

# mruby machine

An Operating System for Microcontroller



HASUMI Hitoshi @hasumikin

*September 4-5, 2020*



# Me

---

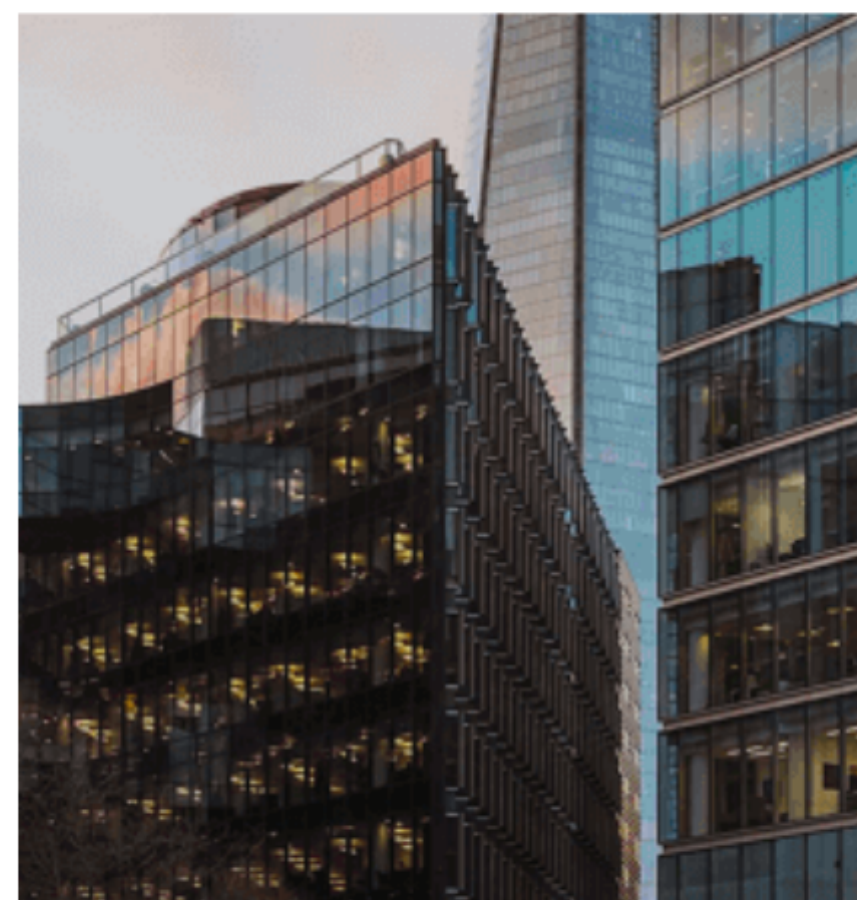
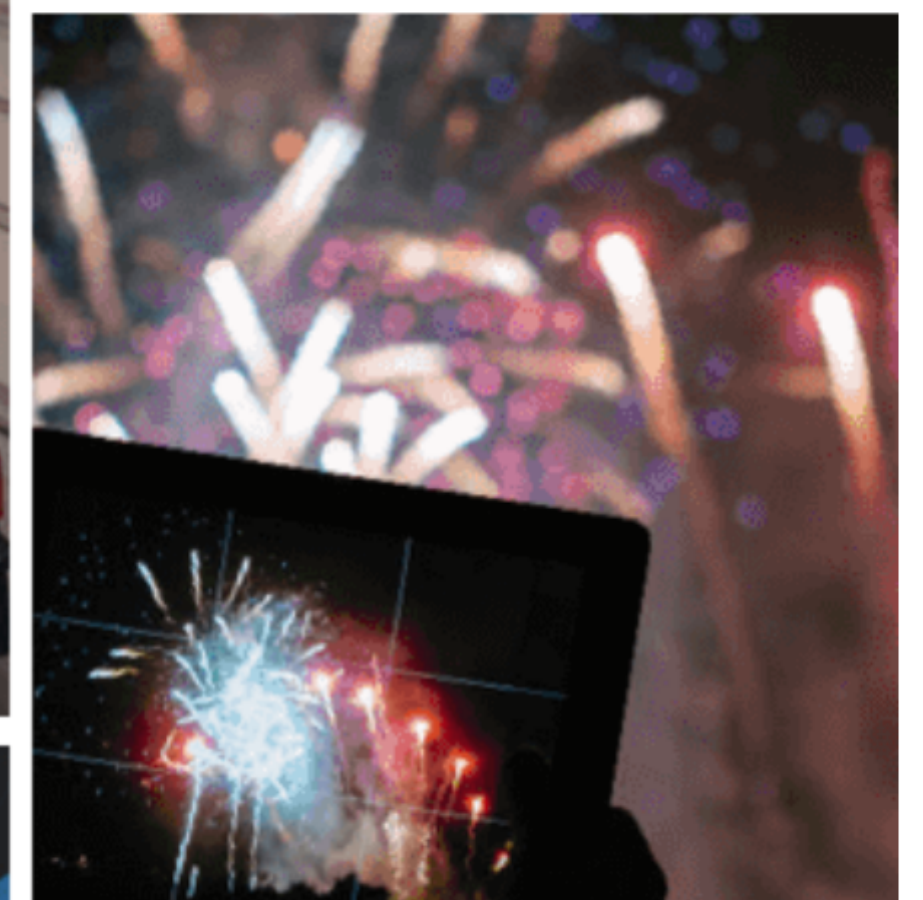
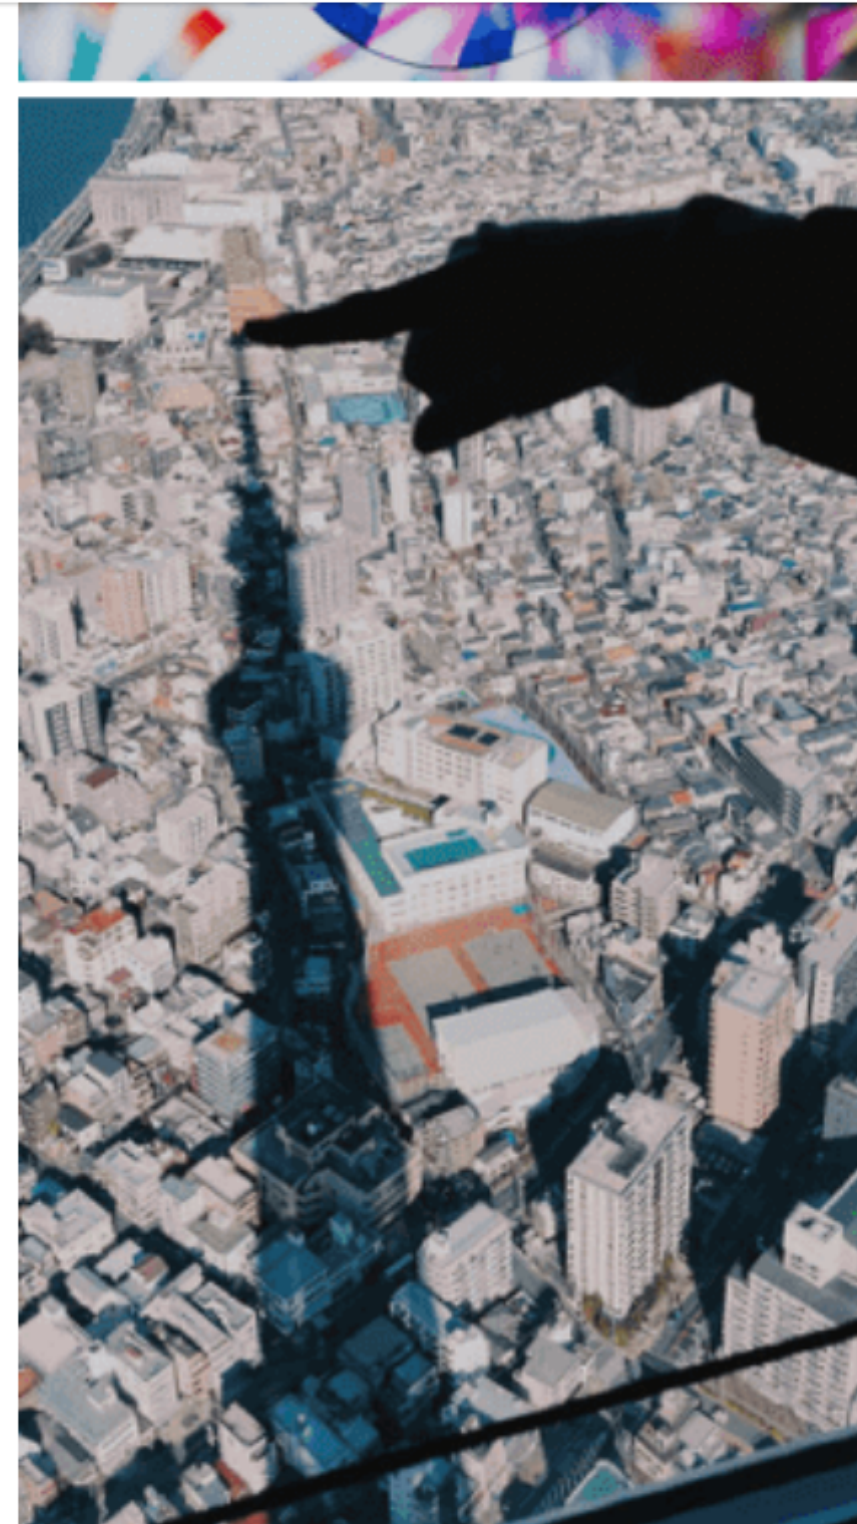
- 🌀 HASUMI Hitoshi (羽角 均)  
@hasumikin (ハスミキン)
- 🌀 From Ruby City Matsue,  
a holy place of 🍷 (熱燗)
- 🌀 Developer on Rails  
(レールに乗った開発者)
- 🌀 Coffee 🍵  
Curry 🍛  
Soba (NoEmojiError)





# With offices in 26 cities around the world, we know your market

Tokyo, Osaka, Matsue,  
Chengdu, Qingdao,  
Shanghai, Beijing, Hanoi,  
Danang, Singapore, Dhaka,  
Cebu, Manila, Dubai,  
Copenhagen, Aarhus,  
London, Manchester,  
Amsterdam, Prague,  
Bangkok, Berlin, New York,





## Chapter 1



# Background

# Motivation 🤖 MicroPython

---

*Generally, we keep minimal configuration of MicroPython under **80K** of ARM Thumb2 code (which includes **compiler** and **interactive prompt**, ...). That means a Cortex-M microcontroller with **128KB** of **flash** can host a minimal version together with some hardware drivers.*

*[cited from ``https://github.com/micropython/micropython/wiki/FAQ'`]*

# Motivation 😐 MicroPython

---

*Regarding **RAM** usage, (...) **8KB** is minimal amount to run simple scripts. As Python is interpreted high-level language, the more memory you have, the more capable applications you can run. The reference MicroPython board, PyBoard, has **128KB** of RAM.*

*[cited from ``https://github.com/micropython/micropython/wiki/FAQ``]*

# Motivation 🙄 mruby and mruby/c

---

- 🌀 Flashing mruby application into 128KB ROM and running it on 128KB RAM are difficult
- 🌀 mruby/c runs on 64KB RAM but it doesn't have mruby compiler
- 🌀 You have to compile Ruby script into VM code before embedding

# Motivation 🤨

---

MicroPython:  
"Overwhelming, we are!"

