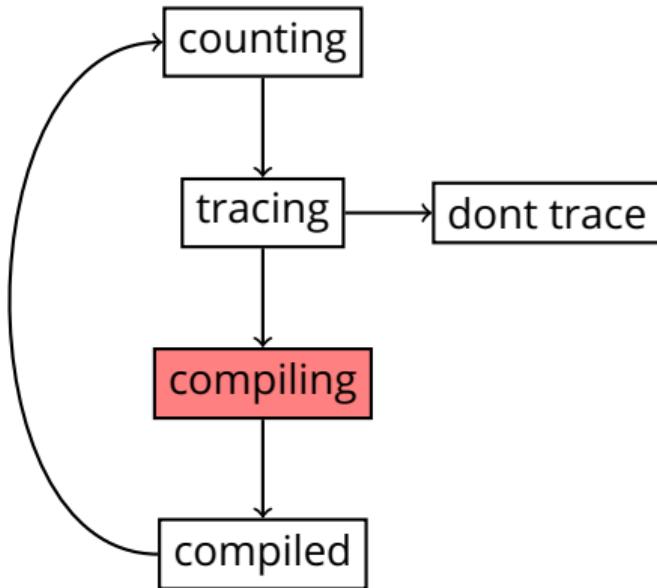
↓ map to LLVM block IDs

`(main<bb6>, do_stuff<bb0>, do_stuff<bb1>, ...)`

Location: state machine



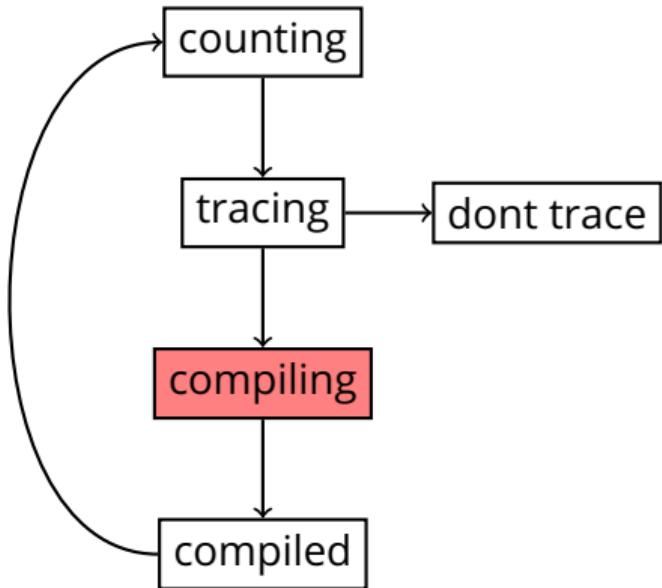
Location: state machine



Compiling

- ▶ Extract blocks from IR.

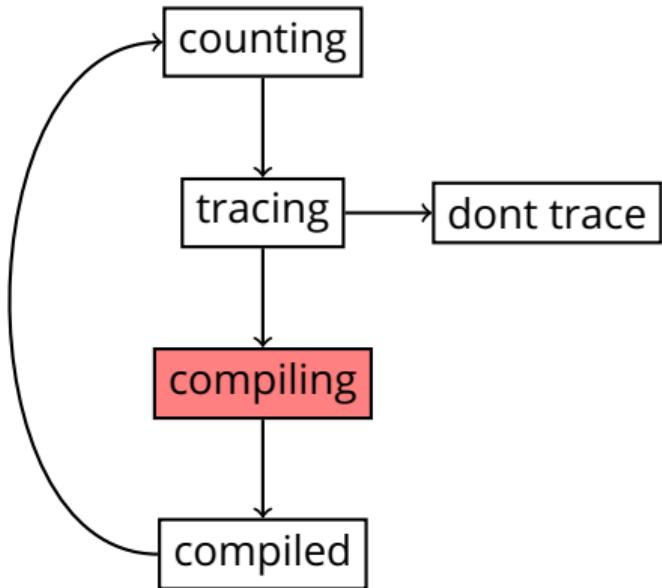
Location: state machine



Compiling

- ▶ Extract blocks from IR.
- ▶ Copy them into a new LLVM module.

