

# Why Aren't More Users More Happy With Our VMs?



Laurence  
Tratt

Warmup work in collaboration with:  
Edd Barrett, Carl Friedrich Bolz, Rebecca Killick, and Sarah Mount



Software Development Team

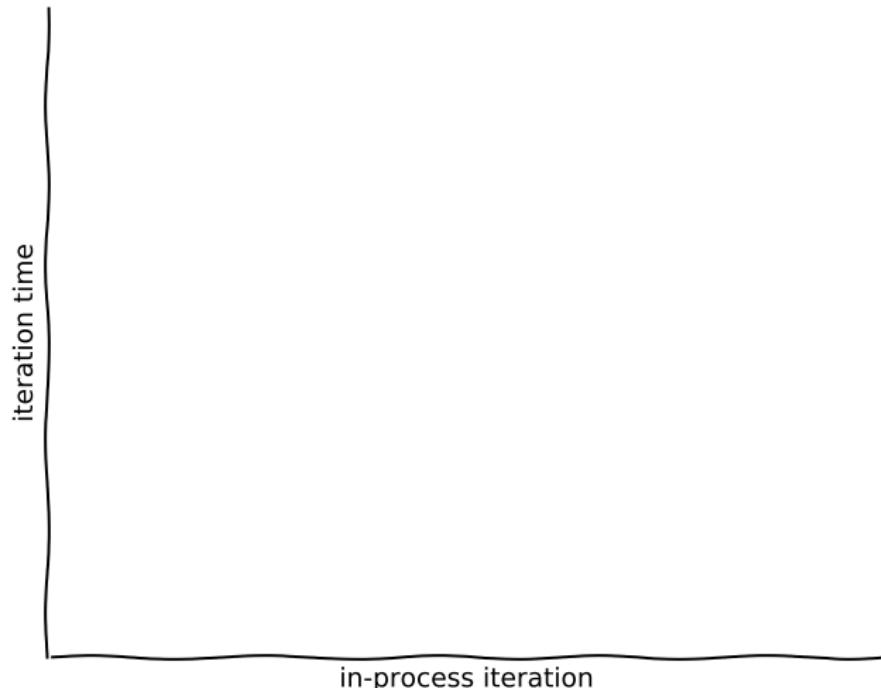
2018-02-05

# What to expect from this talk

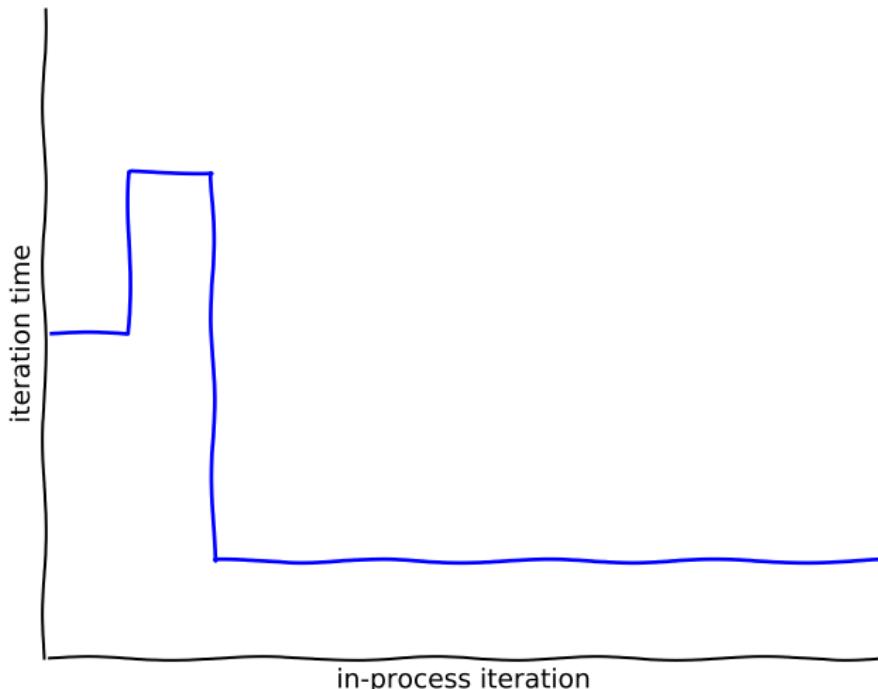
JVMs bring "gcc -O2"  
to the masses

*-Cliff Click: A JVM does that?*

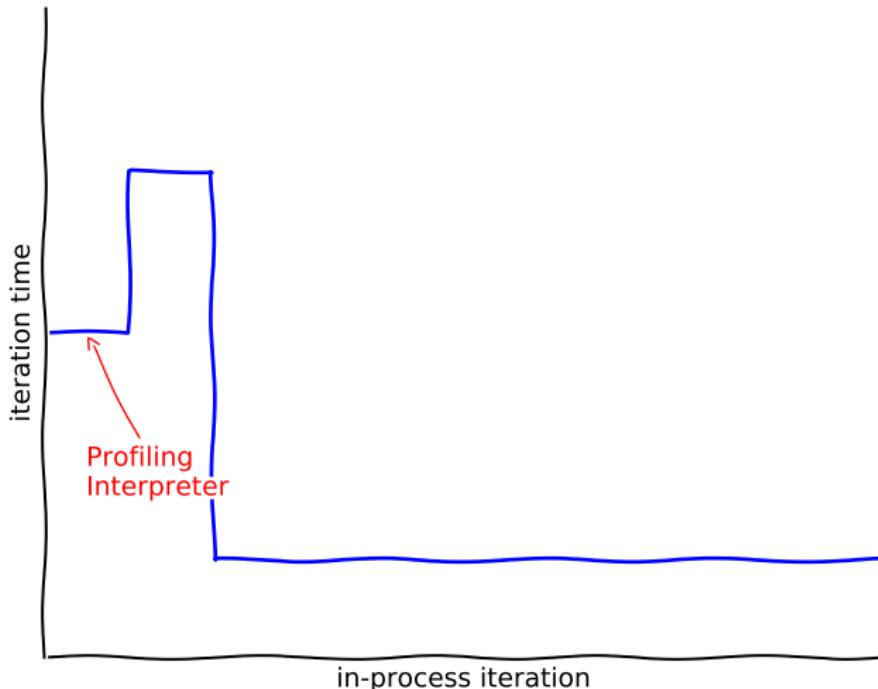
# What do VM claims pertain to?



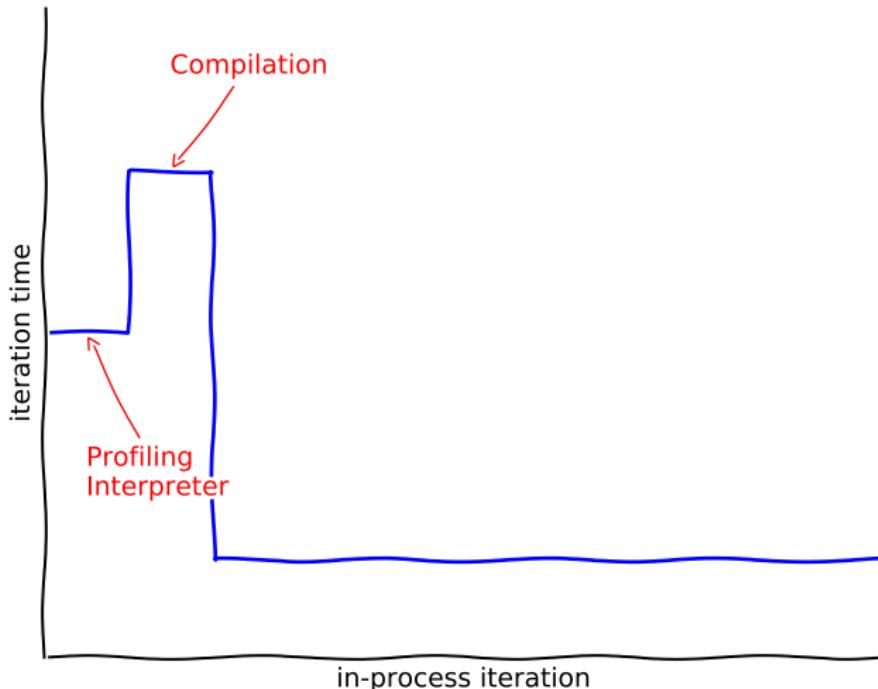
# What do VM claims pertain to?



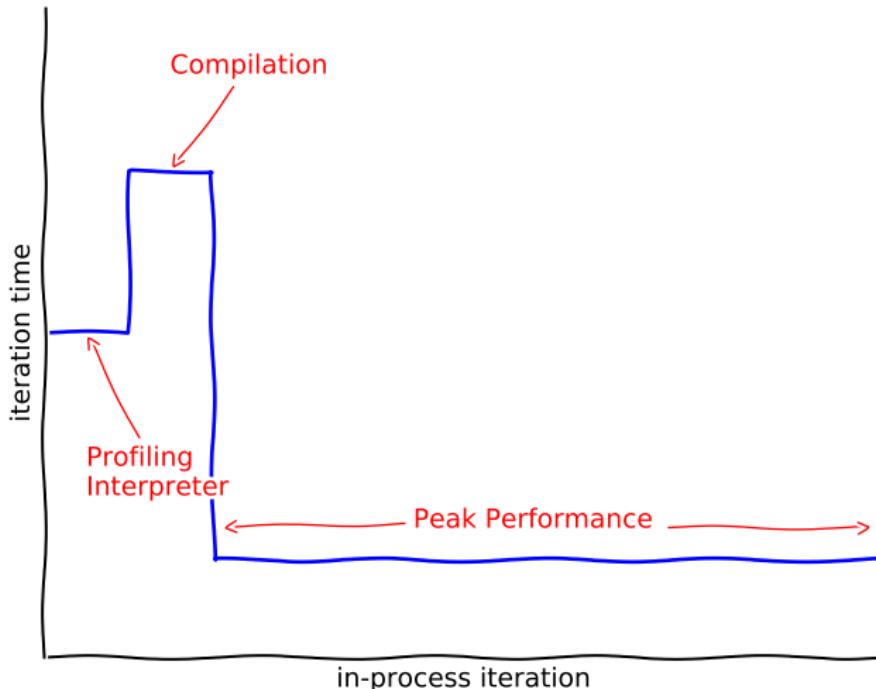
# What do VM claims pertain to?



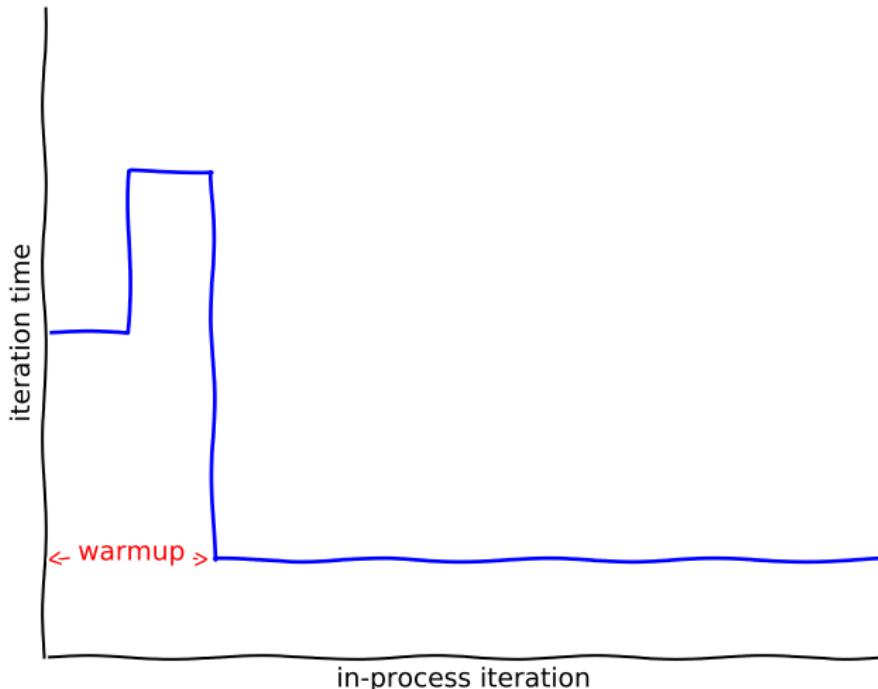
# What do VM claims pertain to?



# What do VM claims pertain to?



# What do VM claims pertain to?



# Warmup

Users *always* perceive warmup

# Warmup

Users *always* perceive warmup

Maybe we should know how long it is?

# The Warmup Experiment

Measure warmup of modern language implementations

# The Warmup Experiment

Measure warmup of modern language implementations

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

## Method 1: Which benchmarks?

The language benchmark games are perfect for us  
(unusually)

## Method 1: Which benchmarks?

The language benchmark games are perfect for us  
(unusually)

We removed any CFG non-determinism

## Method 1: Which benchmarks?

The language benchmark games are perfect for us  
(unusually)

We removed any CFG non-determinism

We added checksums to all benchmarks

## Method 2: How long to run?

2000 *in-process iterations*

## Method 2: How long to run?

2000 *in-process iterations*

30 *process executions*

## Method 3: VMs

- Graal-0.22
- HHVM-3.19.1
- JRuby/Truffle (git #6e9d5d381777)
- Hotspot-8u121b13
- LuaJit-2.0.4
- PyPy-5.7.1
- V8-5.8.283.32
- GCC-4.9.4

Note: same GCC (4.9.4) used for all compilation

## Method 4: Machines

- Linux<sub>4790</sub>, Debian 8, 24GiB RAM
- Linux<sub>E3-1240v5</sub>, Debian 8, 32GiB RAM
- OpenBSD<sub>4790</sub>, OpenBSD 6.0, 32GiB RAM

## Method 4: Machines

- Linux<sub>4790</sub>, Debian 8, 24GiB RAM
  - Linux<sub>E3-1240v5</sub>, Debian 8, 32GiB RAM
  - OpenBSD<sub>4790</sub>, OpenBSD 6.0, 32GiB RAM
- 
- Turbo boost and hyper-threading disabled
  - Network card turned off.
  - Daemons disabled (cron, smtpd)

