

YARV

Past, Present and Future

Koichi “ko1” Sasada

ささだこういち

ko1 at atdot dot net

RUBY MEETS VM

Koichi Sasada

ko1 at atdot dot net

Point of This Presentation

3

Merging YARV is not a goal,
but a start

YARV: Yet Another RubyVM

Notice

4

- I can't speak English well, so I write down all things what I want to say.
 - ▣ Do you get ready for opera glasses?
 - ▣ Unfortunately, some slides are written in Japanese
- You can ask questions with Japanese, C, Ruby, ..., or slow/short English.
- "How to impl. Ruby", not "How to use Ruby"
- "x50" is too big mouth
 - ▣ Maybe x20

Self Introduction

5

- Koichi (given name) Sasada (family name)
 - ▣ ささだ (family name) こういち (given name)
 - ▣ 笹田 (family name) 耕一 (given name)
- Lecturer @ University of Tokyo (Feb. 2008-)
- Only VM developer
 - ▣ Don't have compatibility with Matz
 - ▣ Please call me "ko-i-chi"

Agenda

6

- History of YARV
- Advanced VM Topics
 - ▣ Performance
 - Parallel Thread Execution
 - Embedding Float Value
 - JIT Compiler
 - Pre-Compiler
 - “Ruby to Compiled file” Compiler
 - “Ruby to C” Compiler
 - ▣ New Feature
 - Multi-VM Creation
 - Customizable Ruby Core
 - Debug/Profile support feature
- Summary

History of YARV

7

- 4 Years
 - **1, Jan 2004 Project Start**
 - 2004-2005 VM Core, Optimization
 - Supported by MITO youth Project (IPA)
 - 2005-2006 Thread, etc
 - Supported by MITO Project (IPA)
 - 1, Apr 2006 Got a Job (Assistant on U-Tokyo)
 - 2006-2007 etc, etc
 - Supported by MITO Project (IPA)
 - 25, Dec 2007 Got a Ph.D
 - **25, Dec 2007 (GMT) 1.9 Release**

FYI

8

- Ruby 2.0 – since 2003 3/31
- Perl 6 – since 2003 4/1

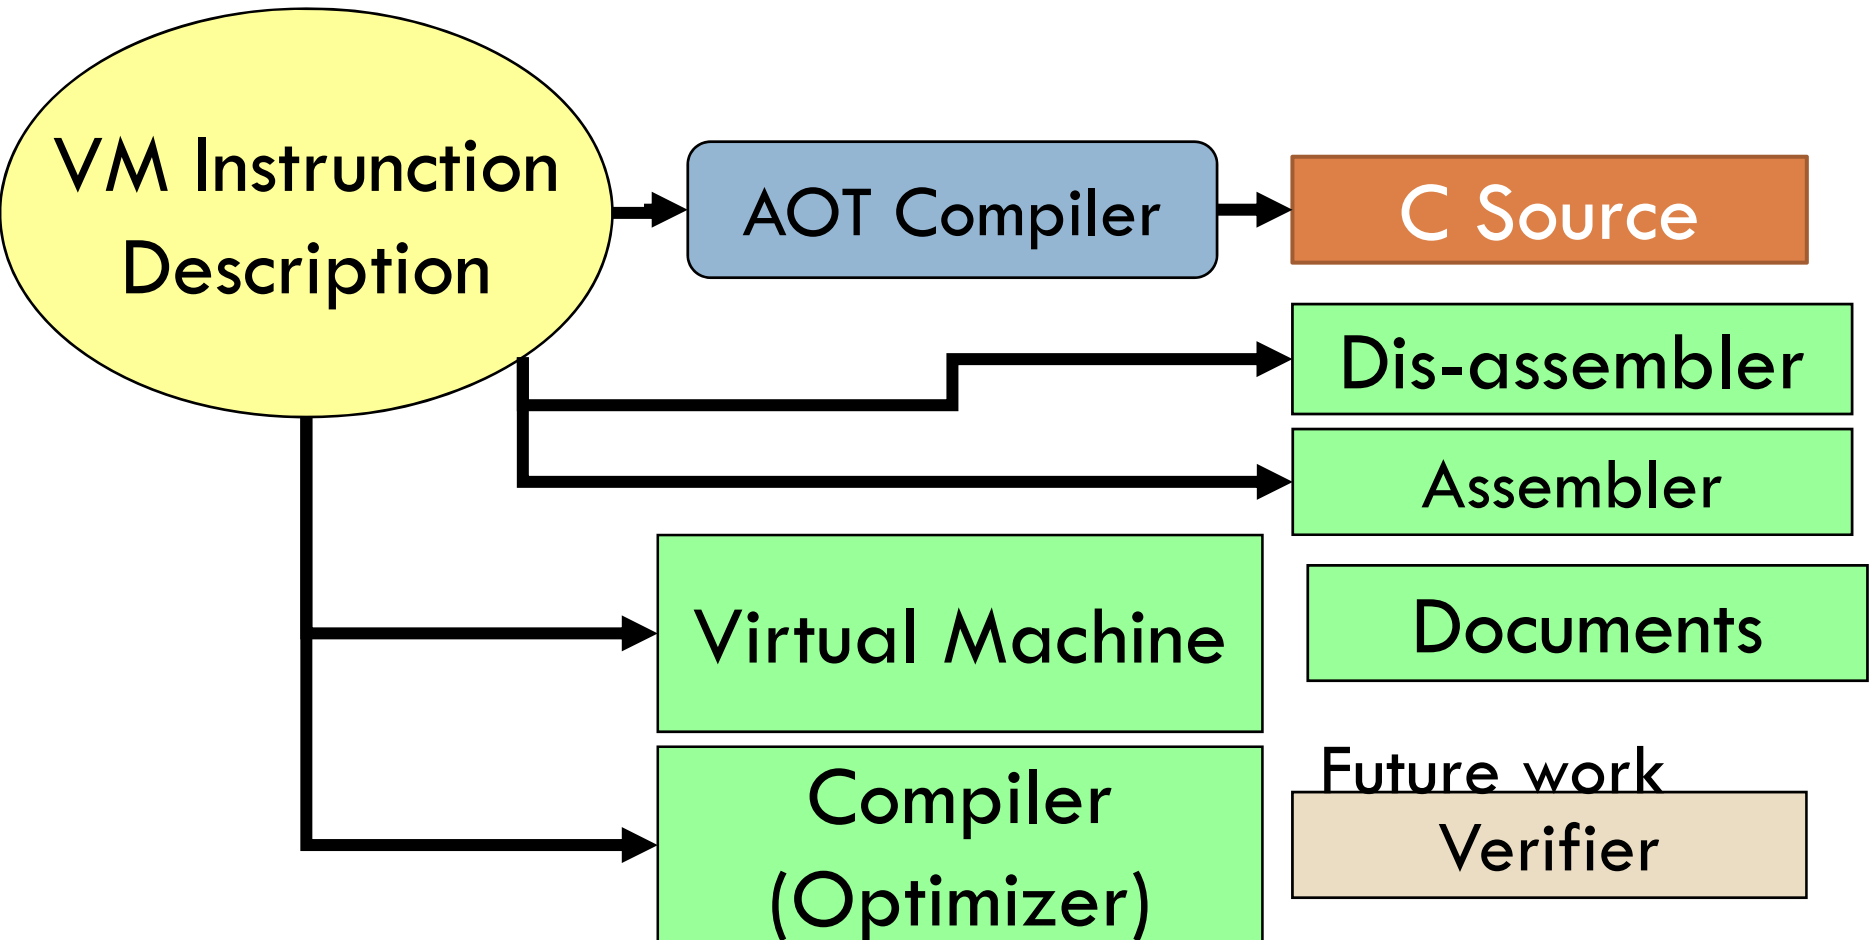
YARV Policy

9

- Performance
 - ▣ Speed, Speed, Speed
 - ▣ Applied many many many optimization Tech.
- Compatibility
 - ▣ C extension API
 - ▣ Not language compatibility
- Auto-generation
 - ▣ VM description to Concrete VM source code

VM Generator

11




Enemies of YARV

12

- Ruby Specification ≐ Matz
 - ▣ Ruby Spec kills many optimization techs
 - ▣ We love “Dynamic” “Meta” Programming, but...
- Changing Spec is also Nightmare
- Portability
 - ▣ We can't use system depending techs.
- Rivals (**not Enemy**)
 - ▣ Jruby, Rubinius, IronRuby, ...
- Peggy work on my Job


