



Nokogiri のお母さん

@flavorjones

ZOMG!!!!!!

**HAPPY
THURSDAY!!**

**WELCOME TO
RubyConf X!**

X-TREME RUBYCONF!!!!!!

Aaron Patterson



at&t



Interactive

AT&T, AT&T logo and all AT&T related marks are trademarks of AT&T Intellectual Property and/or AT&T affiliated companies.

@tenderlove

aaron.patterson@gmail.com



ruby committer

rails committer


Committer HOWTO

Ruby






Rails





RubyConf







WWFMD?

RubyConf 5k

RUBYCONF
FIVE THOUSAND




**ZOMG WHY IS THIS
CODE SO SLOW?**

Performance

Code Analysis

Story Form




IRL

IRL

IRL

Tools

Theory





ARel

What is it?

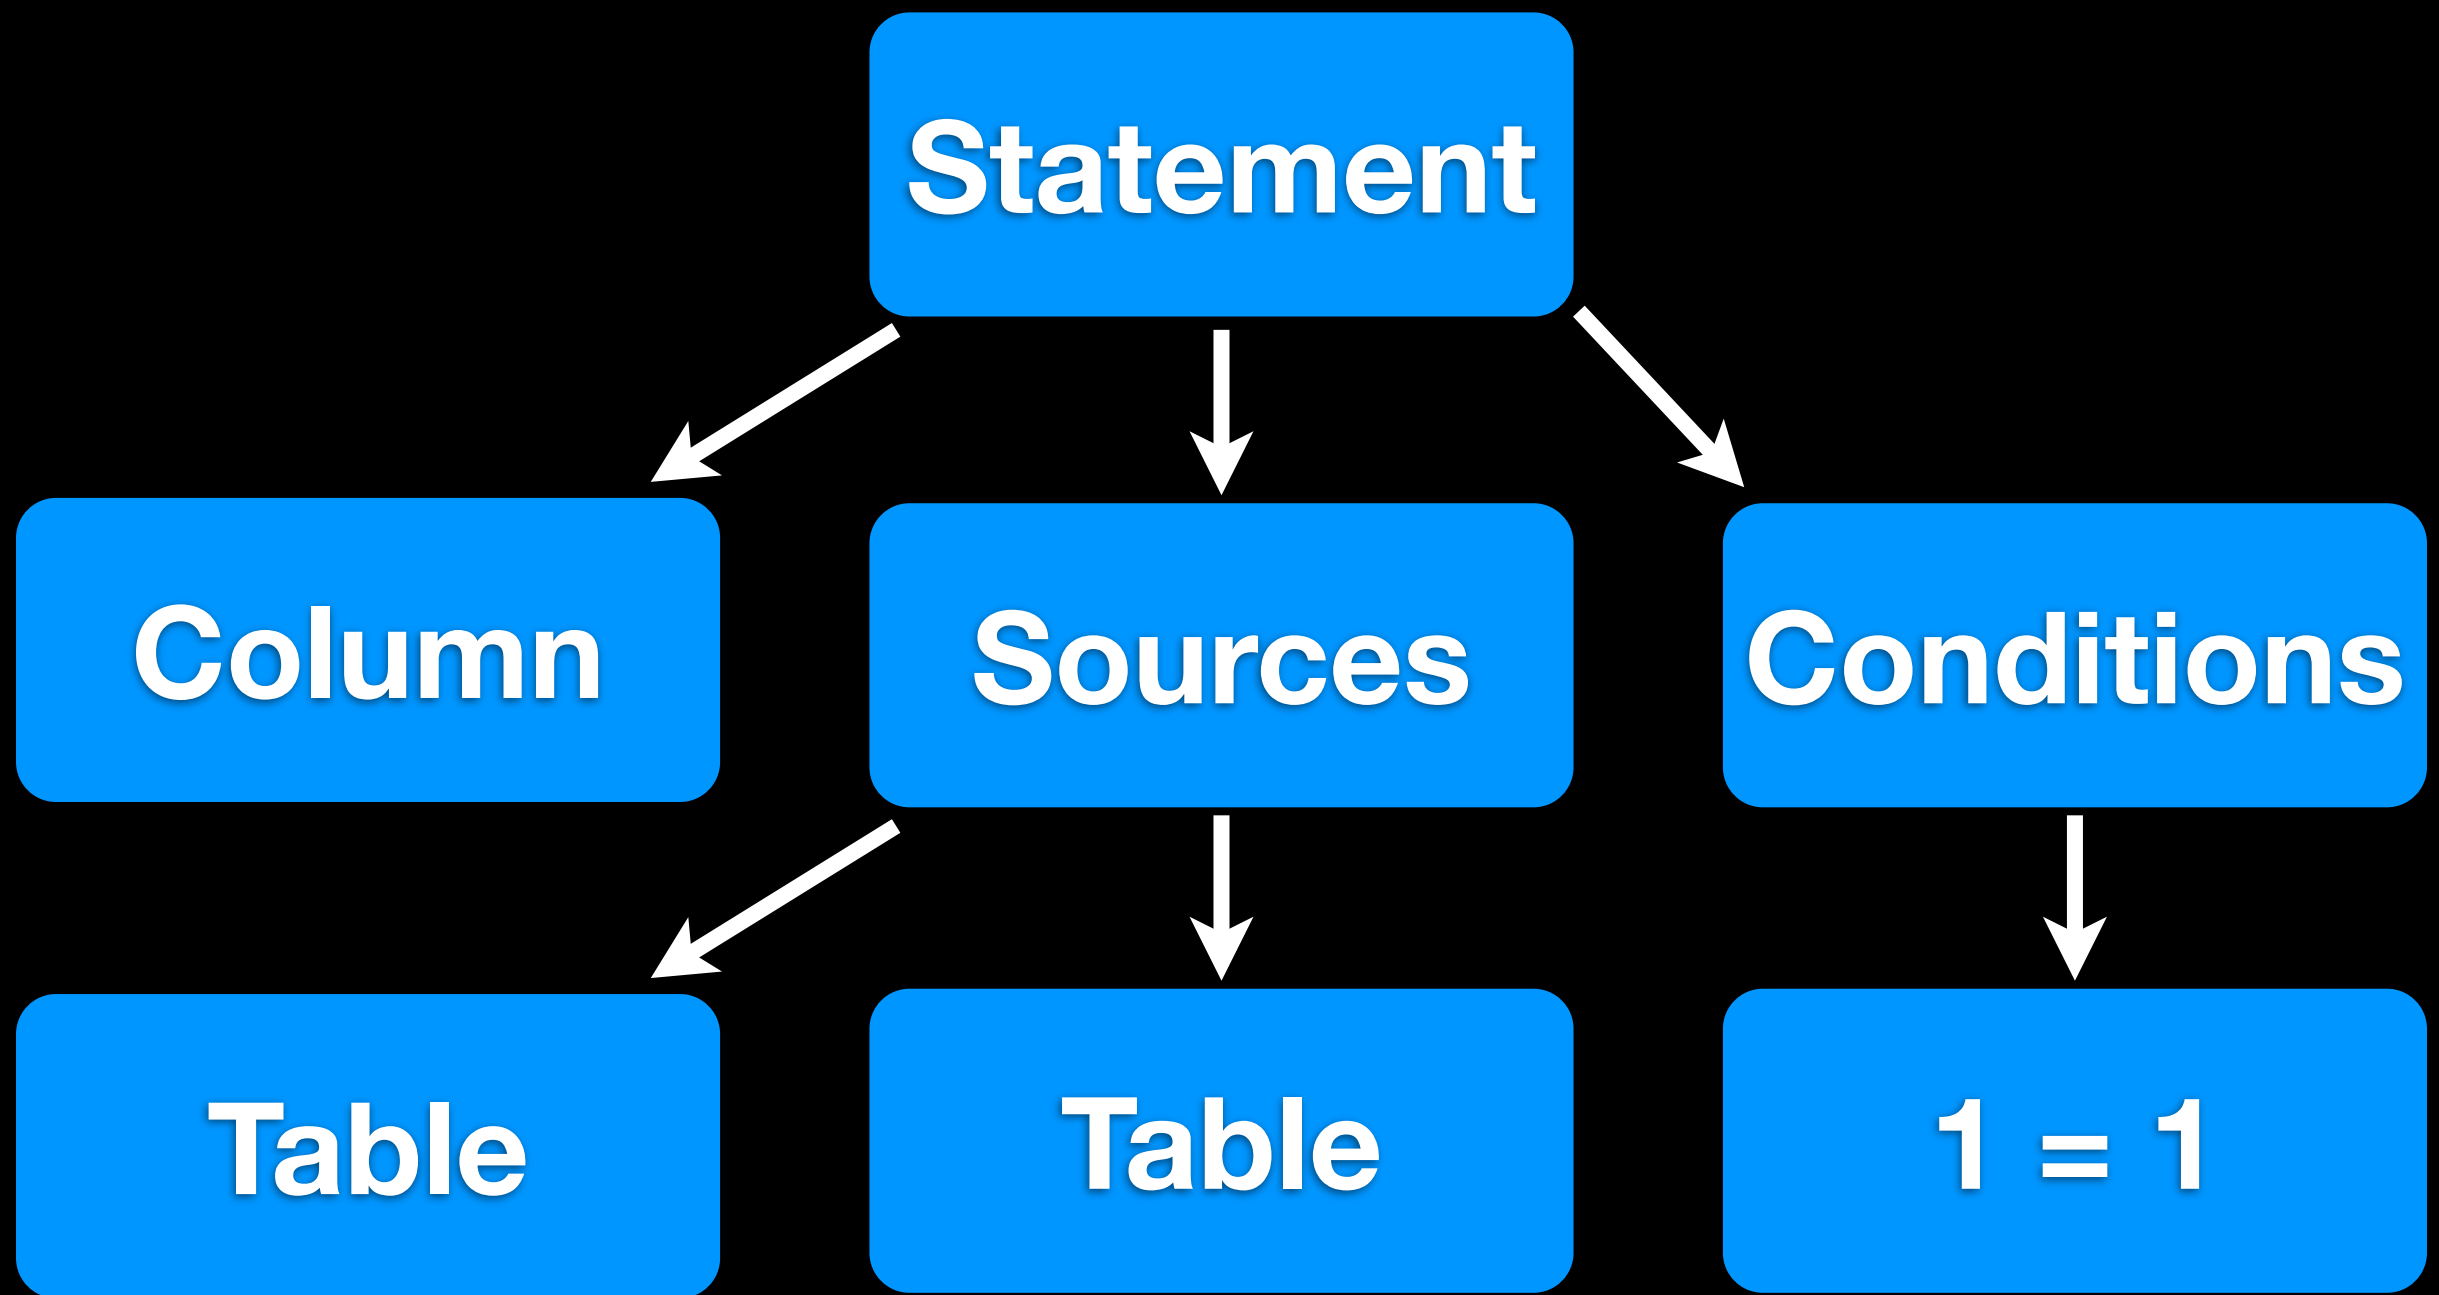
"Relational Algebra"