「圧倒的じゃないか、我が軍は」



# Motivation 🤔

---

Would it be great if we have a  
smaller mruby compiler?

# Target Microcontroller

---

- ⑨ PSoC5LP: Reference board of mruby/c
  - ⑨ Arm Cortex-M3 32bit processor
  - ⑨ 64KB RAM
  - ⑨ 256KB ROM
- ⑨ It will be comparable to MicroPython if Ruby runs on PSoC5LP with compiler



Demo

## Chapter 2



# Terms and Tools



# Terms

---

- ⑨ mruby
  - ⑨ A Ruby implementation for general embedded purpose
- ⑨ mrbc
  - ⑨ The original mruby compiler bundled with mruby
- ⑨ mruby/c
  - ⑨ Another implementation of mruby VM for one-chip microcontroller
  - ⑨ It doesn't have compiler

# Terms

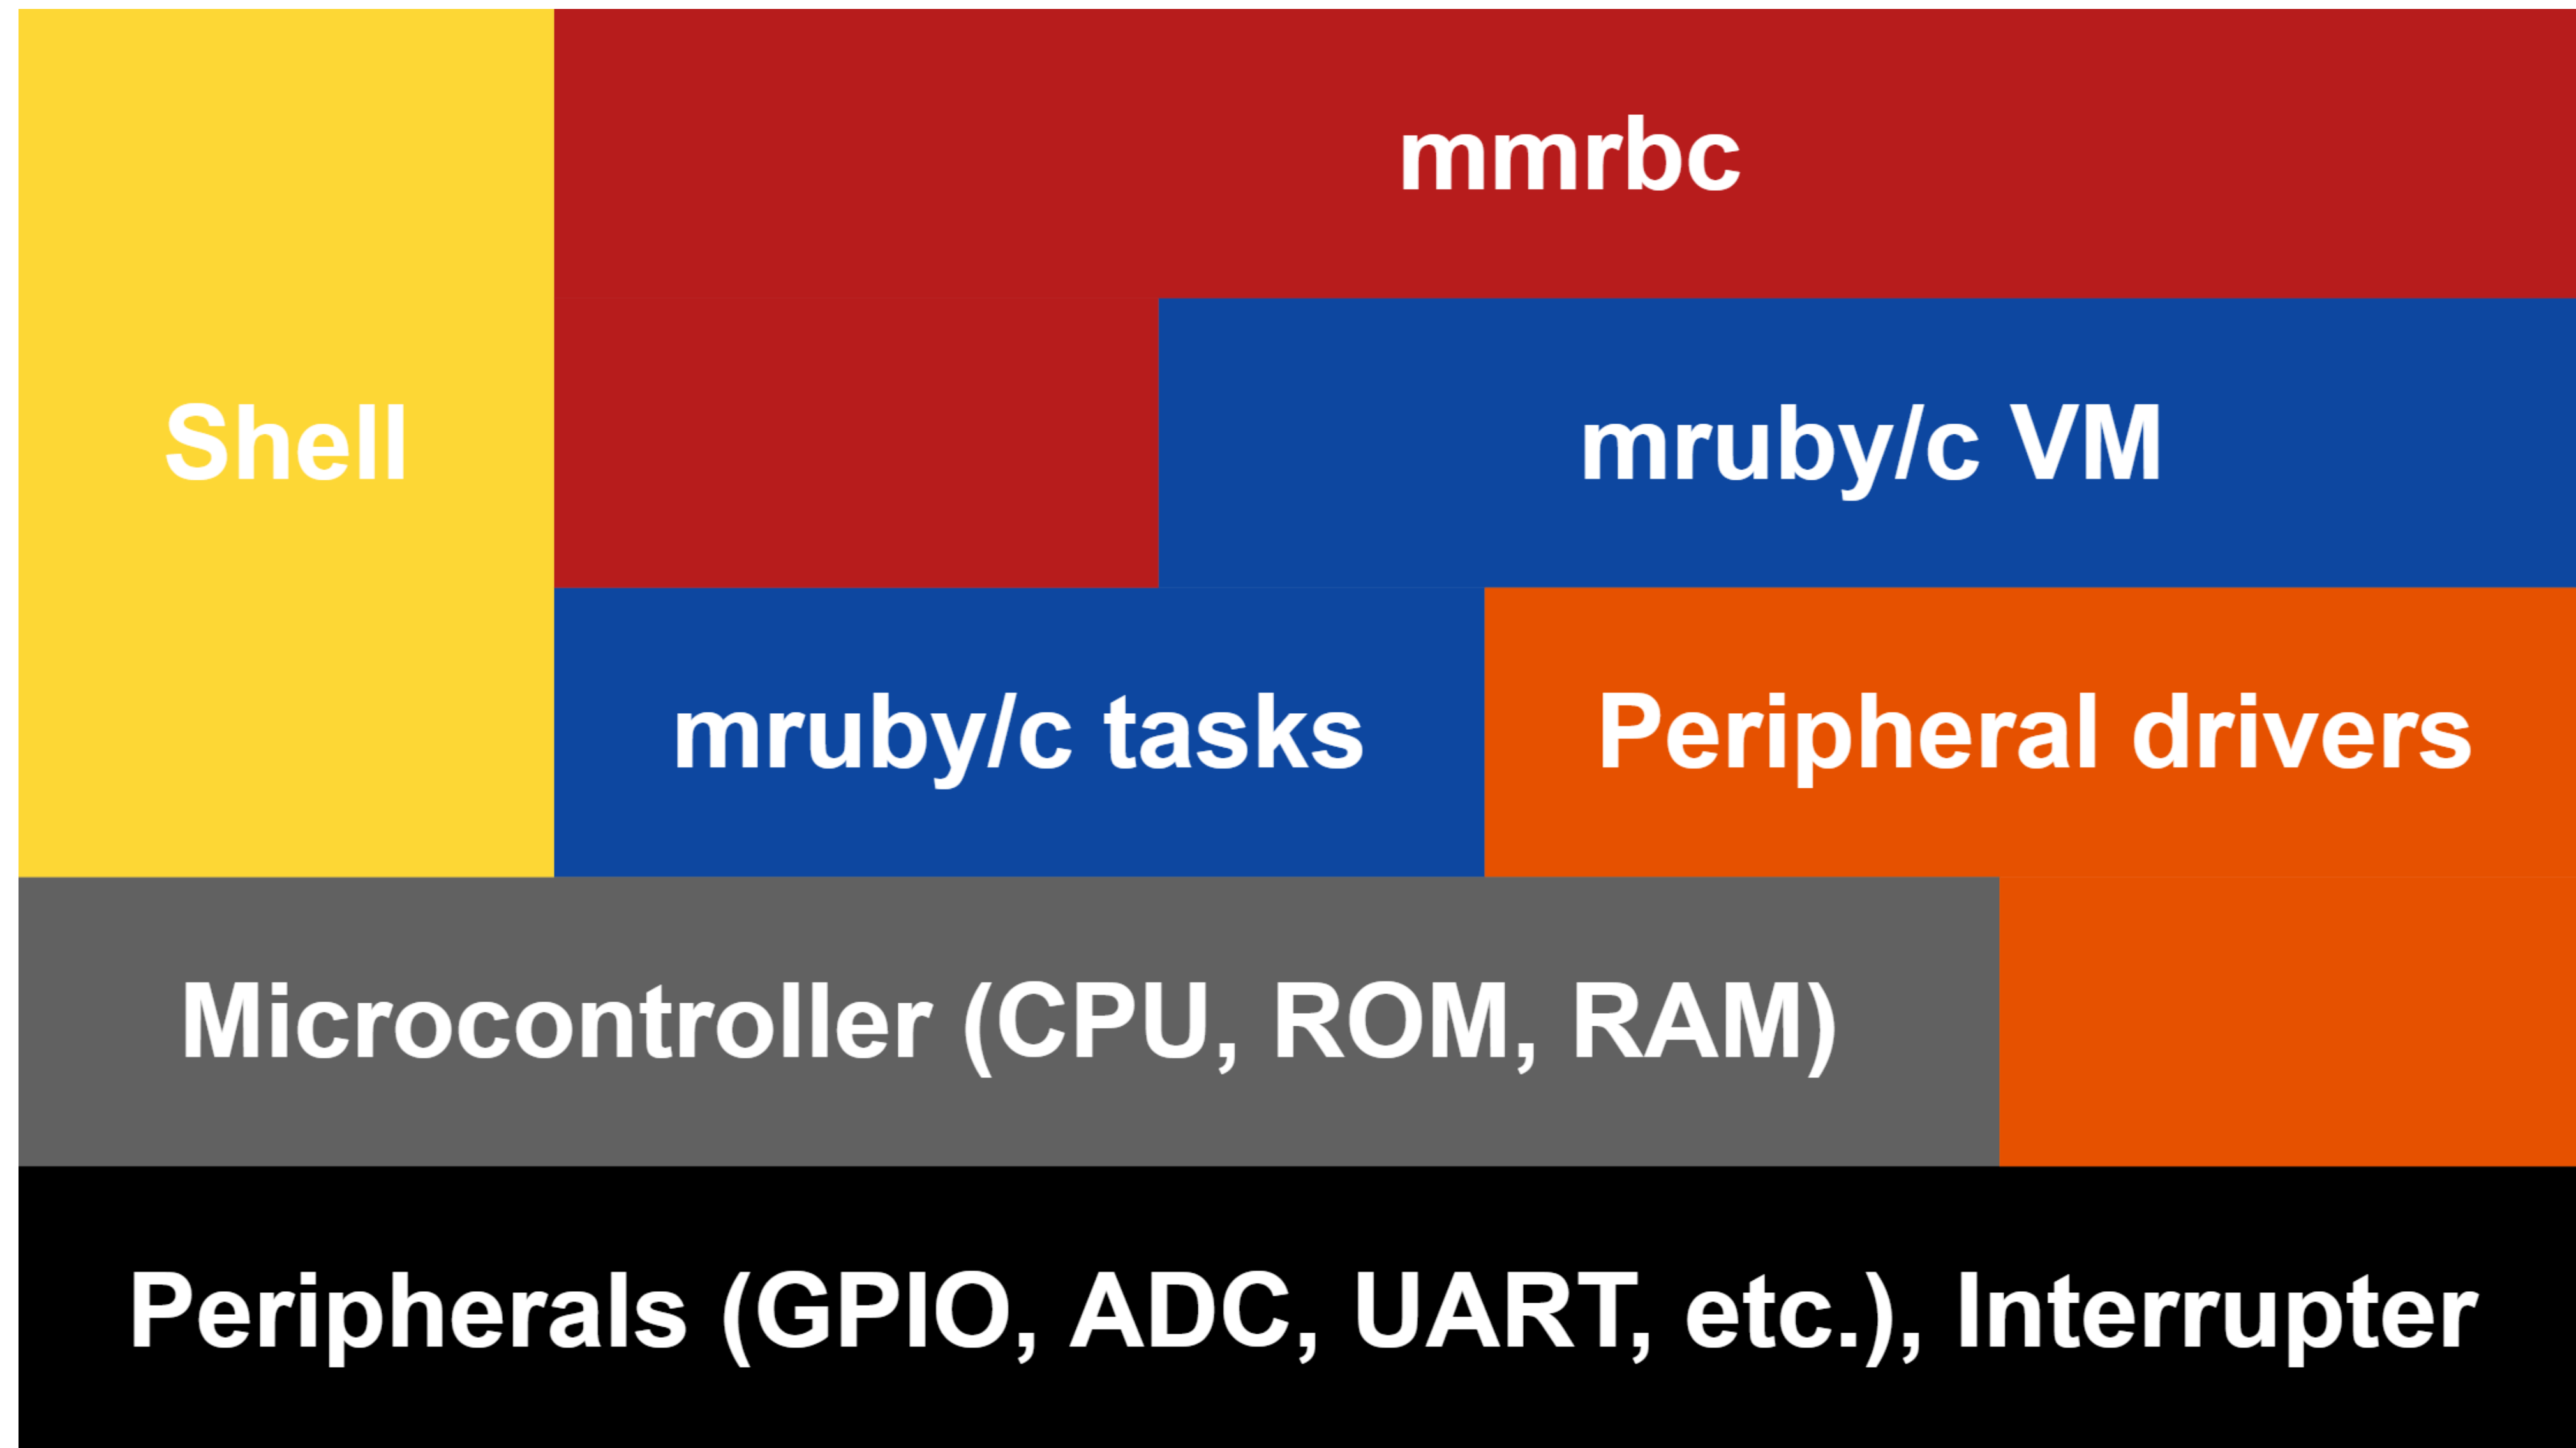
---

- ⑨ mmrbc (mini mruby compiler)
  - ⑨ Another mruby compiler for one-chip microcontroller
- ⑨ mmruby (mini mruby) = mmrbc + mruby/c
- ⑨ mruby machine
  - ⑨ A microcontroller embedded with integration of mmruby, CUI shell and peripheral drivers



# mruby machine

---



# Cross compilation tools for bare metal Arm

---

- ⑨ arm-none-eabi-gcc, arm-none-eabi-ar, arm-none-eabi-ld, arm-none-eabi-strip, etc.
- ⑨ What is "arm-none-eabi-" ?
  - ⑨ arm: Arm architecture
  - ⑨ none: Without OS (Bare metal)
  - ⑨ eabi: Embedded Application Binary Interface



# Three tool-chains in my laptop (Intel 64)

---

1. (x86\_64-linux-gnu-)gcc

Usual compilation in daily development

2. arm-linux-gnueabi-gcc

Binary runs on QEMU (machine emulator)

3. arm-none-eabi-gcc (--with-newlib)

To make a final library archive (libmmrbc.a) for bare metal Arm

# Docker image

---

(('tag:center'))

[https://hub.docker.com/r/  
hasumikin/arm-none-eabi-tools](https://hub.docker.com/r/hasumikin/arm-none-eabi-tools)



# C Library

---

- ⑨ glibc (GNU C Library)
- ⑨ Newlib
  - ⑨ Another C library especially for embedded system
  - ⑨ Some implementation is missed. eg) regex.c
- ⑨ Newlib-nano
  - ⑨ Even smaller version of Newlib for Arm processor
- ⑨ Yet other libraries
  - ⑨ EGLIBC, Clibc, sglibc, etc.

# Chapter 3



mmrbc

# Repo

---

<https://github.com/hasumikin/mmruby>

# Features of mmrbc

---

- ⑨ Portable and Less dependent
  - ⑨ Doesn't depend on mruby nor mruby/c
  - ⑨ (To the contrary, mrbc depends on mruby itself)
  - ⑨ Depends on only standard C library including Newlib(-nano)
  - ⑨ Can be integrated with both mruby and mruby/c



# Features of mmrbc

---

- ⑨ Small (by linking to Newlib-nano)
- ⑨ Fits in less than 256KB ROM (including mruby/c VM and some application code)
- ⑨ Runs enough on less than 128KB RAM