## Method 5: Krun

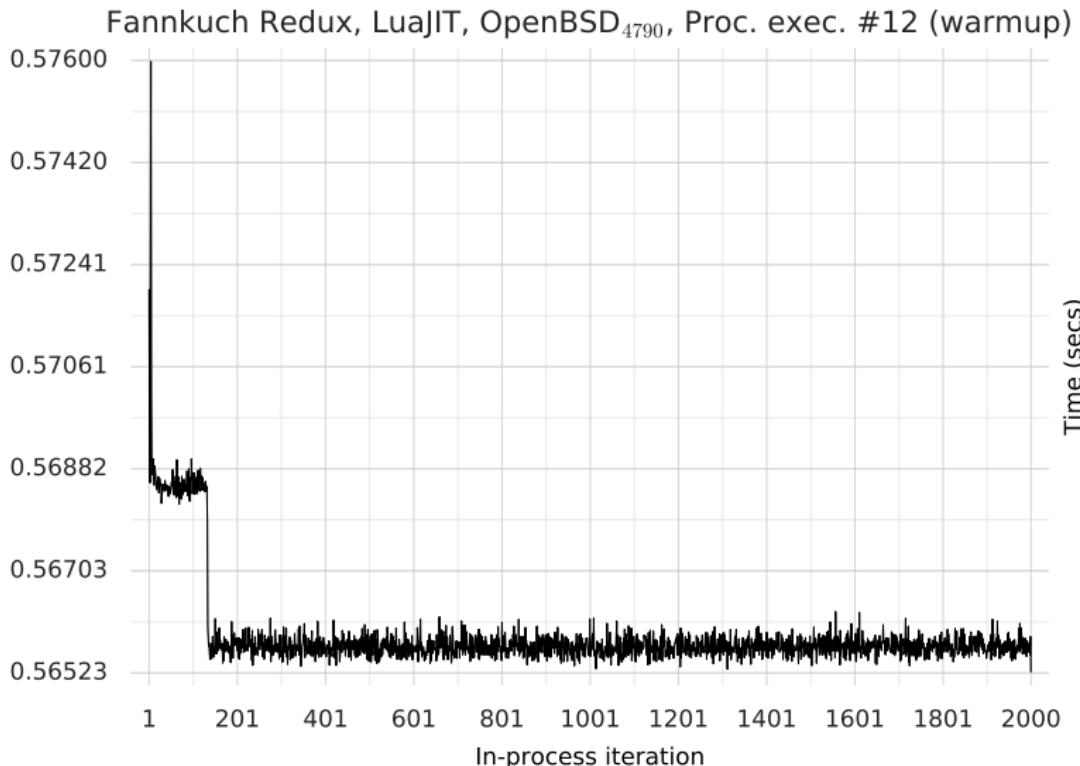
Benchmark runner: tries to control as many confounding variables as possible

## Method 5: Krun

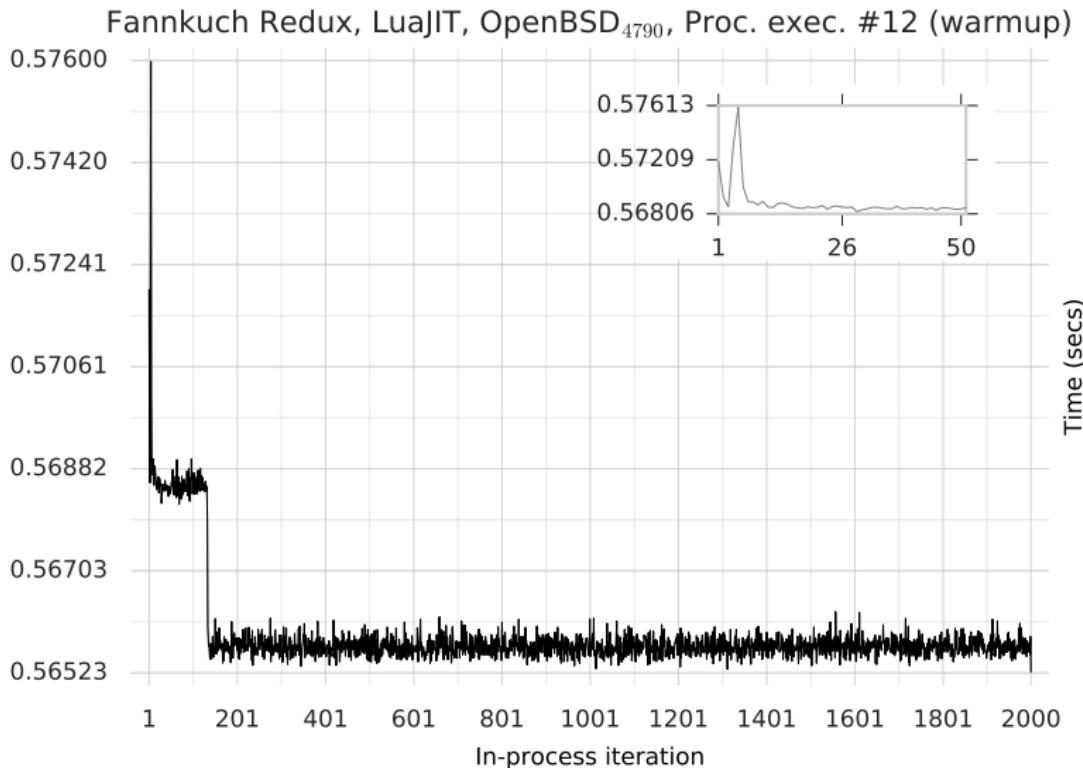
Benchmark runner: tries to control as many confounding variables as possible e.g.:

- Minimises I/O
- Sets fixed heap and stack ulimits
- Drops privileges to a 'clean' user account
- Automatically reboots the system prior to each proc. exec
- Checks `dmesg` for changes after each proc. exec
- Checks system at (roughly) same temperature for proc. execs
- Enforces kernel settings (tickless mode, CPU governors, ...)

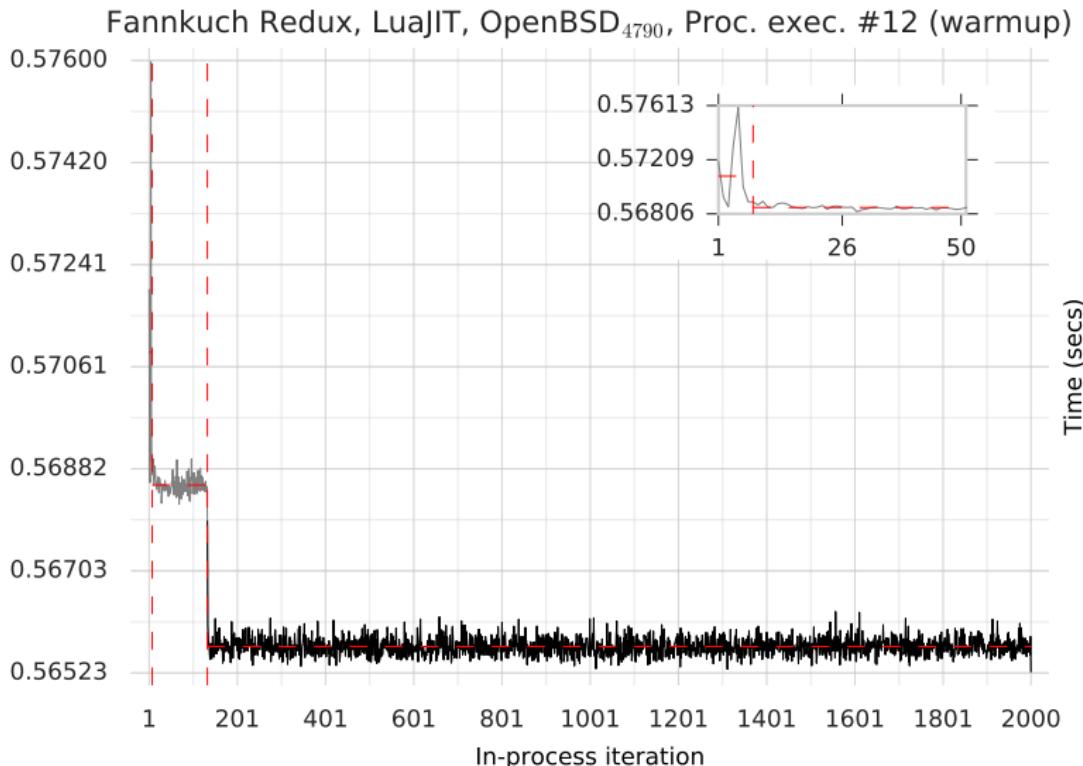
# Warmup & flat (1)



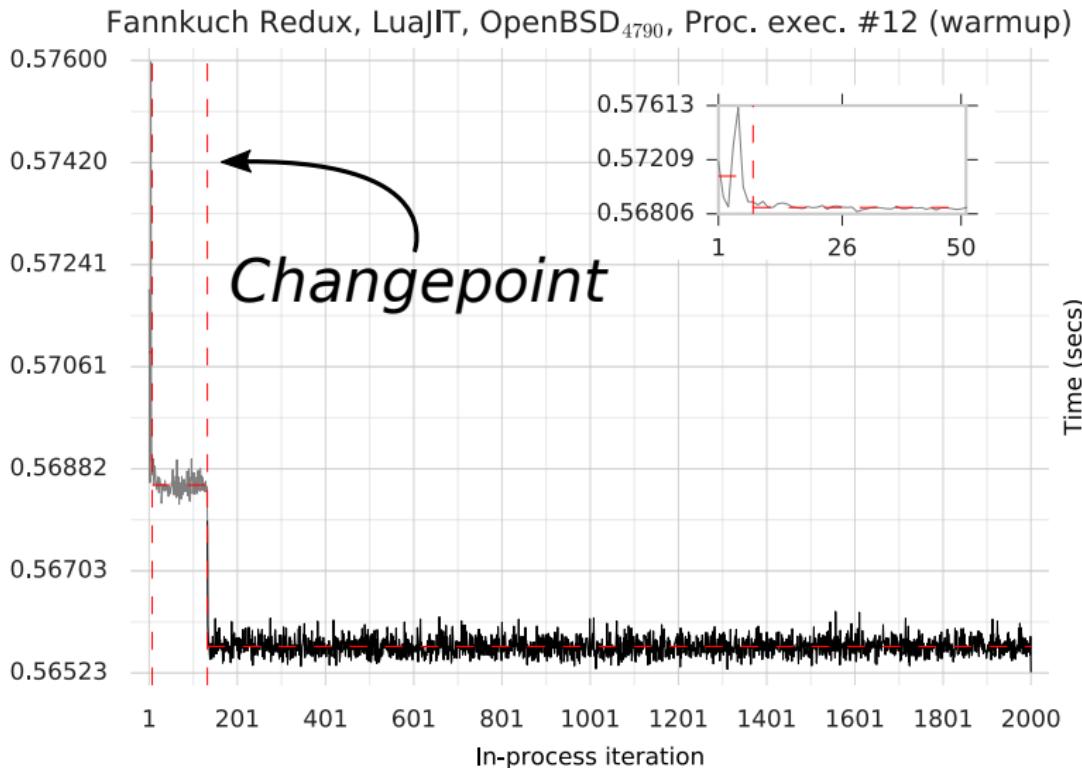
# Warmup & flat (1)



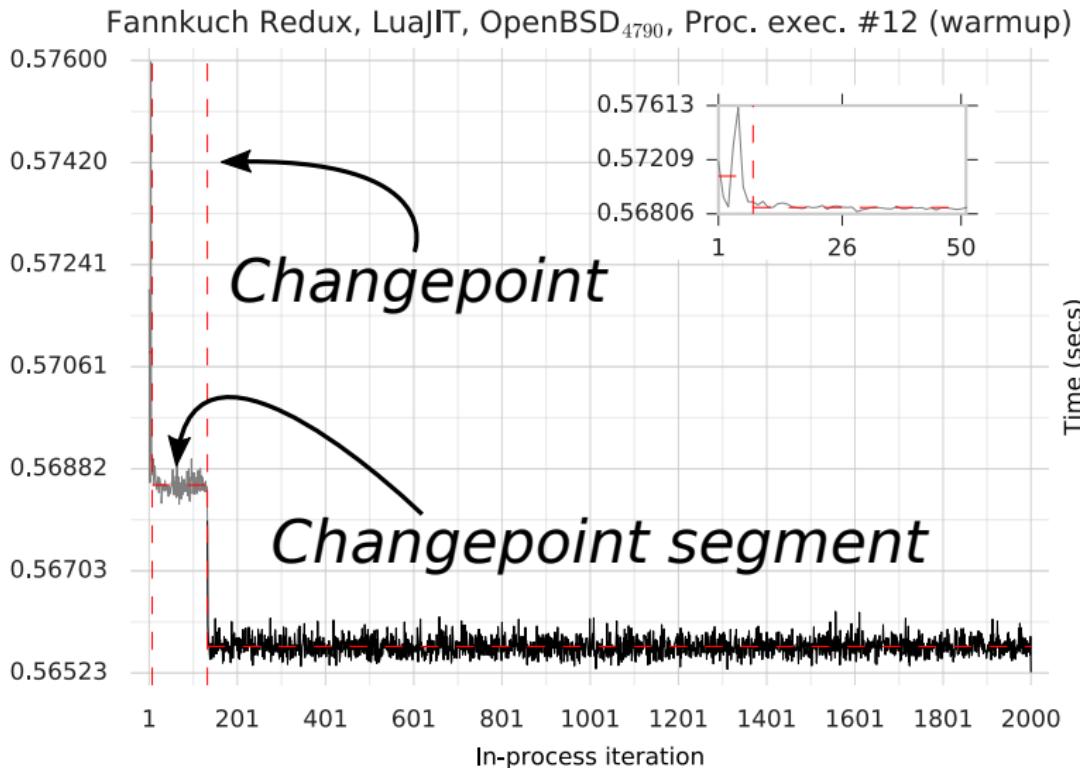
# Warmup & flat (1)



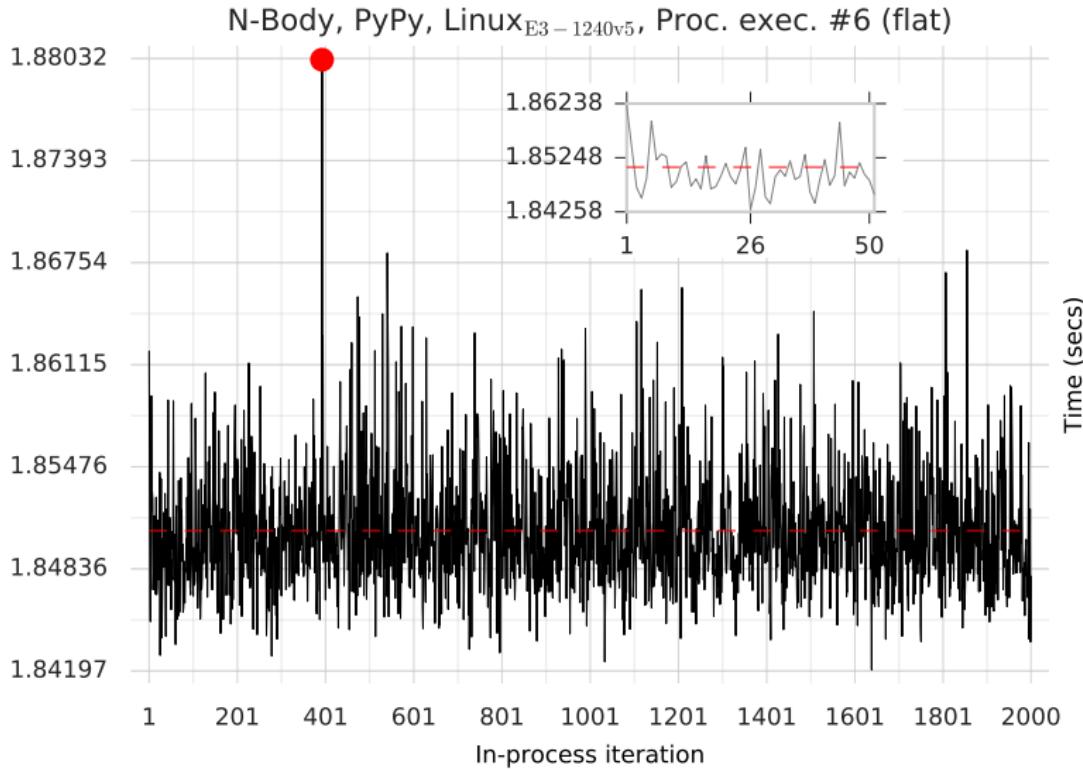
# Warmup & flat (1)