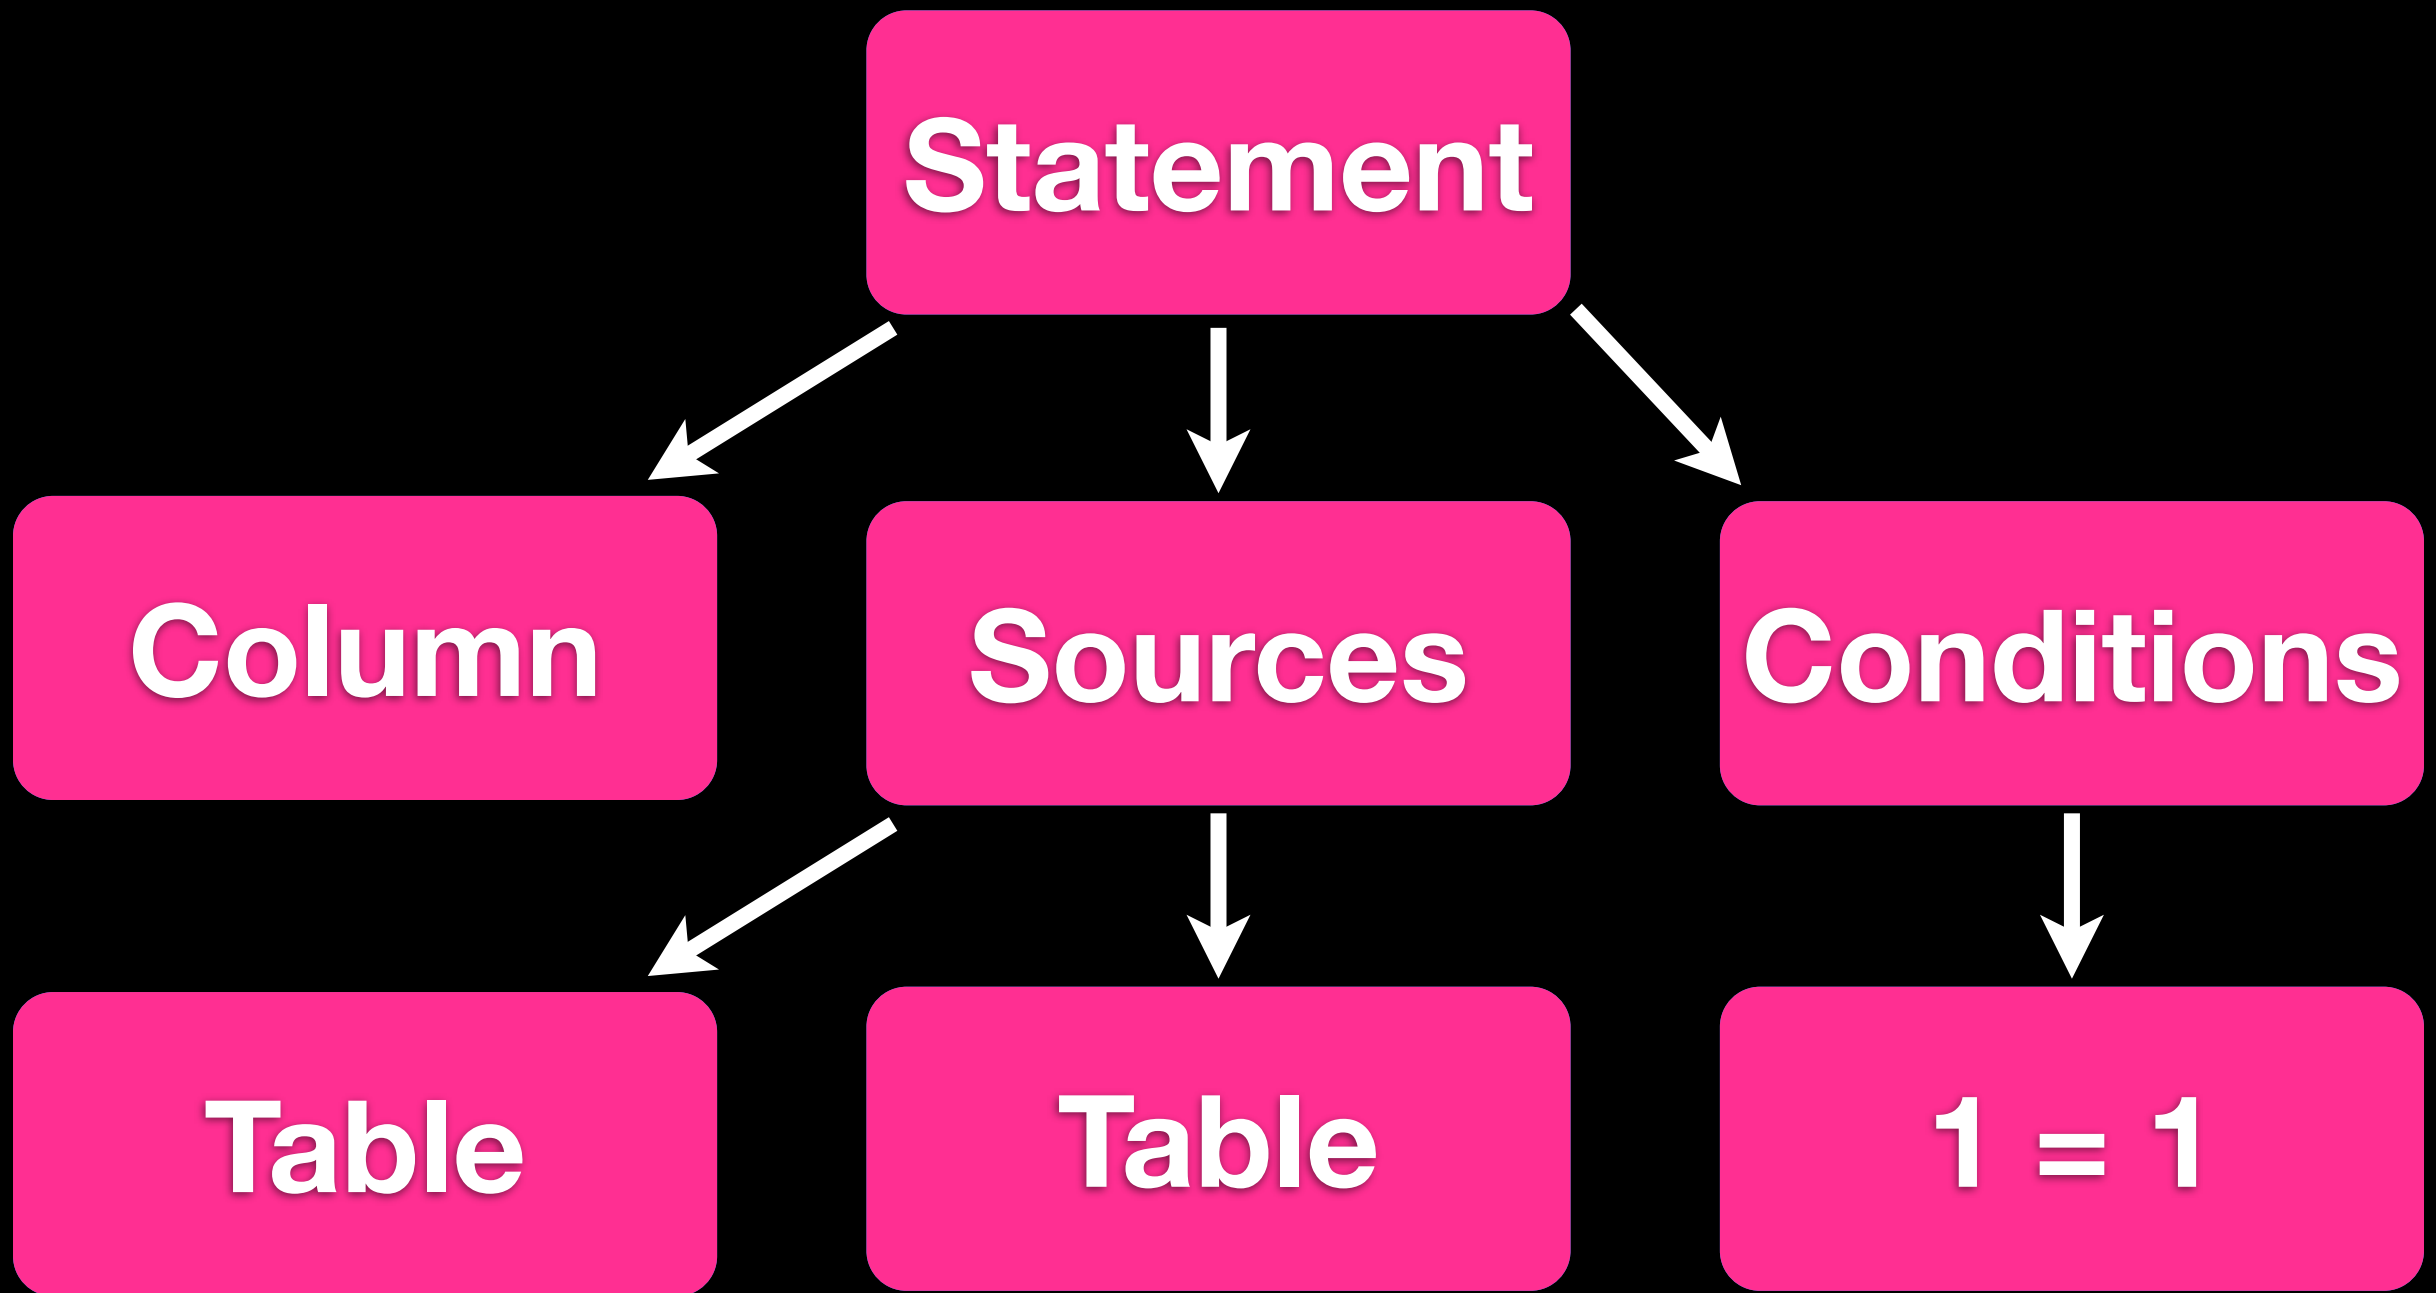
AST Manipulation





AST Translation





SELECT COLUMN
FROM TABLE, TABLE
WHERE 1 = 1


Relationship with Rails

User (YOU!)

ActiveRecord


ARel

User (YOU!)



Records PLZ

ActiveRecord



SQL STMT

ARel

The More You Know™

Getting Started



at&t



Interactive

AT&T, AT&T logo and all AT&T related marks are trademarks of AT&T Intellectual Property and/or AT&T affiliated companies.

Rails

Prepared Statement Caching

DEEPER
UNDERSTANDING
REQUIRED

**ActiveRecord 5x
slower than Rails
2.3.5**

<http://bit.ly/omgsLow>

5x Slower?!?

WTF?



Yes, 5x Slower

**What could possibly
go wrong?**

Motivation

**Why do you care
about speed?**

Scaling Ruby


```
    arel.join(join)
  end

  arel.joins(arel)
end

def build_arel
  arel = table

  arel = build_joins(arel, @joins_values) unless @joins_values.empty?

  (@where_values - []).uniq.each do |where|
    where = Arel.sql(where) if String === where
    arel = arel.where(Arel::Nodes::Grouping.new(where))
  end

  arel = arel.having(*@having_values.uniq.reject{|h| h.blank?}) unless @having_values.empty?

  arel = arel.take(@limit_value) if @limit_value
  arel = arel.skip(@offset_value) if @offset_value

  arel = arel.group(*@group_values.uniq.reject{|g| g.blank?}) unless @group_values.empty?


  arel = arel.order(*@order_values.uniq.reject{|o| o.blank?}) unless @order_values.empty?

  arel = build_select(arel, @select_values.uniq)

  arel = arel.from(@from_value) if @from_value
  arel = arel.lock(@lock_value) if @lock_value

  arel
end

private
```



**When should I make
my code faster?**

**When it isn't fast
enough.**

**What is
"fast enough"?**

Do people notice it?

In comparison to?

**Finishes in a
reasonable amount
of time.**

**What code should I
improve?**

***Only* the code that
matters.**

Don't believe me.

Think Critically

Discovery

What to measure?

Breakdown

Post.find(1)

ARel....

find_by_sql()

execute()

log()

find_by_sql()

execute()

log()

Work per Time

Performance Degraded

Benchmarking

Our Enemies

Time

Space

For Performance:


Things To Reduce

- **Method calls**
- **Branching and looping**
- **Objects (memory consumption)**

For Clean Code:

Things To Reduce

- **Method calls**
- **Branching and looping**
- **Objects (memory consumption)**



(therefore)

Clean Code

==

Performant code

**Measurement is
Paramount**

require 'benchmark'


```
require 'benchmark'

def fib n
  a = 0
  b = 1
  n.times { a, b = b, a + b }
  a
end
```

```
Benchmark.bm(7) do |x|
  x.report("fib") do
    3000.times do |i|
      fib(1000)
    end
  end
end
```


	user	system	total	real
fib	1.570000	0.000000	1.570000	(1.570726)

user	1.570000
system	0.000000
total	1.570000
real	1.570726

```
Benchmark.bm(10) do |x|  
  [1, 10, 100, 1000, 10000].each do |n|  
    x.report("fib #{n}") do  
      n.times { fib(1000) }  
    end  
  end  
end  
end
```

	user	system	total	real
fib 1	0.000000	0.000000	0.000000 (0.000671)
fib 10	0.010000	0.000000	0.010000 (0.008352)
fib 100	0.070000	0.000000	0.070000 (0.074577)
fib 1000	0.740000	0.000000	0.740000 (0.734922)
fib 10000	7.330000	0.000000	7.330000 (7.370046)

○ fib (n, l - 10000)



minitest/benchmark

```
require 'rubygems'
require 'minitest/autorun'
require 'minitest/benchmark'

class BenchFib < MiniTest::Unit::TestCase
  def fib n
    a = 0
    b = 1
    n.times { a, b = b, a + b }
    a
  end

  def bench_fib
    assert_performance_linear 0.99 do |n|
      n.times { fib(1000) }
    end
  end
end
```

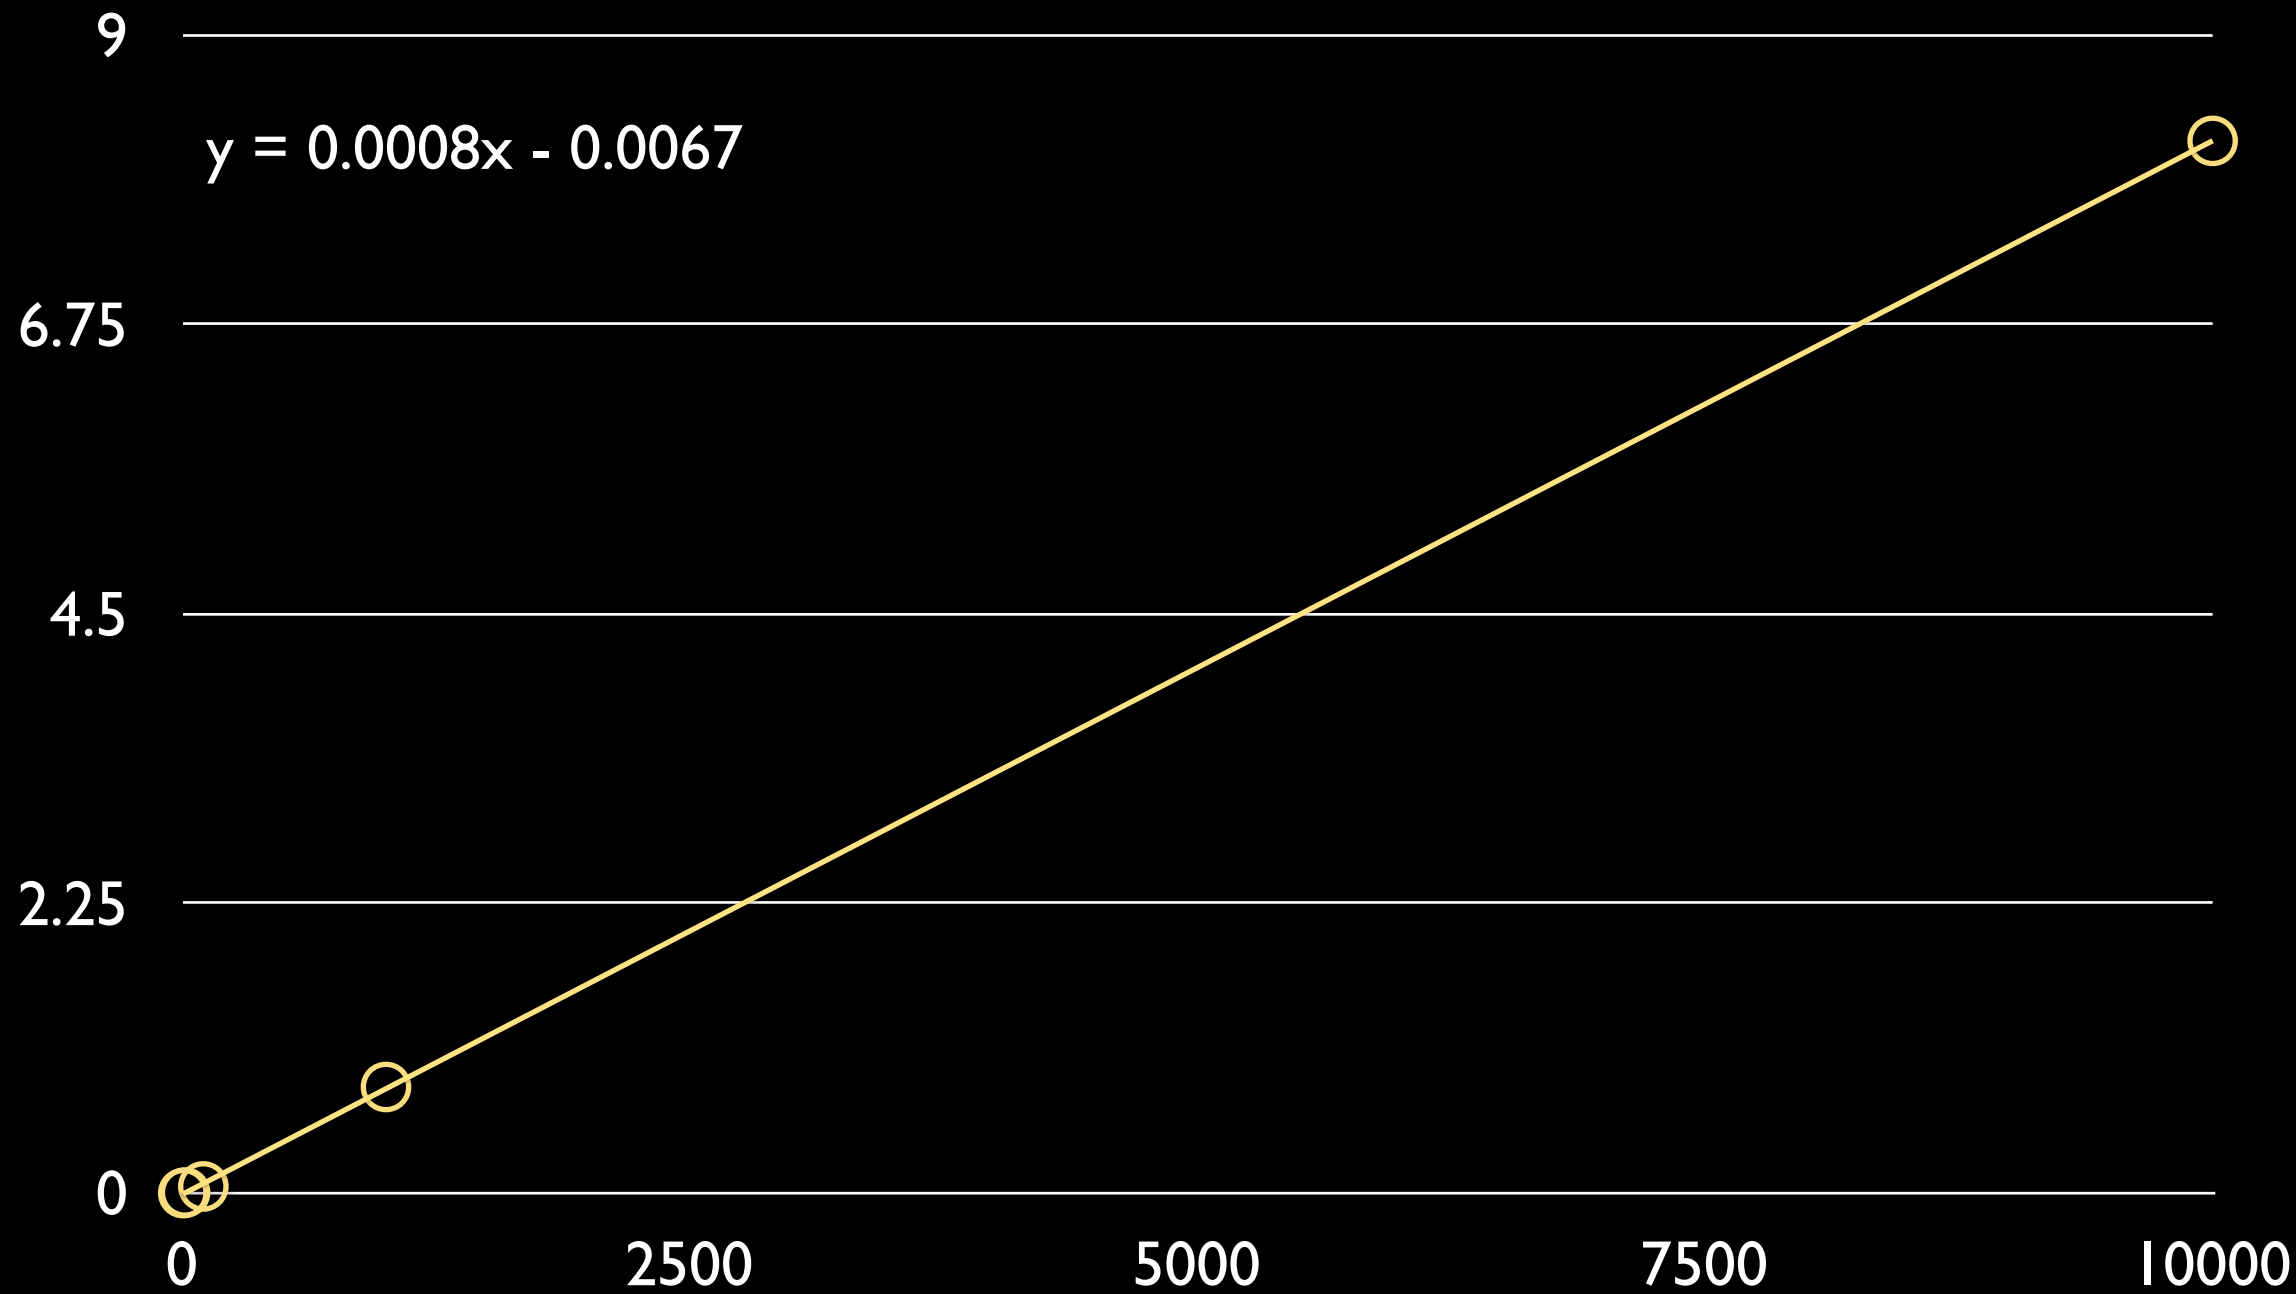


```
Benchmark.bm(10) do |x|  
  [1, 10, 100, 1000, 10000].each do |n|  
    x.report("fib #{n}") do  
      n.times { fib(1000) }  
    end  
  end  
end  
end
```

```
assert_performance_linear 0.99 do |n|  
  n.times { fib(1000) }  
end
```

