

Fast enough VMs in fast enough time



Laurence
Tratt



Software Development Team
2016-02-04

Language designers dilemma

Implementation Effort


Language designers dilemma



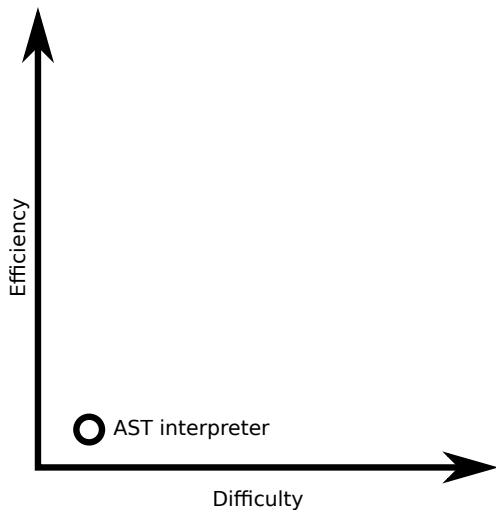
Language designers dilemma




The traditional routes




The traditional routes



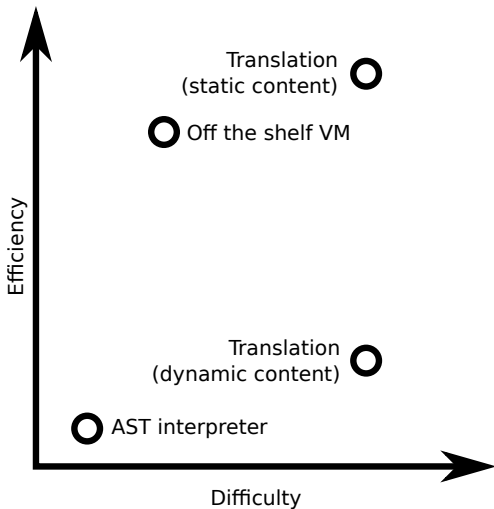
The traditional routes




The traditional routes




The traditional routes



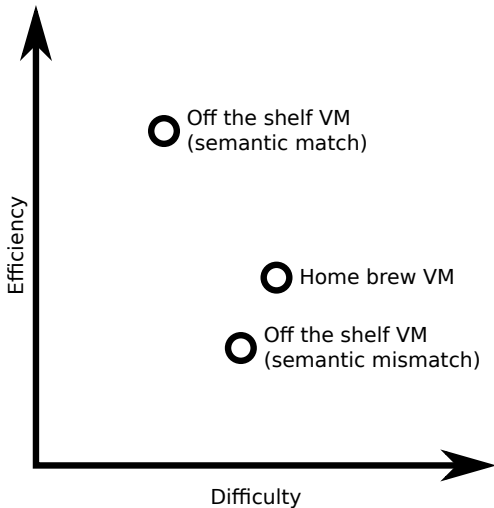
The traditional routes



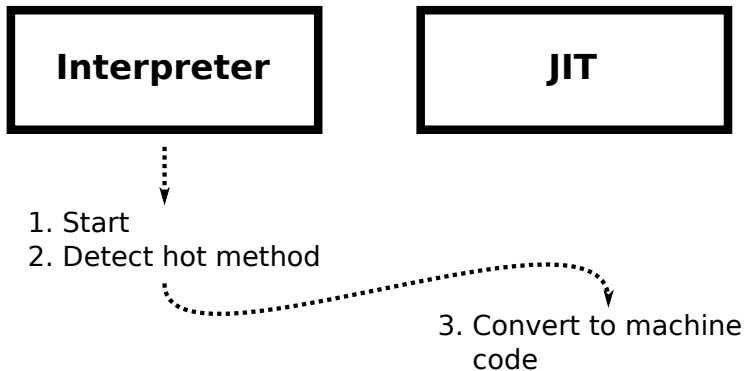
The traditional routes



The traditional routes



JIT compiled VM components



JIT compiled VM issues

Interpreter

JIT

JIT compiled VM issues

Interpreter

- + Easy to write
- Slow

JIT

JIT compiled VM issues

Interpreter

- + Easy to write
- Slow

JIT

- + Fast
- Difficult to write

JIT compiled VM issues

Interpreter

- + Easy to write
- Slow

JIT

- + Fast
- Difficult to write



Interpreter and JIT must be kept in sync


Building JIT compiled VMs is *hard*

Meta-tracing translation with RPython




Interpreter


Meta-tracing translation with RPython




Meta-tracing translation with RPython



Meta-tracing translation with RPython



Meta-tracing translation with RPython



Adding a JIT to an RPython interpreter

```
...
pc := 0
while 1:

    instr := load_next_instruction(pc)
    if instr == POP:
        stack.pop()
        pc += 1
    elif instr == BRANCH:
        off = load_branch_jump(pc)

        pc += off
    elif ...:
        ...
```


Observation: interpreters are big loops.

Adding a JIT to an RPython interpreter


```
...
pc := 0
while 1:
    jit_merge_point(pc)
    instr := load_next_instruction(pc)
    if instr == POP:
        stack.pop()
        pc += 1
    elif instr == BRANCH:
        off = load_branch_jump(pc)
        if off < 0: can_enter_jit(pc)
        pc += off
    elif ...:
        ...
```

Observation: interpreters are big loops.

RPython translation



RPython translation



User program (lang *FL*)

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

Tracing JIT compilers

User program (lang *FL*)

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


Trace when x is set to 6

```
guard_type(x, int)  
guard_not_less_than(x, 0)  
guard_type(x, int)  
x = int_add(x, 2)  
guard_type(x, int)  
x = int_add(x, 3)
```