# Warmup & flat (1)



# Warmup & flat (1)



## Method 7: Classification

Classification algorithm (steps in order):

All segs are equivalent: *flat*

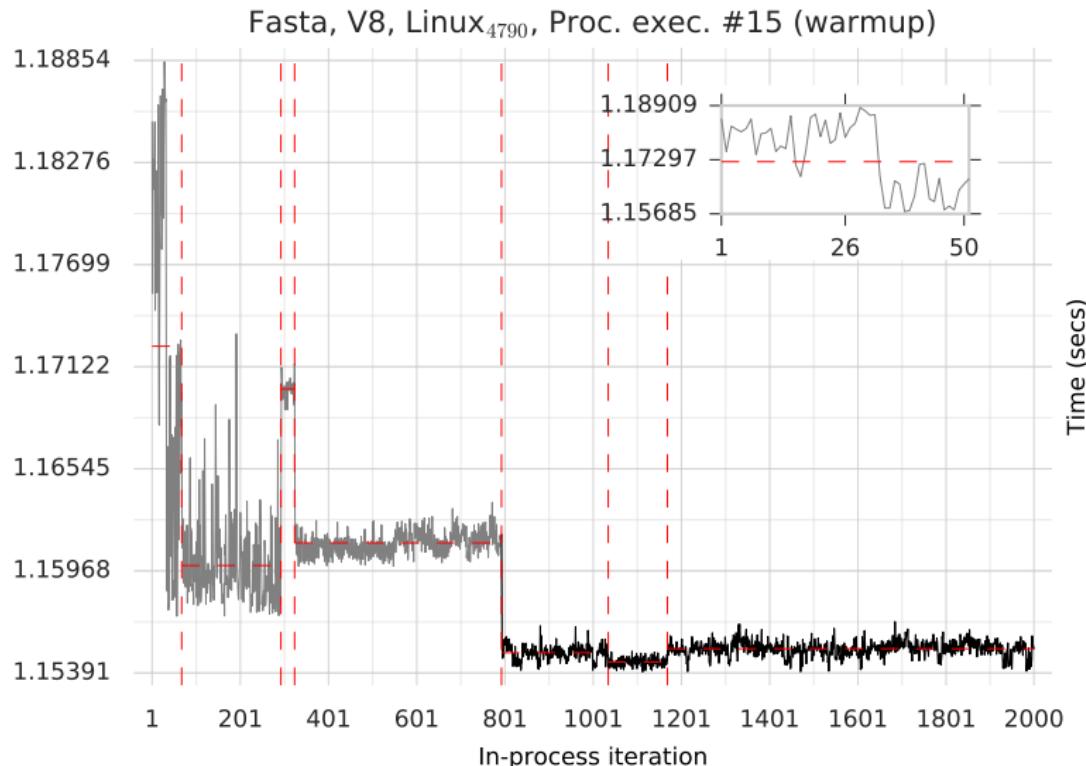
## Method 7: Classification

Classification algorithm (steps in order):

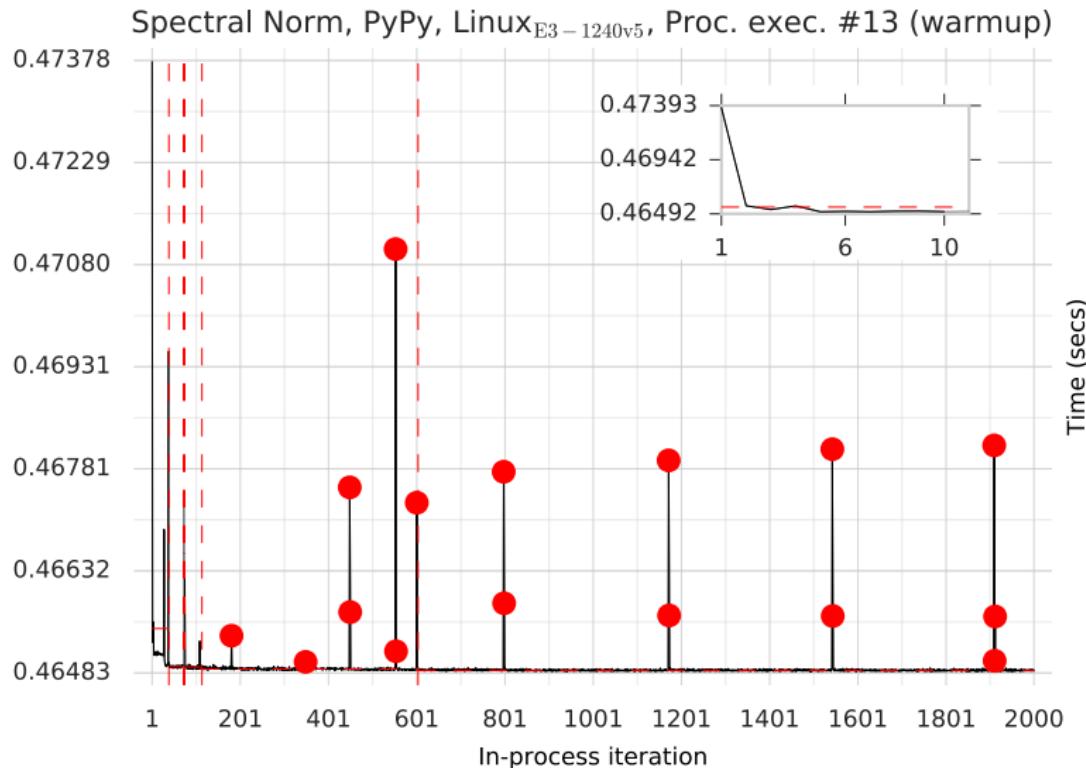
All segs are equivalent: *flat*

Final seg is in fastest set: *warmup*

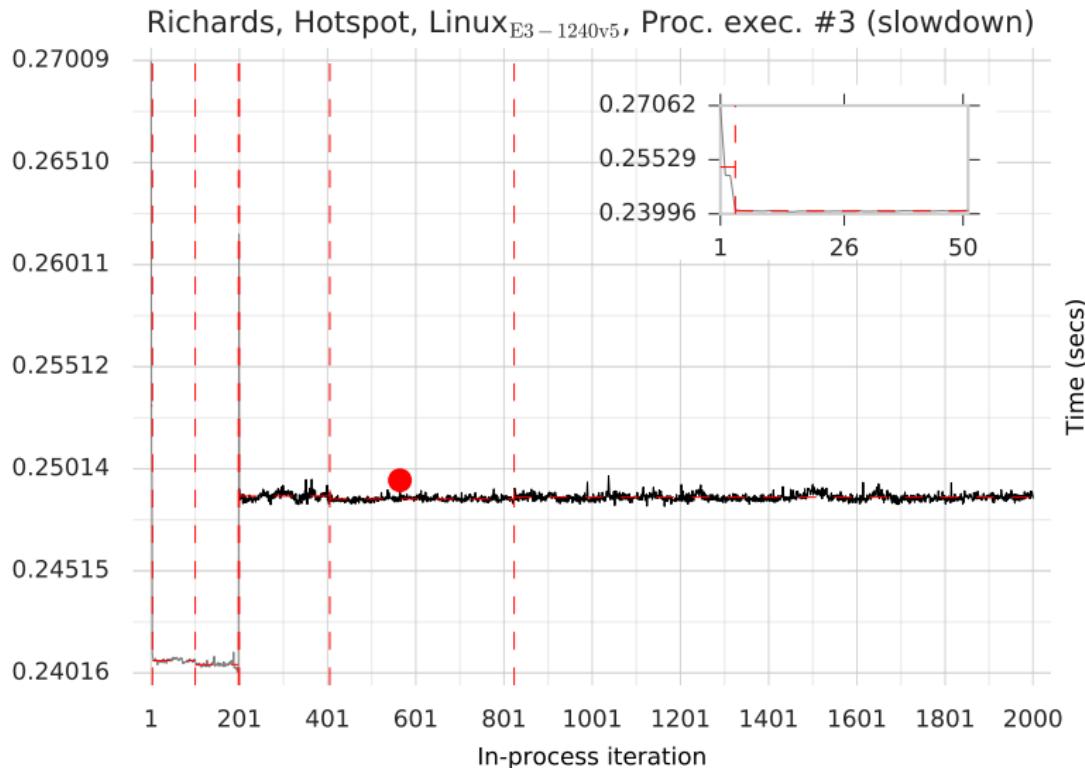
# Warmup & flat (2)



# Warmup & flat (2)



# Slowdown (1)



## Method 7: Classification

Classification algorithm (steps in order):

All segs are equivalent: *flat*

Final seg is in fastest set: *warmup*

## Method 7: Classification

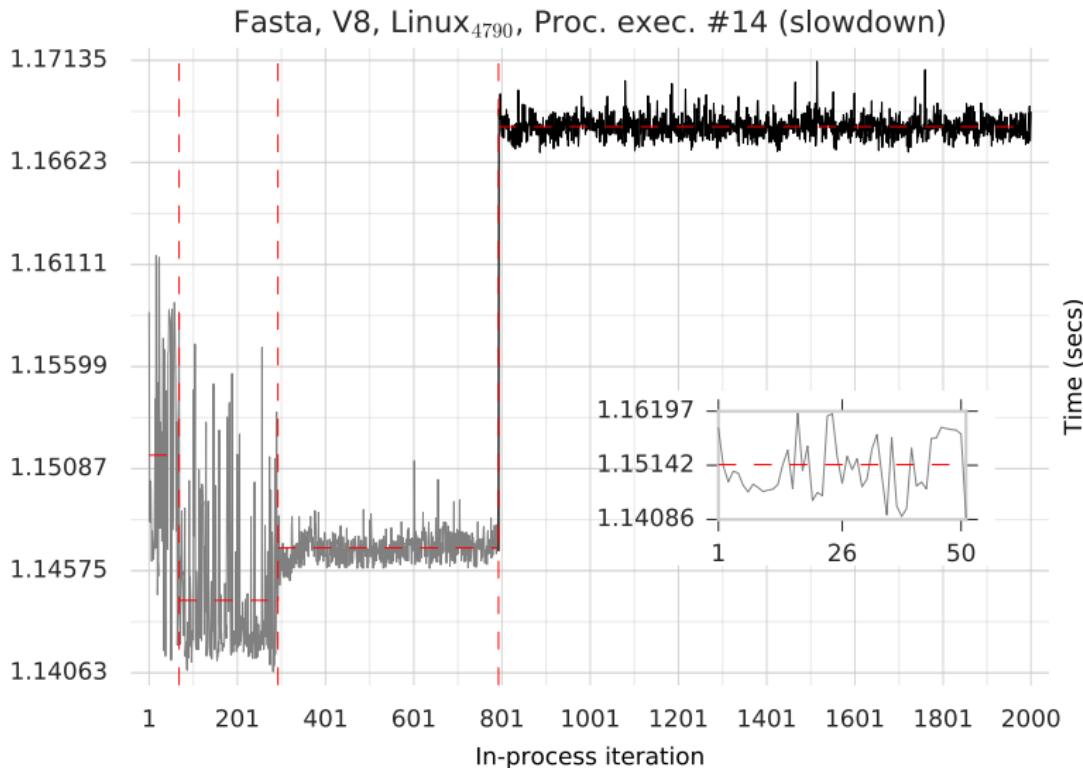
Classification algorithm (steps in order):

All segs are equivalent: *flat*

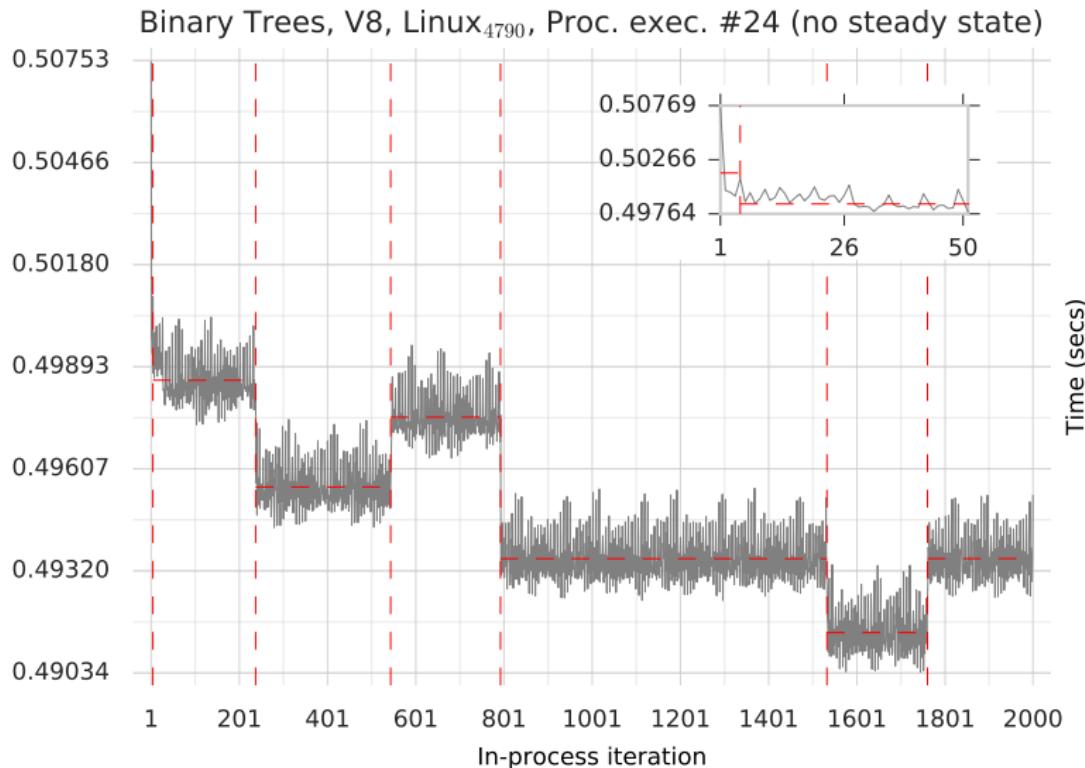
Final seg is in fastest set: *warmup*

Final seg is not in fastest set: *slowdown*

# Slowdown (2)



# No steady state (1)



## Classification (3)

Classification algorithm (steps in order):