# Features of mmrbc

---

- ⑨ **Lemon** as a parser generator instead of Yacc/Bison
- ⑨ Originally a parser generator of SQLite

# Yacc/Bison

---

- ⑨ LALR(1) parser generator, typical choice of parser generator
- ⑨ Yacc (Barkley Yacc)
  - ⑨ Non-reentrant
  - ⑨ Parser-Calls-Lexer style (CRuby and mruby use it)
- ⑨ Bison (Gnu Bison. Modernized version of Yacc)
  - ⑨ Reentrant, Thread-Safe and Lexer-Calls-Parser style
  - ⑨ (You can also select old Yacc style)

# Lemon

---

- ⑨ Reentrant
- ⑨ Lexer-Calls-Parser style



# Parser-Calls-Lexer style (Yacc)

---

```
yyparse(p);

/* in y.tab.c (generated by Yacc) */
int yyparse(parser_state *p) {
    yytoken = yylex(yylval, p);
    /* makes AST according as you defined */
    ...
    goto yyhoge fuga;
    ...
}

/* yylex() is called by parser */
int yylex(void *lval, parser_state *p) {
    # Find next token and return token type
    return tokenType;
}

GenerateVMCode(p->ast);
```

# Lexer-Calls-Parser style (Lemon)

---

```
/* in your own Lexer */
while (hasMoreToken(rubyScript)) {
    token = getToken(rubyScript);
    Parse(parser, tokenType, token);
}
Parse(parser, 0, "");

/* in parse.c (generated by Lemon) */
void Parse(void *yyp, int yymajor, ...) {
    /* makes AST according as you defined */
    ...
}

GenerateVMCode(p->ast);
```

# Lemon

---

- ⑨ LALR(1) parser generator as well
- ⑨ Reentrant
- ⑨ Lexer-Calls-Parser style
- ⑨ Generated parser is, I'm still not sure though,
  - ⑨ Faster than Yacc/Bison (according to its document)
  - ⑨ Smaller code (according to someone)
- ⑨ Grammar syntax of `parse.y` is different from Yacc/-Bison

# Lemon

---

<https://sqlite.org/src/doc/trunk/doc/lemon.html>



## Chapter 4:



# How smaller than mrbc?

# RAM consumption

---

- ⑨ mruby's branch
  - ⑨ `mruby3` as of August 14th

# RAM consumption

---

- ⑨ To make things simple, both mrbc and mmrbc are
  - ⑨ built for x86\_64 architecture
  - ⑨ using `malloc()` and `free()` of glibc
    - ⑨ (mmrbc usually uses `mrbc_alloc()` to manage heap area)
  - ⑨ Linux file system to load Ruby script
    - ⑨ (microcontroller doesn't have file system)
- ⑨ As a result, we can simply compare with **valgrind**

# RAM consumption

---

- ⑨ Analysis command looks like:

```
valgrind \  
  --tool=massif \  
  --stacks=yes \  
  path/to/(mrbc|mmrbc) path/to/script.rb
```

- ⑨ Showing the result:

```
ms_print massif.out.xxxx | less
```



# RAM consumption: Case 1

---

Ruby code:

```
# hello_world.rb  
puts "Hello World!"
```

Output:

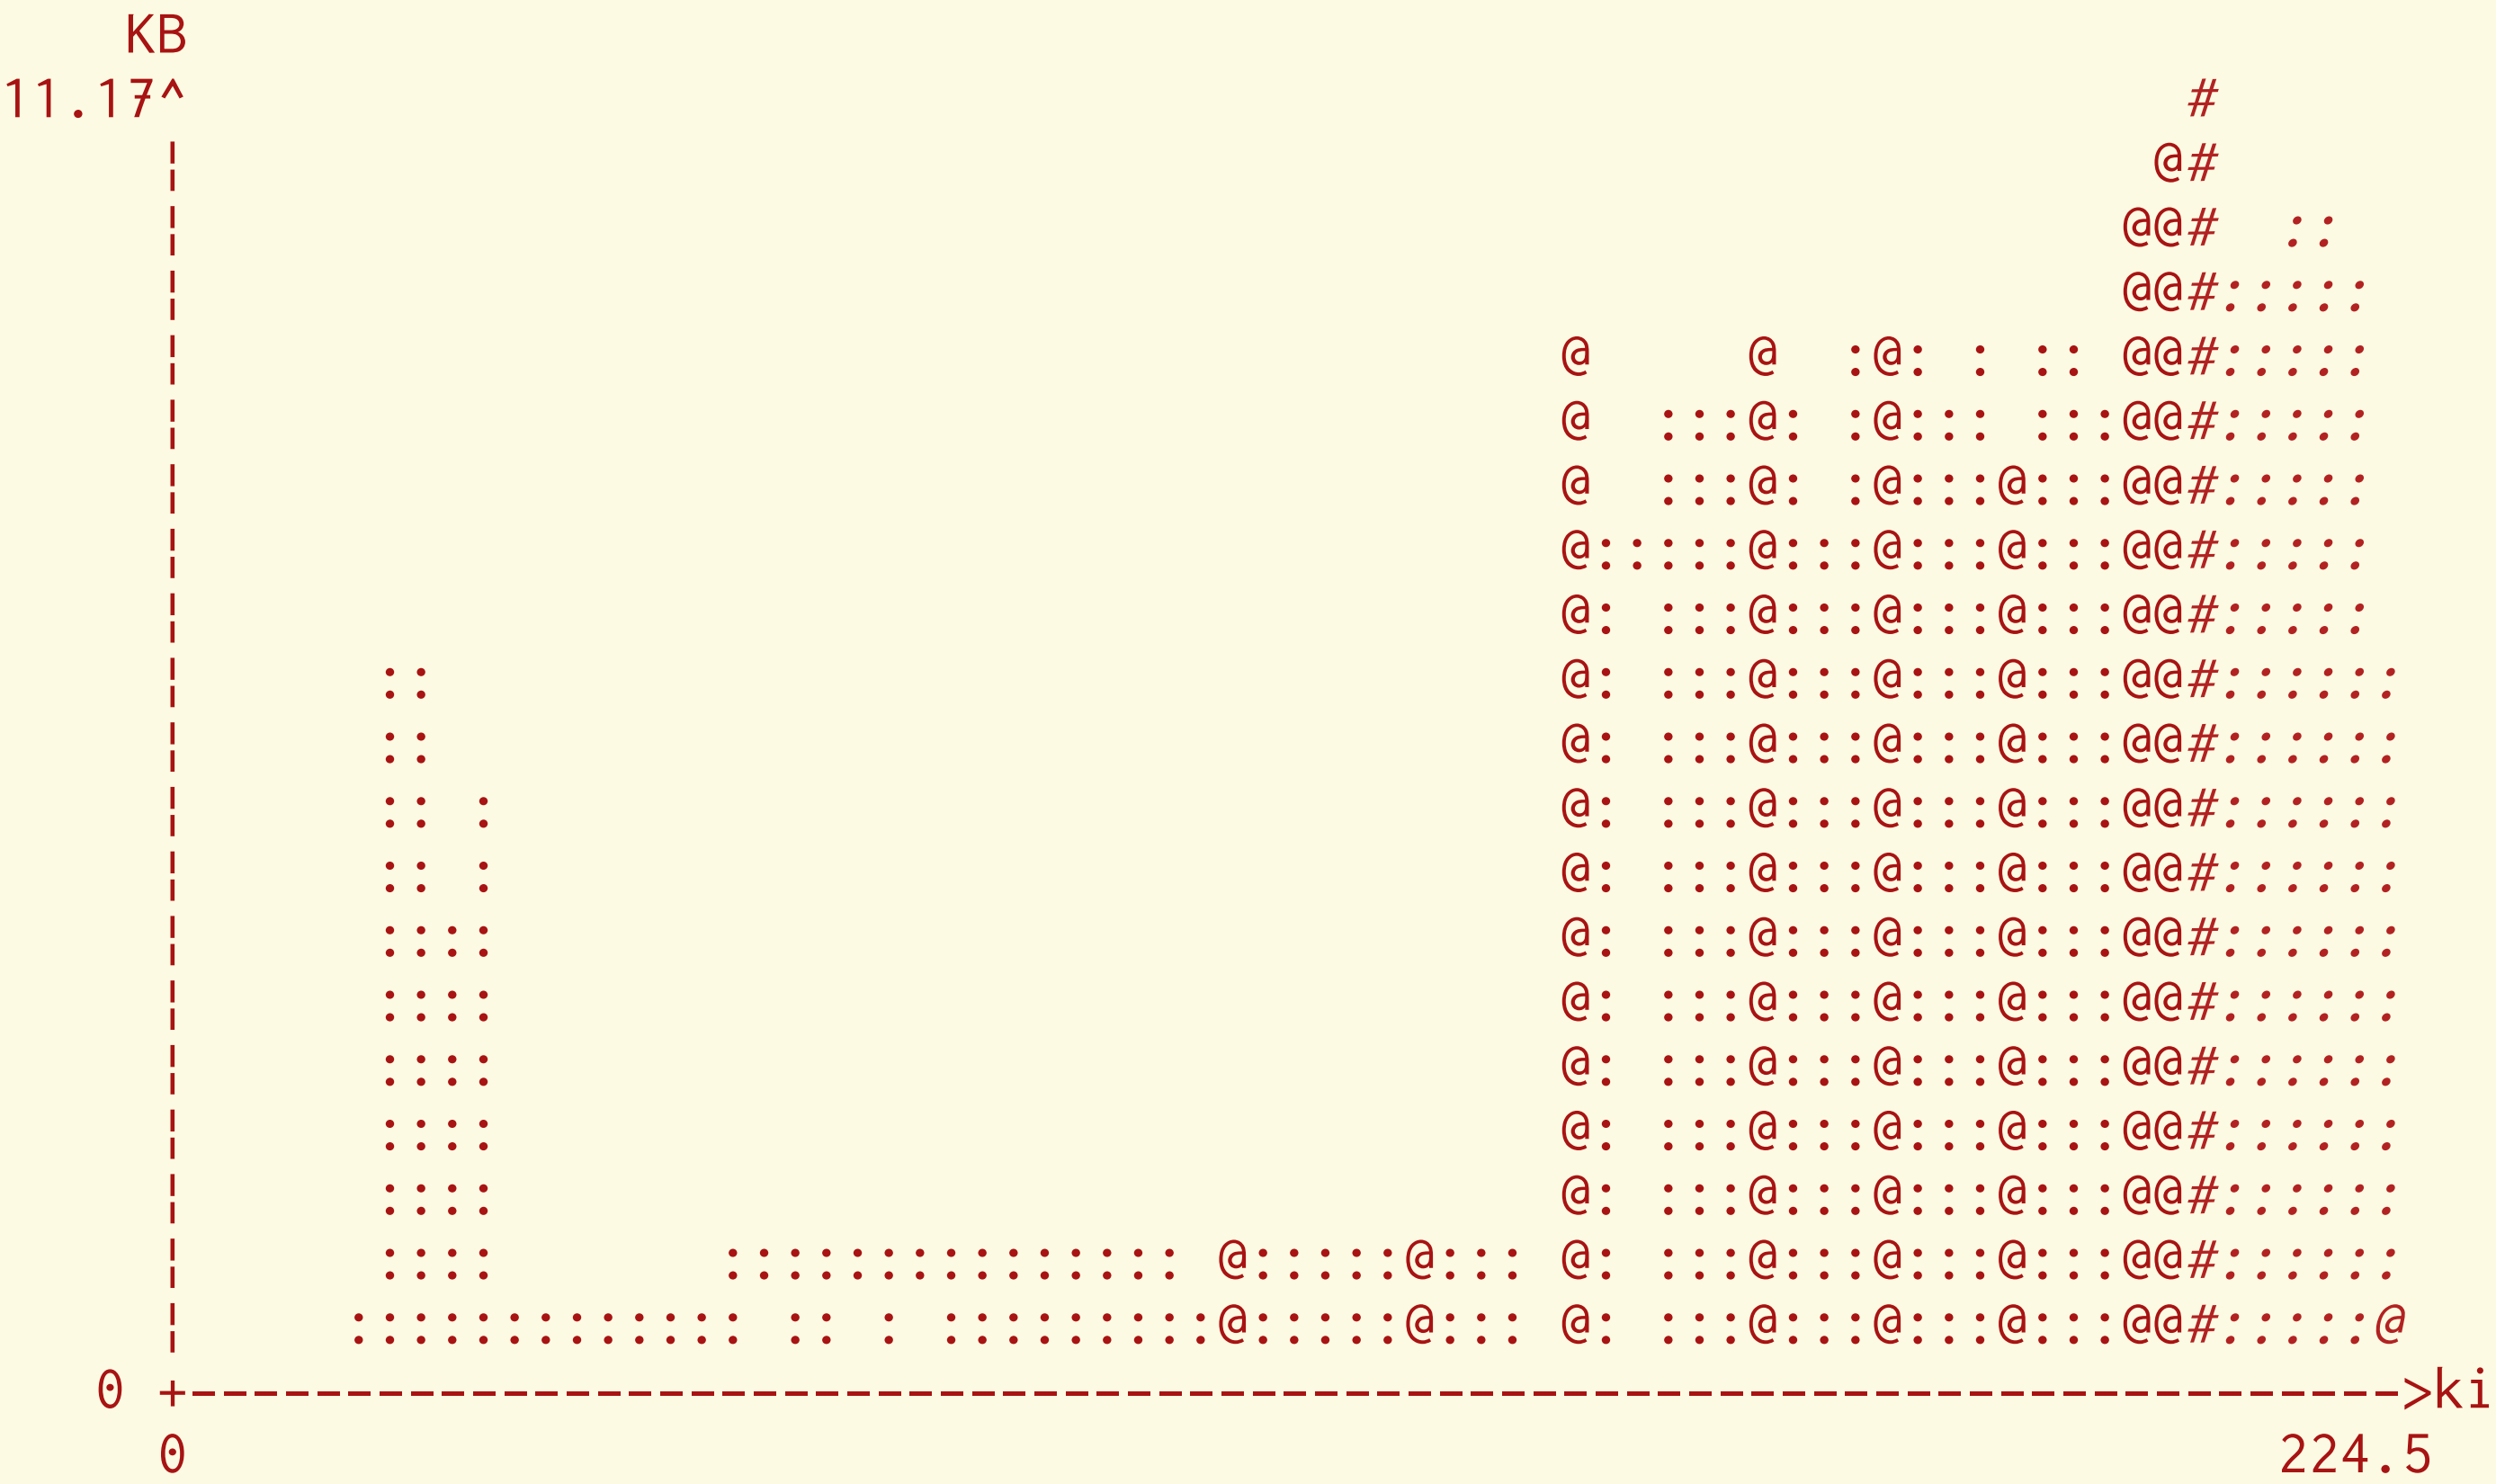
```
Hello World!
```

# RAM consumption: Case 1 / mrbc



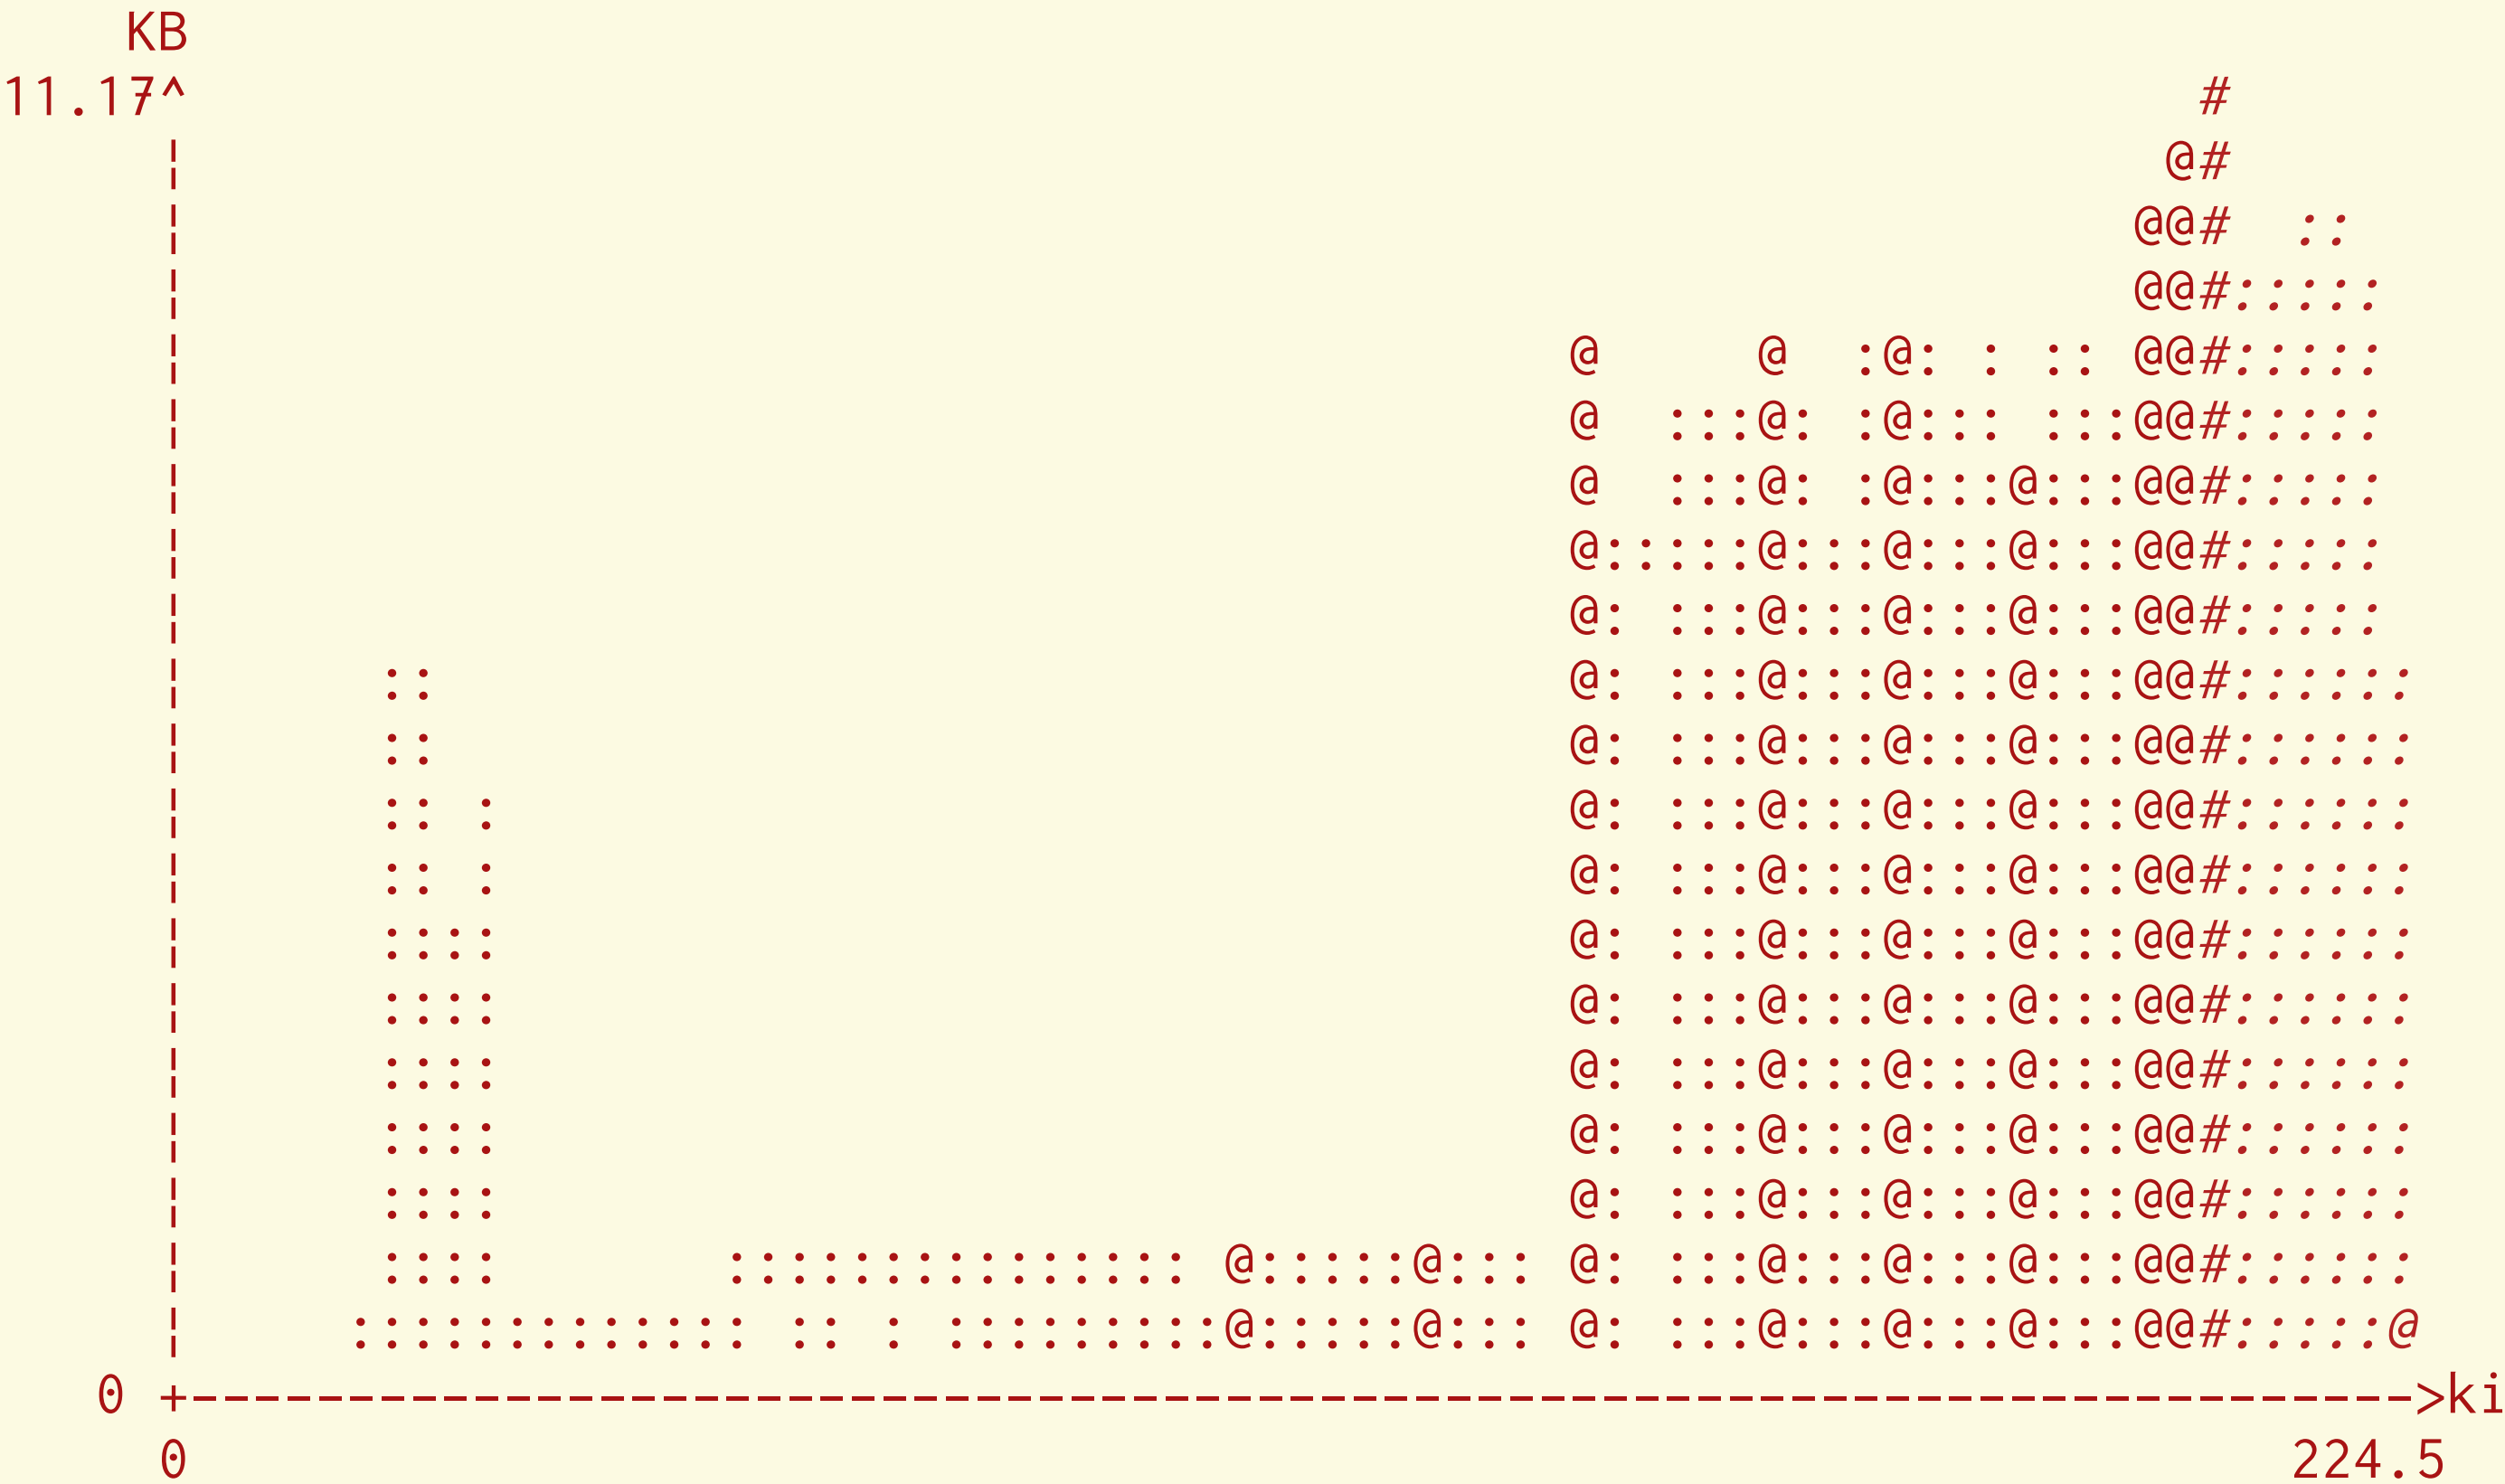
# RAM consumption: Case 1 / mmrbc

```
-----  
Command:      ../mmruby/build/host-production/bin/mmrbc ./hello_world.rb  
Massif arguments:  --stacks=yes  
ms_print arguments: massif.out.1652  
-----
```