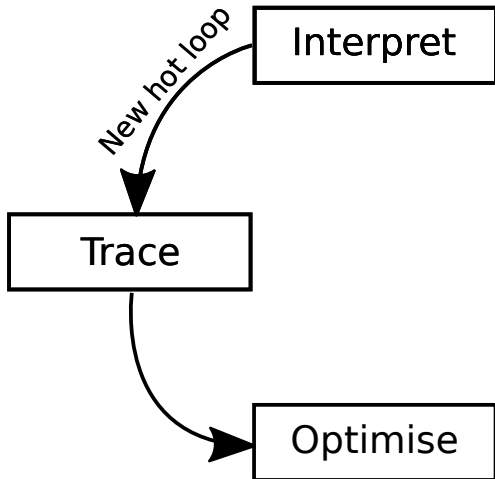
A tracing system's states

Interpret


A tracing system's states



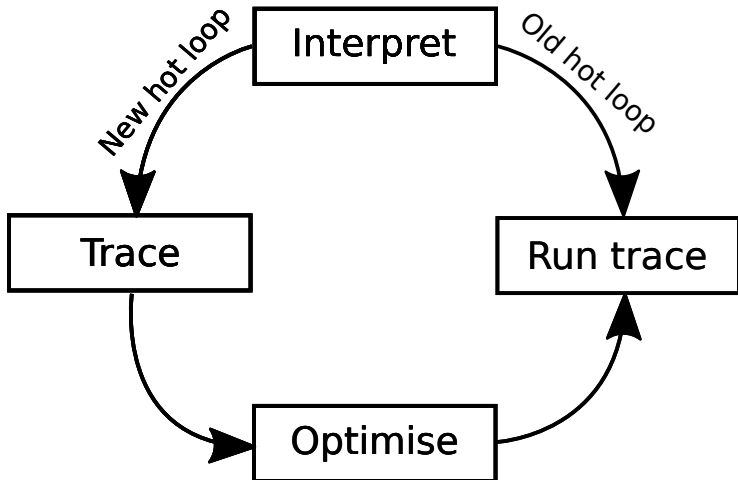
A tracing system's states



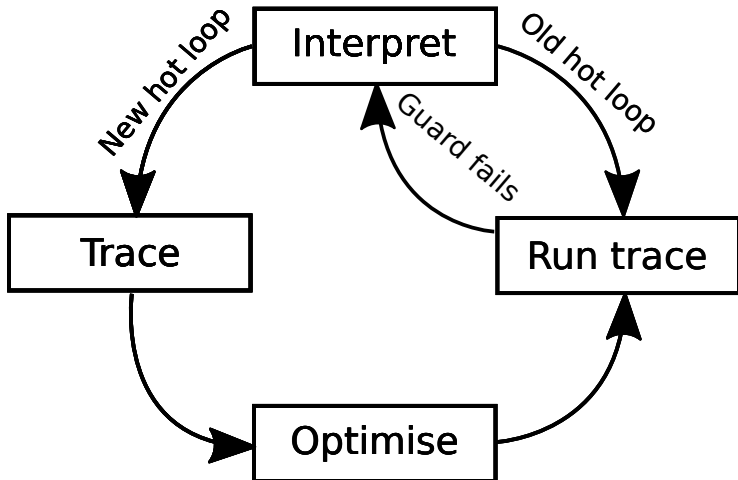
A tracing system's states



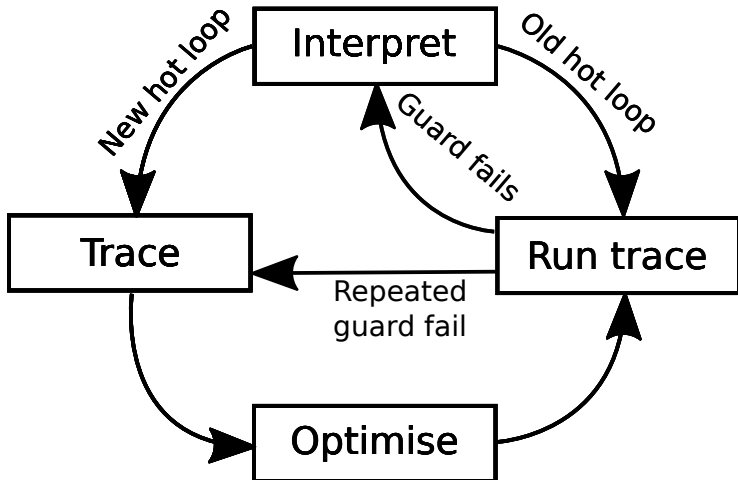
A tracing system's states



A tracing system's states



A tracing system's states



Trace optimisations (1)

- 1 Reuse typical compiler optimisations.

Tracing JIT compilers

User program (lang *FL*)

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

Trace when x is set to 6

```
guard_type(x, int)  
guard_not_less_than(x, 0)  
guard_type(x, int)  
x = int_add(x, 2)  
guard_type(x, int)  
x = int_add(x, 3)
```

Tracing JIT compilers


User program (lang *FL*)

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

Optimised trace

```
guard_type(x, int)  
guard_not_less_than(x, 0)  
x = int_add(x, 5)
```

Meta-tracing VM components



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
```

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)
...
```

Initial trace in full

```