Evaluation: Improve case

13



Evaluation: Macro-Benchmark

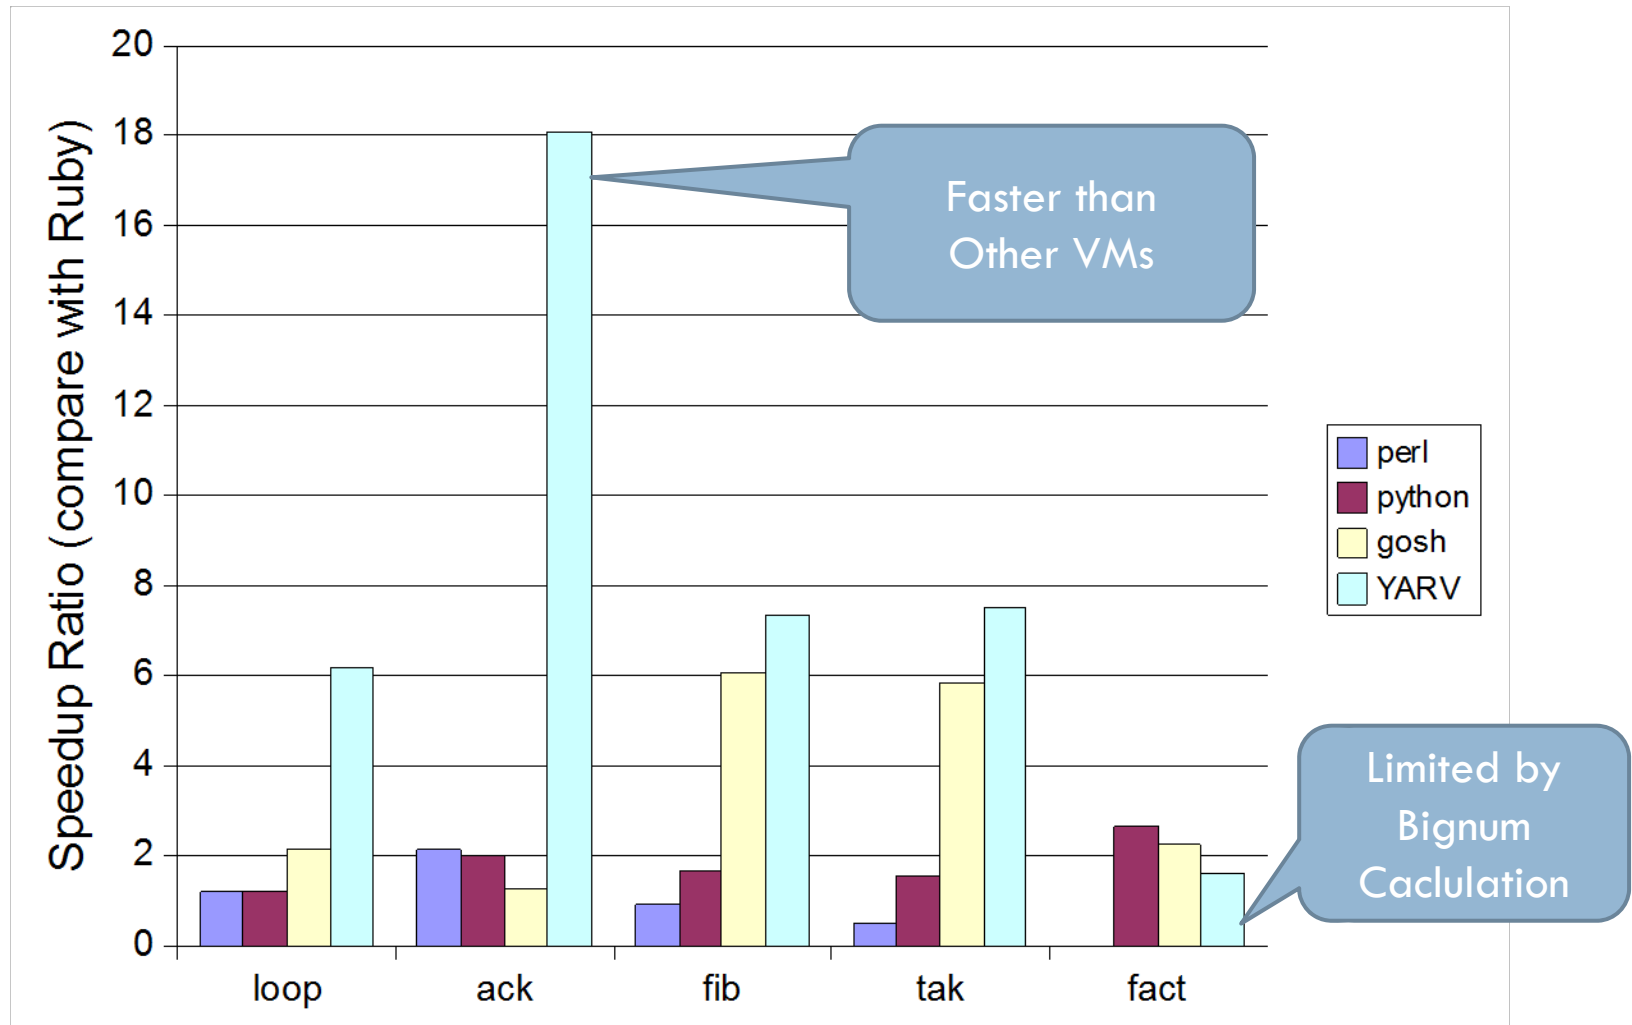
14



Evaluation:

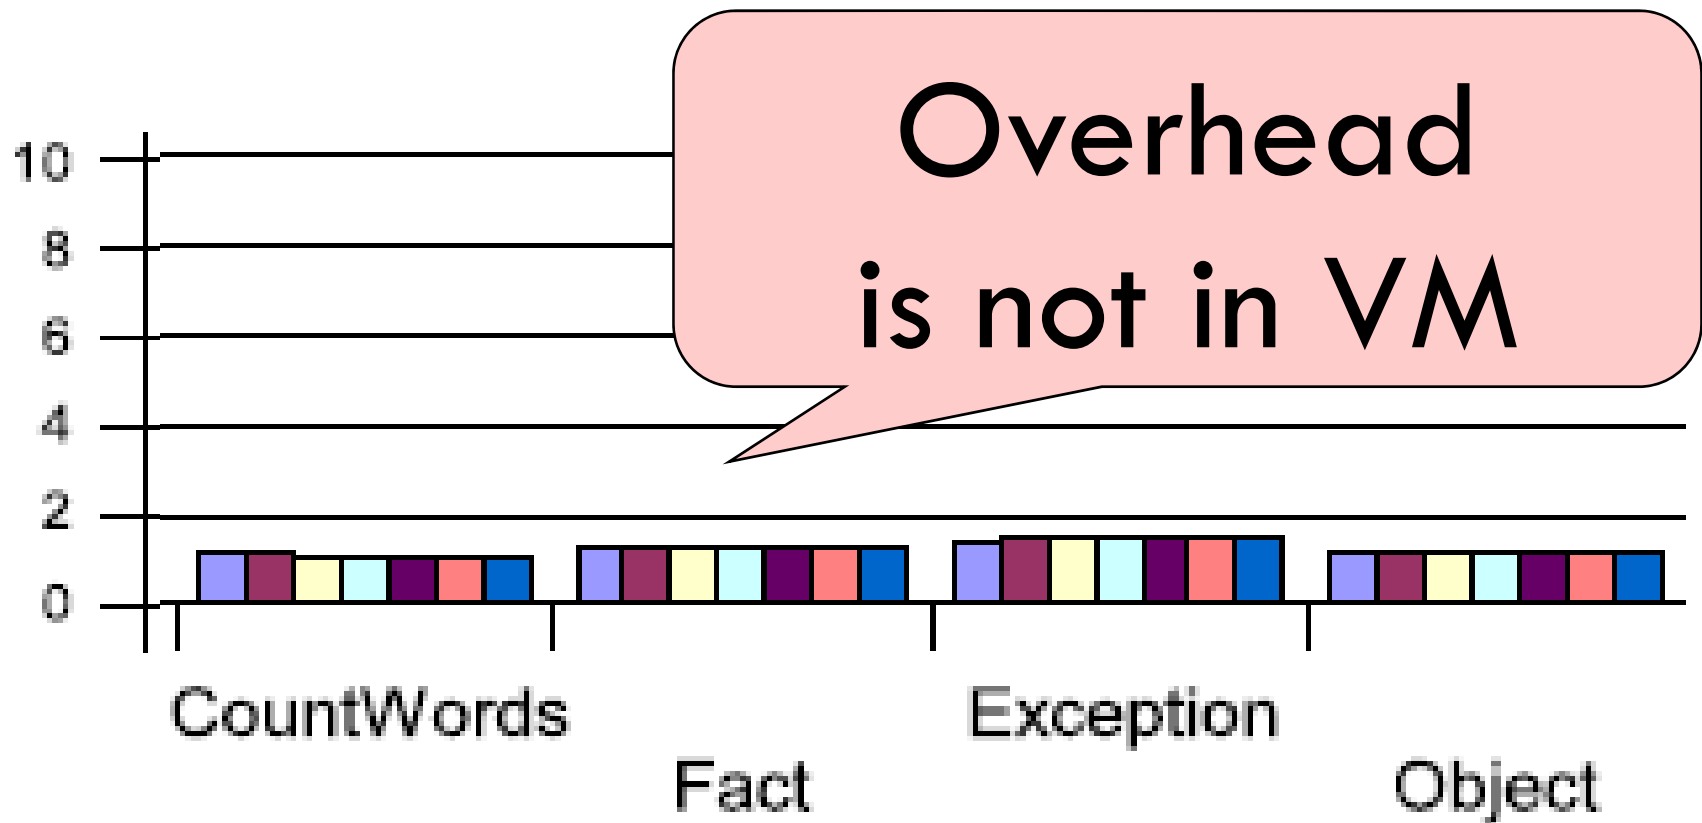
Compare with Other Languages

15



Evaluation: VM doesn't affect

16



Advanced VM Topics

17

- Performance
 - ▣ Parallel Thread Execution
 - ▣ Embedding Float Value
 - ▣ JIT Compiler
 - ▣ Pre-Compiler
 - “Ruby to Compiled file” Compiler
 - “Ruby to C” Compiler
- New Feature
 - ▣ Multi-VM Creation
 - ▣ Customizable Ruby Core
 - ▣ Debug/Profile support feature

Implement a High-Speed Ruby Interpreter

Introduce VM



YARV

(Merged into Ruby 1.9)