# RAM consumption: Case 1 / mmrbc 🙌 8888

```
-----  
Command:          ../mmruby/build/host-production/bin/mmrbc ./hello_world.rb  
Massif arguments:  --stacks=yes  
ms_print arguments: massif.out.1652  
-----
```



# RAM consumption: Case 2

---

## Ruby code:

```
# idiotic_script.rb
ary = Array.new(3)
ary[0] = {a: 123e2}
ary[0][:key] = "string"
ary[1] = %w(abc ABC Hello)
ary[2] = 0x1f
ary[3] = !true
puts "my name is #{self.class}, result is #{ary[2] * ary[0][:a]}"
p ary
```



# RAM consumption: Case 2

---

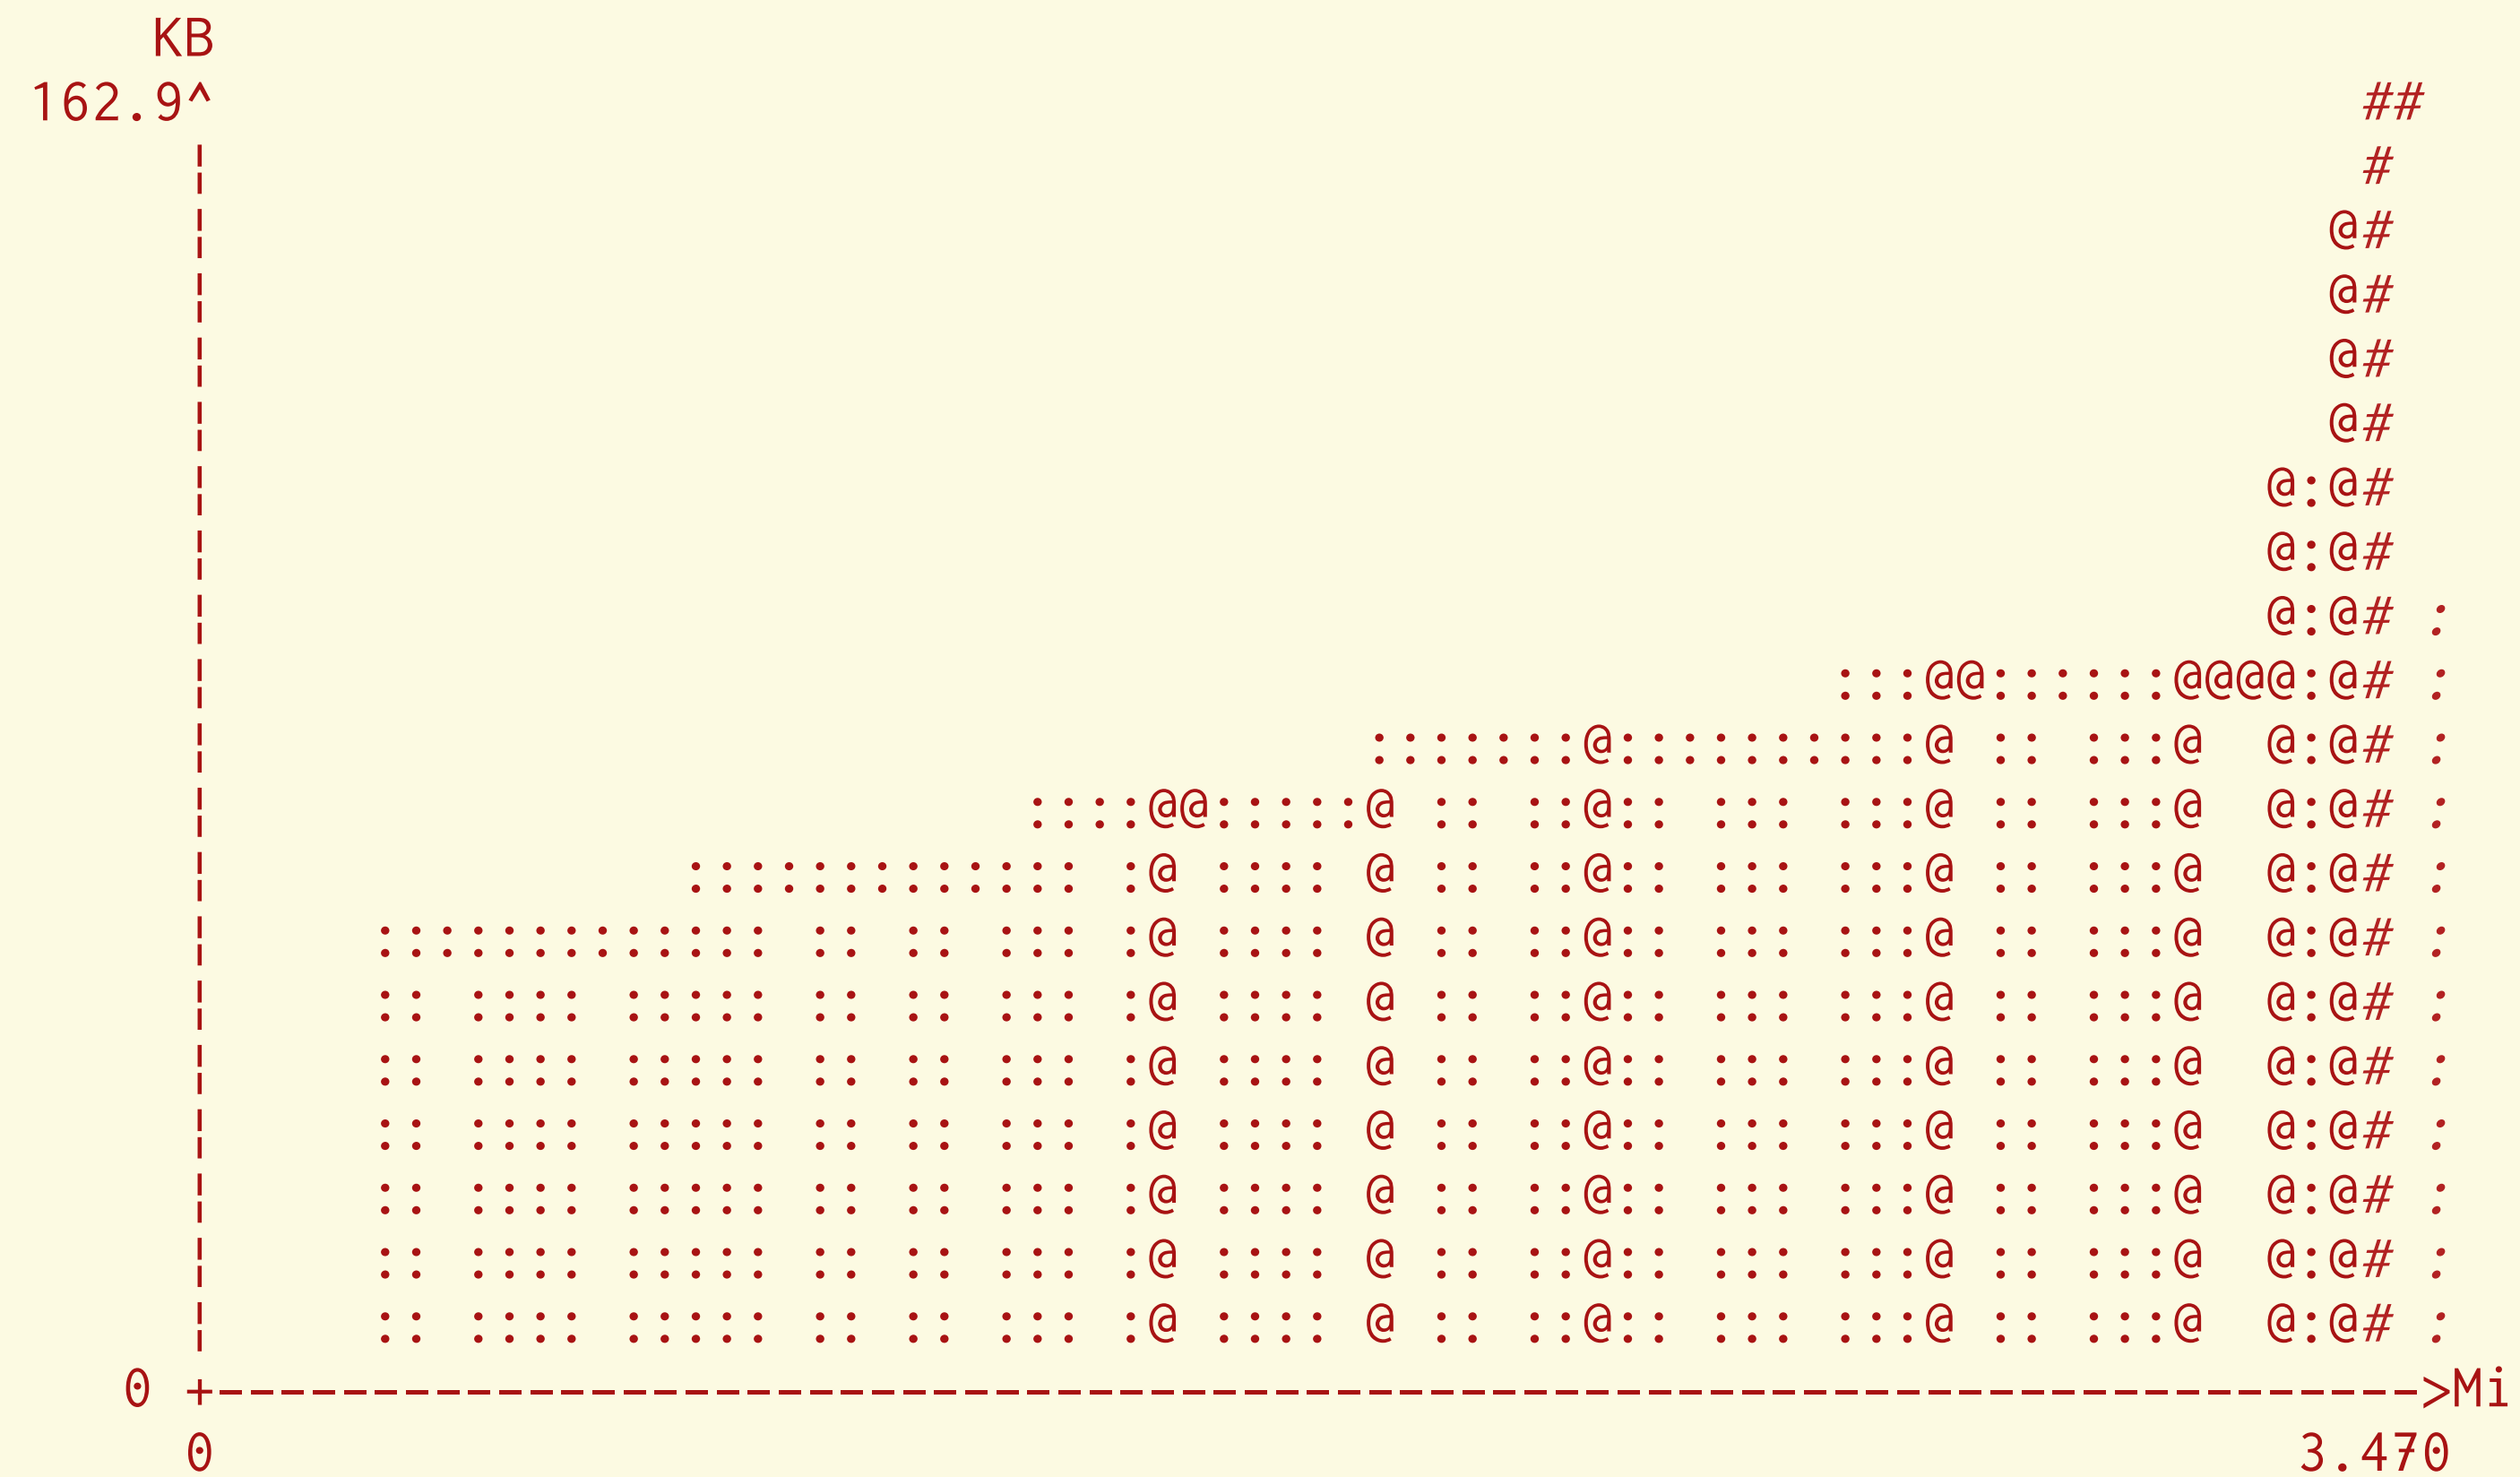
Output:

```
my name is Object, result is 381300
```

```
[{:a=>12300, :key=>"string"}, ["abc", "ABC", "Hello"], 31, false]
```