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)
v14 = add(v11, 2)

v15 = load_instruction(v14)
guard_eq(v15, INSTR_VAR_GET)
v16 = dict_get(v2, "x")
list_append(v1, v16)
v17 = add(v14, 1)
v18 = load_instruction(v17)
guard_eq(v18, INSTR_INT)
list_append(v1, 2)
v19 = add(v17, 1)
v20 = load_instruction(v19)
guard_eq(v20, INSTR_ADD)
v21 = list_pop(v1)
v22 = list_pop(v1)
guard_type(v21, int)
guard_type(v22, int)
v23 = add(v22, v21)
list_append(v1, v23)
v24 = add(v19, 1)
v25 = load_instruction(v24)
guard_eq(v25, INSTR_VAR_SET)
v26 = list_pop(v1)
dict_set(v2, "x", v26)
v27 = add(v24, 1)
v28 = load_instruction(v27)
guard_eq(v28, INSTR_VAR_GET)
v29 = dict_get(v2, "x")

list_append(v1, v29)
v30 = add(v27, 1)
v31 = load_instruction(v30)
guard_eq(v31, INSTR_INT)
list_append(v1, 3)
v32 = add(v30, 1)
v33 = load_instruction(v32)
guard_eq(v33, INSTR_ADD)
v34 = list_pop(v1)
v35 = list_pop(v1)
guard_type(v34, int)
guard_type(v35, int)
v36 = add(v35, v34)
list_append(v1, v36)
v37 = add(v32, 1)
v38 = load_instruction(v37)
guard_eq(v38, INSTR_VAR_SET)
v39 = list_pop(v1)
dict_set(v2, "x", v39)
v40 = add(v37, 1)
```