Parallel Ruby Execution



Parallel Thread Execution

Parallel Thread Execution

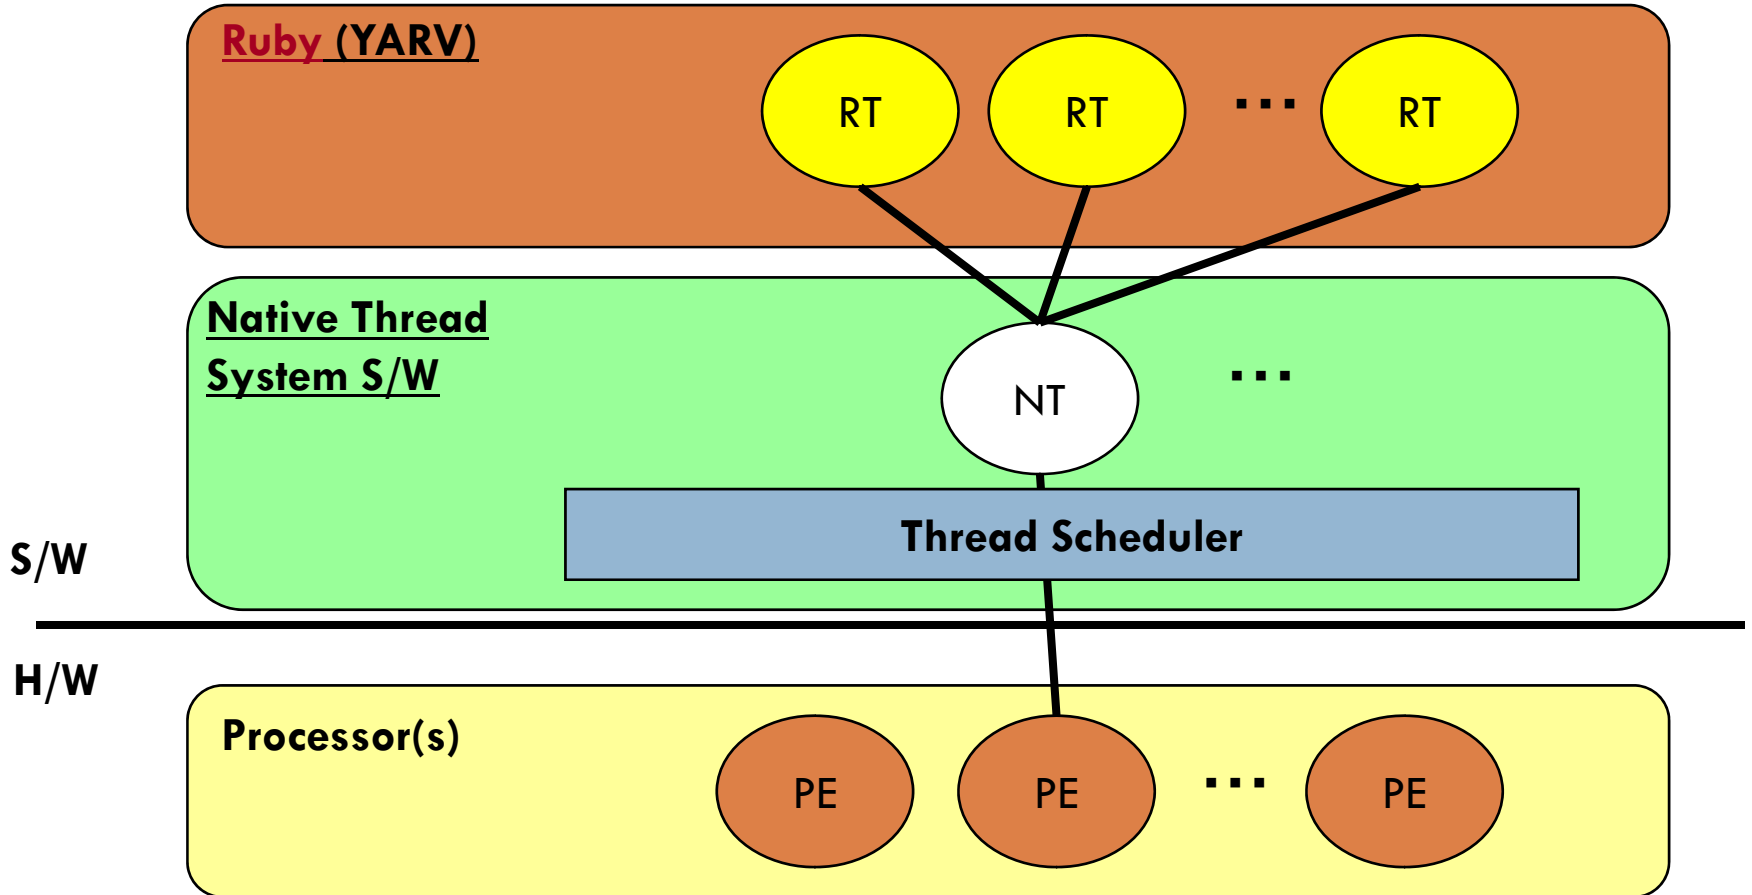
19

- Using Native Thread
- Get rid of Giant VM Lock

Method (1)

Ruby Thread and Native Thread (1:N) a.k.a -1.8 Ruby model

20




PE: Processor Element, UL: User Level, KL: Kernel Level

Method (2)

Ruby Thread and Native Thread (1:1)

21

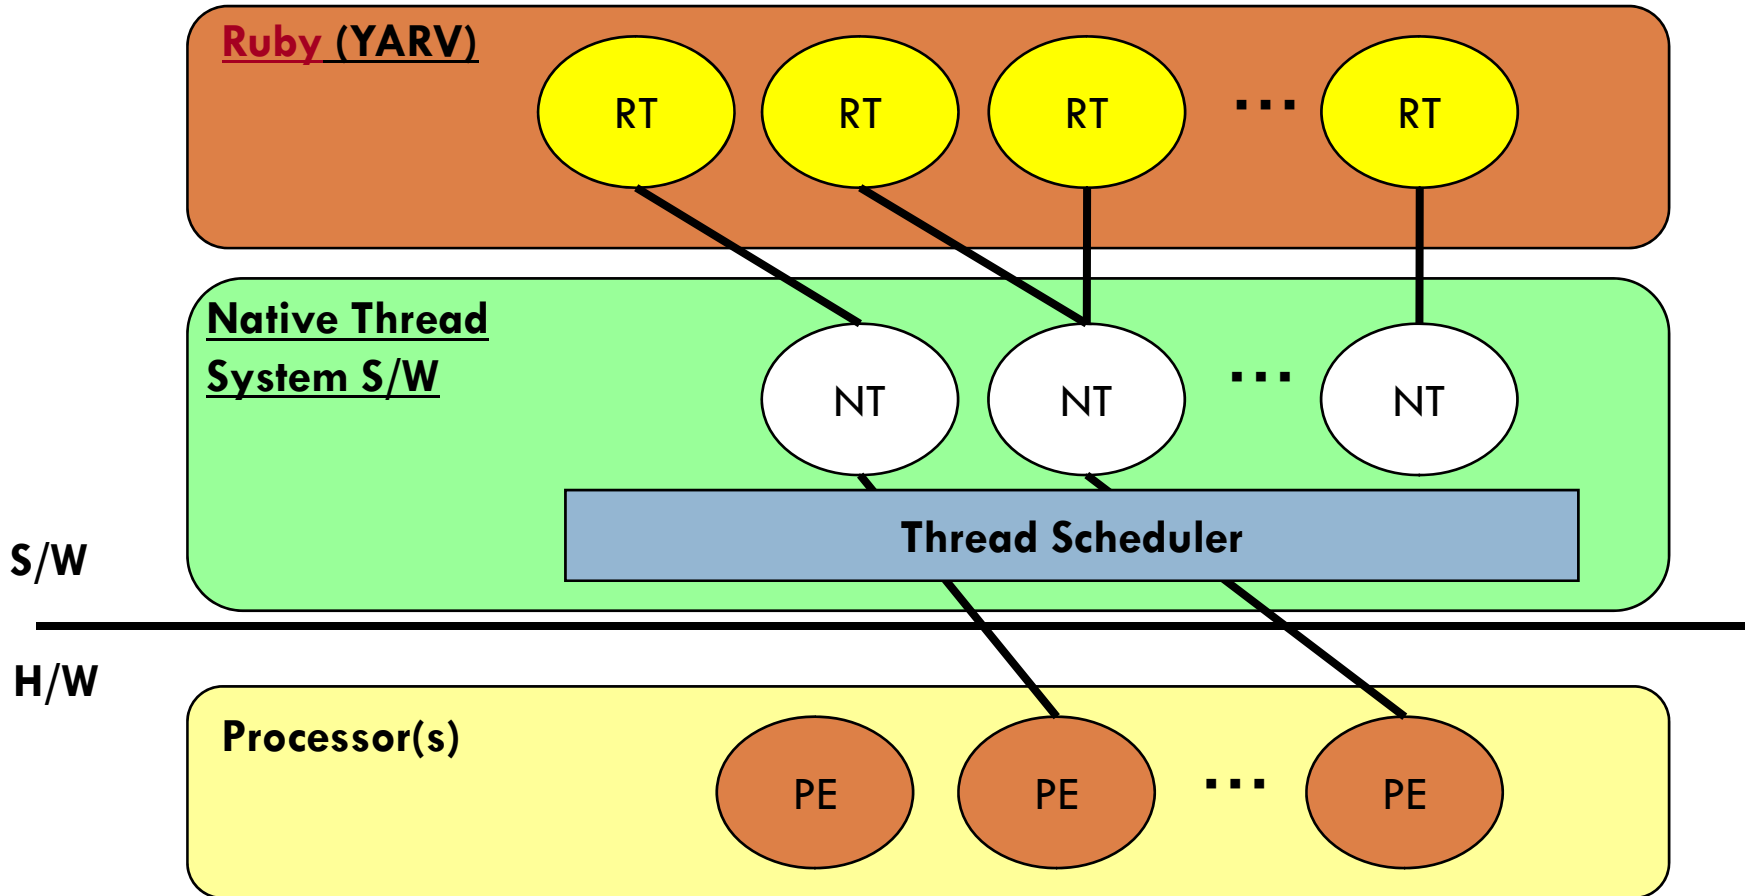


PE: Processor Element, UL: User Level, KL: Kernel Level

Method (3)

Ruby Thread and Native Thread (N:M)

22



PE: Processor Element, UL: User Level, KL: Kernel Level

Discussion

Ruby Thread and Native Thread

23

□ Mapping with Native Thread and Ruby Thread


	Pros	Cons
1:N	Lightweight Thread Control	Can't run in Parallel
1:1	Run in Parallel Simple, Portable	Heavyweight Thread Control (Creation, etc)
N:M	Lightweight Thread Control, Run in Parallel	Complication, Non-Portable

- Accept 1:1 model to make Ruby Simple
 - Depend Performance on Native Thread Libraries

Accepted Method:

Ruby Thread and Native Thread (1:1) ← Ruby 1.9/YARV

24



PE: Processor Element, UL: User Level, KL: Kernel Level

Introduction of Mutual Exclusion

25

- Needed at
 1. Global VM Management Data
 2. Object Management / GC
 3. Inline Cache
 4. Thread Unsafe “C” methods