All segs are equivalent: *flat*

Final seg is in fastest set: *warmup*

Final seg is not in fastest set: *slowdown*

## Classification (3)

Classification algorithm (steps in order):

All segs are equivalent: *flat*

Final seg is in fastest set: *warmup*

Final seg is not in fastest set: *slowdown*

Else: *no steady state*

## Classification (3)

Classification algorithm, in order:

All segs are equivalent: *flat*

Final seg is in fastest set: *warmup*

Final seg is not in fastest set: *slowdown*

Else: *no steady state*

**Good**

## Classification (3)

Classification algorithm, in order:

All segs are equivalent: *flat*

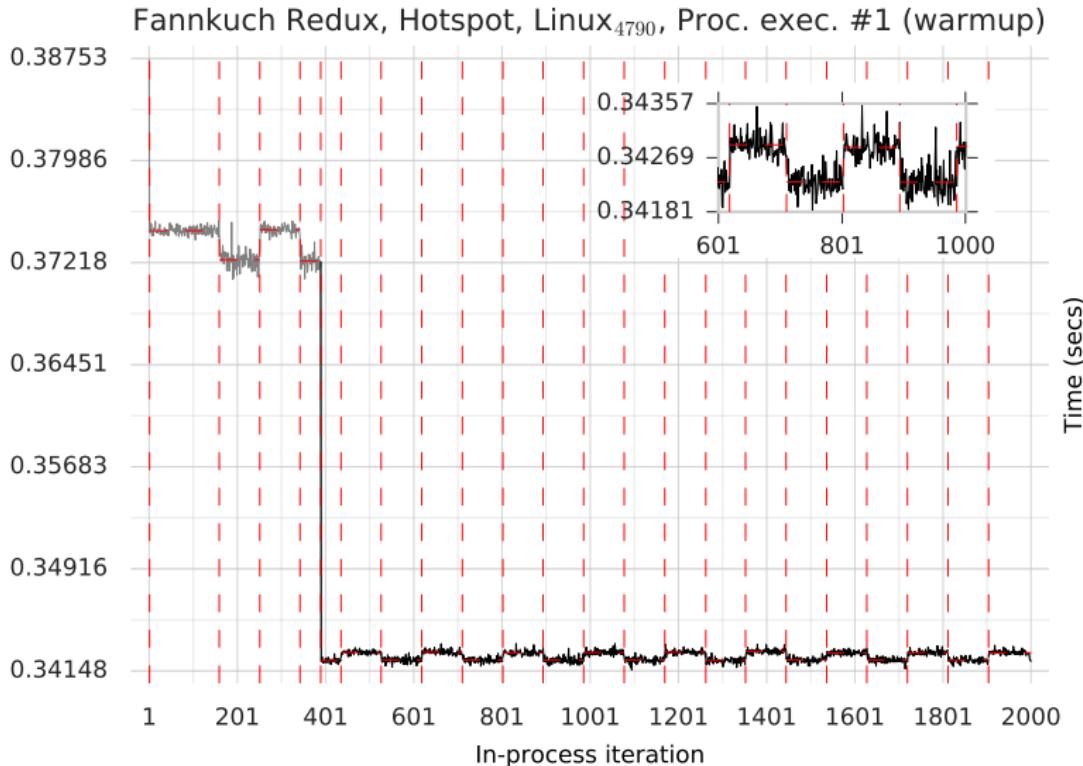
Final seg is in fastest set: *warmup*

Final seg is not in fastest set: *slowdown*

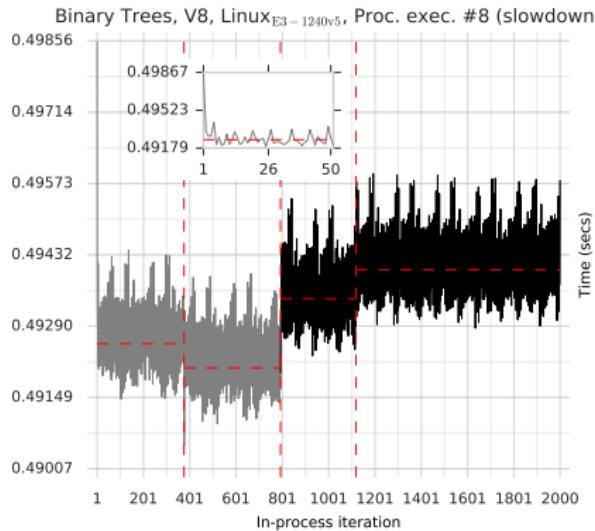
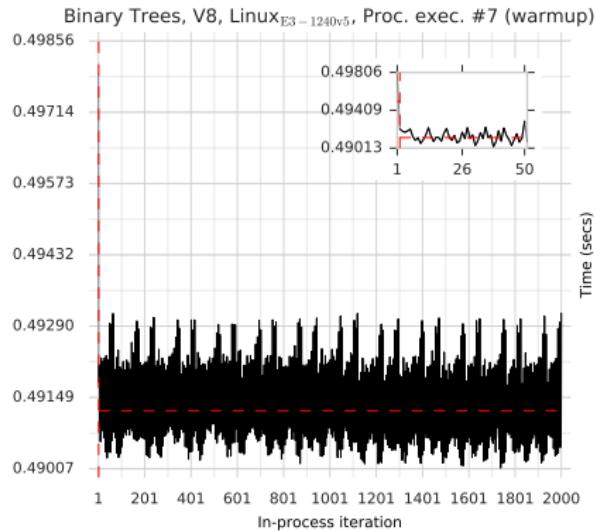
Else: *no steady state*

**Bad**

# Warmup or no steady state?

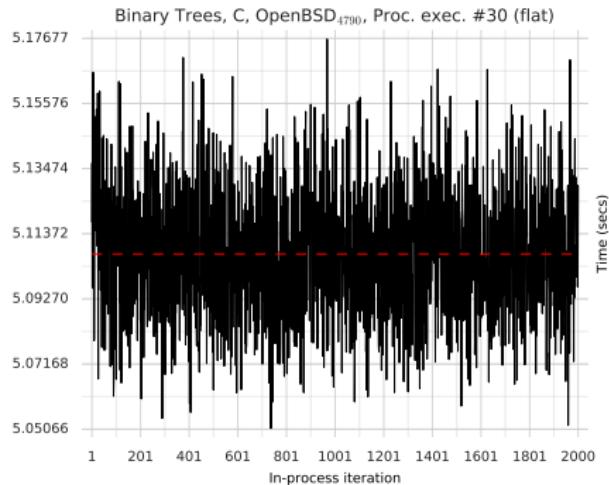
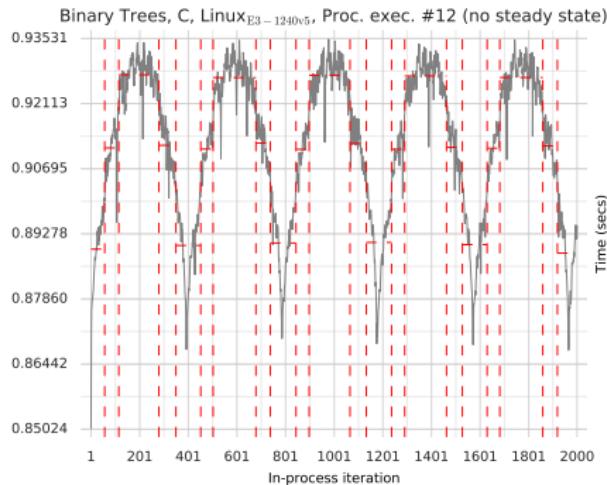


# Inconsistent Process-executions



(Same machine)

# Inconsistent Process-executions



(Different machines. Bouncing ball Linux-specific)

# Individual benchmark stats

	Class.	Steady iter (#)	Steady iter (s)	Steady perf (s)		Class.	Steady iter (#)	Steady iter (s)	Steady perf (s)
C	**	32.0	6.60	0.18594 ±0.00935		-			0.40555 ±0.00510
Graal	✗ (27L, 3F)	(17.0, 193.8)		(3.729, 36.698)		L	8.0	(7.5, 8.5)	1.22 ±0.00045
HHVM	✗ (24L, 4F, 2*)					✗ (16L, 11F, 3*)			0.13334 ±0.00045
HotSpot	✗ (25L, 5F)	7.0	1.19	0.18279 ±0.00116		L	2.0	(2.0, 2.0)	0.13699 ±0.00032
JRuby+Truffle	L	(7.0, 53.5)	(1,182, 9,703)			L	69.0	(69.0, 70.0)	17.95 ±0.00568
LuaJIT	L	1082.0	2219.59	2.05150 ±0.01738		n-body		(17.716, 18.127)	0.20644 ±0.00447
PyPy	✗ (23L, 4F, 2*, 1*)	(999.0, 1232.5)	(2039, 3042, 2516, 021)			-			0.25399 ±0.00447
V8	✗ (27L, 3*)					-			1.85835 ±0.01289
		1.5	0.25	0.49237 ±0.003198		= (25-, 5L)	1.0	(1.0, 361.6)	0.00 ±0.00389
C	✗ (21-, 6L, 2F, 1*)					✗ (19-, 5F, 4*, 2L)			
Graal	✗ (28L, 1*, 1F)					✗ (28L, 1*, 1F)			
HHVM	L	10.0	52.66	1.35779 ±0.01948					
HotSpot	L	390.0	153.70	0.36202 ±0.02767					
JRuby+Truffle	L	(2.0, 390.0)	(407, 155, 254)			✗ (26F, 2*, 2L)			
LuaJIT	L	1016.5	1039.04	1.08833 ±0.038580					
PyPy	✗ (15L, 13-, 2F)	(999.0, 1923.1)	(1014, 290, 1959, 967)			✗ (26F, 2*, 2L)	1021.0	(1014.9, 1027.0)	917.30 ±0.02790
V8	= (19L, 11-)	(1.0, 25.5)	(0.008, 7.525)	0.56285 ±0.00097					0.89509 ±0.02790
		2.0	1.57	1.55442 ±0.026549		✗ (21-, 7F, 2L)	2.0	(2.0, 5.0)	1.12 ±0.01090
		2.0	0.31	0.30401 ±0.000154		L	2.0	(1.114, 4.654)	1.12 ±0.01090
		2.0	0.31	0.30401 ±0.000154		✗ (21-, 7F, 2L)	4.0	(4.0, 16.0)	1.44 ±0.00218
C	-			0.07048 ±0.00020		✗ (28L, 1-, 1F)	997.0	(2.0, 1001.6)	546.38 ±0.00562
Graal	✗ (29L, 1*)					✗ (29F, 1-)	15.0	(0.546, 547.317)	0.89293 ±0.00087
HHVM	✗ (27L, 2-, 1F)					L	35.0	(0.812, 17.404)	12.40 ±0.00087
HotSpot	✗ (18L, 12F)	261.0	30.73	0.11744 ±0.001723		L	35.0	(34.0, 41.0)	139.18 ±0.01113
JRuby+Truffle	L	(6.0, 595.0)	(6.614, 70.021)			L	7.0	(7.0, 8.4)	0.10690 ±0.00029
LuaJIT	**					L	1011.0	(1.901, 2.025)	0.31470 ±0.00029
PyPy	**					L	1007.5	(1007.5, 1014.0)	893.38 ±0.01403
V8	✗ (19L, 10F, 1*)					L	75.0	(888.985, 896.562)	0.83633 ±0.01403
									0.22435 ±0.00020
									0.46489 ±0.00046
									0.24963 ±0.00039
									0.00046 ±0.00039

# Individual benchmark stats

		Class.	Steady iter (#)	Steady iter (s)	Steady perf (s)
C		w			
Graal		⌘ (27L, 3J)	32.0 (17.0, 193.8)	6.60 (3.729, 36.608)	0.18594 ±0.000316
HHVM		⌘ (24L, 4J, 2w)			
HotSpot	binary trees	⌘ (25L, 5J)	7.0 (7.0, 53.5)	1.19 (1.182, 9.703)	0.18279 ±0.000116
JRuby+Truffle		J	1082.0 (999.0, 1232.5)	2219.59 (2039.304, 2516.021)	2.05150 ±0.017737
LuaJIT		⌘ (23L, 4J, 2-, 1w)			
PyPy		⌘ (27J, 3w)			
V8		⌘ (15-, 9L, 6J)	1.5 (1.0, 794.0)	0.25 (0.000, 391.026)	0.49237 ±0.003198

# Overall benchmark stats

Class.	Linux <sub>4790</sub>	Linux <sub>1240v5</sub>	OpenBSD <sub>4790</sub> <sup>†</sup>
⟨VM, benchmark⟩ pairs			
—	8.7%	13.0%	6.7%
⊓	28.3%	23.9%	10.0%
⊓	6.5%	6.5%	0.0%
⊜	4.3%	6.5%	0.0%
=	6.5%	6.5%	13.3%
✗	45.7%	43.5%	70.0%
Process executions			
—	26.4%	20.9%	34.0%
⊓	48.3%	51.5%	52.1%
⊓	16.7%	17.9%	11.1%
⊜	8.7%	9.6%	2.8%

# Overall benchmark stats

Class.	Linux <sub>4790</sub>	Linux <sub>1240v5</sub>	OpenBSD <sub>4790</sub> <sup>†</sup>
⟨VM, benchmark⟩ pairs			
—	8.7%	13.0%	6.7%
⊓	28.3%	23.9%	10.0%
⊓	6.5%	6.5%	0.0%
⊜	4.3%	6.5%	0.0%
=	6.5%	6.5%	13.3%
✗	45.7%	43.5%	70.0%
Process executions			
—	26.4%	20.9%	34.0%
⊓	48.3%	51.5%	52.1%
⊓	16.7%	17.9%	11.1%
⊜	8.7%	9.6%	2.8%

# Summary

Classical warmup occurs for only:

# Summary

Classical warmup occurs for only:

72.4%-74.7% of process executions

## Summary

Classical warmup occurs for only:

72.4%-74.7% of process executions

43.4%-43.5% of (VM, benchmark) pairs

## Summary

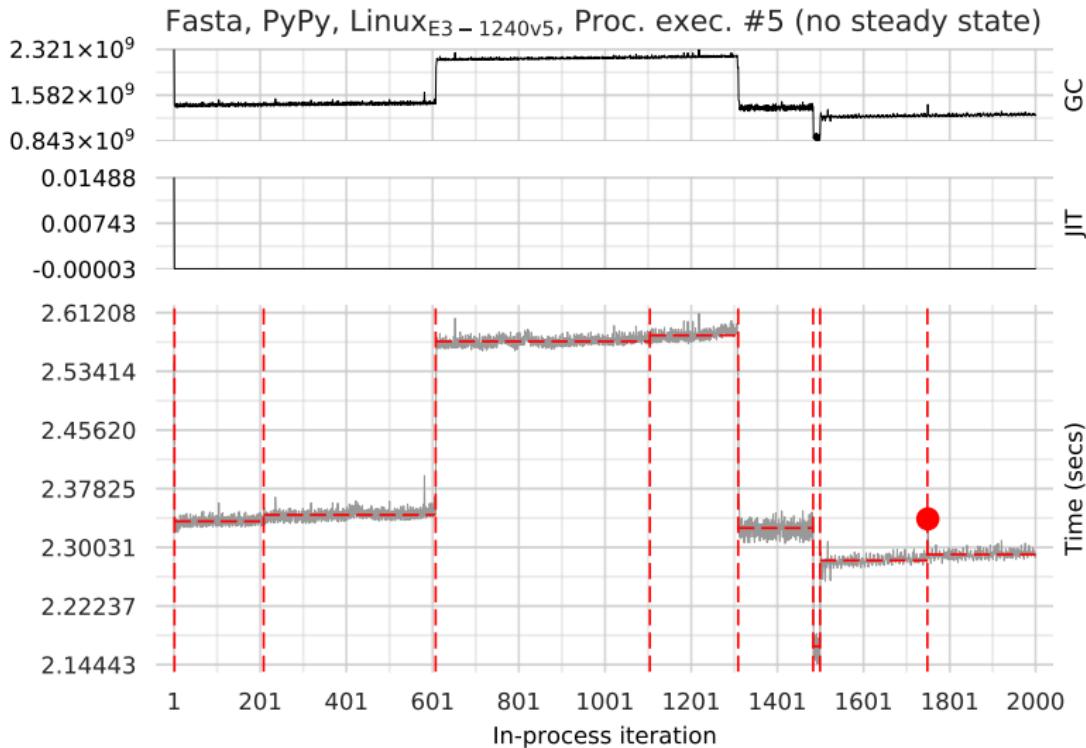
Classical warmup occurs for only:

72.4%-74.7% of process executions

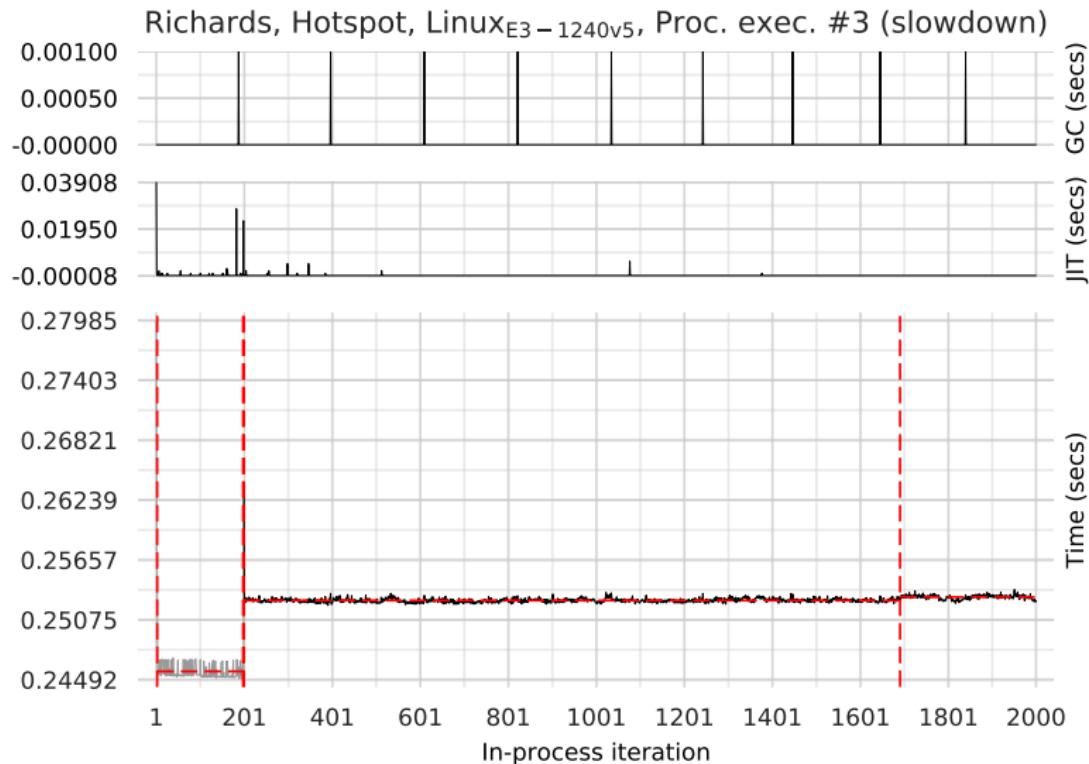
43.4%-43.5% of (VM, benchmark) pairs

0% of benchmarks for (VM, benchmark, machine) triples

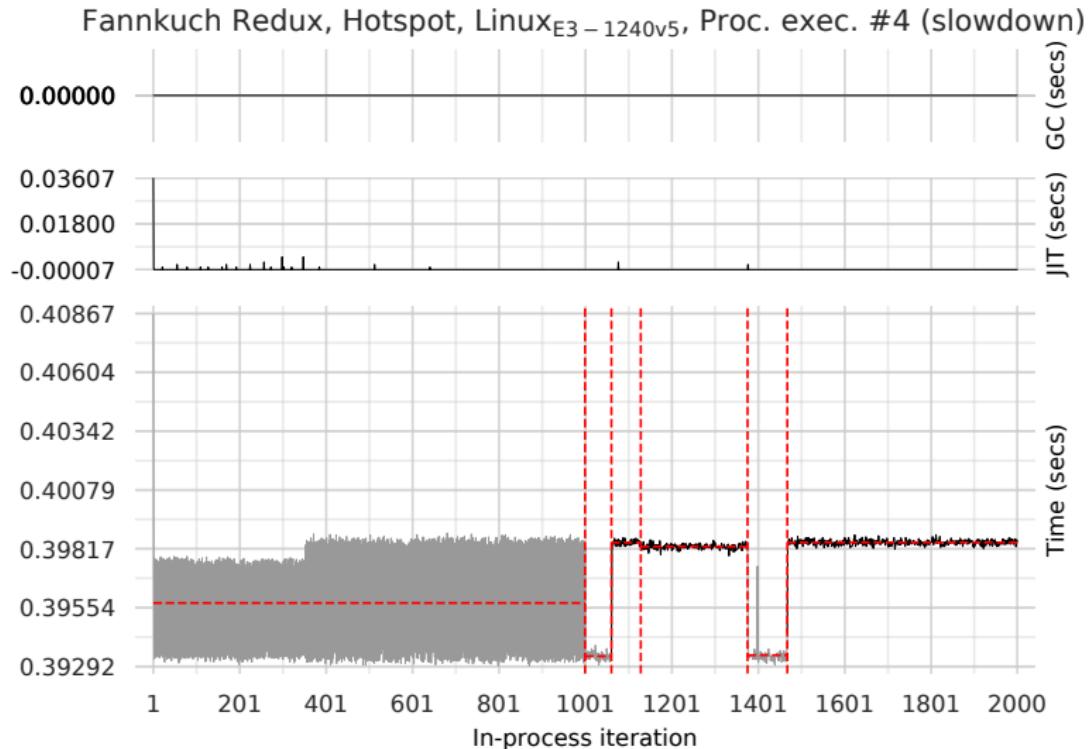
# Are odd effects correlated with compilation and GC?



# Are odd effects correlated with compilation and GC?



# Are odd effects correlated with compilation and GC?



# Benchmark suites

## Benchmark suites

Benchmarks guide our optimisations

## Benchmark suites

Benchmarks guide our optimisations

Are they complete guides?

# A war story

# A war story

Symptom: poor performance of a Pyston  
benchmark on PyPy

# A war story

Symptom: poor performance of a Pyston  
benchmark on PyPy

Cause: RPython traces recursion

# A war story

Symptom: poor performance of a Pyston  
benchmark on PyPy

Cause: RPython traces recursion

Fix: Check for recursion before tracing

# A war story: the basis of a fix

```
diff --git a/rpython/jit/metainterp/pyjitpl.py b/rpython/jit/metainterp/pyjitpl.py
--- a/rpython/jit/metainterp/pyjitpl.py
+++ b/rpython/jit/metainterp/pyjitpl.py
@@ -951,9 +951,31 @@ 
     if warmrunnerstate.inlining:
         if warmrunnerstate.can_inline_callable(greenboxes):
+            # We've found a potentially inlinable function; now we need to
+            # see if it's already on the stack. In other words: are we about
+            # to enter recursion? If so, we don't want to inline the
+            # recursion, which would be equivalent to unrolling a while
+            # loop.
             portal_code = targetjitdriver_sd.mainjitcode
-            return self.metainterp.perform_call(portal_code, allboxes,
-                                                greenkey=greenboxes)
+            inline = True
+            if self.metainterp.is_main_jitcode(portal_code):
+                for gk, _ in self.metainterp.portal_trace_positions:
+                    if gk is None:
+                        continue
+                    assert len(gk) == len(greenboxes)
+                    i = 0
+                    for i in range(len(gk)):
+                        if not gk[i].same_constant(greenboxes[i]):
+                            break
+                    else:
+                        # The greenkey of a trace position on the stack
+                        # matches what we have, which means we're definitely
+                        # about to recurse.
+                        inline = False
+                        break
+            if inline:
+                return self.metainterp.perform_call(portal_code, allboxes,
+                                                    greenkey=greenboxes)
```

# A war story: mixed fortunes

Success: slow benchmark now 13.5x faster

# A war story: mixed fortunes

Success: slow benchmark now 13.5x faster

Failure: some PyPy benchmarks slow down

## A war story: mixed fortunes

Success: slow benchmark now 13.5x faster

Failure: some PyPy benchmarks slow down

Solution: allow *some* tracing into recursion

# A war story: data

#unrollings	1	2	3	5	7	10	
hexiom2	1.3	1.4	1.1	1.0	1.0	1.0	
raytrace-simple	3.3	3.1	2.8	1.4	1.0	1.0	
spectral-norm	3.3	1.0	1.0	1.0	1.0	1.0	
sympy_str	1.5	1.0	1.0	1.0	1.0	1.0	
telco	4	2.5	2.0	1.0	1.0	1.0	
polymorphism	0.07	0.07	0.07	0.07	0.08	0.09	

<http://marc.info/?l=pypy-dev&m=141587744128967&w=2>

## A war story: conclusion

The benchmark suite said 7 levels, so that's what I suggested

## A war story: conclusion

The benchmark suite said 7 levels, so that's what I suggested

*Even though I doubted it was the right global value*

# Benchmark suites (2)

## Benchmark suites (2)

Benchmarks guide our optimisations

Benchmarks guide our optimisations

Are they correct guides?

## 17 JavaScript benchmarks from V8

17 JavaScript benchmarks from V8

Let's make each benchmark run for 2000 iterations

# Octane: pdf.js explodes

```
$ d8 run.js
Richards
DeltaBlue
Encrypt
Decrypt
RayTrace
Earley
Boyer
RegExp
Splay
NavierStokes
PdfJS
```

```
<--- Last few GCs --->
```

```
14907865 ms: Mark-sweep 1093.9 (1434.4) -> 1093.4 (1434.4) MB, 274.8 / 0.0 ms [allocation failure] [GC in old space
14908140 ms: Mark-sweep 1093.4 (1434.4) -> 1093.3 (1434.4) MB, 274.4 / 0.0 ms [allocation failure] [GC in old space
14908421 ms: Mark-sweep 1093.3 (1434.4) -> 1100.5 (1418.4) MB, 280.9 / 0.0 ms [last resort gc].
14908703 ms: Mark-sweep 1100.5 (1418.4) -> 1107.8 (1418.4) MB, 282.1 / 0.0 ms [last resort gc].
```

```
<--- JS stacktrace --->
```

```
===== JS stack trace =====
```

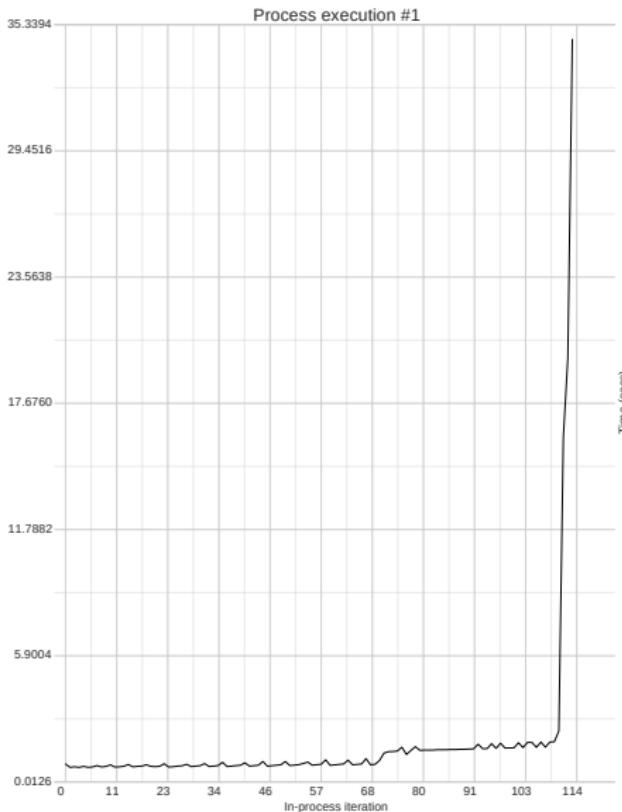
```
Security context: 0x20d333ad3ba9 <JS Object>
```

```
 2: extractFontProgram(aka Type1Parser_extractFontProgram) [pdfjs.js:17004] [pc=0x3a13b275421b] (this=0x3de358283
 3: new Type1Font [pdfjs.js:17216] [pc=0x3a13b2752078] (this=0x4603fbdae9 <a Type1Font with map 0x1f822134f7e1>,
```