Trace optimisation (1)

Removing constants (from jit_merge_point)

```
v1 = <stack>
v2 = <vars>
v4 = dict_get(v2, "x")
list_append(v1, v4)
list_append(v1, 0)
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)
v13 = list_pop(v1)
guard_false(v13)
v16 = dict_get(v2, "x")
list_append(v1, v16)
list_append(v1, 2)
v21 = list_pop(v1)
v22 = list_pop(v1)
guard_type(v21, int)
guard_type(v22, int)
v23 = add(v22, v21)
list_append(v1, v23)
v26 = list_pop(v1)
dict_set(v2, "x", v26)
v29 = dict_get(v2, "x")
list_append(v1, v29)

list_append(v1, 3)
v34 = list_pop(v1)
v35 = list_pop(v1)
guard_type(v34, int)
guard_type(v35, int)
v36 = add(v35, v34)
list_append(v1, v36)
v39 = list_pop(v1)
dict_set(v2, "x", v39)
```

Optimisation #2 & #3

List folded trace

```
v1 = <stack>
v2 = <vars>
v4 = dict_get(v2, "x")
guard_type(v4, int)
guard_not_less_than(v4, 0)
v16 = dict_get(v2, "x")
guard_type(v16, int)
v23 = add(v16, 2)
dict_set(v2, "x", v23)
v29 = dict_get(v2, "x")
guard_type(v29, int)
v36 = add(v29, 3)
dict_set(v2, "x", v36)
```

Optimisation #2 & #3

List folded trace

```
v1 = <stack>
v2 = <vars>
v4 = dict_get(v2, "x")
guard_type(v4, int)
guard_not_less_than(v4, 0)
v16 = dict_get(v2, "x")
guard_type(v16, int)
v23 = add(v16, 2)
dict_set(v2, "x", v23)
v29 = dict_get(v2, "x")
guard_type(v29, int)
v36 = add(v29, 3)
dict_set(v2, "x", v36)
```

Dict folded trace

```
v1 = <stack>
v2 = <vars>
v4 = dict_get(v2, "x")
guard_type(v4, int)
guard_not_less_than(v4, 0)
v23 = add(v4, 2)
guard_type(v23, int)
v36 = add(v23, 3)
dict_set(v2, "x", v36)
```

Optimisation #4 & #5

Type folded trace

```
v1 = <stack>
v2 = <vars>
v4 = dict_get(v2, "x")
guard_type(v4, int)
guard_not_less_than(v4, 0)
v23 = add(v4, 2)
v36 = add(v23, 3)
dict_set(v2, "x", v36)
```

Optimisation #4 & #5

Type folded trace

```
v1 = <stack>
v2 = <vars>
v4 = dict_get(v2, "x")
guard_type(v4, int)
guard_not_less_than(v4, 0)
v23 = add(v4, 2)
v36 = add(v23, 3)
dict_set(v2, "x", v36)
```

Arithmetic folded trace

```
v1 = <stack>
v2 = <vars>
v4 = dict_get(v2, "x")
guard_type(v4, int)
guard_not_less_than(v4, 0)
v23 = add(v4, 5)
dict_set(v2, "x", v23)
```

Optimisation #4 & #5

Type folded trace

```
v1 = <stack>
v2 = <vars>
v4 = dict_get(v2, "x")
guard_type(v4, int)
guard_not_less_than(v4, 0)
v23 = add(v4, 2)
v36 = add(v23, 3)
dict_set(v2, "x", v36)
```

Arithmetic folded trace

```
v1 = <stack>
v2 = <vars>
v4 = dict_get(v2, "x")
guard_type(v4, int)
guard_not_less_than(v4, 0)
v23 = add(v4, 5)
dict_set(v2, "x", v23)
```

Trace optimisation: from 72 trace elements to 7.