(1) Global VM Management Data

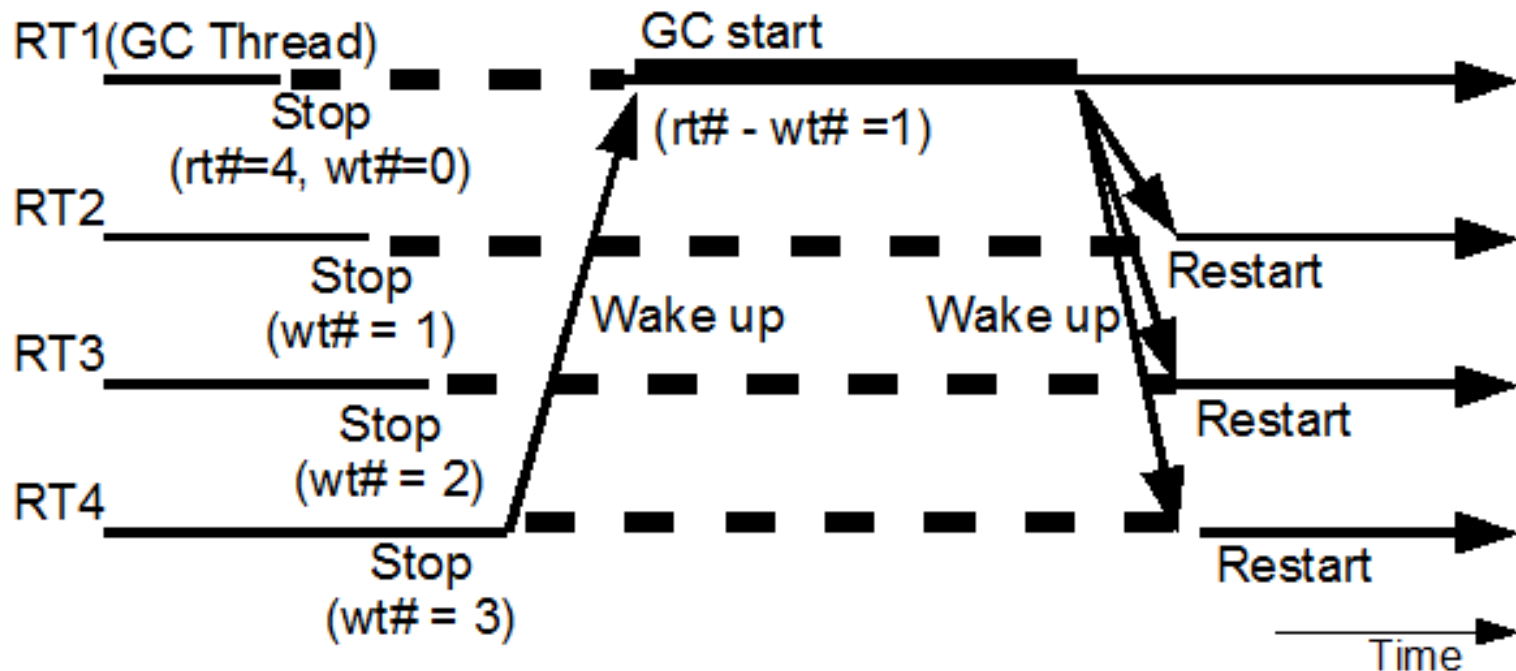
26

- Managed by Table
 - ▣ Variable Name → Value
 - ▣ Method Name → Method Body
 - ▣ ...
- Introduce Synchronization at Table Operation
 - ▣ Get/Set
 - ▣ Easy

(2) Object Management/GC

27

□ Synchronous GC

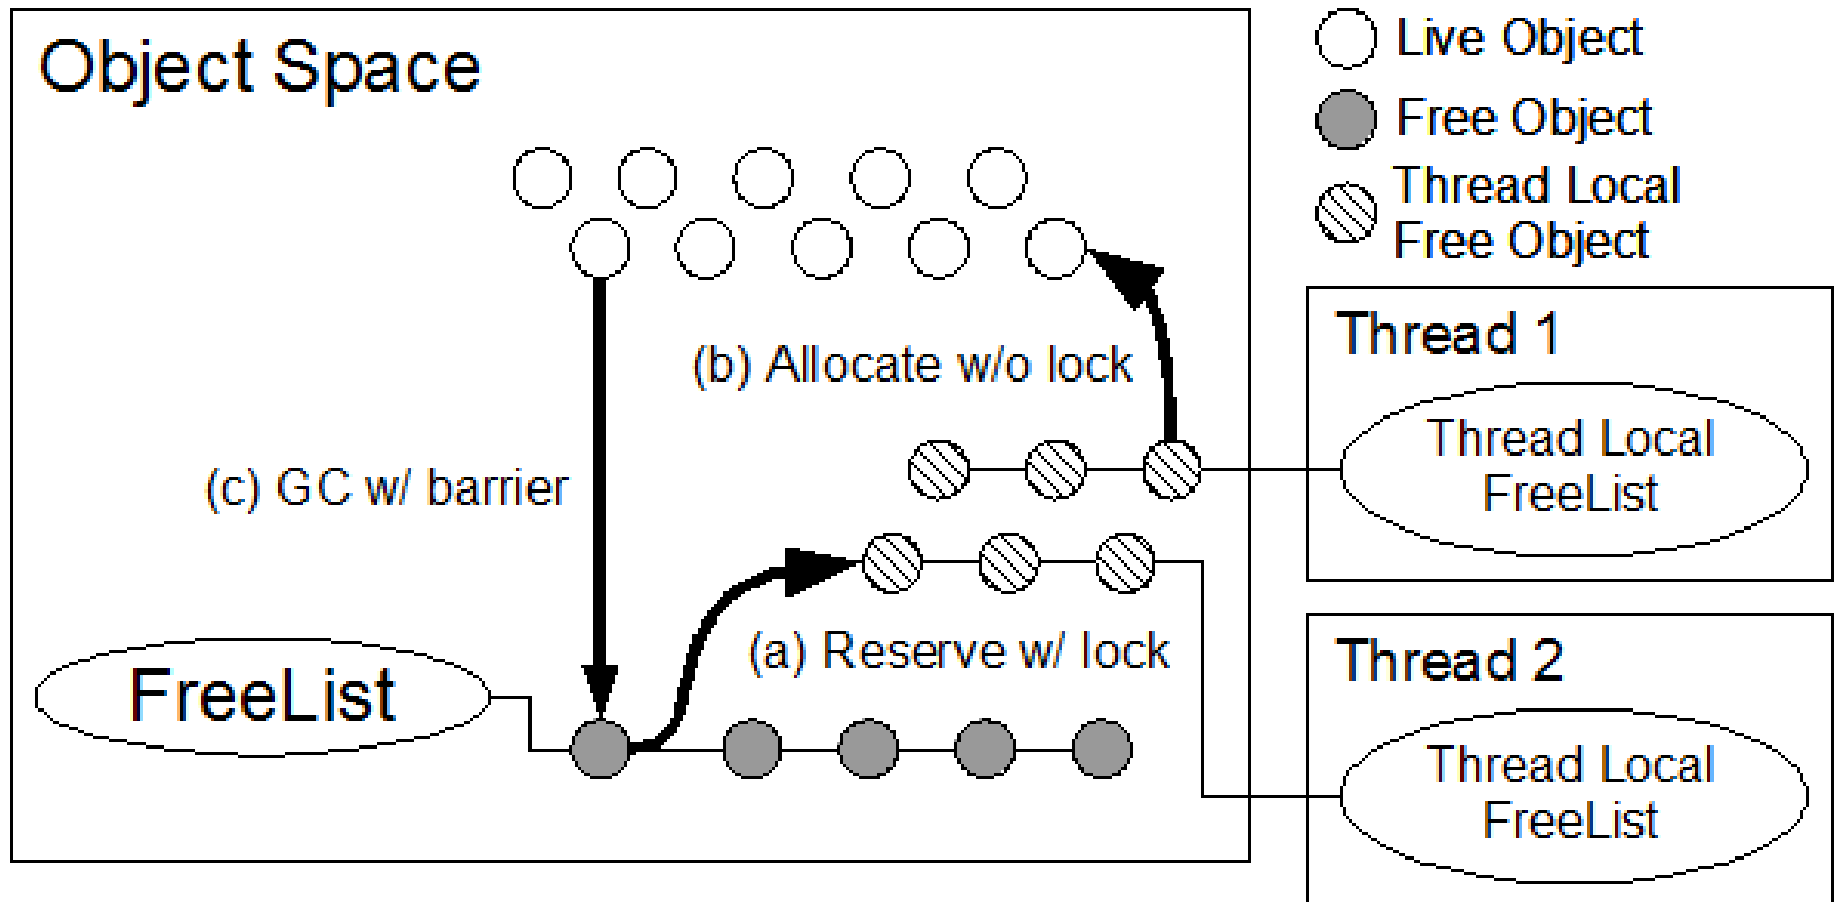


rt#: Ruby Threads Number
wt#: Waiting Threads Number

(2) Object Management

Lock-Free Object Allocation with Thread Local Free List

29



(3) Inline Cache

30

- Using by VM performance
 - ▣ Embed Cache Entries in Instruction Sequence
- Sync. For Coherence -> Performance Problem
 - ▣ Key and Value of Cache Entry
- Sync.-Free Inline Cache
 - ▣ Cache Miss -> Make a new entry
 - GC will clean-up old cache entries
 - Increase Miss-Penalty, but Good for average

(4) Thread Unsafe “C” Methods

31

- CRuby has many many methods implemented by “C”
 - ▣ All of them are “Thread Unsafe”
 - ▣ Because -1.8 doesn’t support parallelization
- Basic Policy: Using Giant-Lock
 - ▣ Invoke old “C” method with Giant-Lock Acquire
 - ▣ Re-write C methods as Thread Safe, this method will be Giant-Lock free