BenchFib	1	10	100	1000	10000		
bench_fib	0.000571	0.005318	0.052582	0.825676	8.180719		

○ BenchFib



find_by_sql

```
def bench_find_by_sql
  assert_performance_linear 0.999 do |n|
    n.times do
      Post.find_by_sql(
        'SELECT * FROM posts WHERE id = 1')
    end
  end
end
```

execute()

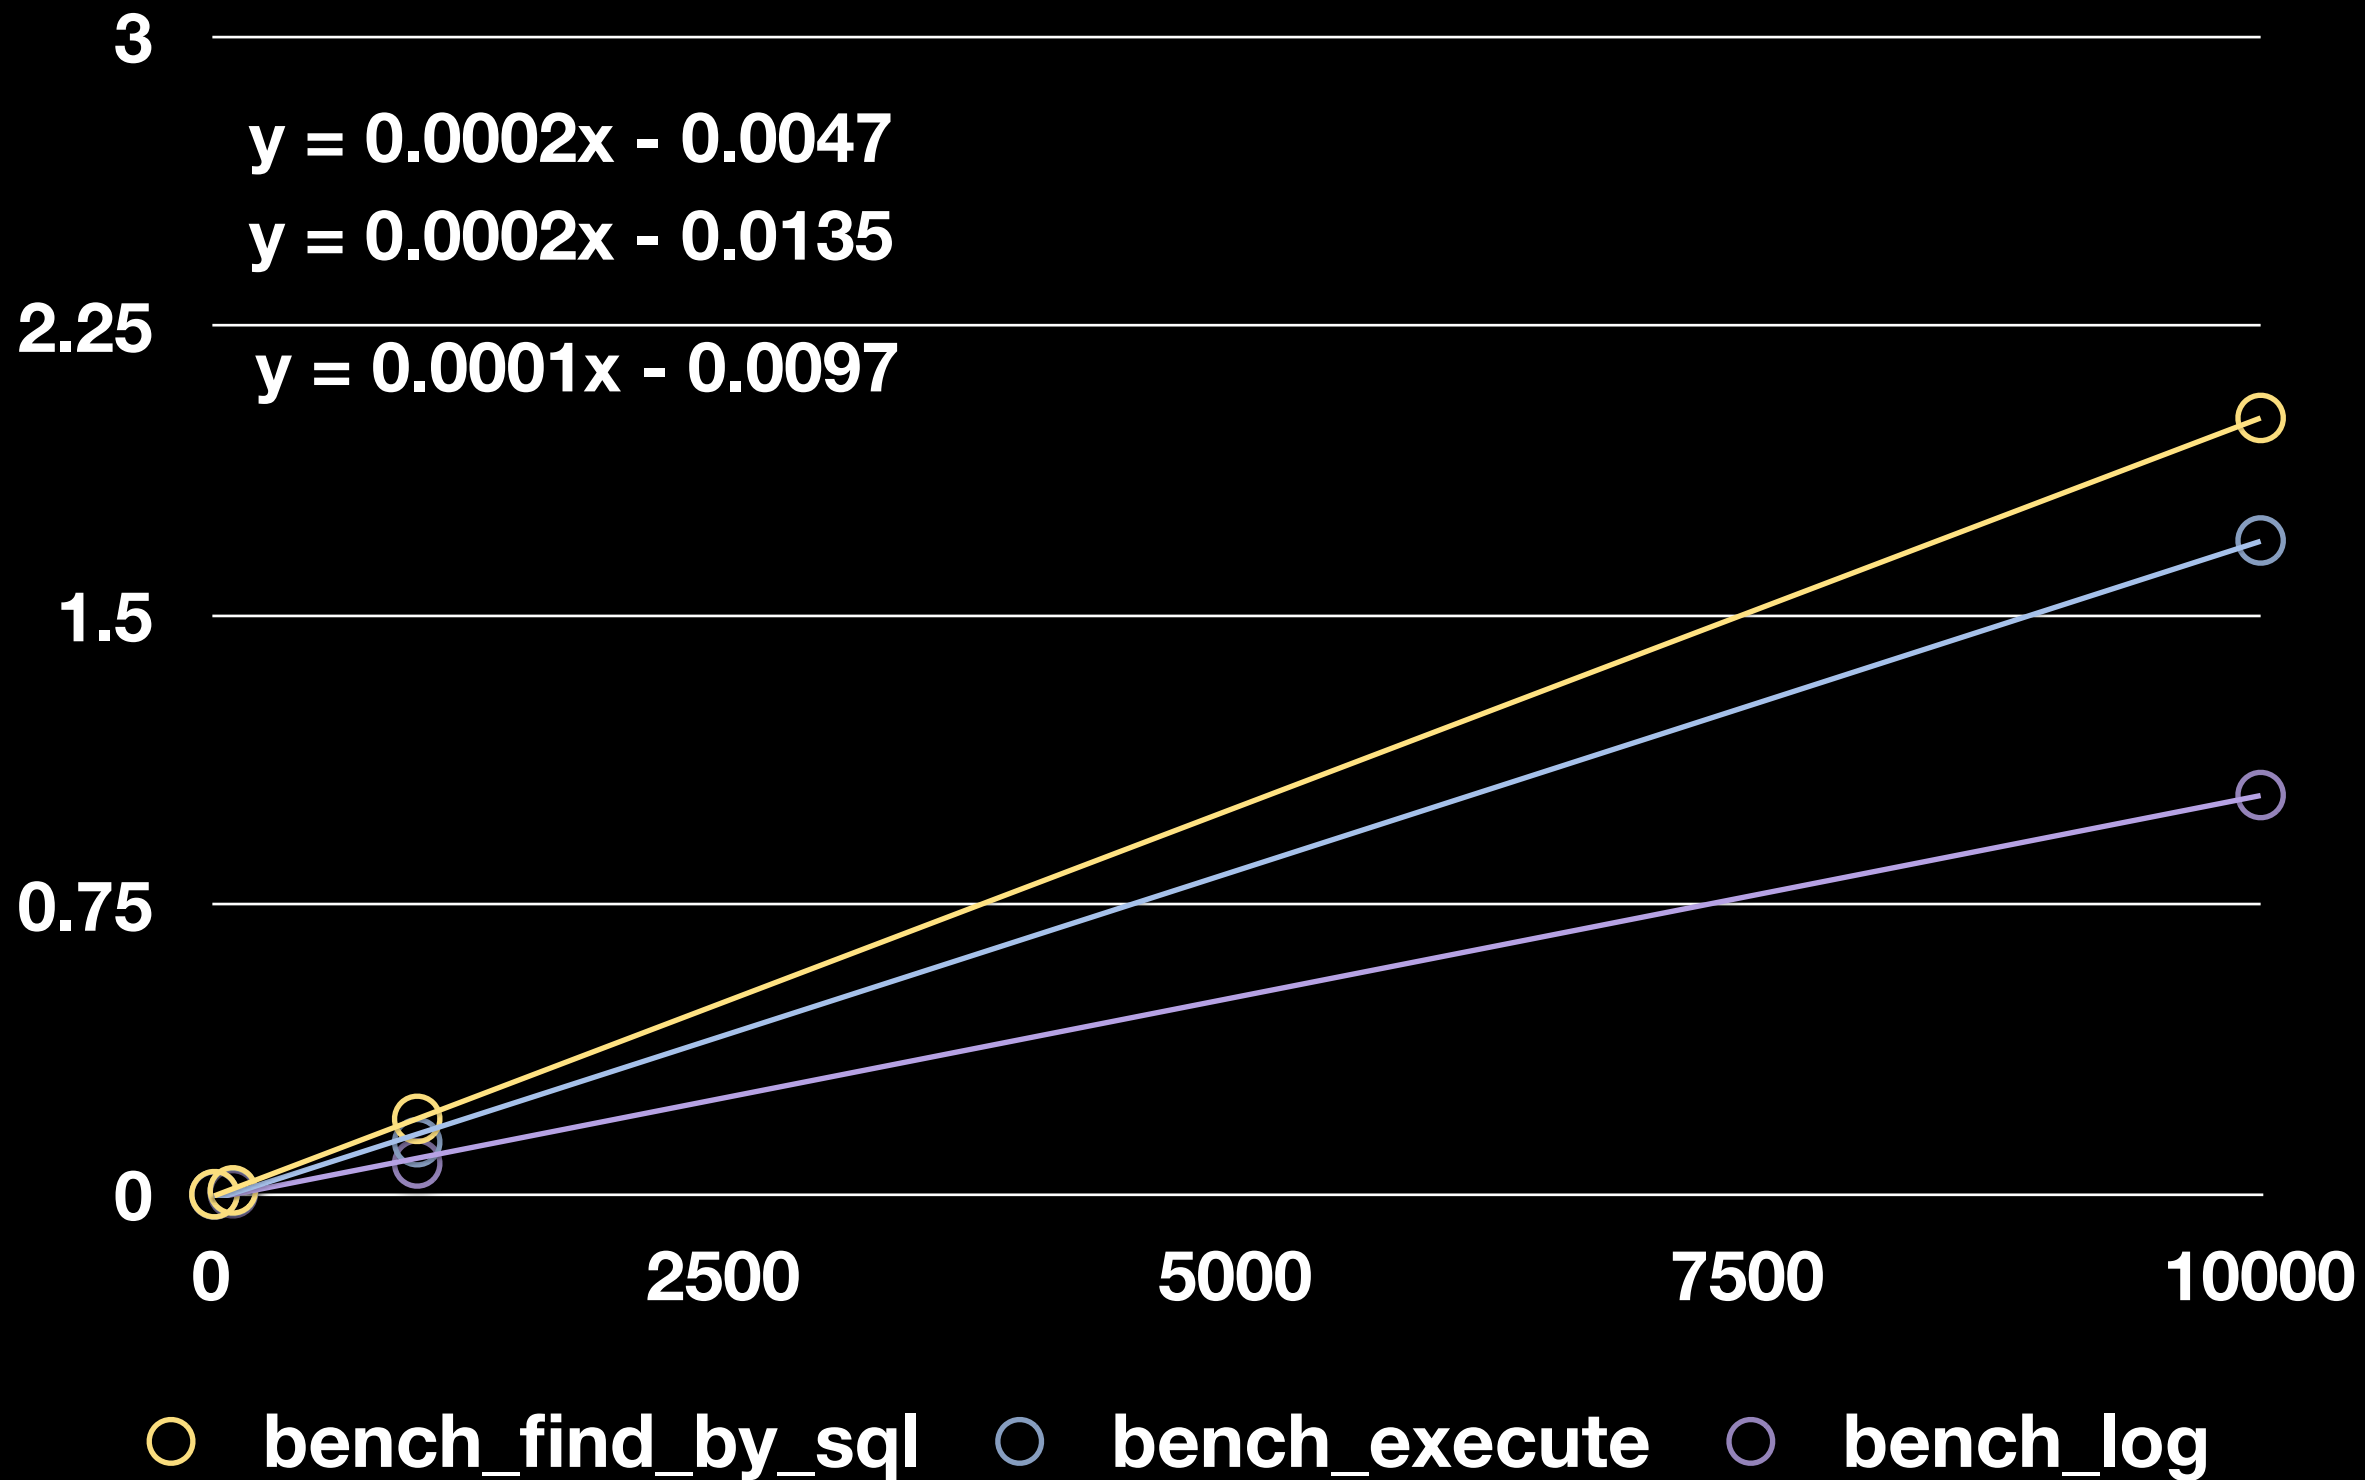
```
def bench_execute
  conn = Post.connection
  assert_performance_linear 0.999 do |n|
    n.times do
      conn.execute(
        'SELECT * FROM posts WHERE id = 1')
    end
  end
end
```

log()


```
def bench_log
  conn = Post.connection
  class << conn
    public :log
  end

  assert_performance_linear 0.999 do |n|
    n.times do
      conn.log('SQL', 'hi mom!') {}
    end
  end
end
```