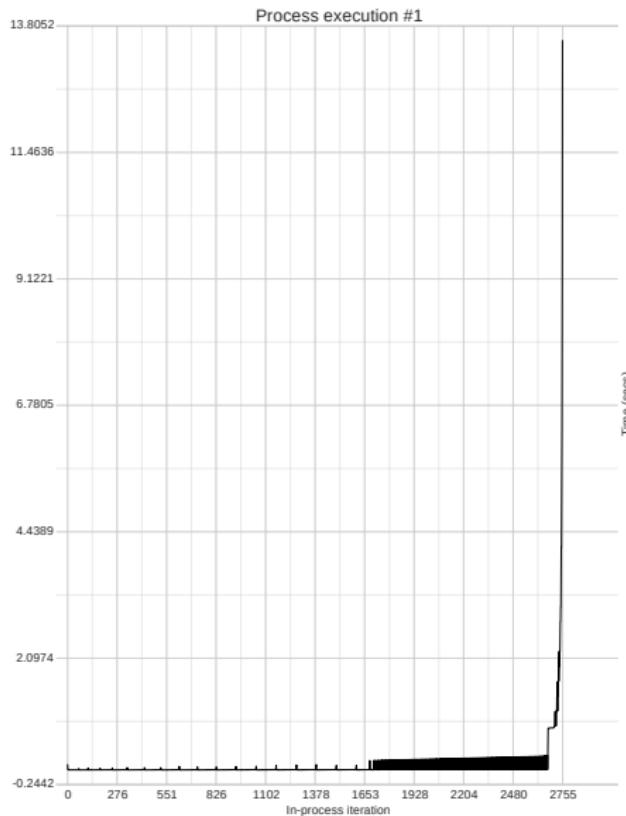
```
# 
# Fatal error in CALL_AND_RETRY_LAST
# Allocation failed - process out of memory
#
```

```
zsh: illegal hardware instruction  d8 run.js
```

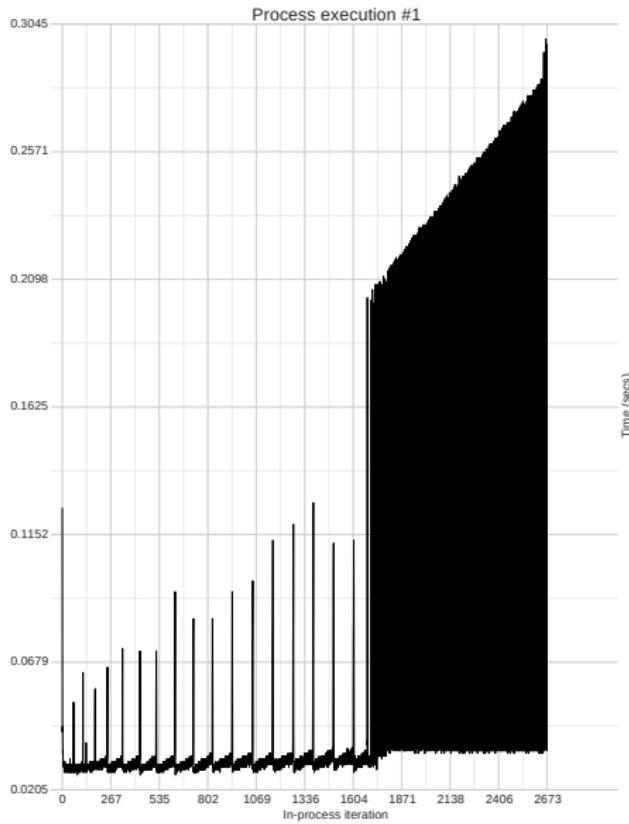
# Octane: analysing pdf.js



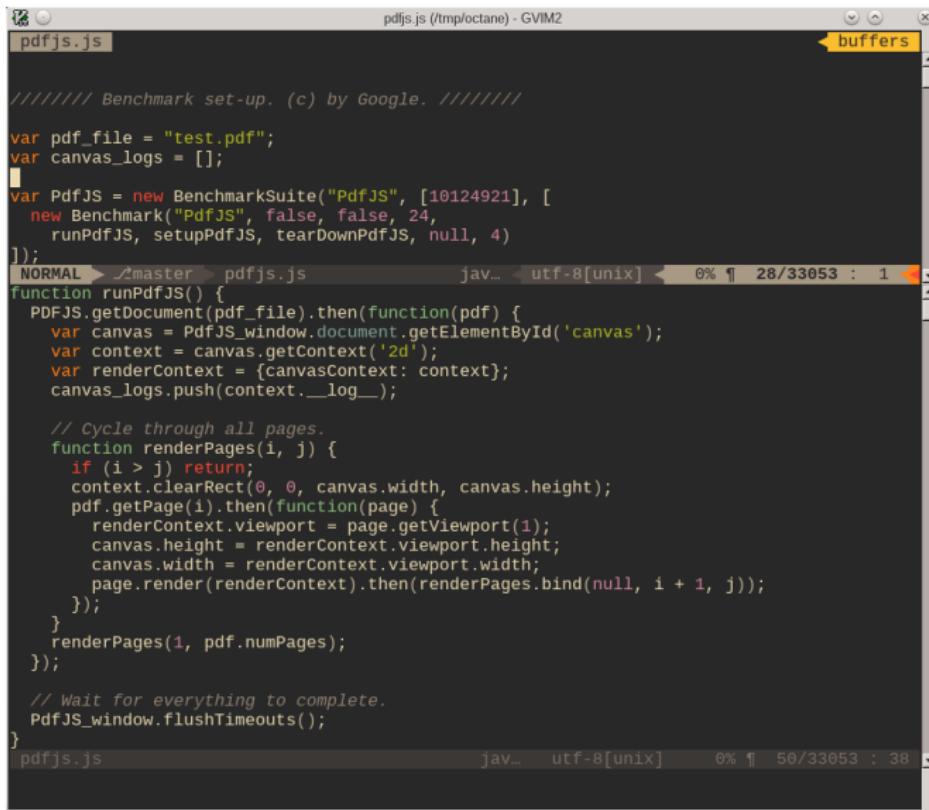
# Octane: analysing pdf.js



# Octane: analysing pdf.js



# Octane: debugging



The screenshot shows a GVIM window with two buffers. The left buffer contains the following JavaScript code:

```
//////// Benchmark set-up. (c) by Google. /////////
var pdf_file = "test.pdf";
var canvas_logs = [];

var PdfJS = new BenchmarkSuite("PdfJS", [10124921], [
  new Benchmark("PdfJS", false, false, 24,
    runPdfJS, setupPdfJS, tearDownPdfJS, null, 4)
]);
function runPdfJS() {
  PDFJS.getDocument(pdf_file).then(function(pdf) {
    var canvas = PdfJS_window.document.getElementById('canvas');
    var context = canvas.getContext('2d');
    var renderContext = {canvasContext: context};
    canvas_logs.push(context.__log__);

    // Cycle through all pages.
    function renderPages(i, j) {
      if (i > j) return;
      context.clearRect(0, 0, canvas.width, canvas.height);
      pdf.getPage(i).then(function(page) {
        renderContext.viewport = page.getViewport(1);
        canvas.height = renderContext.viewport.height;
        canvas.width = renderContext.viewport.width;
        page.render(renderContext).then(renderPages.bind(null, i + 1, j));
      });
    }
    renderPages(1, pdf.numPages);
  });

  // Wait for everything to complete.
  PdfJS_window.flushTimeouts();
}
pdfjs.js
```

The right buffer is empty and labeled "buffers". The status bar at the bottom indicates the file is "pdfjs.js" and the line count is 38.

# Octane: fixing

The screenshot shows a GitHub pull request page for issue #42. The title is "Fix memory leak in pdfjs.js. #42". A green "Open" button indicates the merge status. The commit message is "ltratt wants to merge 2 commits into chromium:master from ltratt:master". Below the title, there are three tabs: "Conversation 5", "Commits 2", and "Files changed 1". The "Files changed" tab is selected, showing a diff for "pdfjs.js". The diff highlights a new line of code: "canvas\_logs.length = 0;".

```
1 pdfjs.js
@@ -43,6 +43,7 @@ function setupPdfJS() {
 43     }
 44
 45     function runPdfJS() {
+    canvas_logs.length = 0;
 46     PDFJS.getDocument(pdf_file).then(function(pdf) {
 47         var canvas = PdfJS_window.document.getElementById('canvas');
 48         var context = canvas.getContext('2d');
```

## Octane: other issues

pdfjs isn't the only problem

## Octane: other issues

pdfjs isn't the only problem

CodeLoadClosure also has a memory leak

## Octane: other issues

pdfjs isn't the only problem

CodeLoadClosure also has a memory leak

zlib complains that Cannot enlarge memory arrays in asm.js (a memory leak? I don't know)

## Octane: other issues

pdfjs isn't the only problem

CodeLoadClosure also has a memory leak

zlib complains that Cannot enlarge memory arrays in asm.js (a memory leak? I don't know)

Timings are made with a non-monotonic microsecond timer

# Summary

# Summary

Why aren't more users more happy with our VMs?

## Summary

Why aren't more users more happy with our VMs?

My thesis: benchmarking *and* benchmarks are performance destiny.

## Summary

Why aren't more users more happy with our VMs?

My thesis: benchmarking *and* benchmarks are performance destiny.

Ours have misled us.

# How to benchmark a bit better

# How to benchmark a bit better

- 1 Run benchmarks for longer to uncover issues.

## How to benchmark a bit better

- 1 Run benchmarks for longer to uncover issues.
- 2 Accept that peak performance may not occur.

## How to benchmark a bit better

- 1 Run benchmarks for longer to uncover issues.
- 2 Accept that peak performance may not occur.
- 3 Always report warmup time.

## How to benchmark a bit better

- 1 Run benchmarks for longer to uncover issues.
- 2 Accept that peak performance may not occur.
- 3 Always report warmup time.
- 4 Stop over-training on small benchmark suites.

## How to benchmark a bit better

- 1 Run benchmarks for longer to uncover issues.
- 2 Accept that peak performance may not occur.
- 3 Always report warmup time.
- 4 Stop over-training on small benchmark suites.
- 5 Collect more benchmarks.

## How to benchmark a bit better

- 1 Run benchmarks for longer to uncover issues.
- 2 Accept that peak performance may not occur.
- 3 Always report warmup time.
- 4 Stop over-training on small benchmark suites.
- 5 Collect more benchmarks.
- 6 Focus on predictable performance.

# The big question

# The big question

Can we fix existing VMs?

# The big question

Can we fix existing VMs?

At least a bit... but a lot? Unclear.

# The big question

Can we fix existing VMs?

At least a bit... but a lot? Unclear.

In case we can't, I have an idea...

# Meta-tracing JITs

## FL Interpreter

```
program_counter = 0; stack = []
vars = {...}
while True:
    jit_merge_point(program_counter)
    instr = load_instruction(program_counter)
    if instr == INSTR_VAR_GET:
        stack.push(
            vars[read_var_name_from_instruction()])
        program_counter += 1
    elif instr == INSTR_VAR_SET:
        vars[read_var_name_from_instruction()] =
            stack.pop()
        program_counter += 1
    elif instr == INSTR_INT:
        stack.push(read_int_from_instruction())
        program_counter += 1
    elif instr == INSTR_LESS_THAN:
        rhs = stack.pop()
        lhs = stack.pop()
        if isinstance(lhs, int) and isinstance(rhs, int):
            if lhs < rhs:
                stack.push(True)
            else:
                stack.push(False)
        else: ...
        program_counter += 1
```

```
elif instr == INSTR_IF:
    result = stack.pop()
    if result == True:
        program_counter += 1
    else:
        program_counter +=
            read_jump_if_instruction()
elif instr == INSTR_ADD:
    lhs = stack.pop()
    rhs = stack.pop()
    if isinstance(lhs, int) and isinstance(rhs, int):
        stack.push(lhs + rhs)
    else: ...
    program_counter += 1
```

# Meta-tracing JITs

---

## *FL* Interpreter

---

```
program_counter = 0; stack = []
vars = {...}
while True:
    jit_merge_point(program_counter)
    instr = load_instruction(program_counter)
    if instr == INSTR_VAR_GET:
        stack.push(
            vars[read_var_name_from_instruction()])
        program_counter += 1
    elif instr == INSTR_VAR_SET:
        vars[read_var_name_from_instruction()] =
            stack.pop()
        program_counter += 1
    elif instr == INSTR_INT:
        stack.push(read_int_from_instruction())
        program_counter += 1
    elif instr == INSTR_LESS_THAN:
        rhs = stack.pop()
        lhs = stack.pop()
        if isinstance(lhs, int) and isinstance(rhs, int):
            if lhs < rhs:
                stack.push(True)
            else:
                stack.push(False)
        else: ...
        program_counter += 1
```

---

# Meta-tracing JITs

## *FL* Interpreter

```
program_counter = 0; stack = []
vars = {...}
while True:
    jit_merge_point(program_counter)
    instr = load_instruction(program_counter)
    if instr == INSTR_VAR_GET:
        stack.push(
            vars[read_var_name_from_instruction()])
        program_counter += 1
    elif instr == INSTR_VAR_SET:
        vars[read_var_name_from_instruction()] =
            stack.pop()
        program_counter += 1
    elif instr == INSTR_INT:
        stack.push(read_int_from_instruction())
        program_counter += 1
    elif instr == INSTR_LESS_THAN:
        rhs = stack.pop()
        lhs = stack.pop()
        if isinstance(lhs, int) and isinstance(rhs, int):
            if lhs < rhs:
                stack.push(True)
            else:
                stack.push(False)
        else: ...
        program_counter += 1
```

## User program (lang *FL*)

```
if x < 0:
    x = x + 1
else:
    x = x + 2
x = x + 3
```

# Meta-tracing JITs

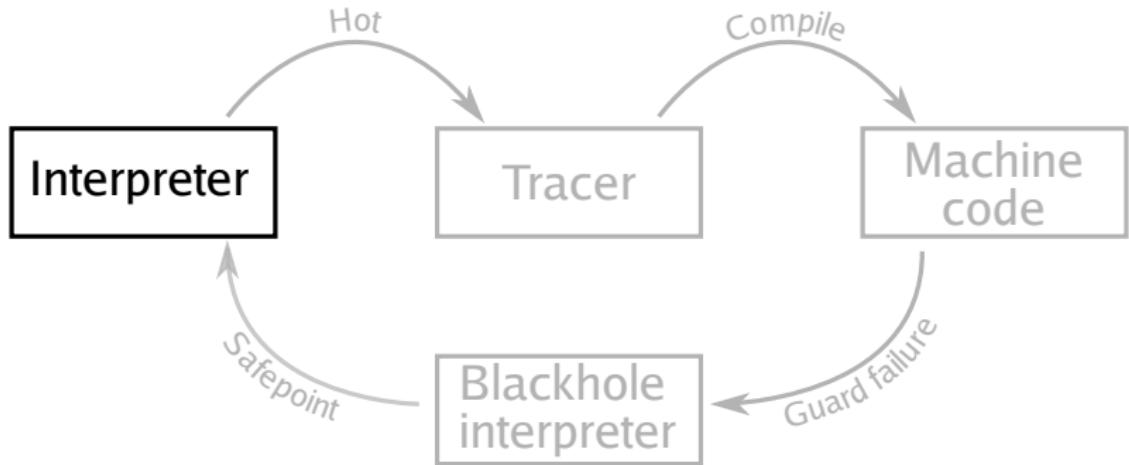
## FL Interpreter

```
program_counter = 0; stack = []
vars = {...}
while True:
    jit_merge_point(program_counter)
    instr = load_instruction(program_counter)
    if instr == INSTR_VAR_GET:
        stack.push(
            vars[read_var_name_from_instruction()])
        program_counter += 1
    elif instr == INSTR_VAR_SET:
        vars[read_var_name_from_instruction()] =
            stack.pop()
        program_counter += 1
    elif instr == INSTR_INT:
        stack.push(read_int_from_instruction())
        program_counter += 1
    elif instr == INSTR_LESS_THAN:
        rhs = stack.pop()
        lhs = stack.pop()
        if isinstance(lhs, int) and isinstance(rhs, int):
            if lhs < rhs:
                stack.push(True)
            else:
                stack.push(False)
        else: ...
        program_counter += 1
```