Problem

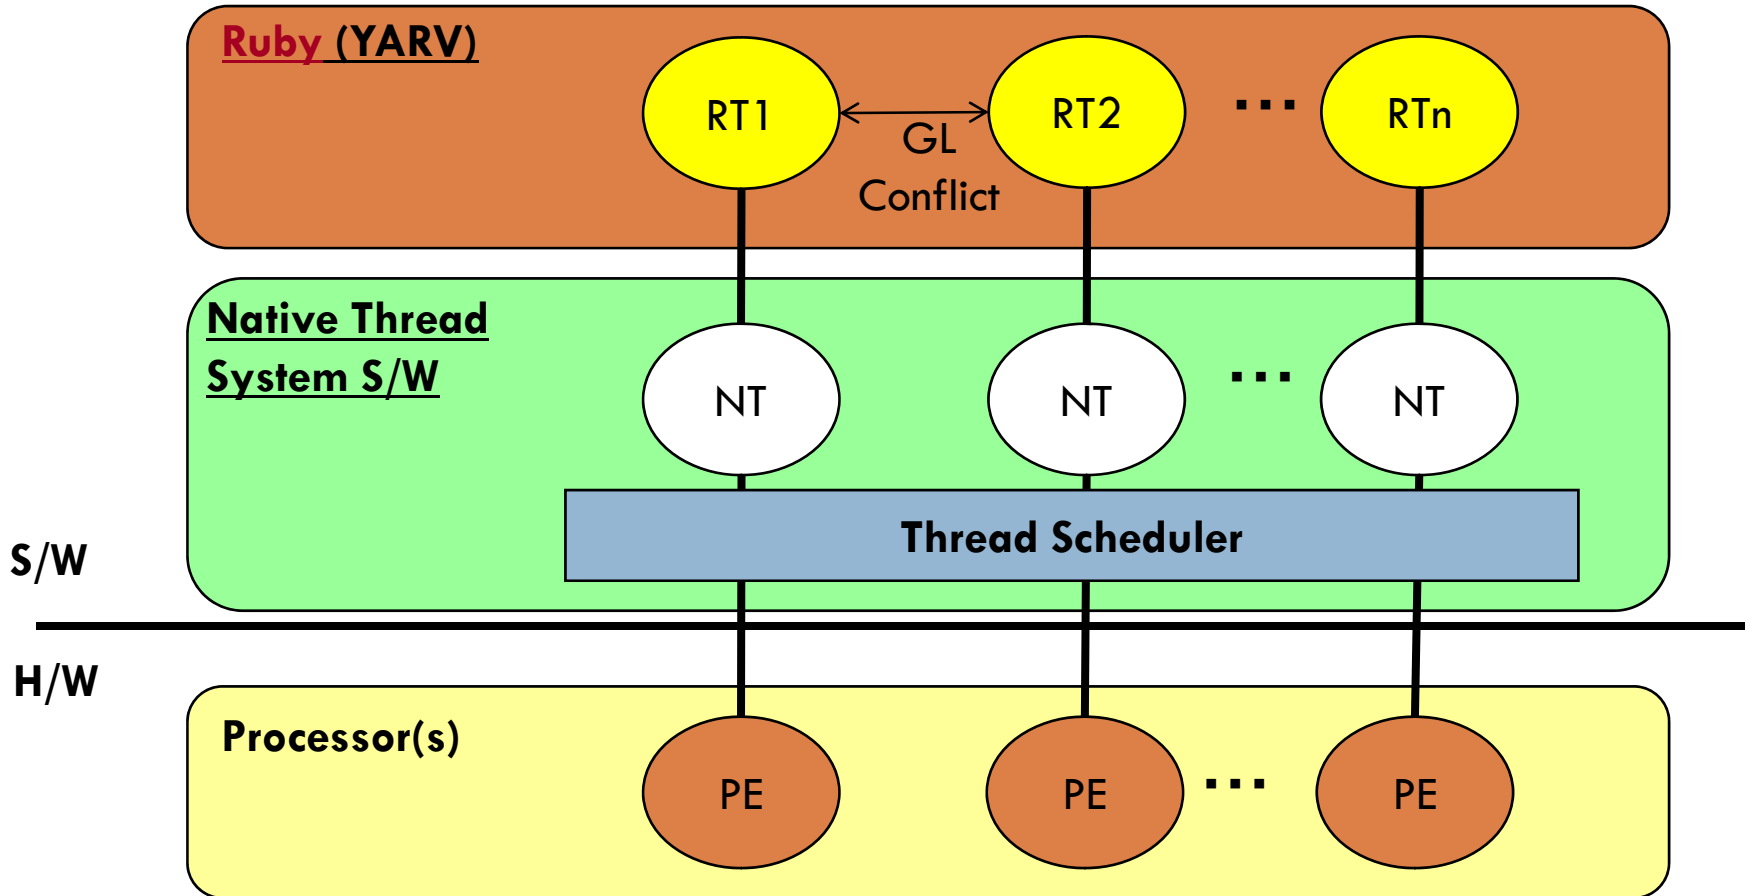
Conflict of Giant-Lock Acquisition

32

- GL Conflict →
Performance Decrement
- Limit Running CPU
 - ▣ Check GL Conflict Periodically
 - ▣ Limit their CPU
 - ▣ Using `pthread_setaffinity_np` on NTPL
 - `SetThreadAffinityMask` on Windows

Running CPU Limitation

33

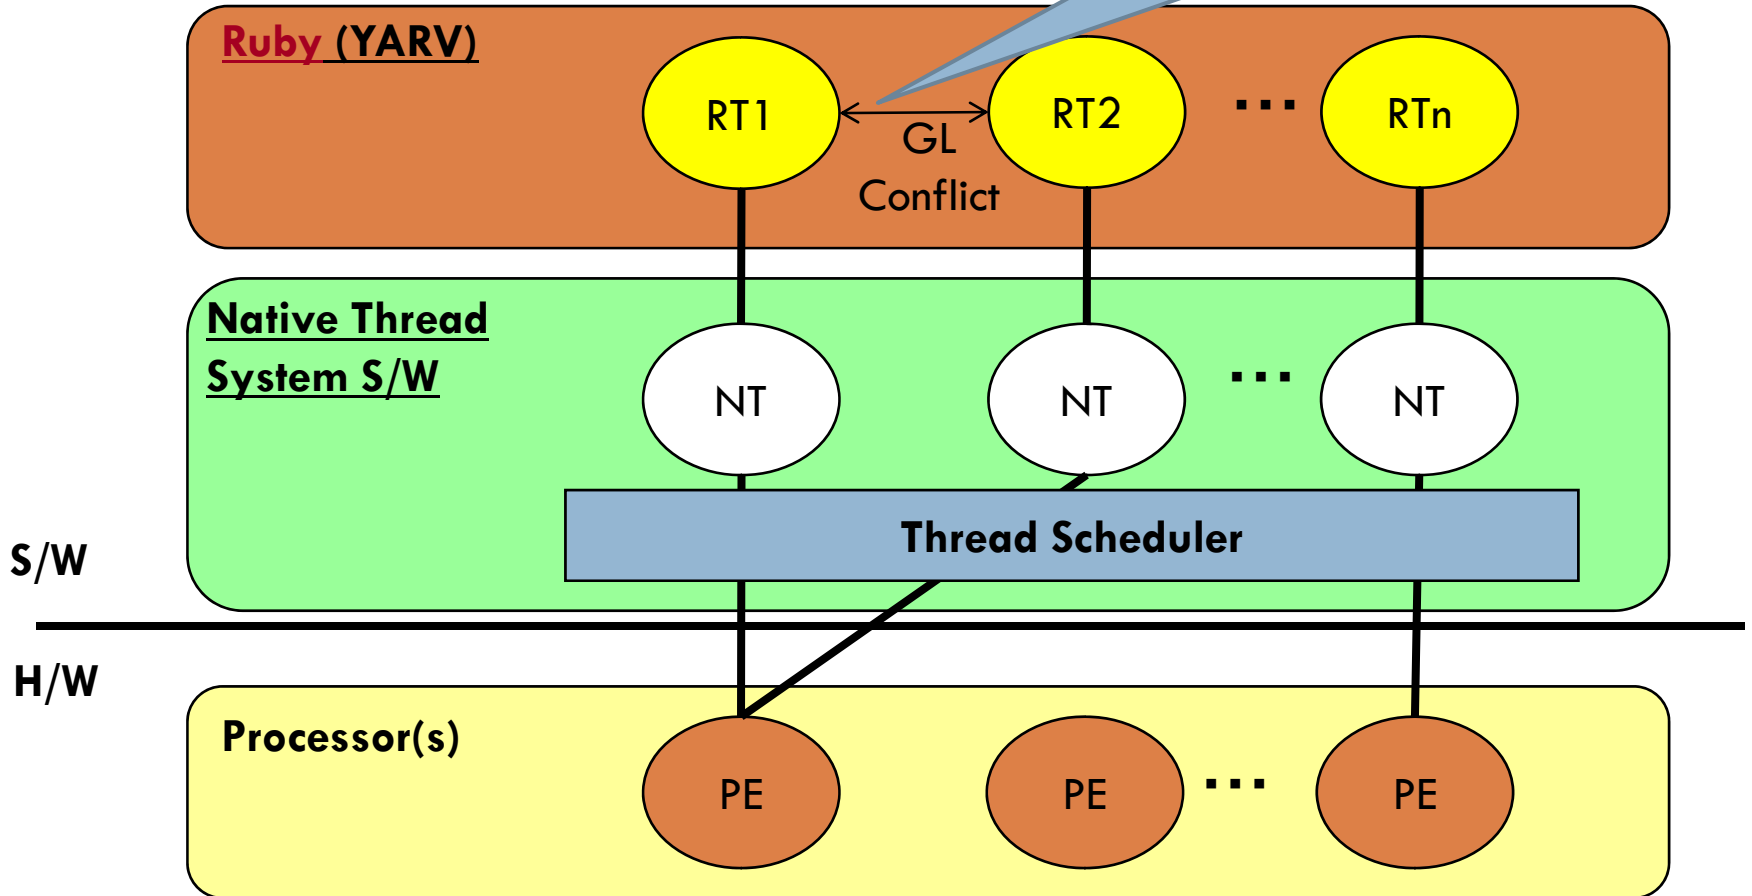


PE: Processor Element, UL: User Level, KL: Kernel Level

Running CPU Limitation

34

Force to run RT1 and RT2 in Serial
→ Avoid Conflict



PE: Processor Element, UL: User Level, KL: Kernel Level

Performance Evaluation Environment

35

- Evaluation Environment
 - ▣ CPU: Intel Xeon CPU E5335 2.0GHz
Quad core x 2 = 8 core
 - ▣ OS: GNU/Linux 2.6.18 x86_64 SMP / NPTL
 - ▣ Compiler: gcc version 4.1.2
- Ruby
 - ▣ `ruby` 1.8.6 (2007-11-02) [x86_64-linux]
- YARV Optimization
 - ▣ All except Unification, Stack caching

Evaluation

Thread control Primitives

36


- Creation, Switch: 0.1 M, Mutual Exclusion: 1 M
- Low Performance for Creation/Join because of Native Thread
 - ▣ Native Thread Overhead
 - ▣ Memory Allocation Overhead
- High Performance Synchronization
- High Performance Thread Context Switch
 - ▣ Independent Stack-Depth (1.8 depends on depth)

	Ruby (sec)	YARV (sec)	Ratio	NTPL (sec)
Creation	0.89	1.95	0.46	0.59
Mutual Exclusion	0.67	0.38	1.76	-
Switch (depth:1)	6.01	0.06	100.17	-
Switch (depth:16)	11.55	0.06	192.5	-

Evaluation

Result (Micro-benchmark)

37




- fib: fib(N) (Make new Thread if $N > 30$)
- hetero: fib + concat (1 thread)
- mandel: Mandelbrot (Big GC overhead)
- concat: String Concatenate (No Parallelism)

