

## Ruby - Bug #5217

### lineno is broken when source code has about 7000 lines

08/23/2011 11:42 PM - mame (Yusuke Endoh)

<b>Status:</b> Closed	
<b>Priority:</b> Normal	
<b>Assignee:</b> ko1 (Koichi Sasada)	
<b>Target version:</b>	
<b>ruby -v:</b> -	<b>Backport:</b>
<b>Description</b>	
<pre>asakusa.rb (simplecov) soap4r simplecov [BUG] bug  \$ ruby -e 'puts "p\n" * 7000; puts "p([1"; puts ")'" &gt; t.rb  \$ cat t.rb p p p  ... 7000 p ...  p p  \$ ./ruby -rcoverage -e 'Coverage.start; load "t.rb"'  diff --git a/thread.c b/thread.c index 6970d8f..57a6962 100644 --- a/thread.c +++ b/thread.c @@ -4764,7 +4764,7 @@ update_coverage(rb_event_flag_t event, VALUE proc, VALUE self, ID id, VALUE klas long line = rb_sourceline() - 1; long count; if (RARRAY_PTR(coverage)[line] == Qnil) {     •     • } count = FIX2LONG(RARRAY_PTR(coverage)[line]) + 1; if (POSFIXABLE(count)) {  iseq unsigned short 65536 16 65536 7000  rb_vm_get_sourceline coverage</pre>	

```
set_trace_func(proc {|type, file, line,| p line if type == "line" })
```

```
p
p
p
```

```
... 7000   p ...
```

```
p
p
p
```

```
XXXXXXXXXX
```

```
2
3
4
5
...
```

```
6550
6551
6552
6553
7001
7001
7001
7001
7001
7001
...
```

```
XXXXXXXXXXXXXXXXXXXXXXXXXXXX
```

```
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX short XXXXX
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
XXXXX > XXXXX
```

```
diff --git a/iseq.h b/iseq.h
index beeacbb..9c19501 100644
--- a/iseq.h
+++ b/iseq.h
@@ -44,9 +44,9 @@ struct rb_compile_option_struct {
};
```

```
struct iseq_insn_info_entry {
```

- unsigned short position;
  - unsigned short line\_no;
  - unsigned short sp;
  
  - unsigned long position;
  - unsigned long line\_no;
  - unsigned long sp;
- ```
};
```

```
struct iseq_catch_table_entry {
```

```
--
```

```
Yuskue Endoh mame@tsg.ne.jp
```

## Associated revisions

Revision 7049d9c8 - 08/24/2011 06:31 AM - ko1 (Koichi Sasada)

- iseq.h, iseq.c, compile.c: Change the line number data structure to solve an issue reported at [ruby-dev:44413] [Ruby 1.9 - Bug #5217]. Before this fix, each instruction has an information including line number (iseq:iseq\_insn\_info\_table). Instead of this data

structure, recording only line number changing places (iseq::iseq\_line\_info\_table).  
The order of entries in iseq\_line\_info\_table is ascending order of iseq\_line\_info\_table\_entry::position. You can get a line number by an iseq and a program counter with this data structure.  
This fix reduces memory consumption of iseq (bytecode).  
On my measurement, a rails application consumes 21.8MB for iseq with this fix on the 32bit CPU. Without this fix, it consumes 24.7MB for iseq [ruby-dev:44415].

- proc.c: ditto.
- vm\_inshelper.c: ditto.
- vm\_method.c: ditto.
- vm.c (rb\_vm\_get\_sourceline): change to use rb\_iseq\_line\_no().

git-svn-id: svn+ssh://ci.ruby-lang.org/ruby/trunk@33046 b2dd03c8-39d4-4d8f-98ff-823fe69b080e

### History

---

#### #1 - 08/23/2011 11:43 PM - mame (Yusuke Endoh)

- Assignee set to ko1 (Koichi Sasada)

#### #2 - 08/24/2011 07:23 AM - ko1 (Koichi Sasada)

- ruby -v changed from ruby 1.9.4dev (2011-08-23 trunk 33019) [i686-linux] to -

#####

(2011/08/23 7:43), Yusuke Endoh wrote:

```
##### short #####
#####
#### > #####
```

#####

```
##### IRC #####
#####
```

```
##### Rails #####
#### (*1)
##### 24,699,850 B 24.7MB
##### 21,835,244 B 21.8MB
##### 32bit: ruby 1.9.4dev (2011-08-22 trunk 33022)
[i686-linux]
```

\*1: objspace #####

```
require 'objspace'
p ObjectSpace.memsize_of_all(RubyVM::InstructionSequence)
exit!
```

```
##### rack-*/lib/rack/handler/webrick.rb self.run
#####
#####
```

```
##### fork #####
#####
#####
```

#####  
<http://www.atdot.net/sp/view/wpheql/readonly?lang=diff>

```
#####trunk #####
####
```

--  
// SASADA Koichi at atdot dot net

#### #3 - 08/24/2011 10:08 AM - nahi (Hiroshi Nakamura)

Koichi Sasada wrote:

```
##### IRC #####
```

00000000000000000000000000000000

000000000000

000000000000000000000000

<http://www.atdot.net/sp/view/wpheql/readonly?lang=diff>

000000000000000000000000trunk 0000000000000000  
0000

00000000000000000000000000000000

**#4 - 08/24/2011 12:53 PM - mame (Yusuke Endoh)**

000000

2011082410:08 Hiroshi Nakamura [nakahiro@gmail.com](mailto:nakahiro@gmail.com):