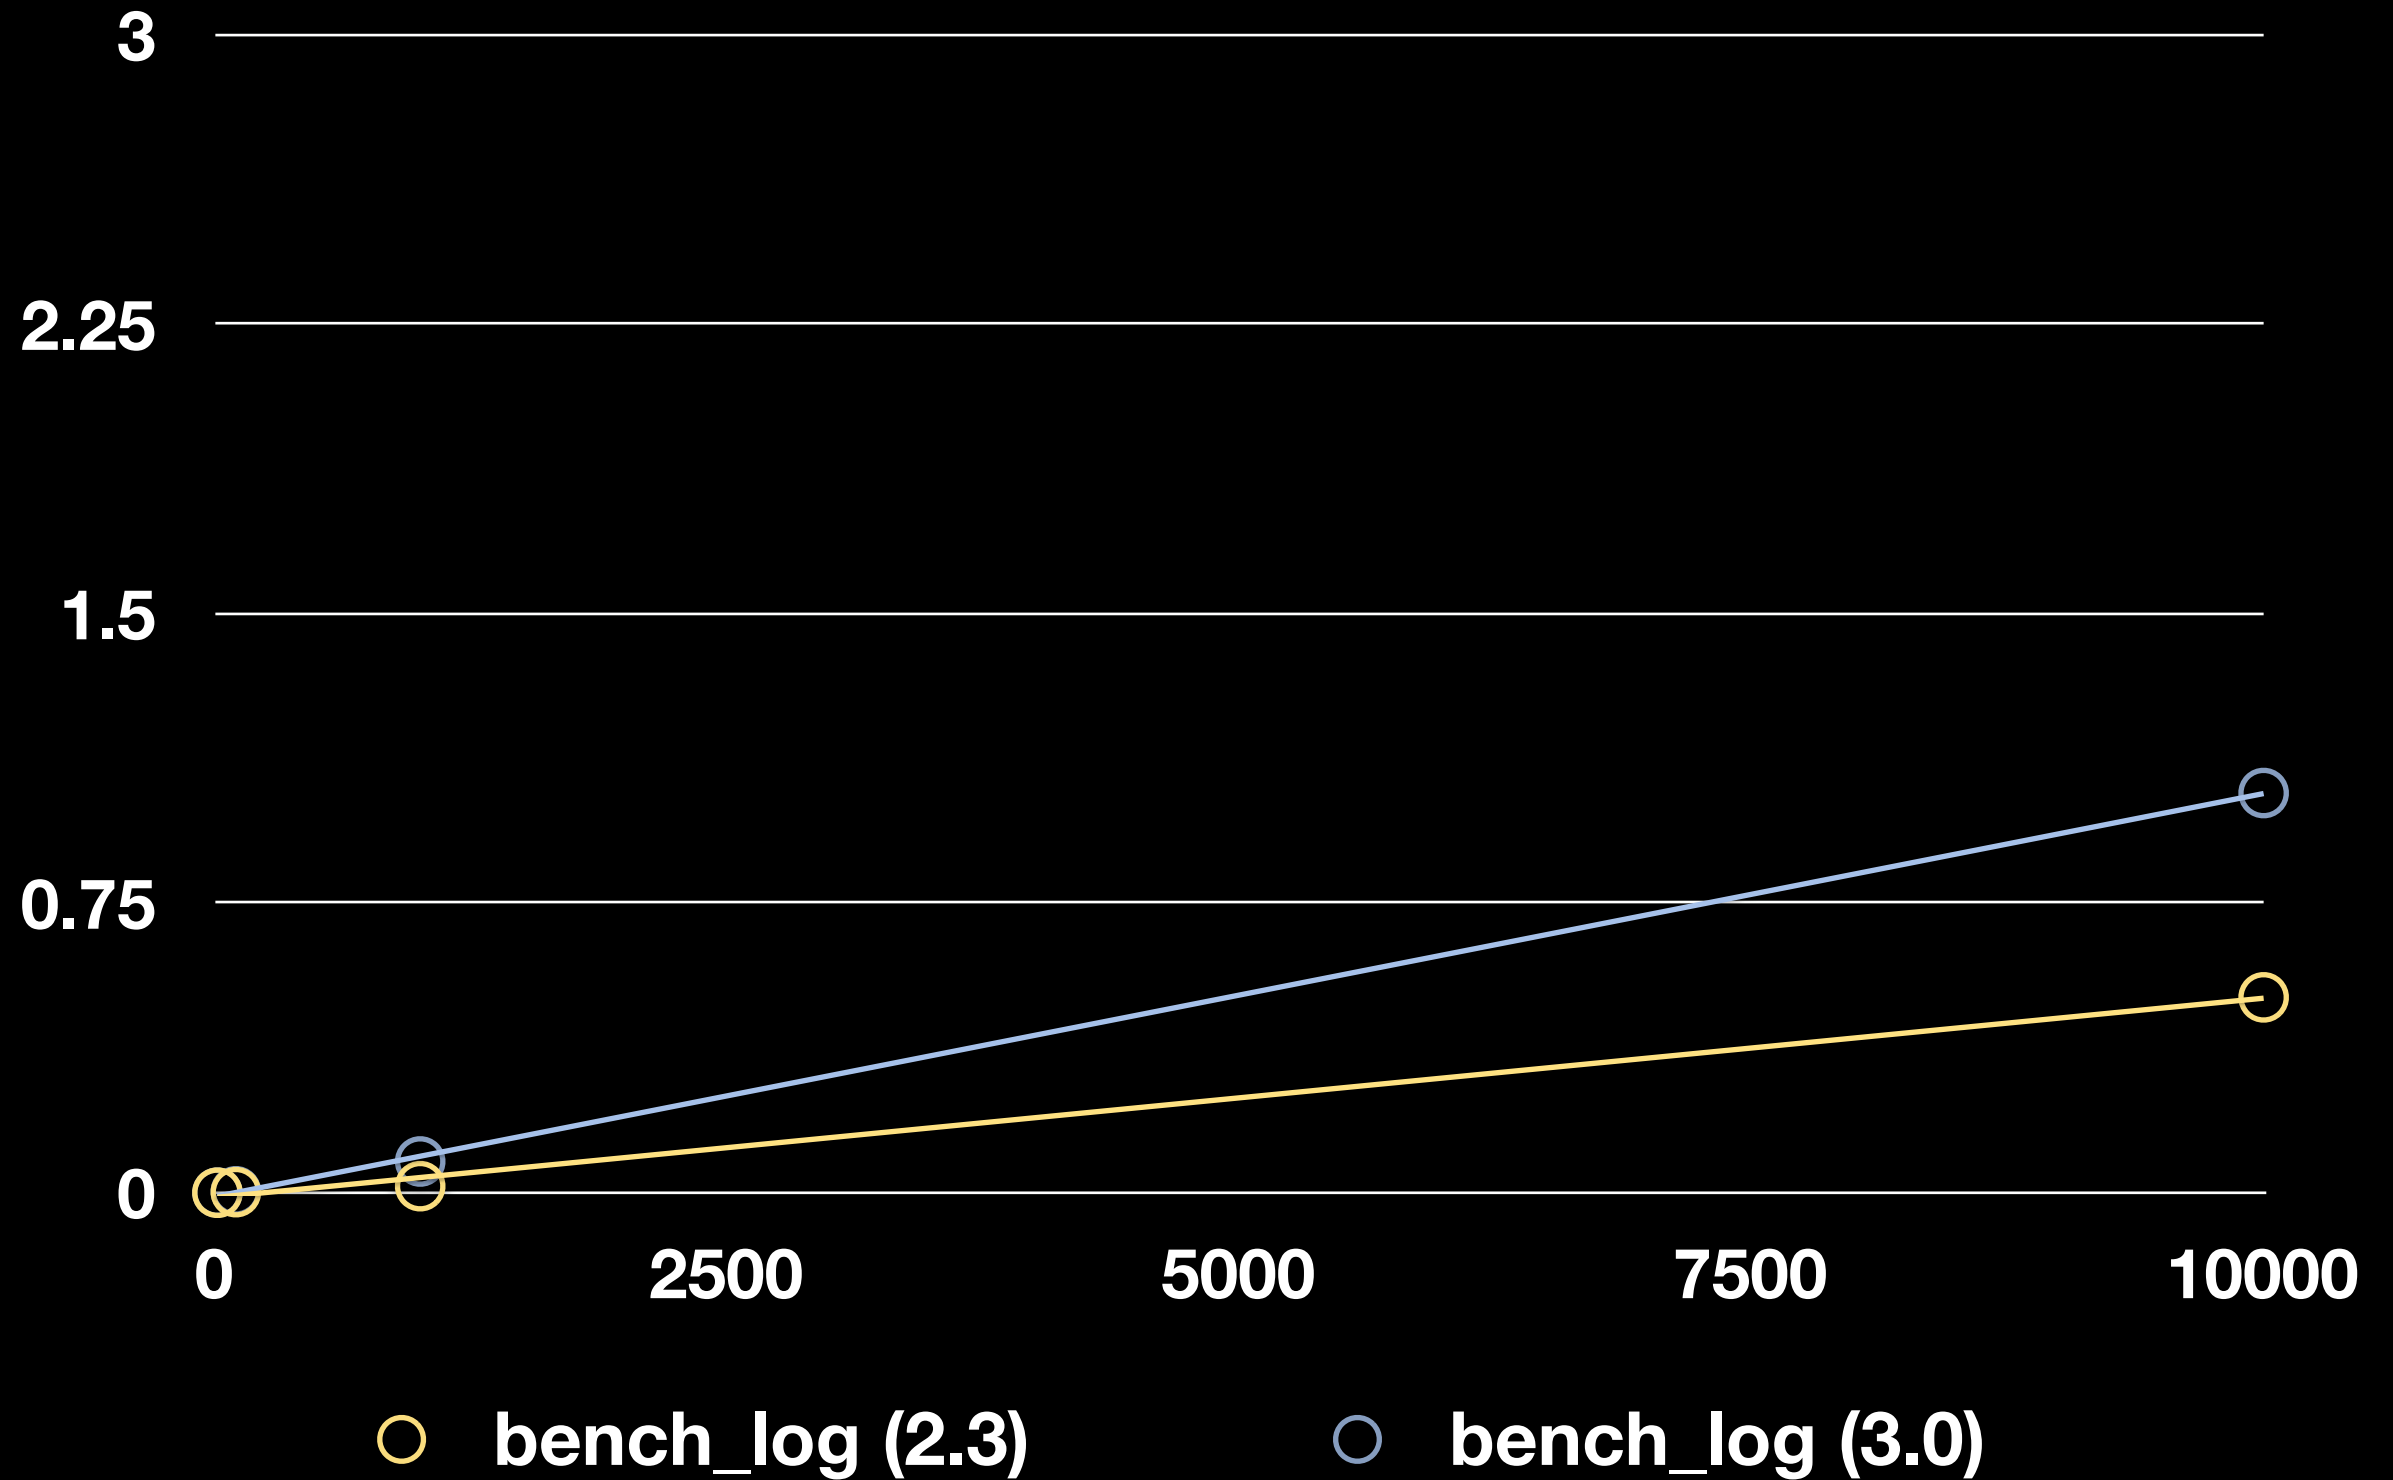
ActiveRecord 3.0 Beta



ActiveRecord 2.3.x



bench_log 2.3.x + 3.0



Δ find_by_sql()

=

Δ execute()

=

Δ log()

$$\Delta \text{execute}() - \Delta \text{log}() = 0$$

Method Call Analysis

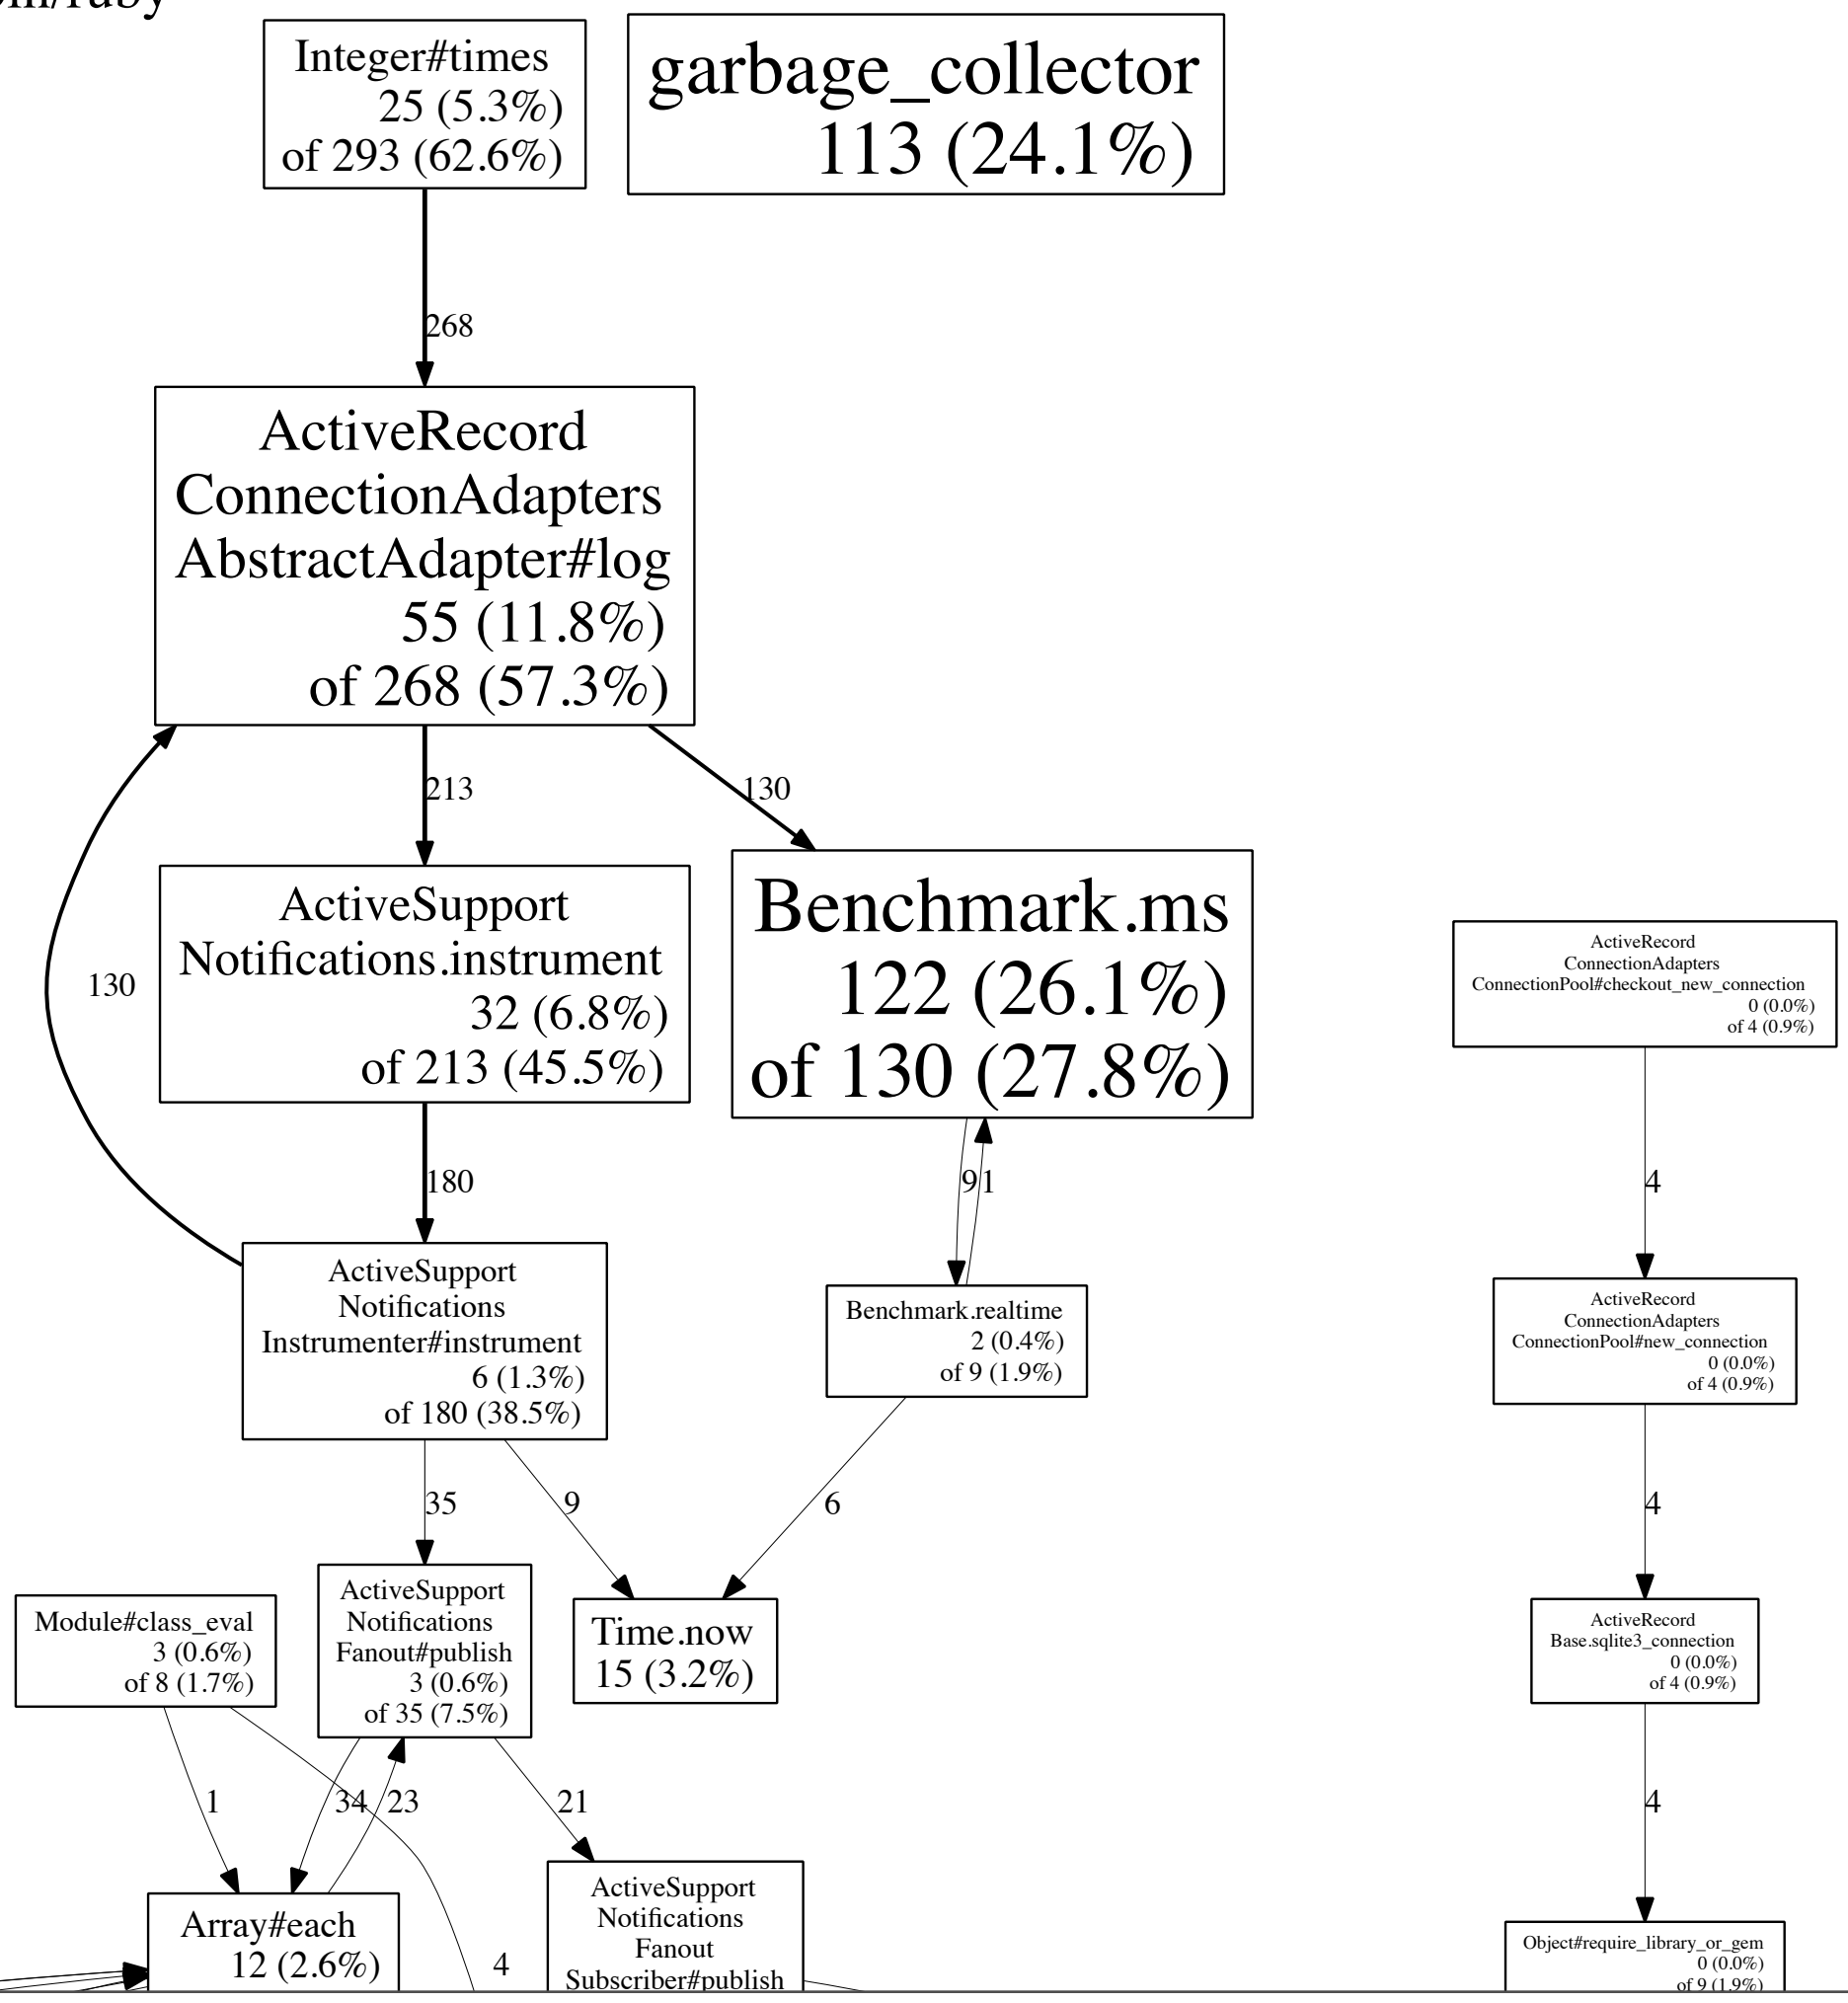
perftools.rb

<http://github.com/tmm1/perftools.rb/>

```
CPUPROFILE=/tmp/  
my_profile RUBYOPT="-  
r`gem which perftools | tail  
-1`" ruby ...
```

```
CPUPROFILE=/tmp/  
my_profile RUBYOPT="-  
r`gem which perftools | tail  
-1`" ruby ...
```


