

Ruby trunk - Feature #9633

[PATCH]Symbol GC

03/13/2014 07:59 AM - authorNari (Narihiro Nakamura)

Status:	Closed
Priority:	Normal
Assignee:	matz (Yukihiro Matsumoto)
Target version:	2.2.0

Description

I've written a patch to collect most symbols.

PATCH: https://github.com/authorNari/ruby/compare/13834fb3c...symbol_gc.patch

Summary

- Most symbols in Ruby level are GC-able generated by #to_sym, #intern, etc..
- Exclude a symbol which is translated ID in C-level from GC-able symbols
- Keep Ruby's C extension compatibility
- Pass make test-all

Benchmark

A benchmark program is here.

```
obj = Object.new
100_000.times do |i|
  obj.respond_to?("sym#{i}".to_sym)
end
GC.start
puts "symbol : #{Symbol.all_symbols.size}"
```

```
% time RBENV_VERSION=ruby-r45059 ruby -v /tmp/a.rb
ruby 2.2.0dev (2014-02-20 trunk 45059) [x86_64-linux]
symbol : 102416
0.24s user 0.01s system 91% cpu 0.272 total
```

```
% time RBENV_VERSION=symgc ruby -v /tmp/a.rb
ruby 2.2.0dev (2014-02-20 trunk 45059) [x86_64-linux]
symbol : 2833
0.21s user 0.01s system 90% cpu 0.247 total
```

The total number of symbols is declined.
The total time of symgc version is improved because Full GC pressure has been reduced.

The result of make benchmark.

<https://gist.github.com/authorNari/9359704>

There is no significant slowdown.

(I would welcome to try an additional benchmark and report)

Implementation Detail

I classify Dynamic symbol and Static symbol.

- Static symbol