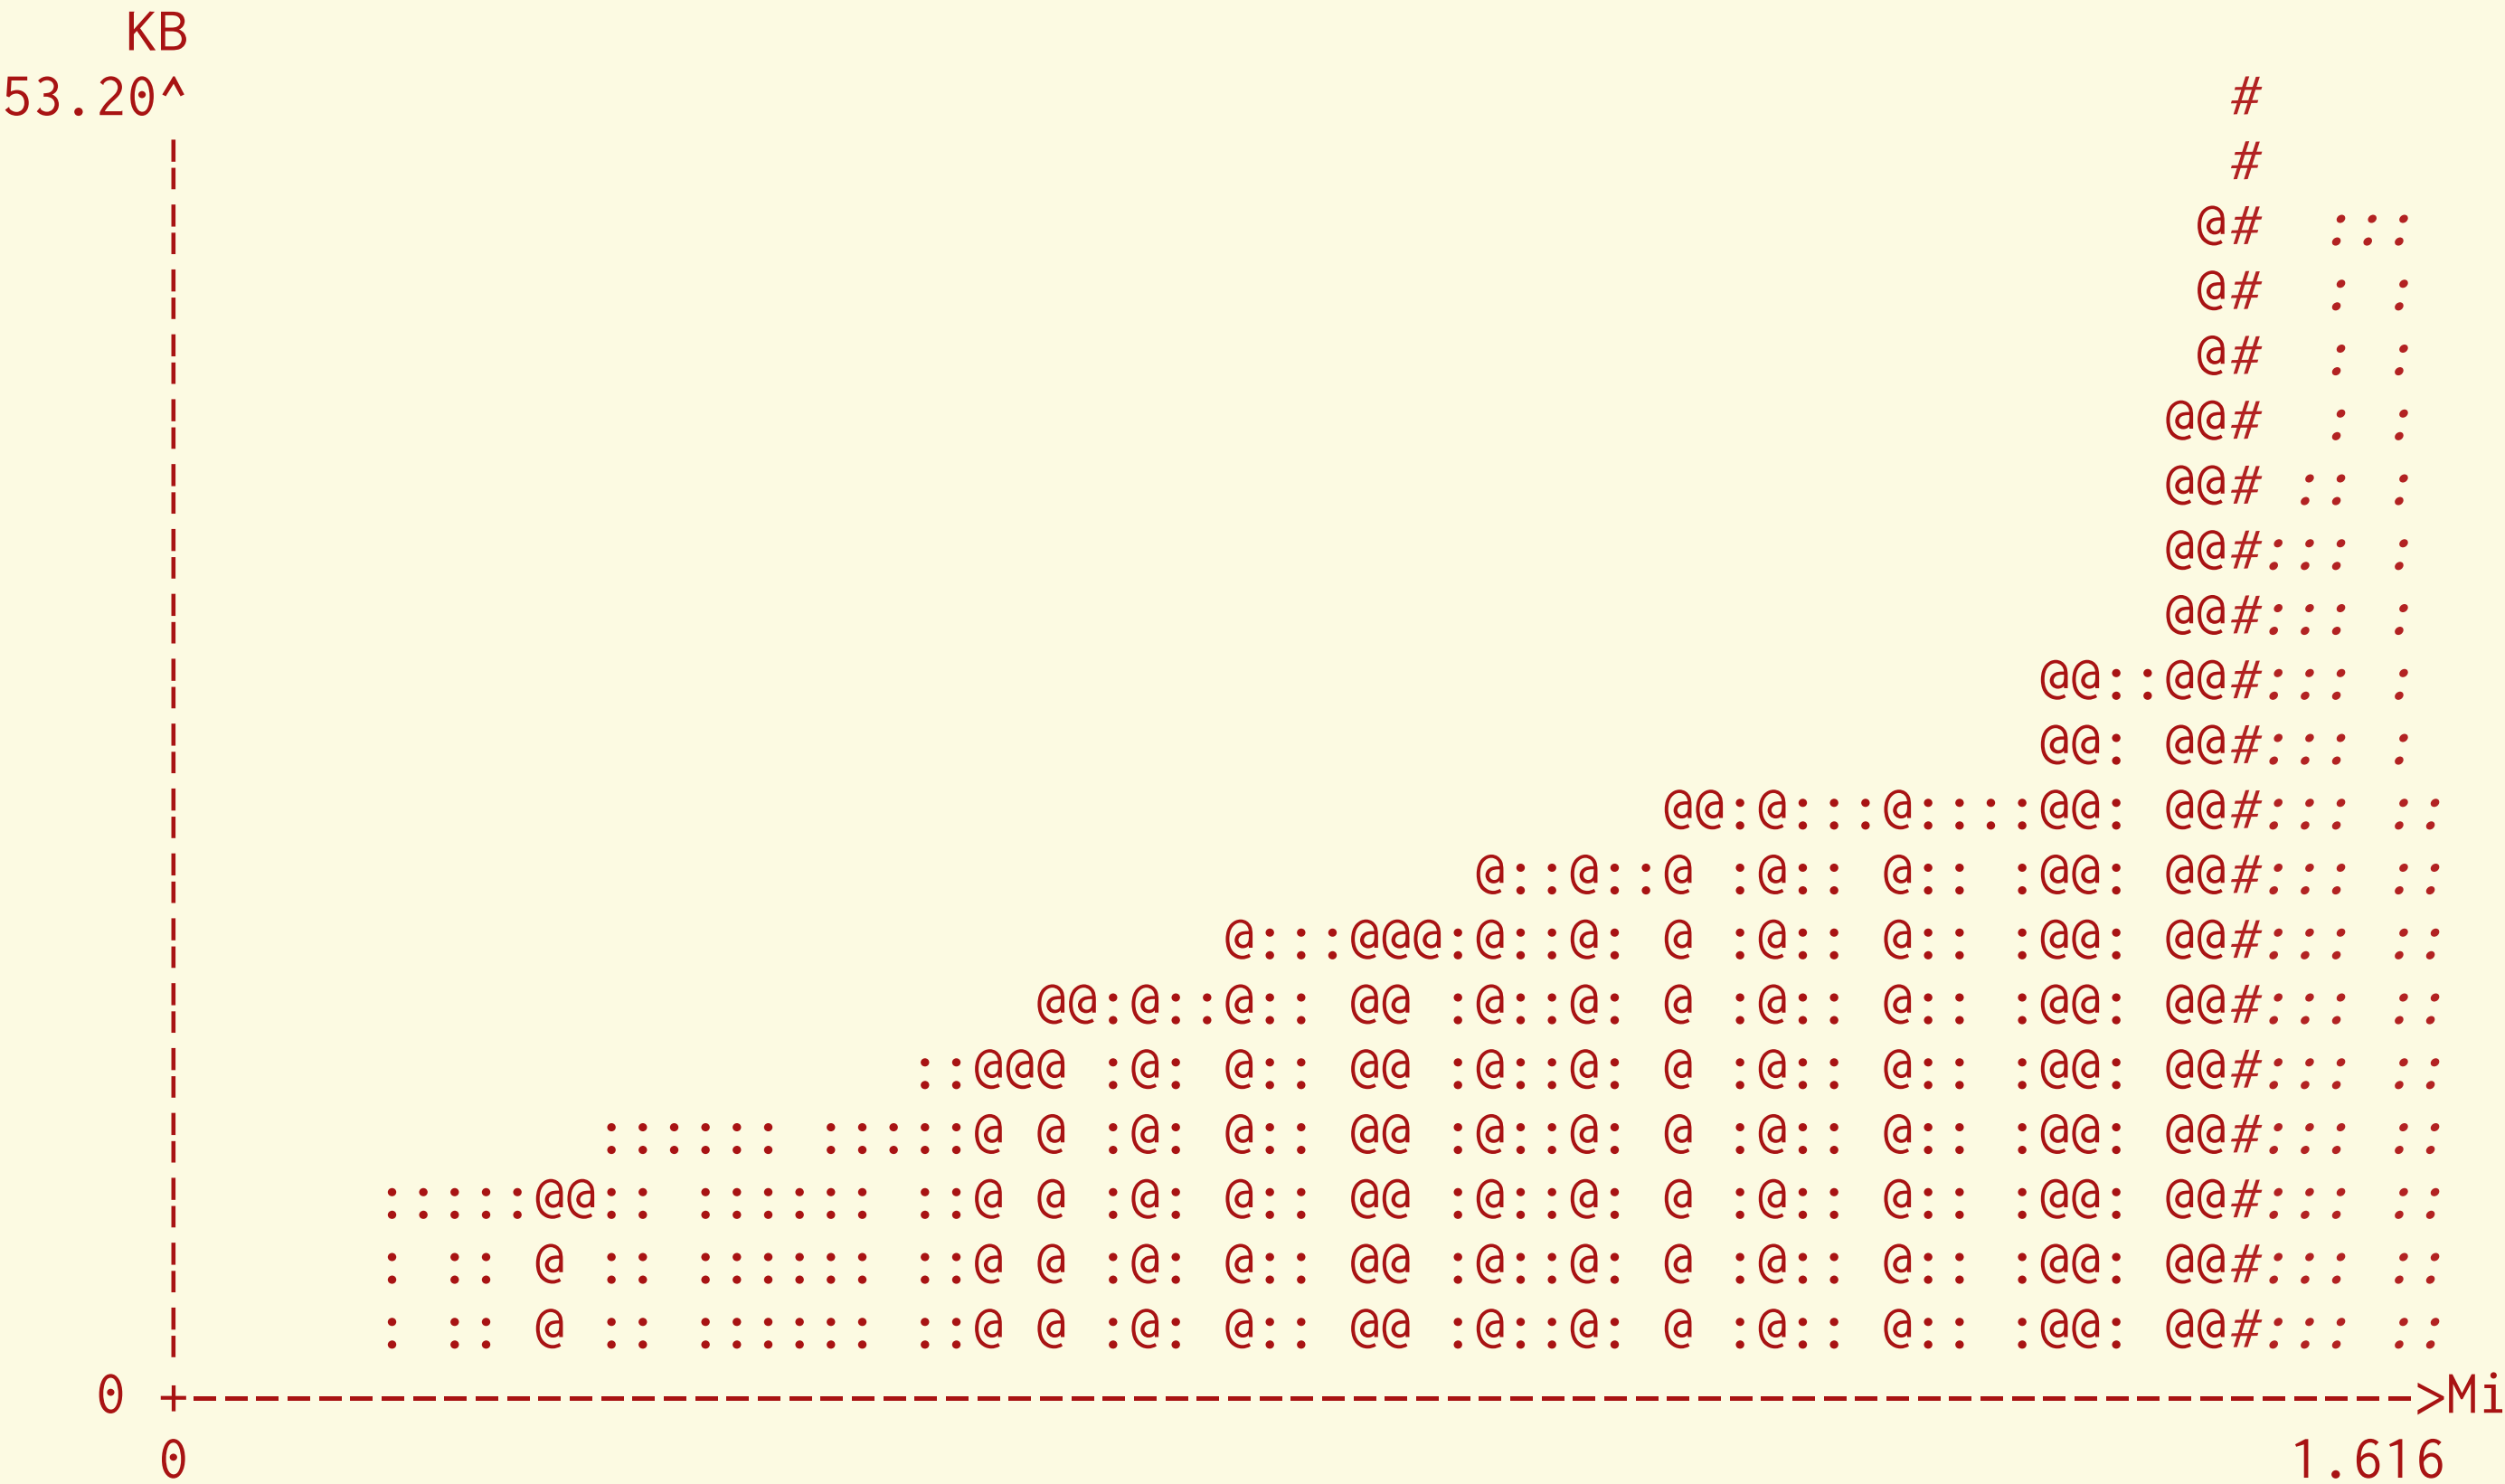
# RAM consumption: Case 2 / mrbc

```
Command:      ../mruby/build/host/bin/mrbc ./idiotic_script.rb
Massif arguments:  --stacks=yes
ms_print arguments:  massif.out.4133
```