Rails 3.0 Beta

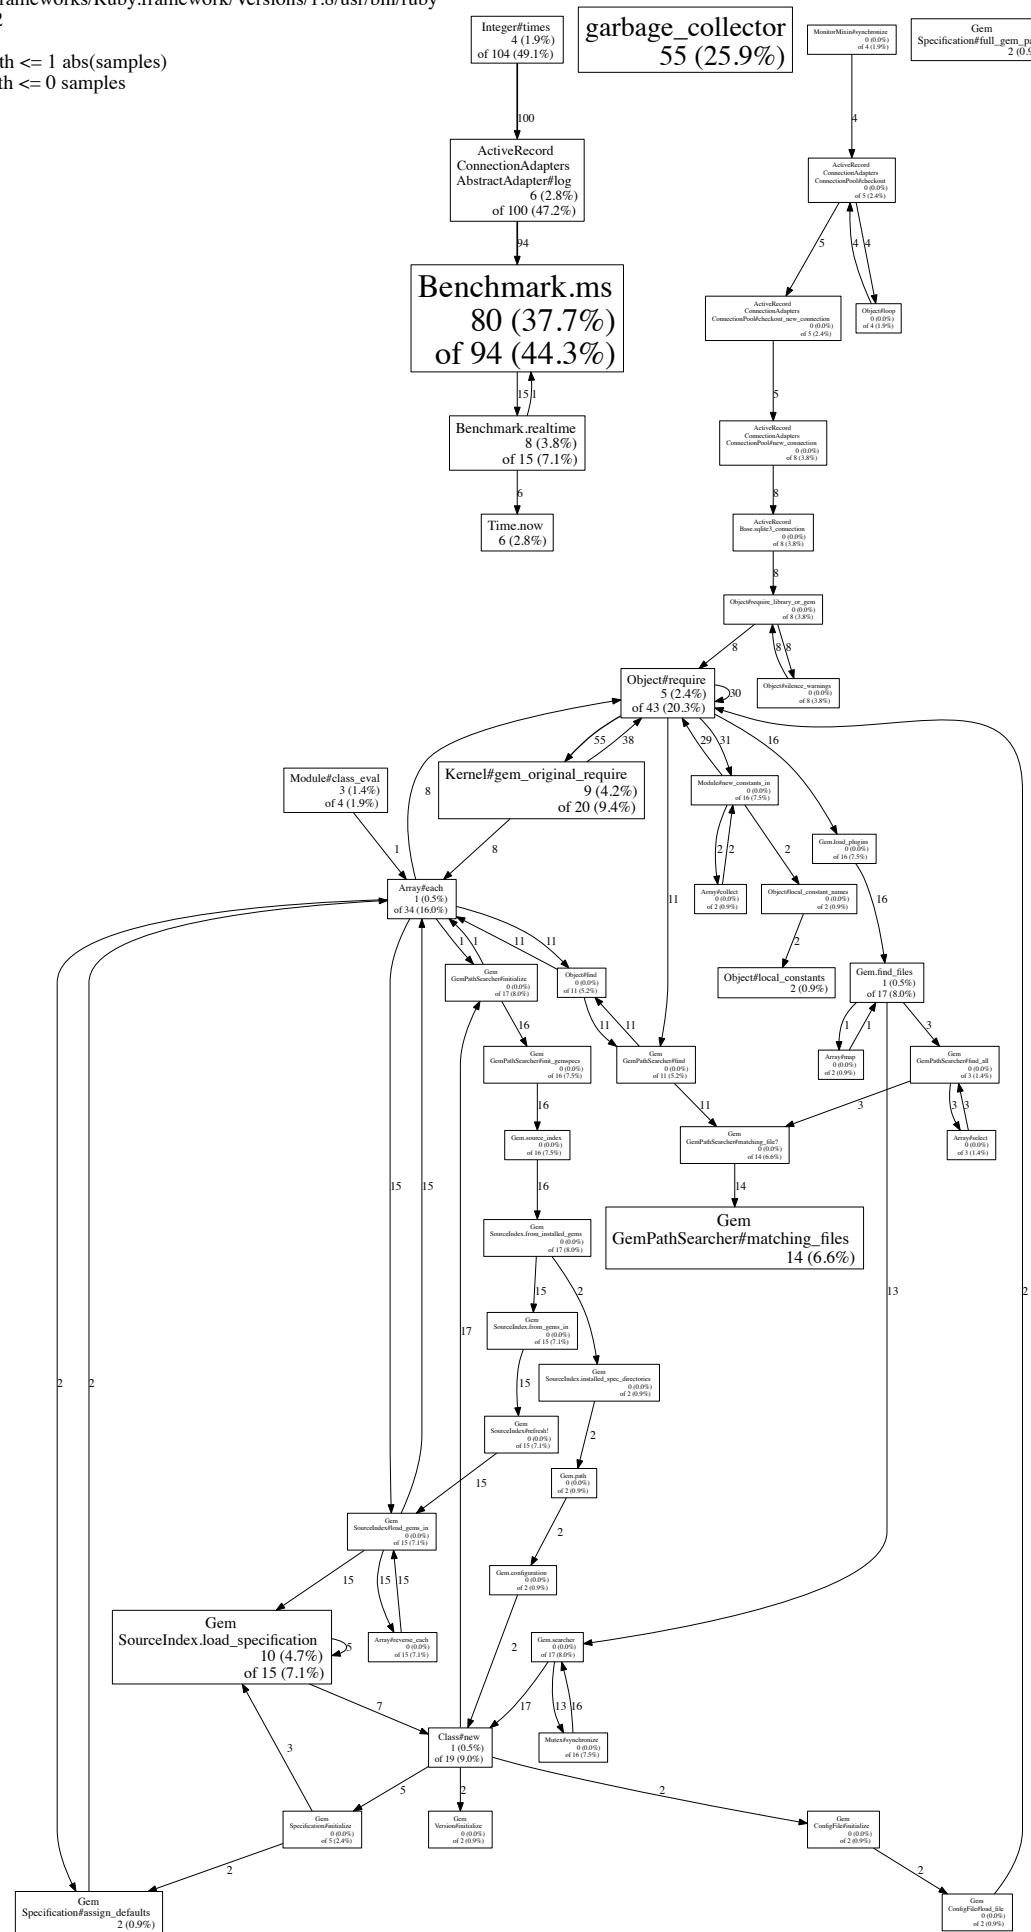


Total: 468 samples

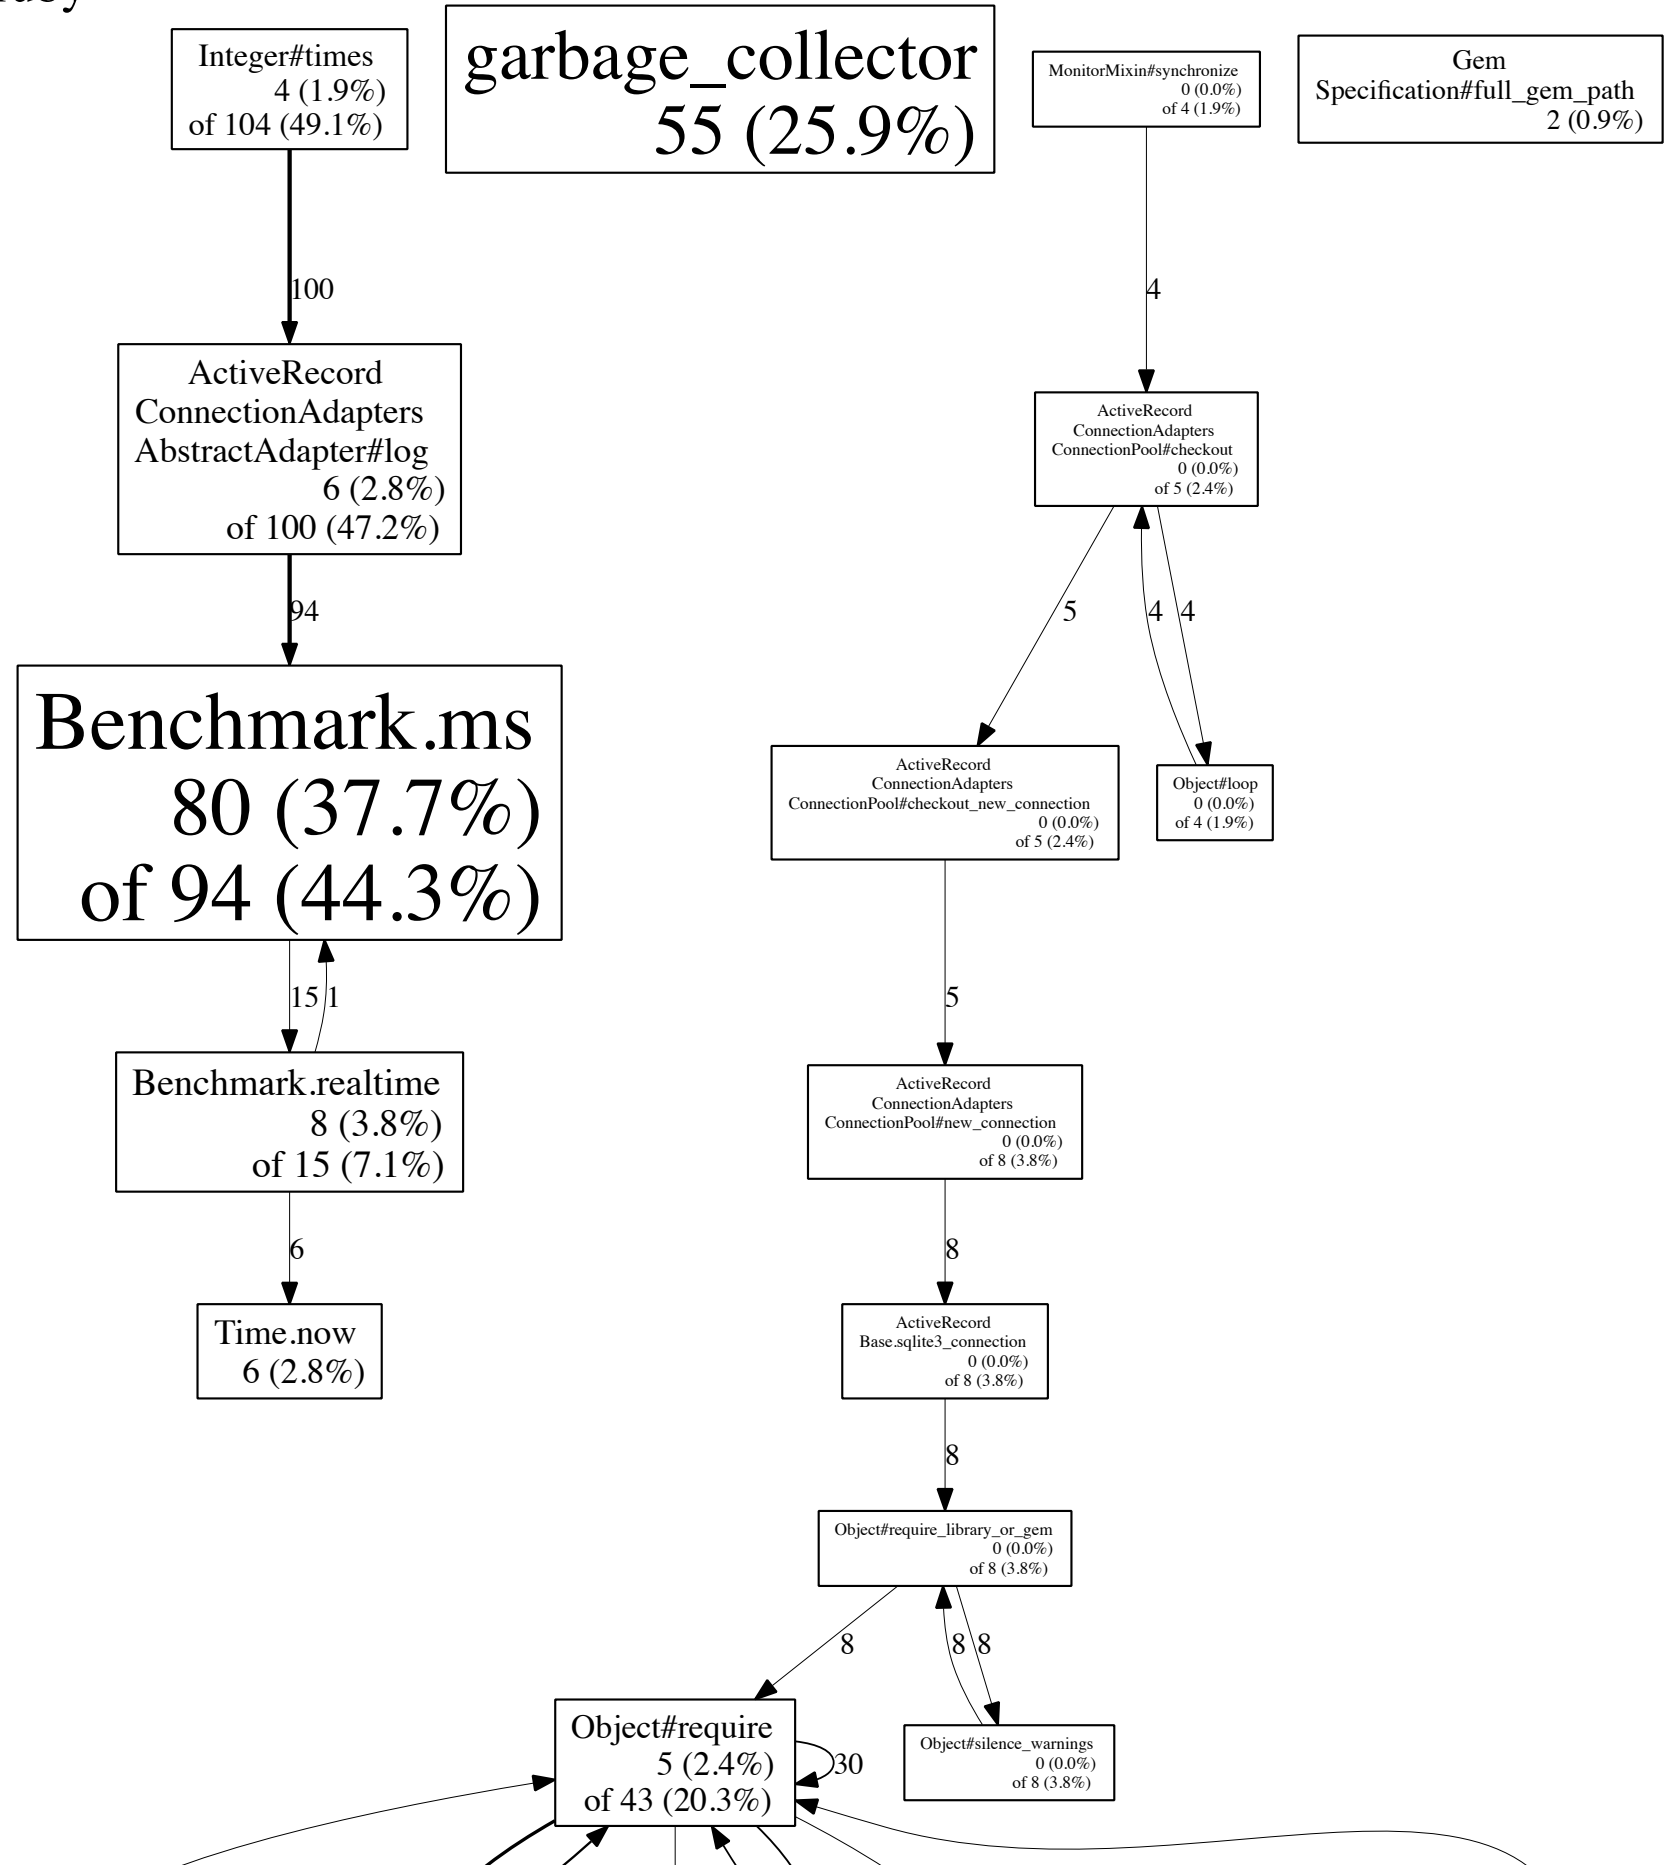
122	26.1%	26.1%	130	27.8%	Benchmark.ms
113	24.1%	50.2%	113	24.1%	garbage_collector
55	11.8%	62.0%	268	57.3%	::AbstractAdapter#log
32	6.8%	68.8%	213	45.5%	::Notifications.instrument
25	5.3%	74.1%	293	62.6%	Integer#times
18	3.8%	78.0%	18	3.8%	::LogSubscriber#call
15	3.2%	81.2%	15	3.2%	Time.now
13	2.8%	84.0%	13	2.8%	GemPathSearcher#matching_files
12	2.6%	86.5%	70	15.0%	Array#each
9	1.9%	88.5%	15	3.2%	SourceIndex.load_specification
9	1.9%	90.4%	51	10.9%	Object#require