Evaluation

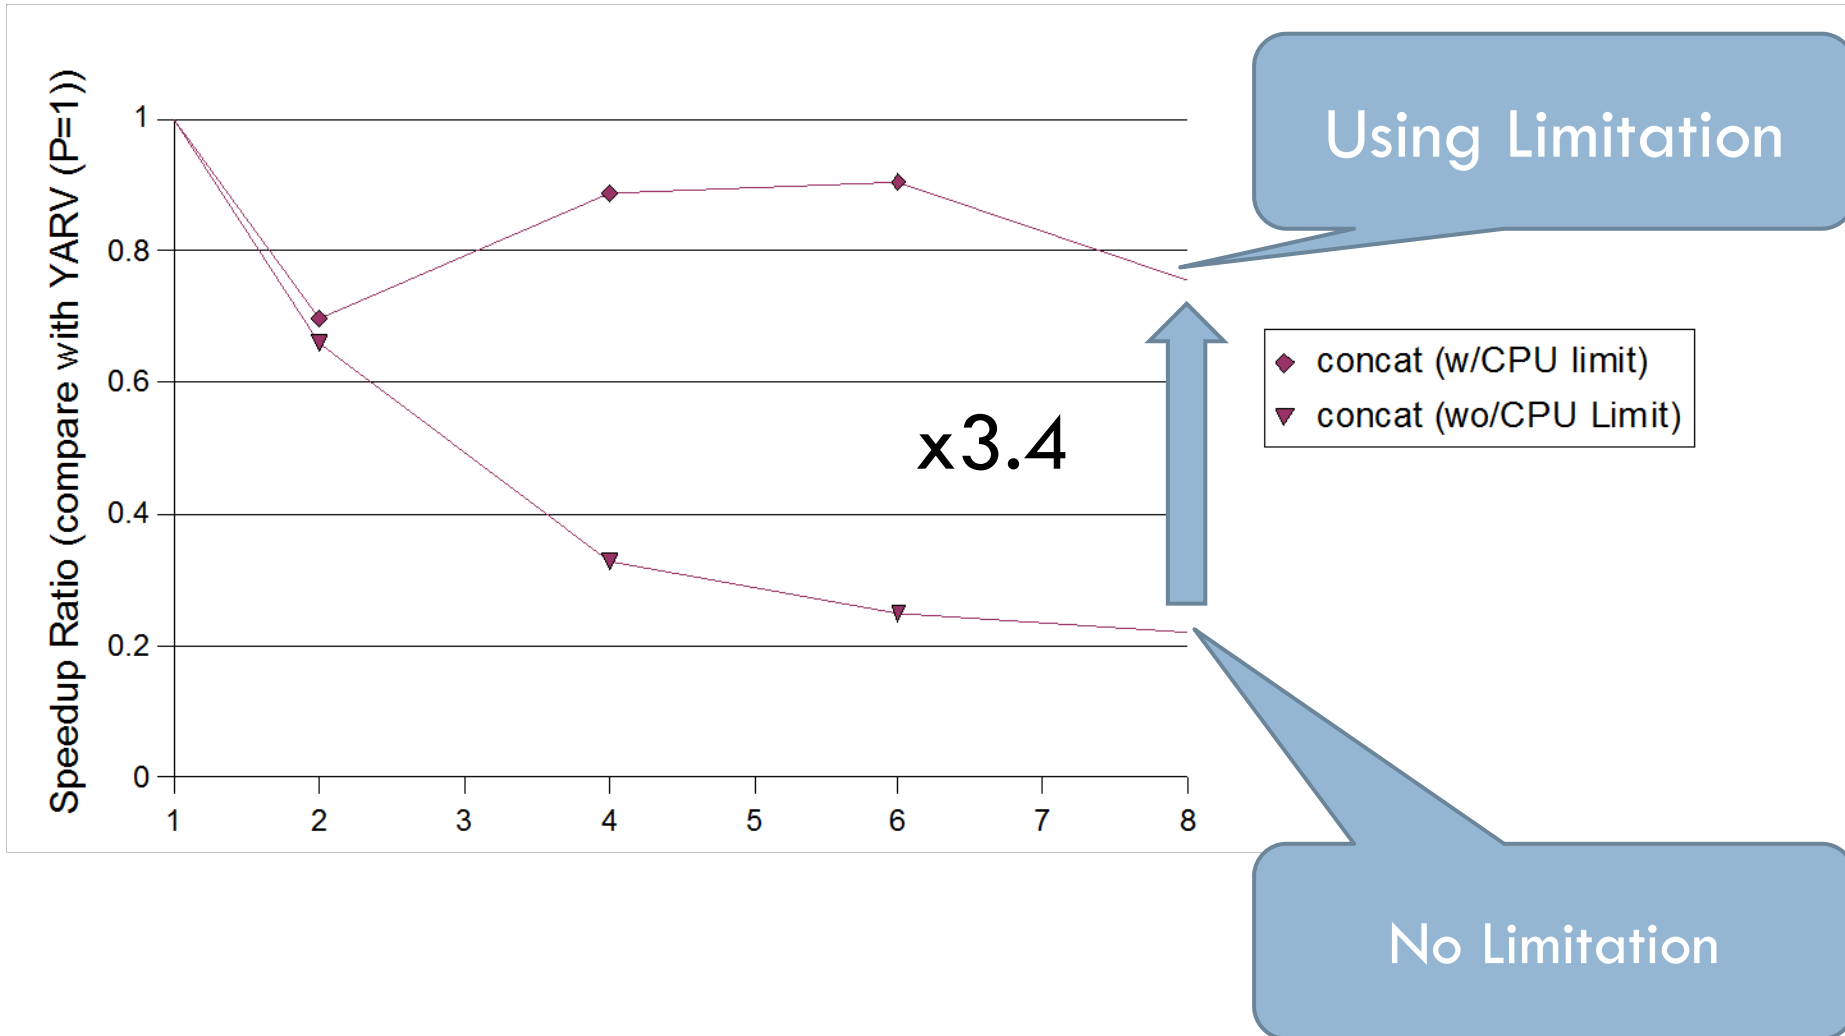
Result (Micro-benchmark)

38



CPU Limitation

39



Parallel Thread Execution

Problem

40

- Unsafe “Old C” methods
 - ▣ Replacing all is not easy task
 - ▣ Man Power Problem?
- Programming Model
 - ▣ Is Parallel Thread Application easy to write?
- Ruby 1.9
 - ▣ 1.9 support Native Thread
 - ▣ 1.9 doesn't support Parallel Thread Execution

Embedded Float Representation

41

- Float object is not immediate type
 - ▣ This means that Float is “allocated” each time
 - ▣ Ex) Fixnum, Symbol, etc
- Half execution time of Float calculation is Memory management

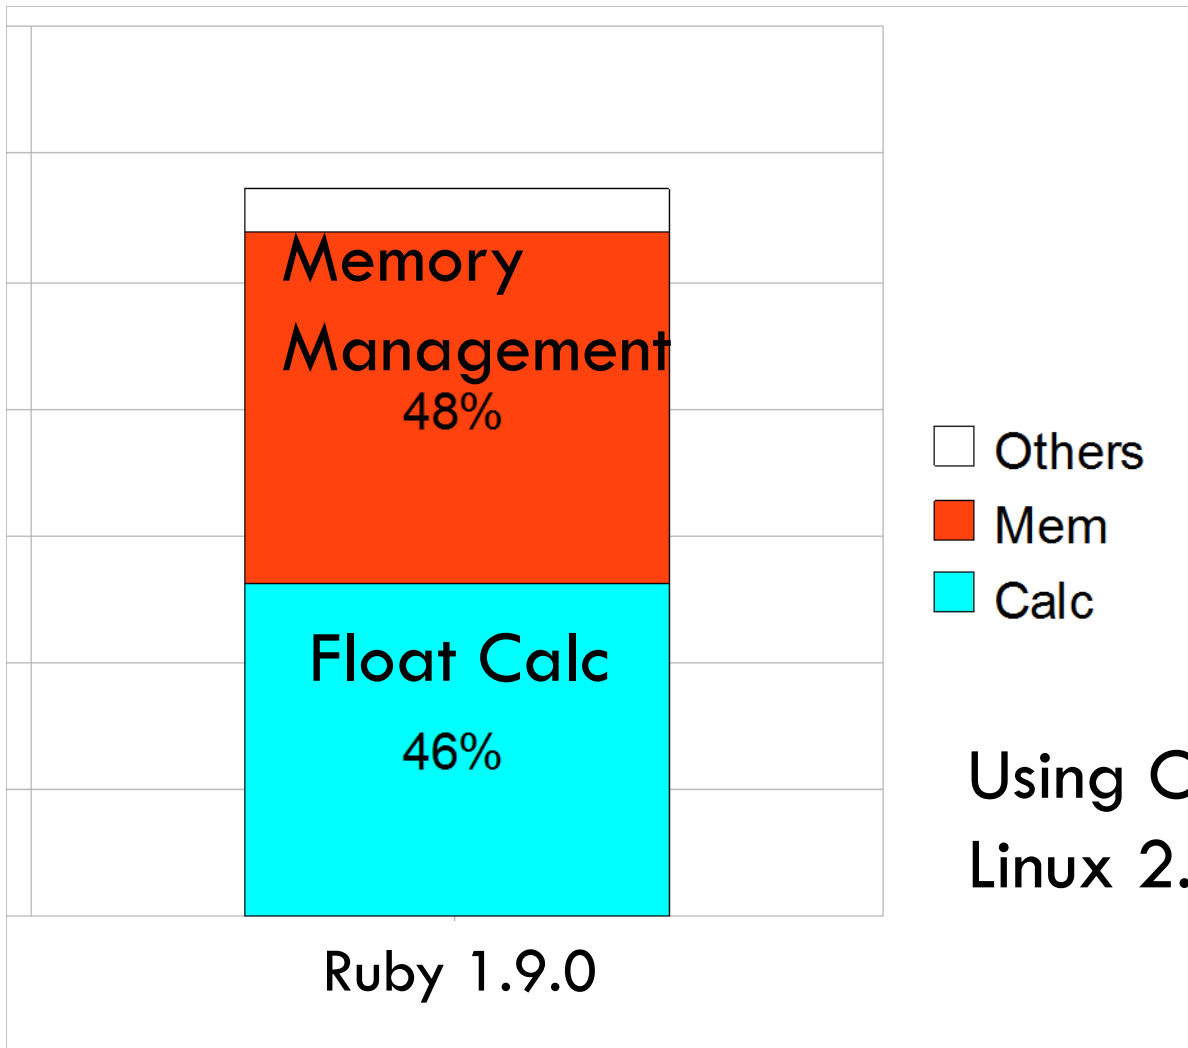
Toy-Program

42

```
i = 0; f = 0.0
while i < 30_000_000
  i += 1
  f += 0.1; f -= 0.1
  f += 0.1; f -= 0.1
  f += 0.1; f -= 0.1
  f += 0.1; f -= 0.1
end
```

List of Execution Time Toy-Program

43



- Others
- Mem
- Calc

Using OProfile
Linux 2.6 (x86_64), Xeon

Embed Float Object as Fixnum


44

- Solution: Embed 64bit Float value to 64bit CPU Pointer type!