Koichi Sasada wrote:

0 000000000000 IRC 000000000000000000000000  
00000000000000000000000000000000

000000000000

000000000000000000000000

0 000000000000000000000000

<http://www.atdot.net/sp/view/wpheql/readonly?lang=diff>

0 000000000000000000000000trunk 0000000000000000  
0000

trunk 0000000000000000  
1.9.3 0000000000000000

00000000000000000000000000000000

000 (BUG) bug) 00000000trunk 000000000000  
000000000000 (r33030) 000000000000000000000000  
00 (BUG) 000000000000000000000000

Yugui 00000000 (r33030) 0 1.9.3 0000000000000000

--  
Yusuke Endoh [mame@tsg.ne.jp](mailto:mame@tsg.ne.jp)

**#5 - 08/24/2011 01:53 PM - ko1 (Koichi Sasada)**

(2011/08/23 20:43), Yusuke ENDOH wrote:

000000000000000000000000

000000000000000000000000

--  
// SASADA Koichi at atdot dot net

**#6 - 08/24/2011 03:31 PM - ko1 (Koichi Sasada)**

- Status changed from Open to Closed

- % Done changed from 0 to 100

This issue was solved with changeset r33046.  
Yusuke, thank you for reporting this issue.  
Your contribution to Ruby is greatly appreciated.

- iseq.h, iseq.c, compile.c: Change the line number data structure to solve an issue reported at [\[ruby-dev:44413\]](#) [Ruby 1.9 - Bug #5217]. Before this fix, each instruction has an information including line number (iseq::iseq\_insn\_info\_table). Instead of this data structure, recording only line number changing places (iseq::iseq\_line\_info\_table). The order of entries in iseq\_line\_info\_table is ascending order of iseq\_line\_info\_table\_entry::position. You can get a line number by an iseq and a program counter with this data structure. This fix reduces memory consumption of iseq (bytecode). On my measurement, a rails application consumes 21.8MB for iseq with this fix on the 32bit CPU. Without this fix, it consumes 24.7MB for iseq [\[ruby-dev:44415\]](#).
- proc.c: ditto.
- vm\_insnhelper.c: ditto.
- vm\_method.c: ditto.
- vm.c (rb\_vm\_get\_sourceline): change to use rb\_iseq\_line\_no().

#7 - 08/24/2011 06:23 PM - kosaki (Motohiro KOSAKI)

```

[BUG] bug trunk
(r33030)
[BUG]

Yugui (r33030) 1.9.3

```

```

backport project 1.9.3

```

#8 - 08/24/2011 07:53 PM - ursm (Keita Urashima)

```

Asakusa.rb

[BUG] r33046


```

```

diff --git a/thread.c b/thread.c
index 57a6962..880e5f8 100644
--- a/thread.c
+++ b/thread.c
@@ -4764,7 +4764,7 @@ update_coverage(rb_event_flag_t event, VALUE
proc, VALUE self, ID id, VALUE klas
long line = rb_sourceline() - 1;
long count;
if (RARRAY_PTR(coverage)[line] == Qnil) {

```

- [REDACTED]
  - [REDACTED]
- ```

}
count = FIX2LONG(RARRAY_PTR(coverage)[line]) + 1;
if (POSFIXABLE(count)) {

```

#9 - 08/24/2011 09:53 PM - mame (Yusuke Endoh)

20110824 18:09 KOSAKI Motohiro [kosaki.motohiro@gmail.com](mailto:kosaki.motohiro@gmail.com):

```

[BUG] bug trunk
(r33030)
[BUG]

Yugui (r33030) 1.9.3

```

backport project 1.9.3  
Asakusa.rb

--  
Yusuke Endoh [mame@tsg.ne.jp](mailto:mame@tsg.ne.jp)

#10 - 08/24/2011 09:59 PM - mame (Yusuke Endoh)

2011 08 24 19:30 Keita Urashima [urasm@urasm.jp](mailto:urasm@urasm.jp):

Asakusa.rb

[BUG] r33046

coverage  
rb\_bug return

rb\_bug

--  
Yusuke Endoh [mame@tsg.ne.jp](mailto:mame@tsg.ne.jp)

#11 - 08/24/2011 11:36 PM - mame (Yusuke Endoh)

2011 08 24 10:08 Hiroshi Nakamura [nakahiro@gmail.com](mailto:nakahiro@gmail.com):

Koichi Sasada wrote:

IRC

<http://www.atdot.net/sp/view/wpheql/readonly?lang=diff>

trunk

trunk 1.9.3

(BUG) bug trunk r33030

Yugui (r33030) 1.9.3

--

Yusuke Endoh [mame@tsg.ne.jp](mailto:mame@tsg.ne.jp)

**#12 - 08/24/2011 11:53 PM - yugui (Yuki Sonoda)**

2011/8/24 KOSAKI Motohiro [kosaki.motohiro@gmail.com](mailto:kosaki.motohiro@gmail.com):

backport project 1.9.3

done

--

Yuki Sonoda (Yugui)

[yugui@yugui.jp](mailto:yugui@yugui.jp)

<http://yugui.jp>