# RAM consumption: Case 2 / mmrbc

```
-----  
Command:          ../mmruby/build/host-production/bin/mmrbc idiotic_script.rb  
Massif arguments:  --stacks=yes  
ms_print arguments: massif.out.13414  
-----
```



# RAM consumption: Summary

---

Ruby script	mrbc	mmrbc
hello_world.rb	157.8KB	11.17KB
idiotic_script.rb	162.9KB	53.20KB

Note: These results are larger than possible results on 32bit architecture because of pointer size

# Idiotic script

---

```
# idiotic_script.rb
ary = Array.new(3)
ary[0] = {a: 123e2}
ary[0][:key] = "string"
ary[1] = %w(abc ABC Hello)
ary[2] = 0x1f
ary[3] = !true
puts "my name is #{self.class}, result is #{ary[2] * ary[0][:a]}"
p ary

=>
my name is Object, result is 381300
[{:a=>12300, :key=>"string"}, ["abc", "ABC", "Hello"], 31, false]
```

Syntax that mmrbc can compile



# Syntax implemented

---

- ⑨ lvar, \$gvar, @ivar, CONSTANT
- ⑨ assignment, fcall, method chain
- ⑨ Literals: Array, Hash, String, Fixnum, Float, etc.
- ⑨ Keywords: true, false, nil, return, self, etc.
- ⑨ +, -, \*, /, %, \*\*, +=, -=, \*=, /=, %=, &=, |=, ^=, \*\*=, <<=, >>=
- ⑨ !, ~, +, -, not, &, |, ^, <<, >>
- ⑨ Syntax sugar: %w(a b c)
- ⑨ Interpolation: "The answer is #{get\_result(@attr)}!"

# Syntax under construction

---

- ⑨ if then elif else end, unless
- ⑨ case when else then
- ⑨ while, until, break, next, redo, retry, for in
- ⑨ [0,1].each do |var| end
- ⑨ def method\_name(arg); end
- ⑨ class MyClass < ParentClass; end
- ⑨ rescue, ensure
- ⑨ ..., etc.

# Roadmap

---

<https://github.com/hasumikin/mmruby/issues/6>

# Code size (to estimate ROM consumption)

```
% ls -l
(...)
-rwxrwxrwx 609k hasumi 17 Aug  7:03 mruby_machine_PSoC5LP.a
-rwxrwxrwx 903k hasumi 17 Aug  7:03 mruby_machine_PSoC5LP.elf
-rwxrwxrwx 650k hasumi 17 Aug  7:03 mruby_machine_PSoC5LP.hex
-rwxrwxrwx 425k hasumi 17 Aug  7:03 mruby_machine_PSoC5LP.map

% size mruby_machine_PSoC5LP.elf
   text    data     bss     dec      hex filename
 136080     320    63009   199409   30af1 mruby_machine_PSoC5LP.elf
```

# Code size (to estimate ROM consumption)

```
% ls -l
(...)
-rwxrwxrwx 609k hasumi 17 Aug 7:03 mruby_machine_PSoC5LP.a
-rwxrwxrwx 903k hasumi 17 Aug 7:03 mruby_machine_PSoC5LP.elf
-rwxrwxrwx 650k hasumi 17 Aug 7:03 mruby_machine_PSoC5LP.hex
-rwxrwxrwx 425k hasumi 17 Aug 7:03 mruby_machine_PSoC5LP.map

% size mruby_machine_PSoC5LP.elf
   text    data     bss     dec      hex filename
 136080     320    63009   199409   30af1 mruby_machine_PSoC5LP.elf
```

$$136\_080 + 320 = 136\_400 \doteq \mathbf{133.2KB} \text{ 👍}$$



# Summary

---

- ⑨ mmrbc is a mini mruby compiler
  - ⑨ VM code runs on both mruby VM and mruby/c VM
  - ⑨ Still incomplete regarding syntax and memory efficiency
- ⑨ mruby machine
  - = mmrbc + mruby/c + shell + drivers
  - ⑨ Runs on PSoC5LP { RAM: 64KB, ROM: 256KB }

Thank you!