Review

IEEE 754 Double Precision Representation

45



$$-1^s 2^{e-1024} m \quad (m = 1 + \sum_{i=0}^{51} \frac{b_i}{2^{52-i}})$$

Discussion

How to Embed 64 bit Double?

46

- VALUE embed Object doesn't need memory overhead
- **64bit CPU have 64 bit pointer type**
→ **Use 64 bit CPU**
- At least we need 1 bit for TAG bit
 - ▣ From Mantissa?
 - Decrease Precision
 - ▣ From Exponential?
 - Decrease Representation Range

Proposal

47

- From Exponential

But Store Float in Heap if it is Out of Range

→ **Save a Range and Precision**

- ▣ Often used $010000000000b \sim 101111111111b$
($2^{-512} \sim 2^{511}$)

- ▣ If Float is out of range, alloc from Heap
- ▣ Float Out of this Range is Rare Number on Numeric Application -> Practical Solution

Proposal

Real Program

48

- 1) Check the Range of "e" (512~1535)
- 2) IEEE745 double -> Float (Encoding)
- 3) Float -> IEEE745 double (Decoding)



Proposal

Float Representation with Tag

49

- e: 100000000000b ~ 01111111111b
- Note that if "b62 != b61", we can embed.
- On Ruby, Tag is at LSB
→ Left Rotate **3bit**
- b63 b62 b61 b60 ... b0 → 3 bit rotate
b60 ... b0 b63 b62 b61

Proposal

Real Encoding Code

50

```
VALUE rb_float_new(double dbl) {
    VALUE v1 = BIT_DOUBLE2VALUE(dbl);
    VALUE v2 = BIT_ROTL(v1, 3);
    if ((v2 + 1) & 0x02) // check lower 2 bits
        return v2 | 0x01;    // Embed tag
    else {
        if (dbl == 0) // 0.0
            return ruby_float_zero;
        else // alloc from Heap
            return rb_float_new_in_heap(dbl);
    }
}
```

Proposal

Real Decode Code

51

```
double RFLOAT_VALUE(VALUE v) {
    if (v & 1) {
        VALUE v1 = v ^ ((v >> 1) & 1);
        VALUE v2 = BIT_ROTTR(v1, 3);
        return BIT_VALUE2DOUBLE(v2);
    }
    else
        return RFLOAT(v)->float_value;
}
```

Implementation

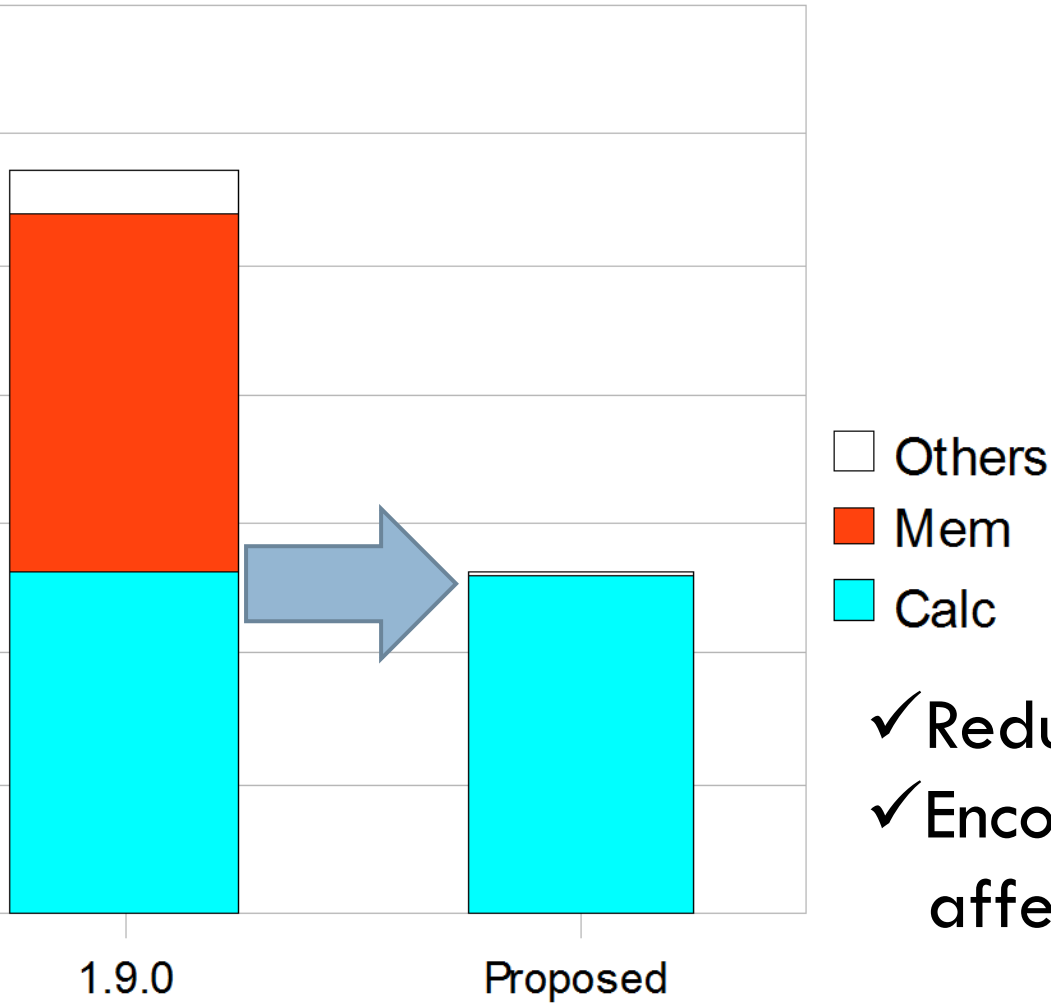
52

- Ruby 1.9.0-0
- Easy to Implementation
- No Spec Changes

Evaluation

Toy-Program

53

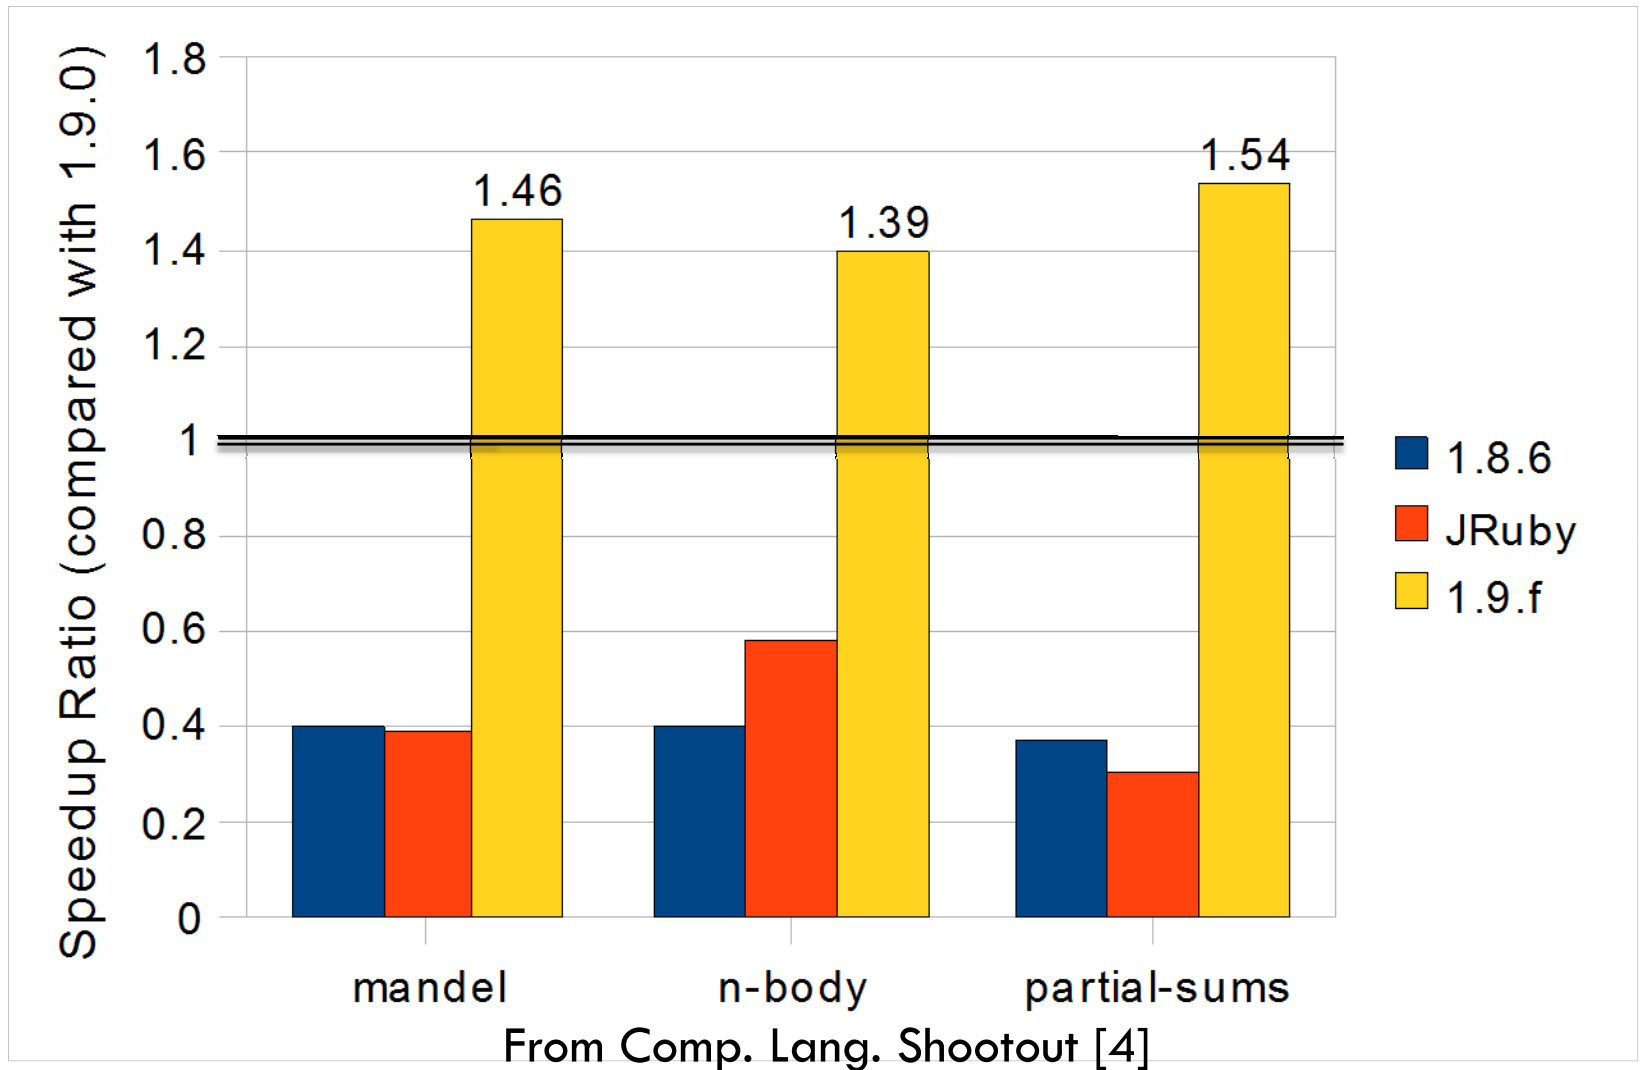


- ✓ Reduce Mem Time
- ✓ Encode/Decode don't affect to Performance

Evaluation

Compared with other Ruby Impl.