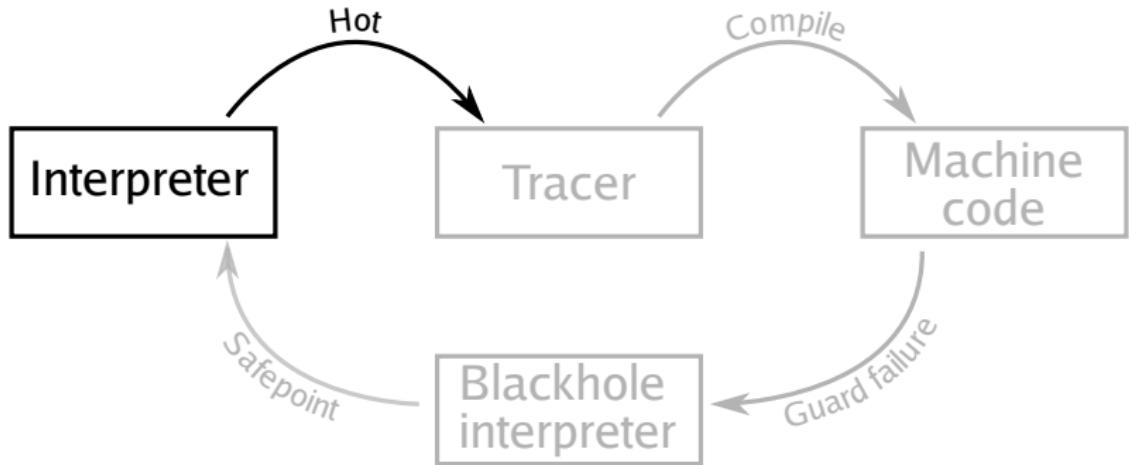
## Initial trace

```
v0 = <program_counter>
v1 = <stack>
v2 = <vars>
v3 = load_instruction(v0)
guard_eq(v3, INSTR_VAR_GET)
v4 = dict_get(v2, "x")
list_append(v1, v4)
v5 = add(v0, 1)
v6 = load_instruction(v5)
guard_eq(v6, INSTR_INT)
list_append(v1, 0)
v7 = add(v5, 1)
v8 = load_instruction(v7)
guard_eq(v8, INSTR_LESS_THAN)
v9 = list_pop(v1)
v10 = list_pop(v1)
guard_type(v9, int)
guard_type(v10, int)
guard_not_less_than(v9, v10)
list_append(v1, False)
v11 = add(v7, 1)
v12 = load_instruction(v11)
guard_eq(v12, INSTR_IF)
v13 = list_pop(v1)
guard_false(v13)
...
```