Trace optimisations (2)

- 1 Reuse typical compiler optimisations.

Trace optimisations (2)

- 1 Reuse typical compiler optimisations.
- 2 Make use of tracings natural tendency to inline.

Trace optimisations (2)

- 1 Reuse typical compiler optimisations.
- 2 Make use of tracings natural tendency to inline.
- 3 Insert *language-specific* hints.

Example: Language-Specific Runtime Feedback

```
def lookup(cls, name):  
    ...  
  
def call_method(obj, func_name, args):  
    cls = obj.get_class()  
  
    func = lookup(cls, func_name)  
    return func.call(obj, args)
```


Example: Language-Specific Runtime Feedback

```
def lookup(cls, name):  
    ...  
  
def call_method(obj, func_name, args):  
    cls = obj.get_class()  
    promote(cls)  
    func = lookup(cls, func_name)  
    return func.call(obj, args)
```

Example: Language-Specific Runtime Feedback

```
@elidable
def lookup(cls, name):
    ...

def call_method(obj, func_name, args):
    cls = obj.get_class()
    promote(cls)
    func = lookup(cls, func_name)
    return func.call(obj, args)
```

An RPython experiment

An RPython experiment

A Converge VM in RPython.

A Converge VM in RPython.

How hard can it be?

Converge 1 vs. Converge 2 VMs

	Converge 1	Converge 2
Size (KLoc)		
Effort (person months)		
Performance		

Converge 1 vs. Converge 2 VMs

	Converge 1	Converge 2
Size (KLoc)	13	
Effort (person months)		
Performance		

Converge 1 vs. Converge 2 VMs

	Converge 1	Converge 2
Size (KLoc)	13	5.5
Effort (person months)		
Performance		

Converge 1 vs. Converge 2 VMs

	Converge 1	Converge 2
Size (KLoc)	13	5.5
Effort (person months)	18	
Performance		

Converge 1 vs. Converge 2 VMs

	Converge 1	Converge 2
Size (KLoc)	13	5.5
Effort (person months)	18	3
Performance		

Converge 1 vs. Converge 2 VMs

	Converge 1	Converge 2
Size (KLoc)	13	5.5
Effort (person months)	18	3
Performance	x	

Converge 1 vs. Converge 2 VMs

	Converge 1	Converge 2
Size (KLoc)	13	5.5
Effort (person months)	18	3
Performance	x	2-150x

PyPy: an industrial strength VM

	Converge 2	PyPy
Size (KLoc)	5.5	
Effort (person months)	3	
Performance	x	

PyPy: an industrial strength VM

	Converge 2	PyPy
Size (KLoc)	5.5	60 (+190 for libraries)
Effort (person months)	3	
Performance	x	

PyPy: an industrial strength VM

	Converge 2	PyPy
Size (KLoc)	5.5	60 (+190 for libraries)
Effort (person months)	3	300
Performance	x	

PyPy: an industrial strength VM

	Converge 2	PyPy
Size (KLoc)	5.5	60 (+190 for libraries)
Effort (person months)	3	300
Performance	x	3-5x

PyPy demo

Where we are now

+ Get a good JIT compiler for *very little effort*.

Where we are now

- + Get a good JIT compiler for *very little effort*.
- RPython is... RPython.

Where we are now

- + Get a good JIT compiler for *very little effort*.
- RPython is... RPython.
- Fairly poor warmup.

Where we are now

- + Get a good JIT compiler for *very little effort*.
- RPython is... RPython.
- Fairly poor warmup.
- Poor multi-threading support.

Summary

What language designers
dilemma?

Thank you for listening

<http://soft-dev.org/>