54



Evaluation

Compared with Other Languages

55


- Evaluate with other languages
- Note that C/Java use “volatile” to avoid optimization

```
i = 0; f = 0.0
while i < 30_000_000
  i += 1
  f += 0.1; f -= 0.1
  f += 0.1; f -= 0.1
  f += 0.1; f -= 0.1
  f += 0.1; f -= 0.1
end
```

Evaluation

Compared with Other Languages

56



[PLAN]

JIT Compiler

57

- I'm re-designing to reduce VM instructions to impl. it easy
 - ▣ Current VM has about 50 instructions
 - ▣ Ex) "definemethod" move to "Method"

[PLAN]

Pre-Compiler

58

- YARV VM Generator helps us
 - ▣ Ruby to “Pre-compiled”
 - ▣ Ruby to “C”
- Purpose
 - ▣ Eliminate Loading-Time
 - ▣ More aggressive optimization
 - ▣ Obfuscation (?)

[PLAN]

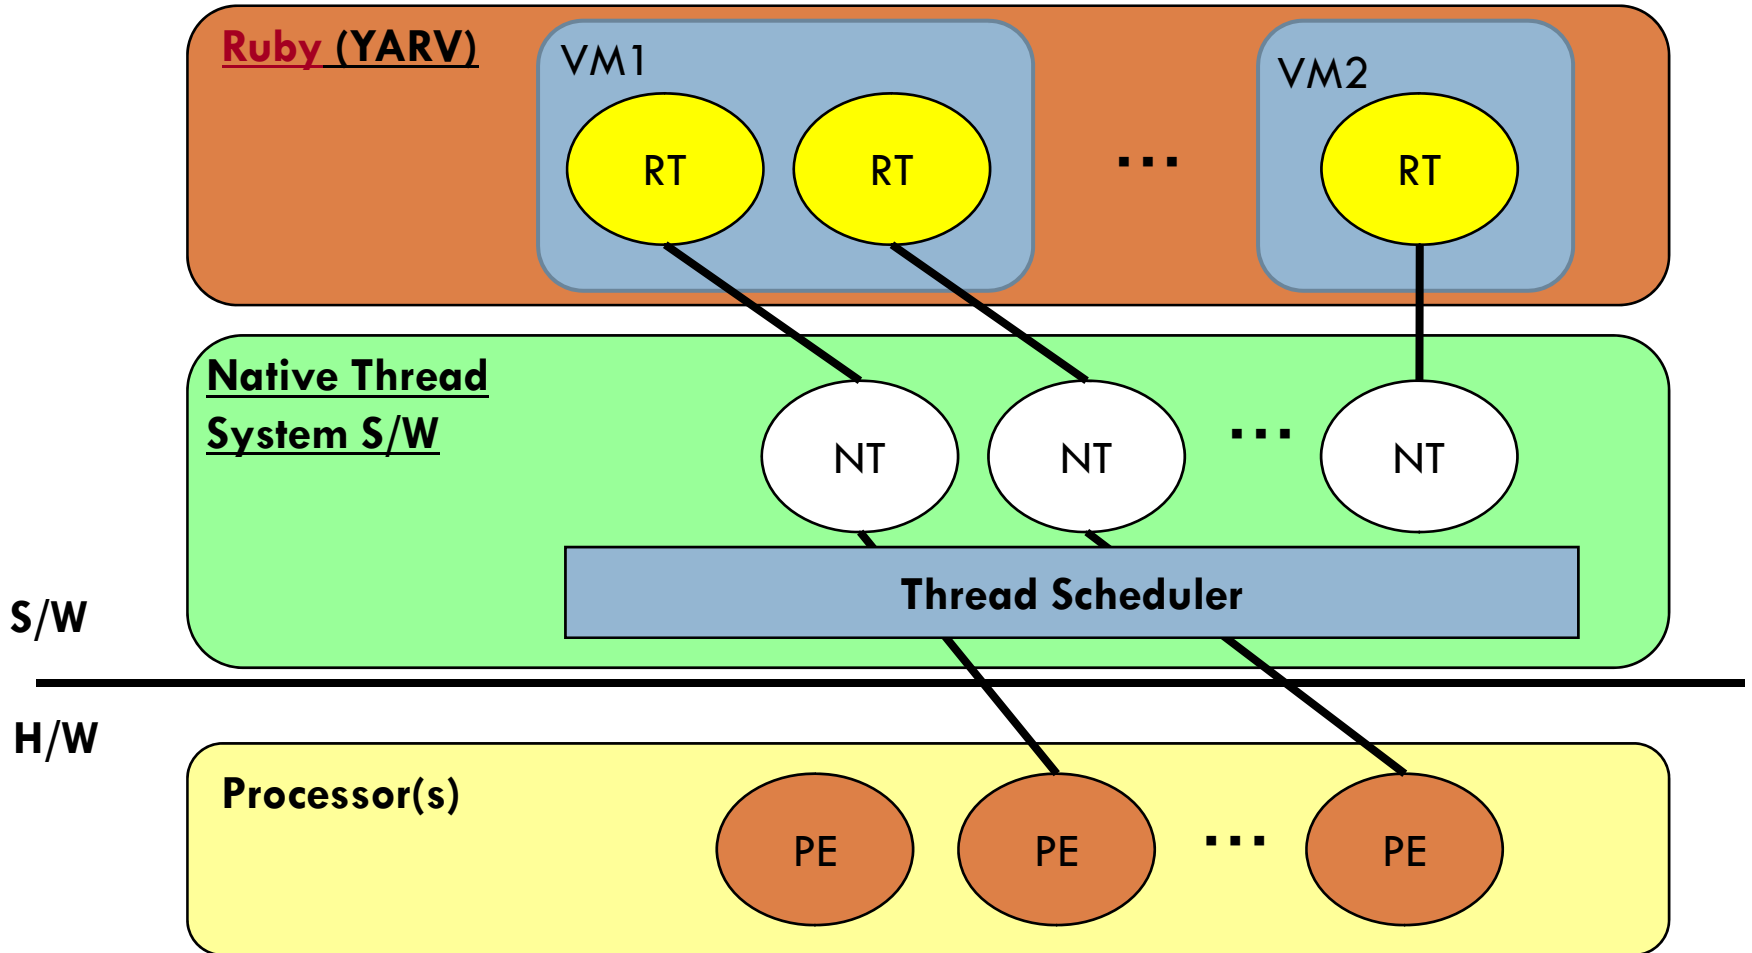
Multi-VM Creation

59

- Purpose
 - ▣ Embed Ruby into Application
 - mod_ruby, ...
 - ▣ Sand-box

Multi-VM Overview

60



PE: Processor Element, UL: User Level, KL: Kernel Level

Multi-VM Points

61

- How to control VMs?
 - ▣ C Level? → Designed with Nobuyoshi Nakada
 - Making new VM is need only 3 lines
 - ▣ Ruby Level?
- How to share environments Inter VM
 - ▣ Trade off between Isolation and Util.

[PLAN]

Customizable VM Core

62

- Ruby is toooooo FAT to use XXX purpose
 - ▣ Many Many Convenience Methods/Feature
- Need Re-design Ruby Core

[PLAN]

Debug/Profile Support Feature

63

- Only cheap Debugger/Profiler API
 - ▣ `set_trace_func`, `Thread#set_trace_func`
- Introduce “break” instruction?

Future Work

Benchmark

64

- Current Benchmark suits is for checking YARV Performance
 - ▣ Focus to YARV optimization
 - ▣ Toy benchmarks
- We need more pragmatic benchmarks

Summary

65

- YARV Merged into Ruby 1.9
- I'm working at Advanced VM Topics
 - ▣ Performance
 - Parallel Thread Execution
 - Embedding Float Value
 - JIT Compiler
 - Pre-Compiler
 - "Ruby to Compiled file" Compiler
 - "Ruby to C" Compiler
 - ▣ New Feature
 - Multi-VM Creation
 - Customizable Ruby Core
 - Debug/Profile support feature

Summary

66

Merging YARV is not a goal,
but a start

VM is a very flexible
infrastructure to hack

One more thing...

Sasada Lab@U-Tokyo

68

- I'll make a laboratory from 2009, Apr
 - ▣ Department of Creative Informatics, Graduate School of Information Science and Technology, The University of Tokyo
 - ▣ Graduate School
 - ▣ Lab is at Akihabara, Tokyo, Japan
- Unfortunately I can't employ you as Research Assistant
 - ▣ There are not enough grants in Japan...

Research Topics

69

- Ruby, Ruby, Ruby, PHP, Ruby, Ruby
- Ruby, Ruby, Ruby, Ruby, Python, Ruby
- Ruby, Perl, Ruby, Ruby, Ruby, Ruby
- Ruby, Ruby, Ruby, Ruby, Lua, Ruby
- Ruby, Ruby, Java, Ruby, Ruby, Ruby
- Implementation of Programming Language
- Operating System / Processor Architecture
- Software development

Sasada Lab.

70

```
□ if you.have_interest(  
  :Japan, Tokyo, :Akihabara,  
  :Japanese,  
  :Ruby, :Research, :Development  
)  
  you.send_mail_to "ko1 at atdot dot net"  
end
```

Thank you for your attention!
Any Questions?

ささだ こういち

Koichi Sasada

Ko1 at atdot dot net