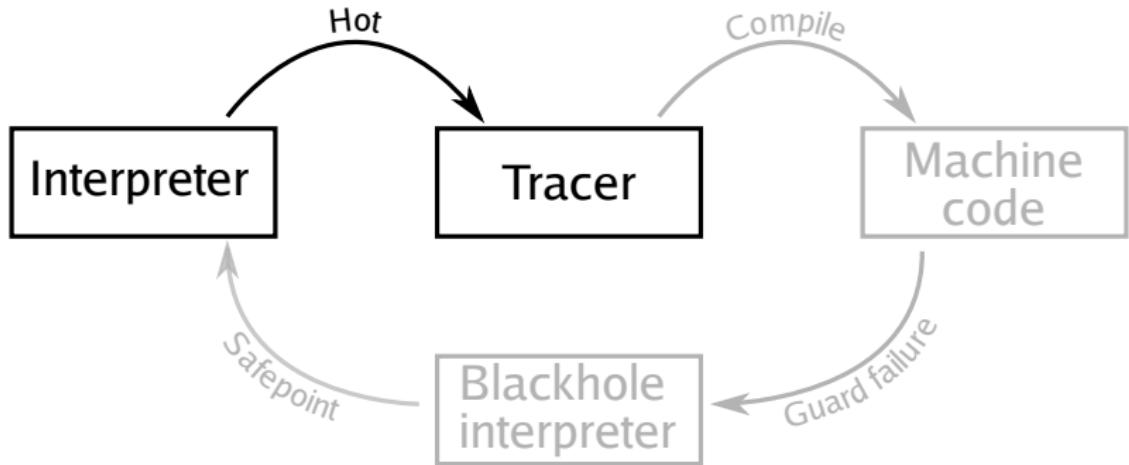
# Meta-tracer states



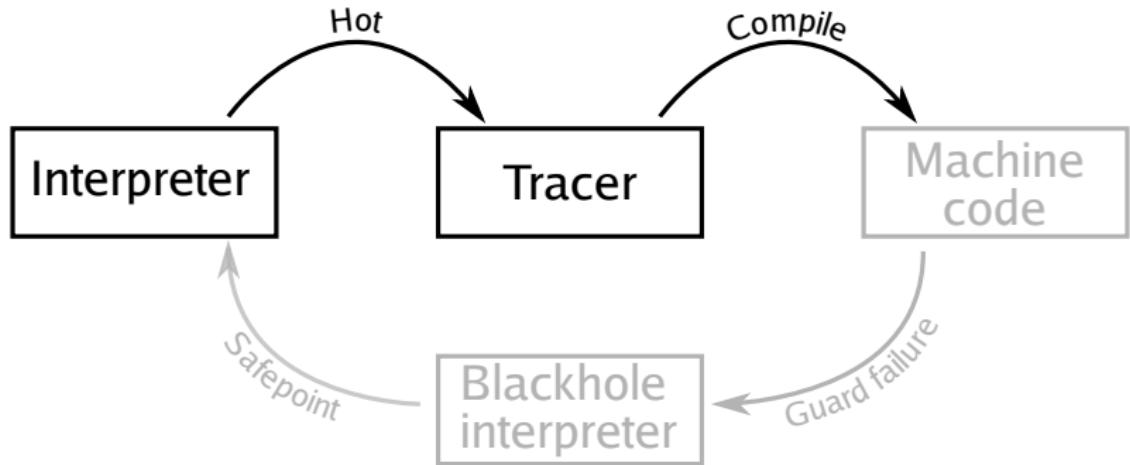
# Meta-tracer states



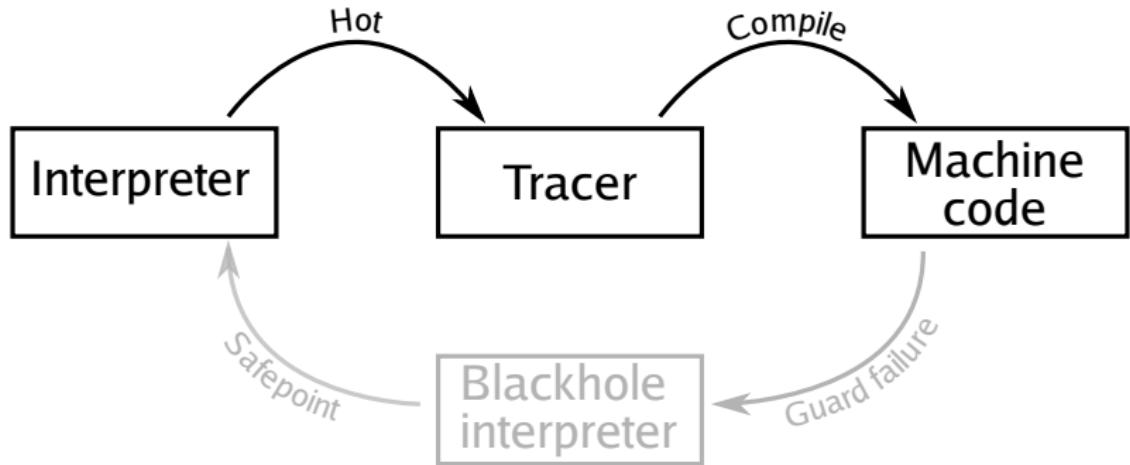
# Meta-tracer states



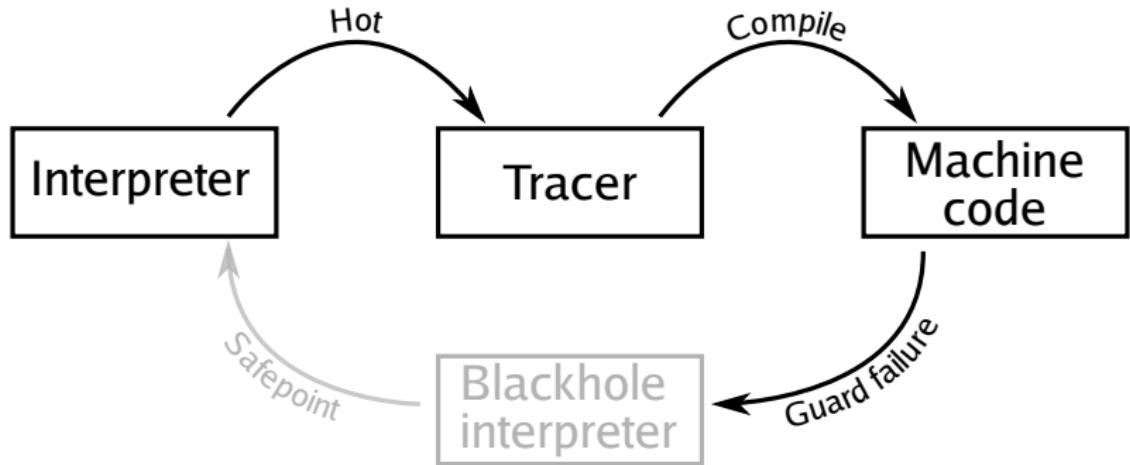
# Meta-tracer states



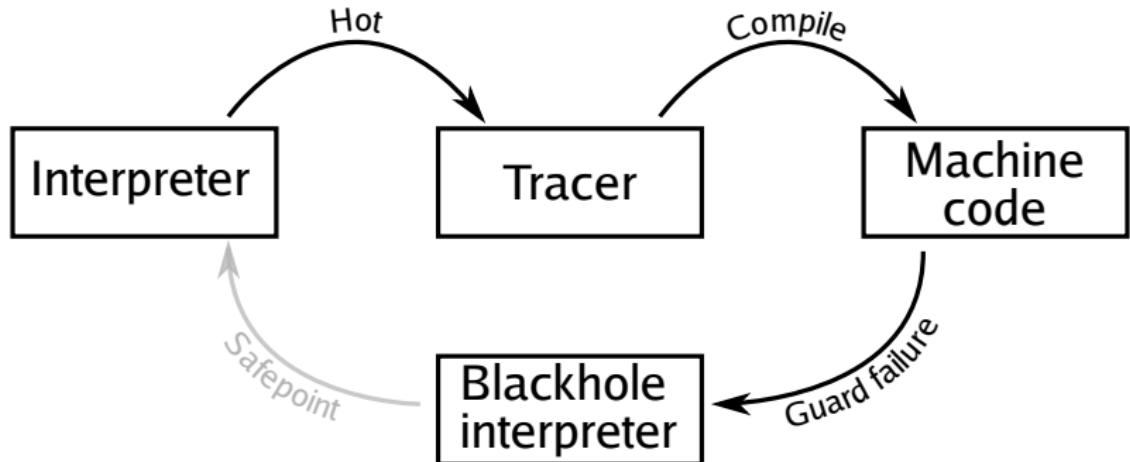
# Meta-tracer states



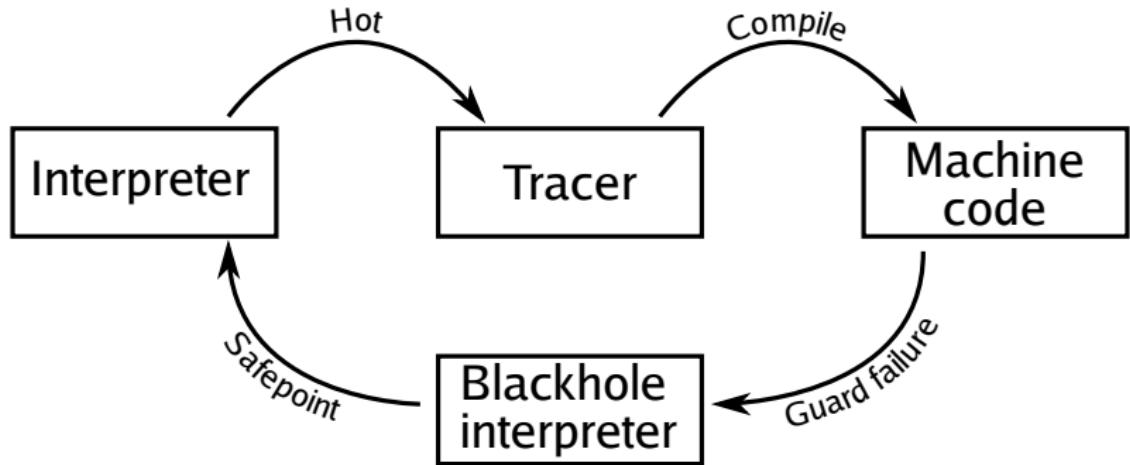
## Meta-tracer states



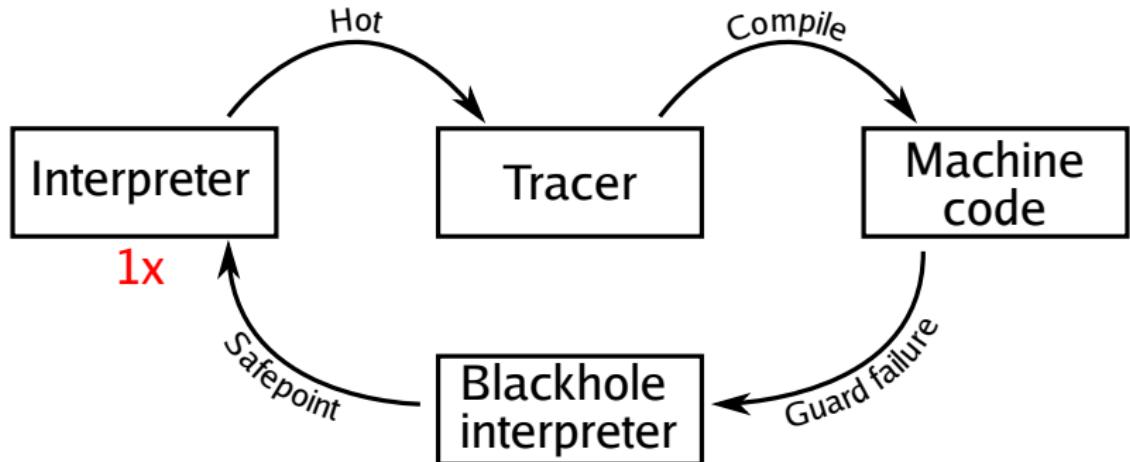
## Meta-tracer states



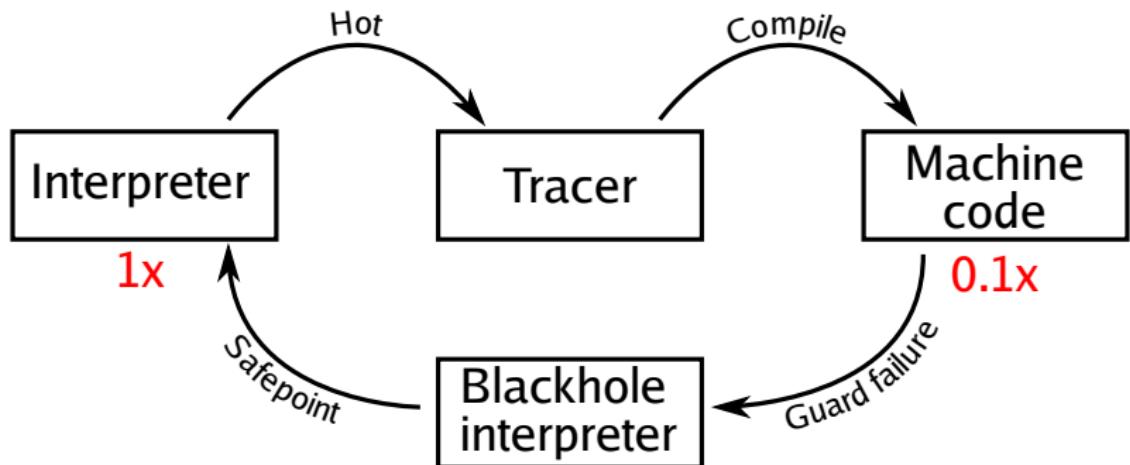
## Meta-tracer states



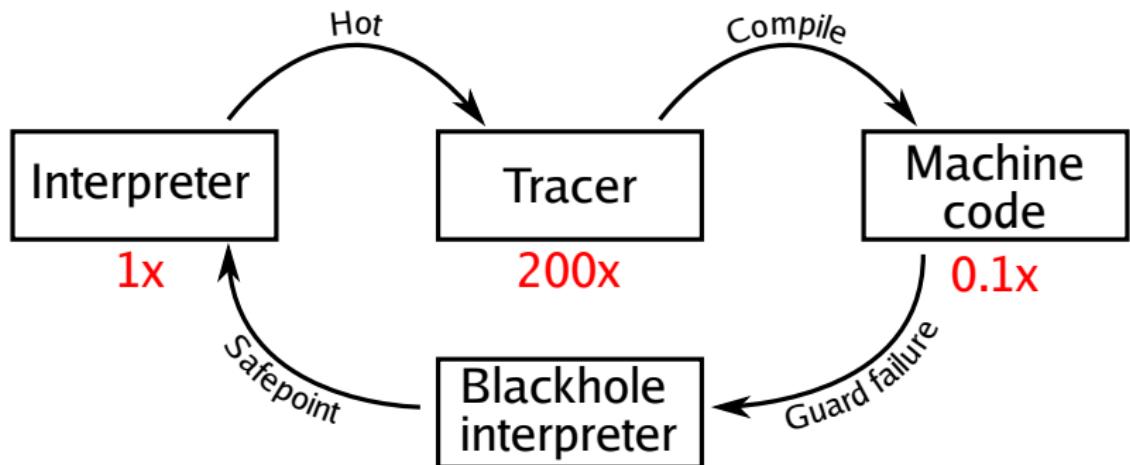
# Meta-tracer performance (now)



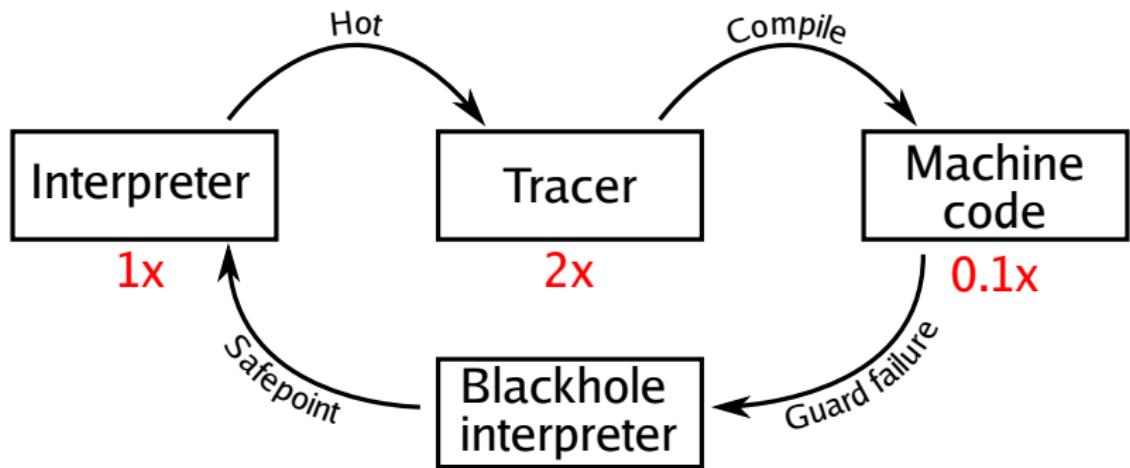
## Meta-tracer performance (now)



## Meta-tracer performance (now)



# Meta-tracer performance (our aim)



# References

## **VM Warmup Blows Hot and Cold**

*E. Barrett, C. F. Bolz, R. Killick, V. Knight, S. Mount and L. Tratt.*

## **Rigorous Benchmarking in Reasonable Time**

*T. Kalibera and R. Jones*

## **Specialising Dynamic Techniques for Implementing the Ruby Programming Language**

*C. Seaton (Chapter 4)*

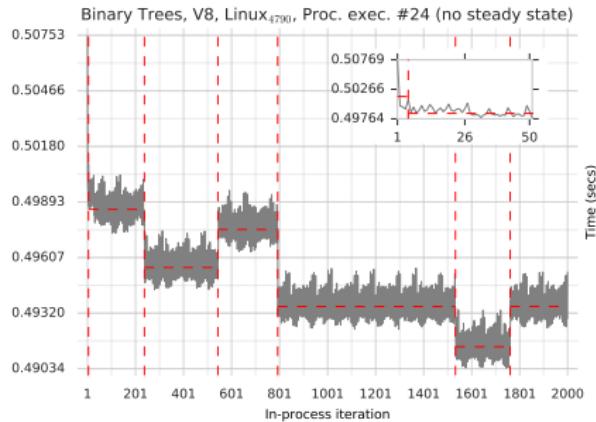
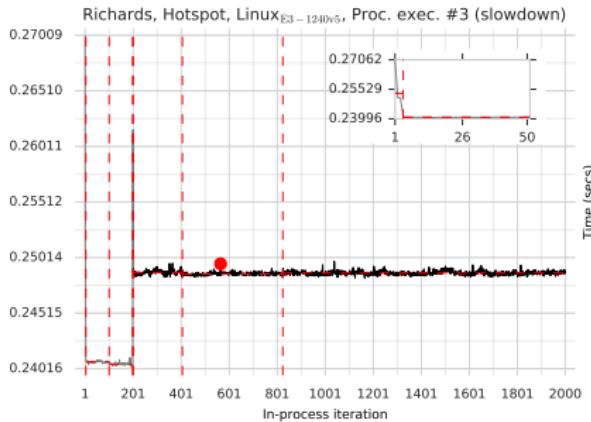
## **Quantifying performance changes with effect size confidence intervals**

*T. Kalibera and R. Jones*

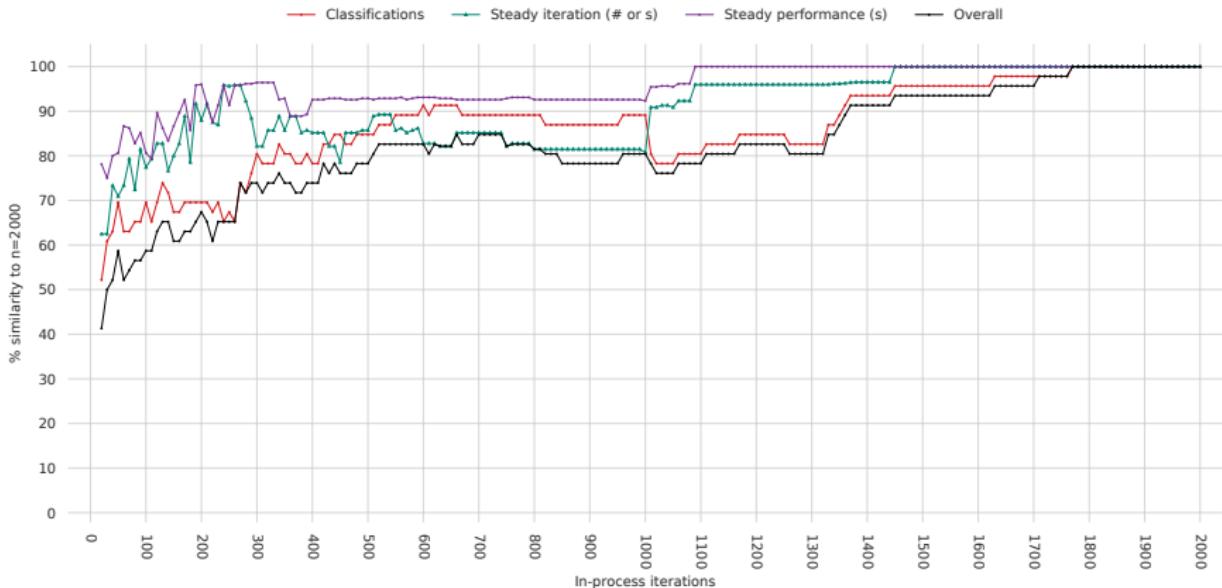
# Thanks

- EPSRC: *COOLER* and *Lecture*.
- Oracle.
- Cloudflare.

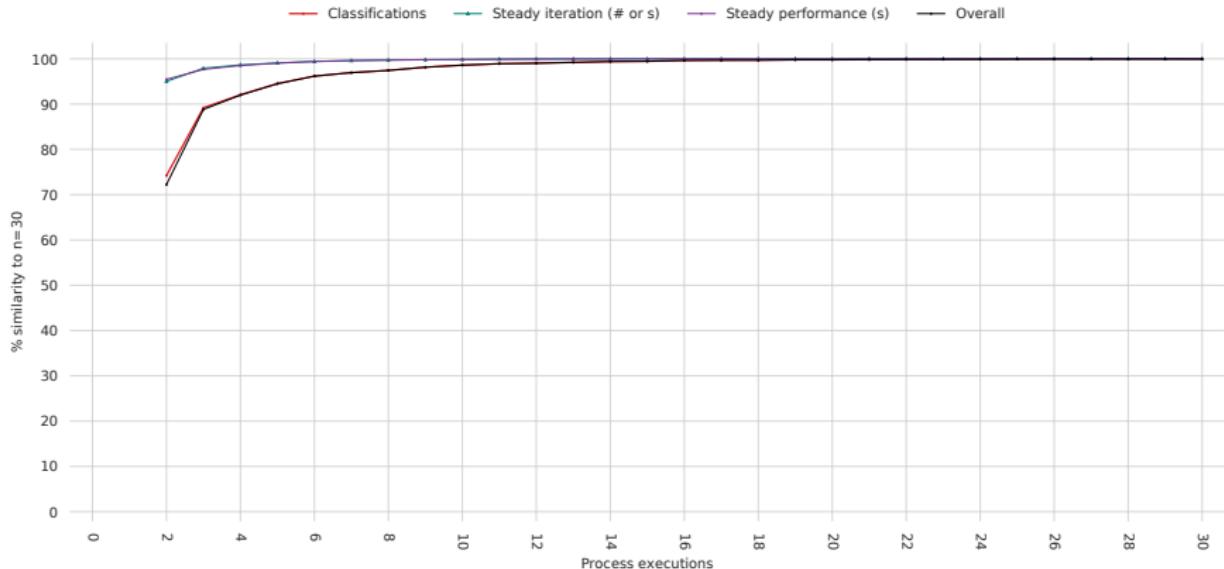
# Thanks for listening



# How long to run things for



# How long to run things for



# Comparing results

	Class.	Steady iter (#)	Steady iter (s)	Steady perf (s)
binarytrees	= (13-, 2L)	1.0 (1.0,701.3)	0.00 (0.000,88.313)	0.12555 ±0.001198
capnproto_decode	-			0.12813 ±0.001797
capnproto_encode	L	4.0 $\delta = -115.5$ (4.0,4.3)	0.11 $\delta = -18.177$ (0.105,0.118)	0.03364 $\delta = -0.087450$ ±0.000527
fannkuch_redux	L	2.0 (2.0,71.3)	0.13 (0.129,0.026)	0.12620 $\delta + 0.006886$ ±0.001135
fasta	Normal	-		0.12122 ±0.000403
jsonlua_decode	-			0.11448 ±0.002019
jsonlua_encode	-			0.13125 ±0.002070
luacheck	-			1.00955 ±0.004522
luacheck_parser		✗ (14-, 1w)		
luafun	= (14L, 1-)	42.0 (29.7,42.3)	6.36 (4.431,6.447)	0.15423 $\delta = -0.013313$ ±0.000905
md5	-			0.11168 $\delta = -0.001118$ ±0.000947
nbody	-			0.18291 $\delta + 0.023855$ ±0.004045
richards	L	2.0 (2.0,2.0)	0.15 (0.146,0.149)	0.14306 ±0.002472
series	L	2.0 (2.0,2.0)	0.35 (0.347,0.347)	0.33403 ±0.000320
spectralnorm	-			0.13988 ±0.000001