Rails 2-3-stable

/System/Library/Frameworks/Ruby.framework/Versions/1.8/usr/bin/ruby
Total samples: 212
Focusing on: 212
Dropped nodes with <= 1 abs(samples)
Dropped edges with <= 0 samples



abs(samples)
samples



Total: 212 samples

80	37.7%	37.7%	94	44.3%	Benchmark.ms
55	25.9%	63.7%	55	25.9%	garbage_collector
14	6.6%	70.3%	14	6.6%	::GemPathSearcher#matching_files
10	4.7%	75.0%	15	7.1%	::SourceIndex.load_specification
9	4.2%	79.2%	20	9.4%	Kernel#gem_original_require
8	3.8%	83.0%	15	7.1%	Benchmark.realtime
6	2.8%	85.8%	100	47.2%	::AbstractAdapter#log
6	2.8%	88.7%	6	2.8%	Time.now
5	2.4%	91.0%	43	20.3%	Object#require
4	1.9%	92.9%	104	49.1%	Integer#times
3	1.4%	94.3%	4	1.9%	Module#class_eval
2	0.9%	95.3%	2	0.9%	::Specification#assign_defaults
2	0.9%	96.2%	2	0.9%	Gem::Specification#full_gem_path
2	0.9%	97.2%	2	0.9%	Object#local_constants

ruby-prof

Usage

```
result = RubyProf.profile do
  ...
end
printer = RubyProf::FlatPrinter.new(result)
printer.print(STDOUT, 0)
```

$$n = 1000$$

Rails 3.0 Beta

Total: 0.160831

%self	total	self	wait	child	calls	name
28.26	0.15	0.05	0.00	0.10	1000	<::Notifications>#instrument
9.36	0.03	0.02	0.00	0.01	1000	<::Benchmark>#realtime
9.16	0.09	0.01	0.00	0.07	1000	::Instrumenter#instrument
7.28	0.02	0.01	0.00	0.01	4000	<Class::Time>#now
5.41	0.16	0.01	0.00	0.15	1000	::AbstractAdapter#log
5.10	0.01	0.01	0.00	0.01	1000	<::Notifications>#instrumenter
2.92	0.00	0.00	0.00	0.00	4000	<Class::Time>#allocate
2.73	0.02	0.00	0.00	0.01	1000	Array#each
2.57	0.03	0.00	0.00	0.03	1000	<Module::Benchmark>#ms
2.57	0.00	0.00	0.00	0.00	4000	Time#initialize
2.55	0.03	0.00	0.00	0.02	1000	Notifications::Fanout#publish
2.45	0.01	0.00	0.00	0.01	1000	LogSubscriber#call
2.37	0.01	0.00	0.00	0.01	1000	::Fanout::Subscriber#publish
2.04	0.01	0.00	0.00	0.00	1000	LogSubscriber#logger
1.84	0.00	0.00	0.00	0.00	1000	::Fanout#listeners_for
1.64	0.16	0.00	0.00	0.16	1	Integer#times

Rails 2-3-stable

Thread ID: 2148237740

Total: 0.051336

%self	total	self	wait	child	calls	name
27.00	0.03	0.01	0.00	0.01	1000	<Module::Benchmark>#realtime
23.35	0.05	0.01	0.00	0.04	1000	::AbstractAdapter#log
11.67	0.01	0.01	0.00	0.00	2000	<Class::Time>#now
7.66	0.03	0.00	0.00	0.03	1000	<Module::Benchmark>#ms
5.37	0.00	0.00	0.00	0.00	1000	::AbstractAdapter#log_info
5.11	0.05	0.00	0.00	0.05	1	Integer#times
4.52	0.00	0.00	0.00	0.00	2000	<Class::Time>#allocate
3.89	0.00	0.00	0.00	0.00	2000	Time#initialize
3.83	0.00	0.00	0.00	0.00	2000	Time#to_f

Methods in Common

- **<Class::Time>#now**
- **<Class::Time>#allocate**

n = 1000

	3.0 Beta	2-3 Stable
Time#now	4000	2000
Time#allocate	4000	2000

<http://bit.ly/omgslow>

"It's all fixed!"

Wait a few hours...

**"It's better, but still
2x slower"**

Post.find(1)

ARel....

find_by_sql()

execute()

log()

ARel....

Side note:
This is when Ryan
told me to rewrite.

Superficial Improvements


**Limited
Domain / System
Knowledge**

VM Tricks

See Results Quickly

Tapers off Over Time

System Impact



attr_*


```
def some_attribute
  @some_attribute
end
```

vs

```
attr_reader :some_attribute
```


+ bench_method

○ bench_reader



attr_reader

```
case VM_METHOD_TYPE_IVAR: {
  if (argc != 0) {
    rb_raise(rb_eArgError,
             "wrong number of arguments (%d for 0)", argc);
  }
  val = rb_attr_get(recv, def->body.attr.id);
  break;
}
```

method call

```
case VM_METHOD_TYPE_ISEQ: {
    rb_control_frame_t *reg_cfp;
    int i;

    rb_vm_set_finish_env(th);
    reg_cfp = th->cfp;

    CHECK_STACK_OVERFLOW(reg_cfp, argc + 1);

    *reg_cfp->sp++ = recv;
    for (i = 0; i < argc; i++) {
        *reg_cfp->sp++ = argv[i];
    }

    vm_setup_method(th, reg_cfp, recv, argc, blockptr, 0 /* flag */, me);
    val = vm_exec(th);
    break;
}
```

vm_setup_method

- **Check for stack overflow**
- **Pushing a stack frame**
- **Copying arguments**

Predicate Methods

```
class Foo
  attr_reader :some_attribute

  def some_attribute?
    @some_attribute
  end
end
```

Predicate Methods

```
class Foo
  attr_reader :some_attribute
  alias :some_attribute? :some_attribute
end
```

Hash[] vs inject({})

inject({})

```
some_list.inject({}) do |hash, val|  
  hash[val] = some_transform(val)  
  hash  
end
```


Hash

```
values = some_list.map { |val|  
  [val, some_transform(val)]  
}  
Hash[values]
```

```
@list.inject({}) do |hash, val|  
  hash[val] = val.length  
  hash  
end
```

vs

```
Hash[@list.map { |val| [val, val.length] }]
```

+ Hash[]

○ inject({})