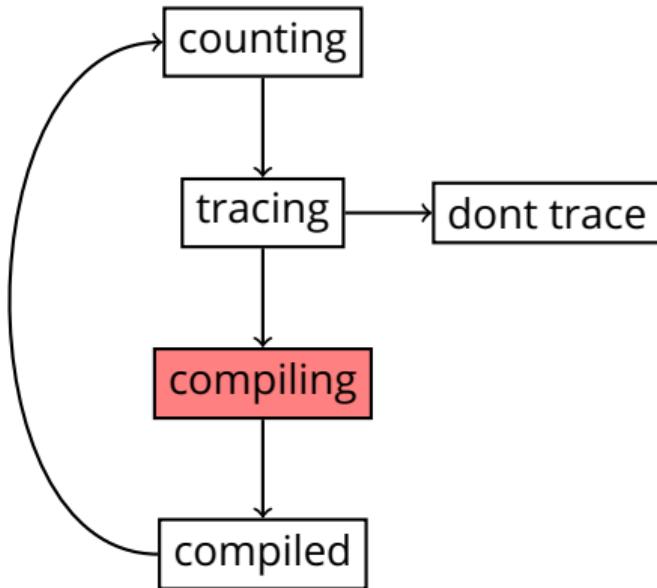
Location: state machine



Compiling

- ▶ Extract blocks from IR.
- ▶ Copy them into a new LLVM module.
- ▶ Rewrite terminators to guards.

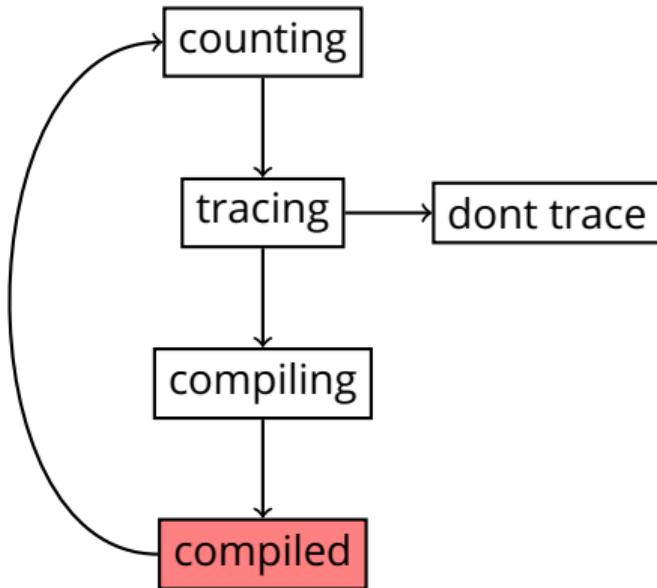
Location: state machine



Compiling

- ▶ Extract blocks from IR.
- ▶ Copy them into a new LLVM module.
- ▶ Rewrite terminators to guards.
- ▶ Compile to executable code.

Location: state machine



Compiled

- ▶ Jump to compiled code instead of interpreting.

- ▶ We can use this for interpreters written in any LLVM language.
 - ▶ Rust, C, C++, Julia, D,
- ▶ You get a tracing JIT for free!
- ▶ We can adapt existing interpreters.
 - ▶ CPython, Lua, Ruby MRI,
- ▶ A lot of the trace compilation work can be done by LLVM.
 - ▶ At least initially...

No performance numbers yet!

Plan of action

- ▶ Get some C interpreters JITting on our system.
 - ▶ Lua and CPython.
- ▶ Implement and evaluate additional ideas to improve warmup.

Thanks for Listening