7

$$y = 5.268E-6x - 0.0075$$

$$y = 6.17E-6x - 0.0108$$

5.25

3.5

1.75

0

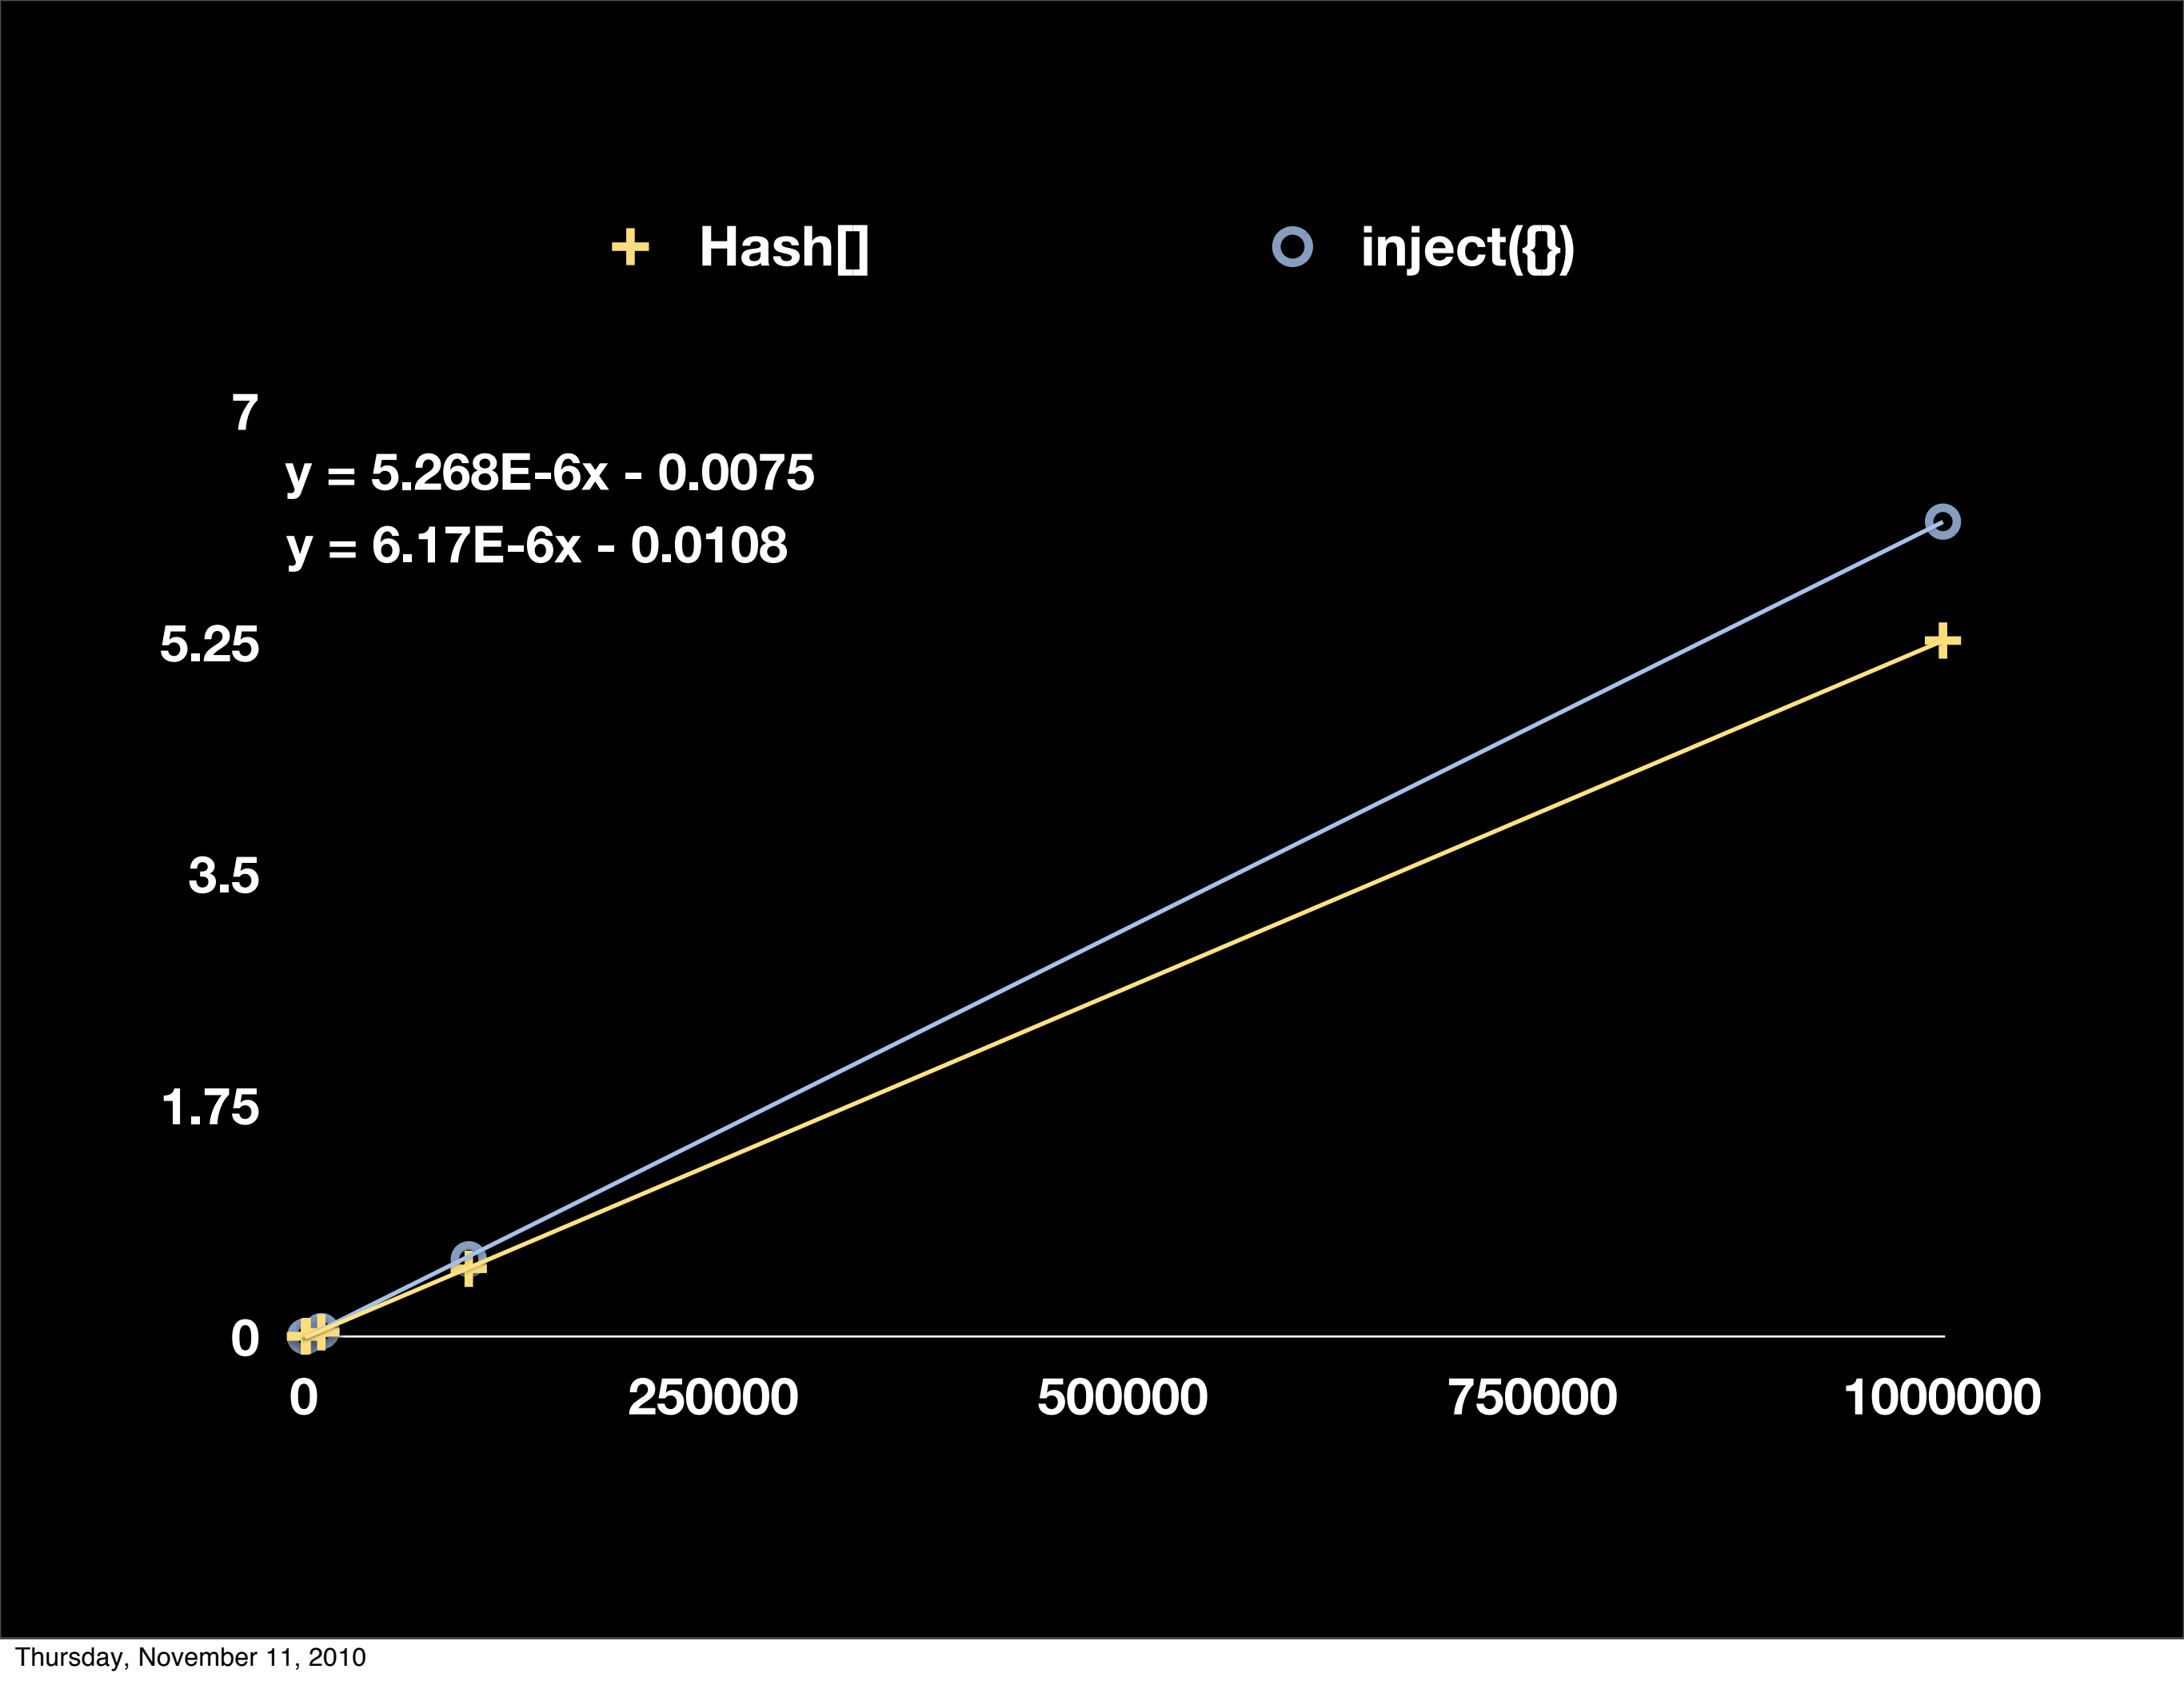
0

250000

500000

750000

1000000



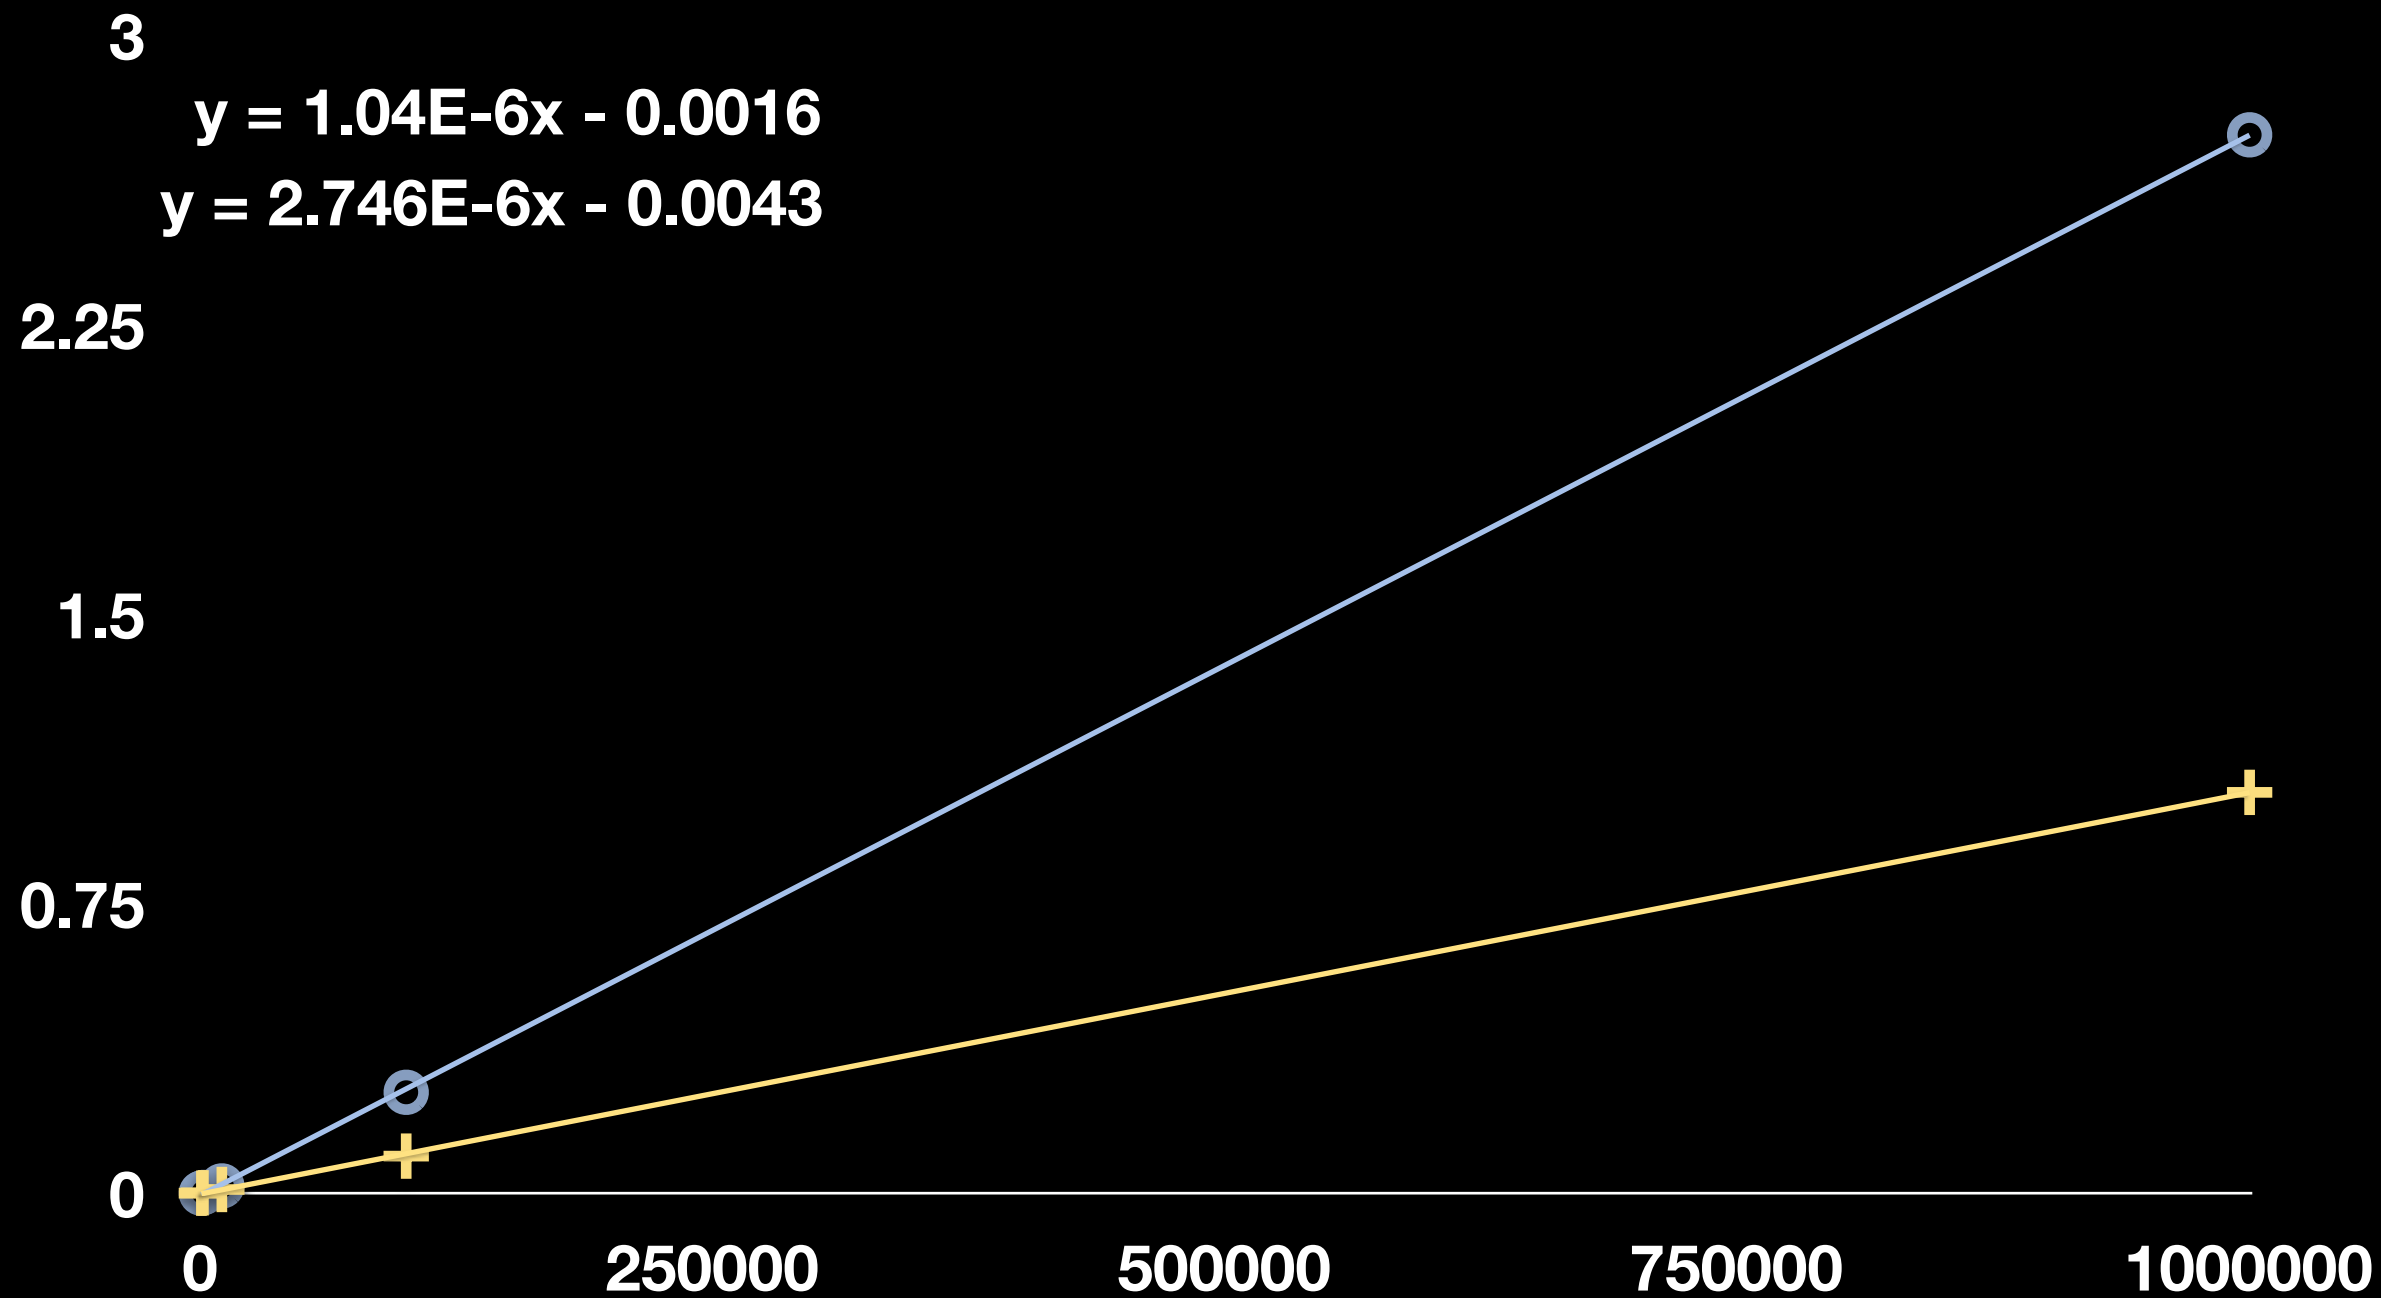
Strangeness

```
def bench_naked_each
  assert_performance_linear 0.999 do |n|
    m = nil
    n.times { @list.each { |v| m = v } }
  end
end
```

```
def bench_naked_inject
  assert_performance_linear 0.999 do |n|
    n.times { @list.inject { |m,v| m = v } }
  end
end
```

+ naked each

○ naked inject



TANGENT

When to use inject()

**When one
calculation depends
on the previous**

```
@list.inject({}) do |hash, val|  
  hash[val] = val.length  
  hash  
end
```

vs

```
Hash[@list.map { |val| [val, val.length] }]
```

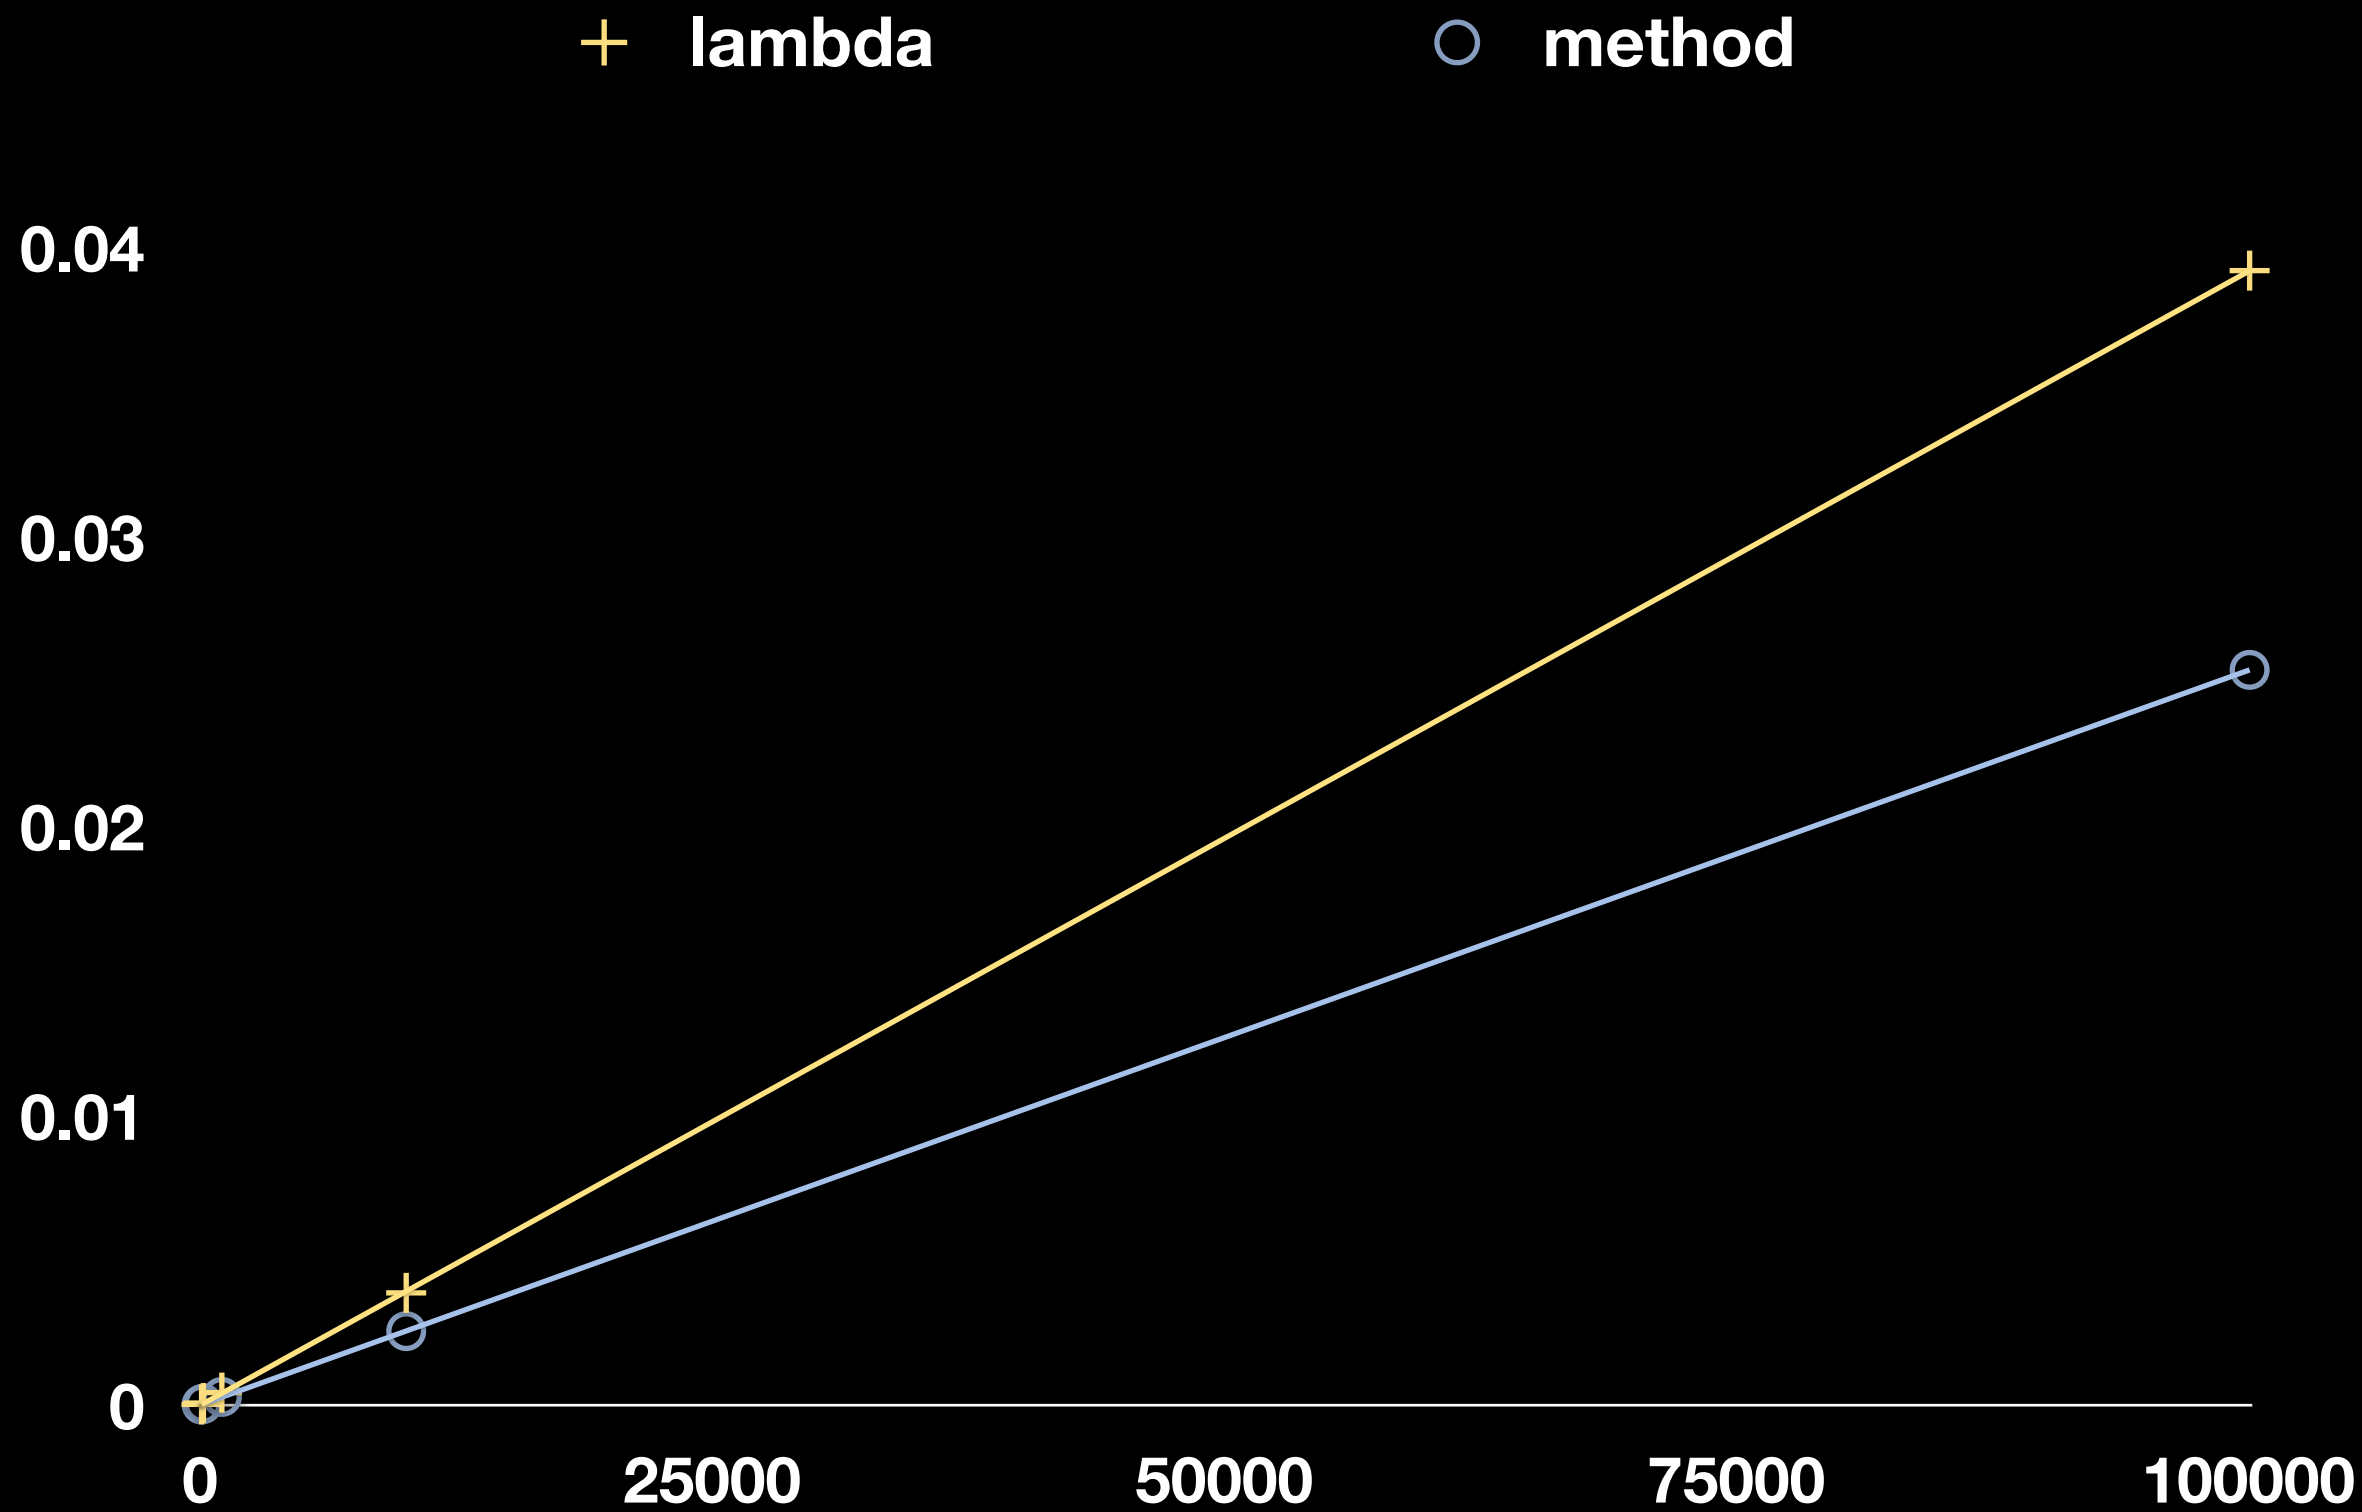
```
%w{  
  Foo  
  Bar  
  Baz  
}.inject(Object) { |klass,string|  
  klass.const_get(string.to_sym)  
}
```

Proc Activation

```
lambda { ... }
```

```
# vs
```

```
class Callable  
  def call; ... end  
end
```



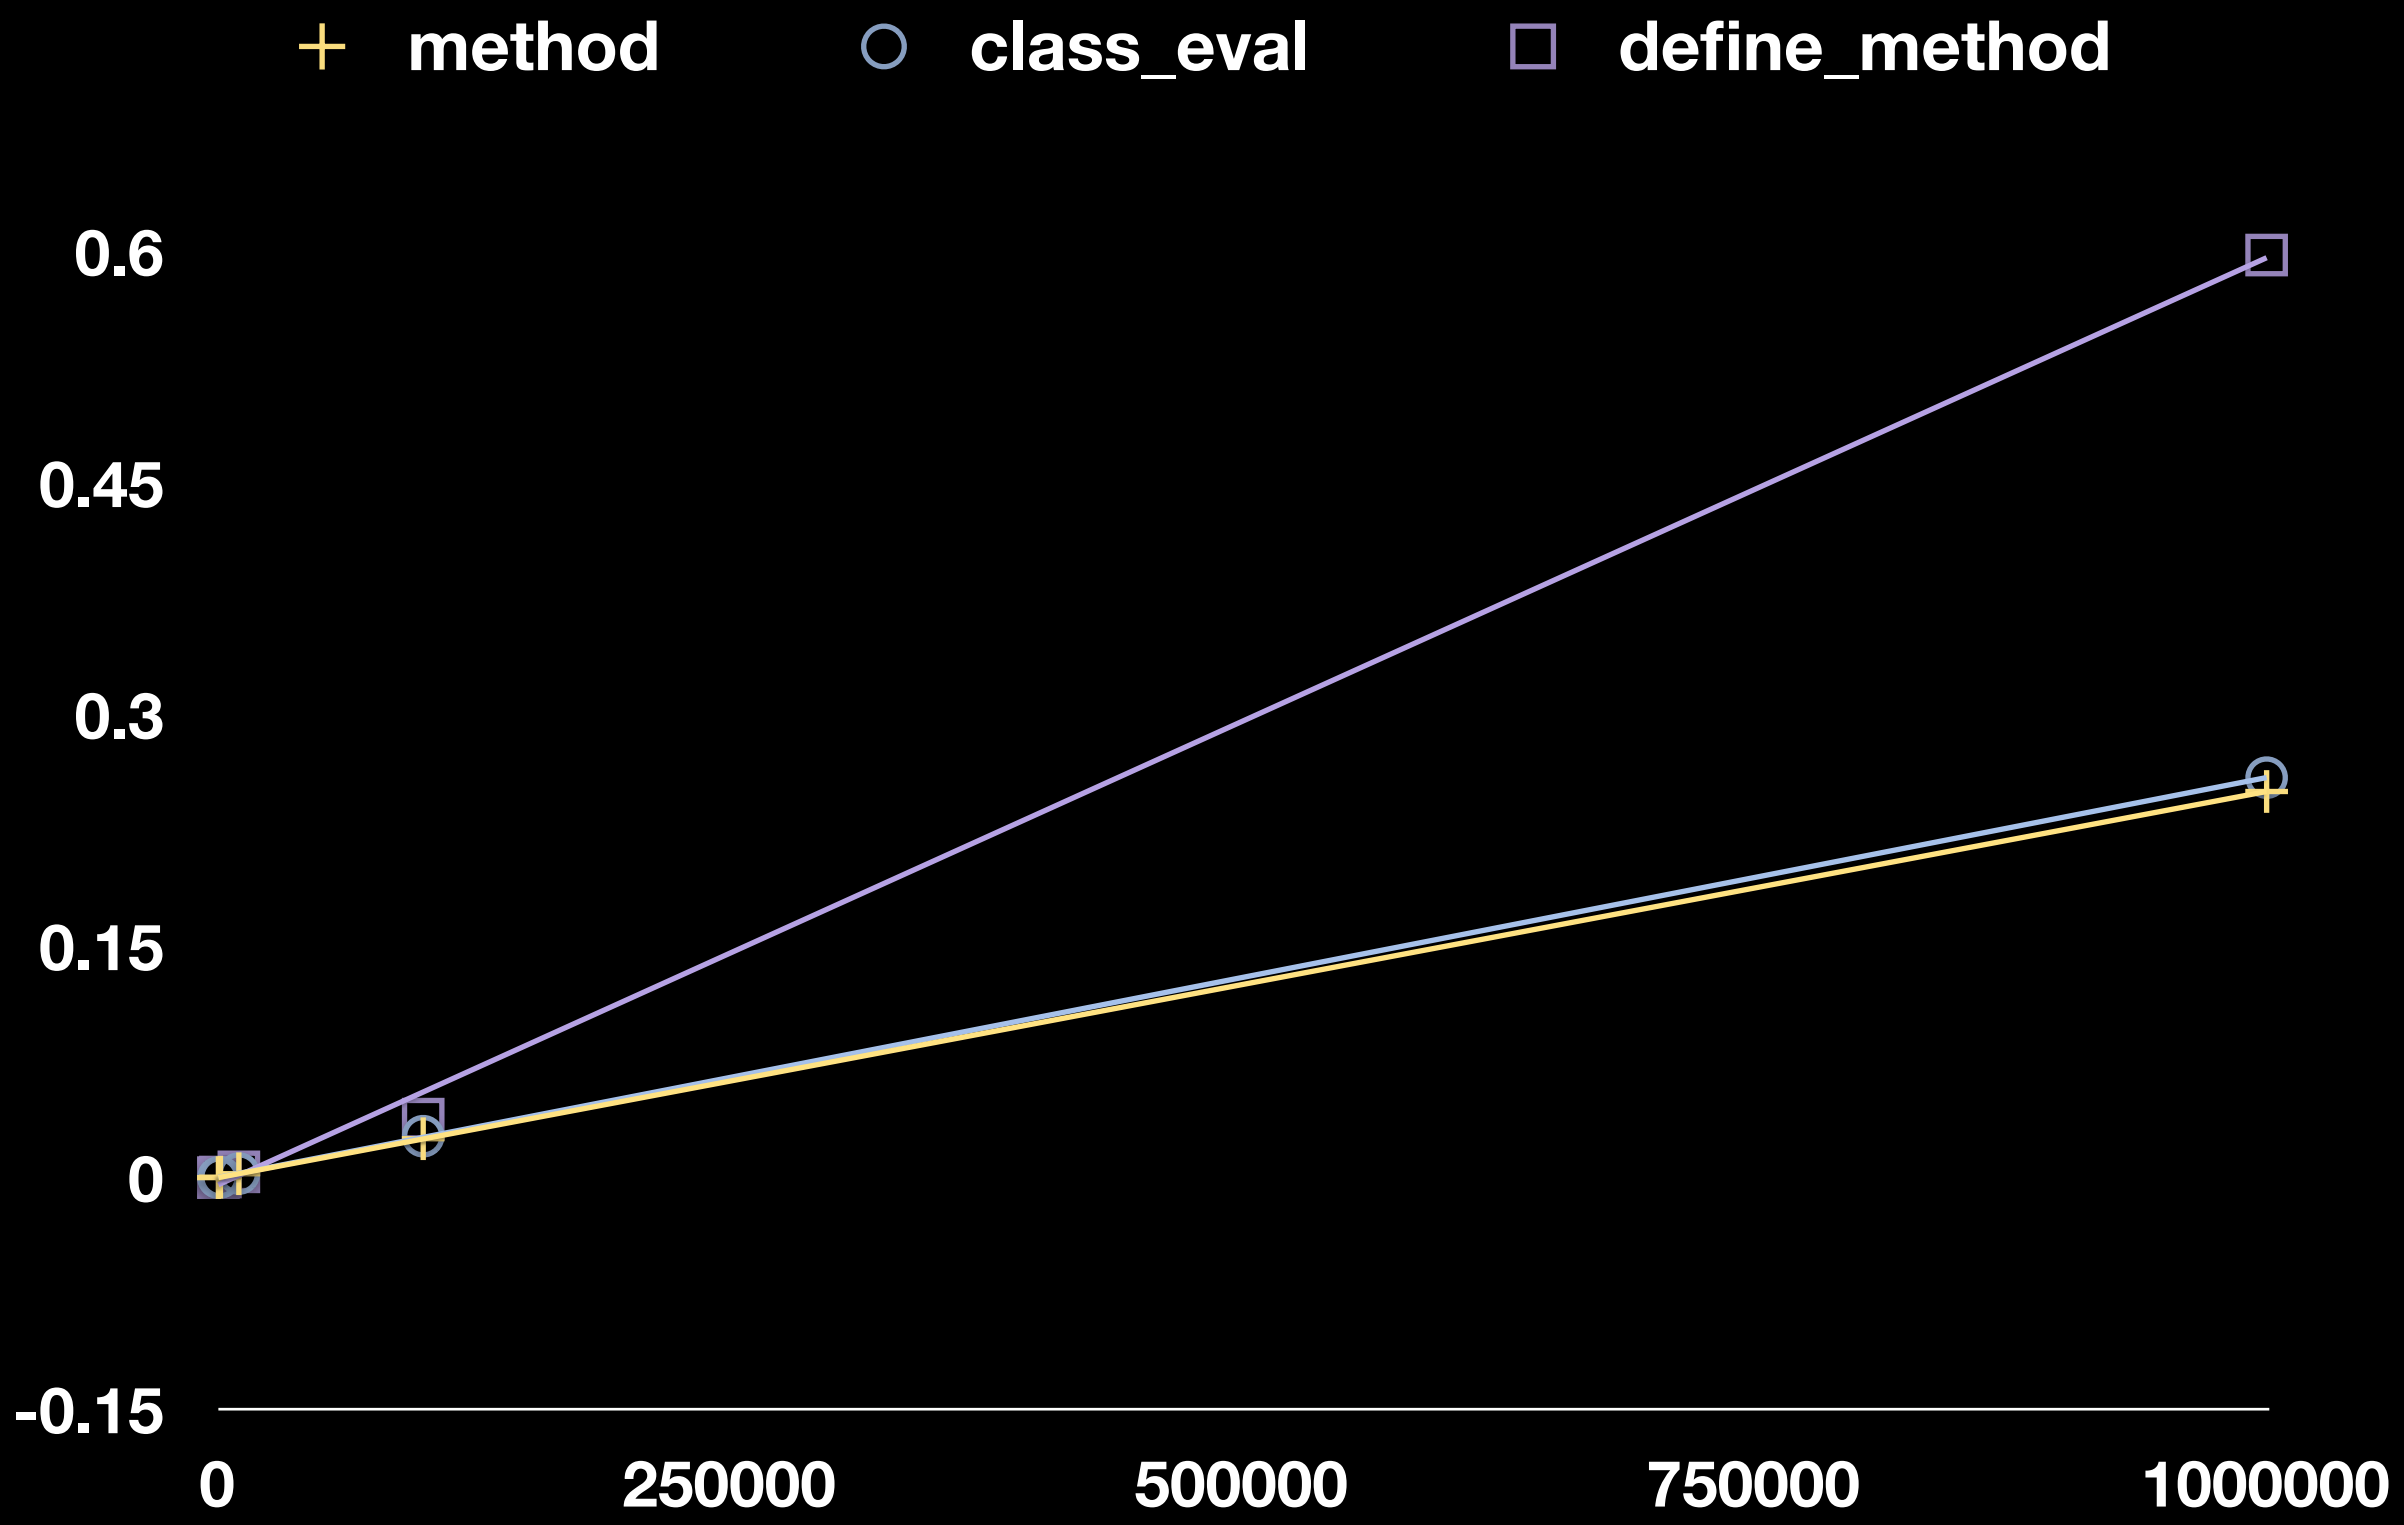
is_a?(Proc)

respond_to?(:call)


```
class Callable
  def call(...)
    ...
  end
end
```

define_method

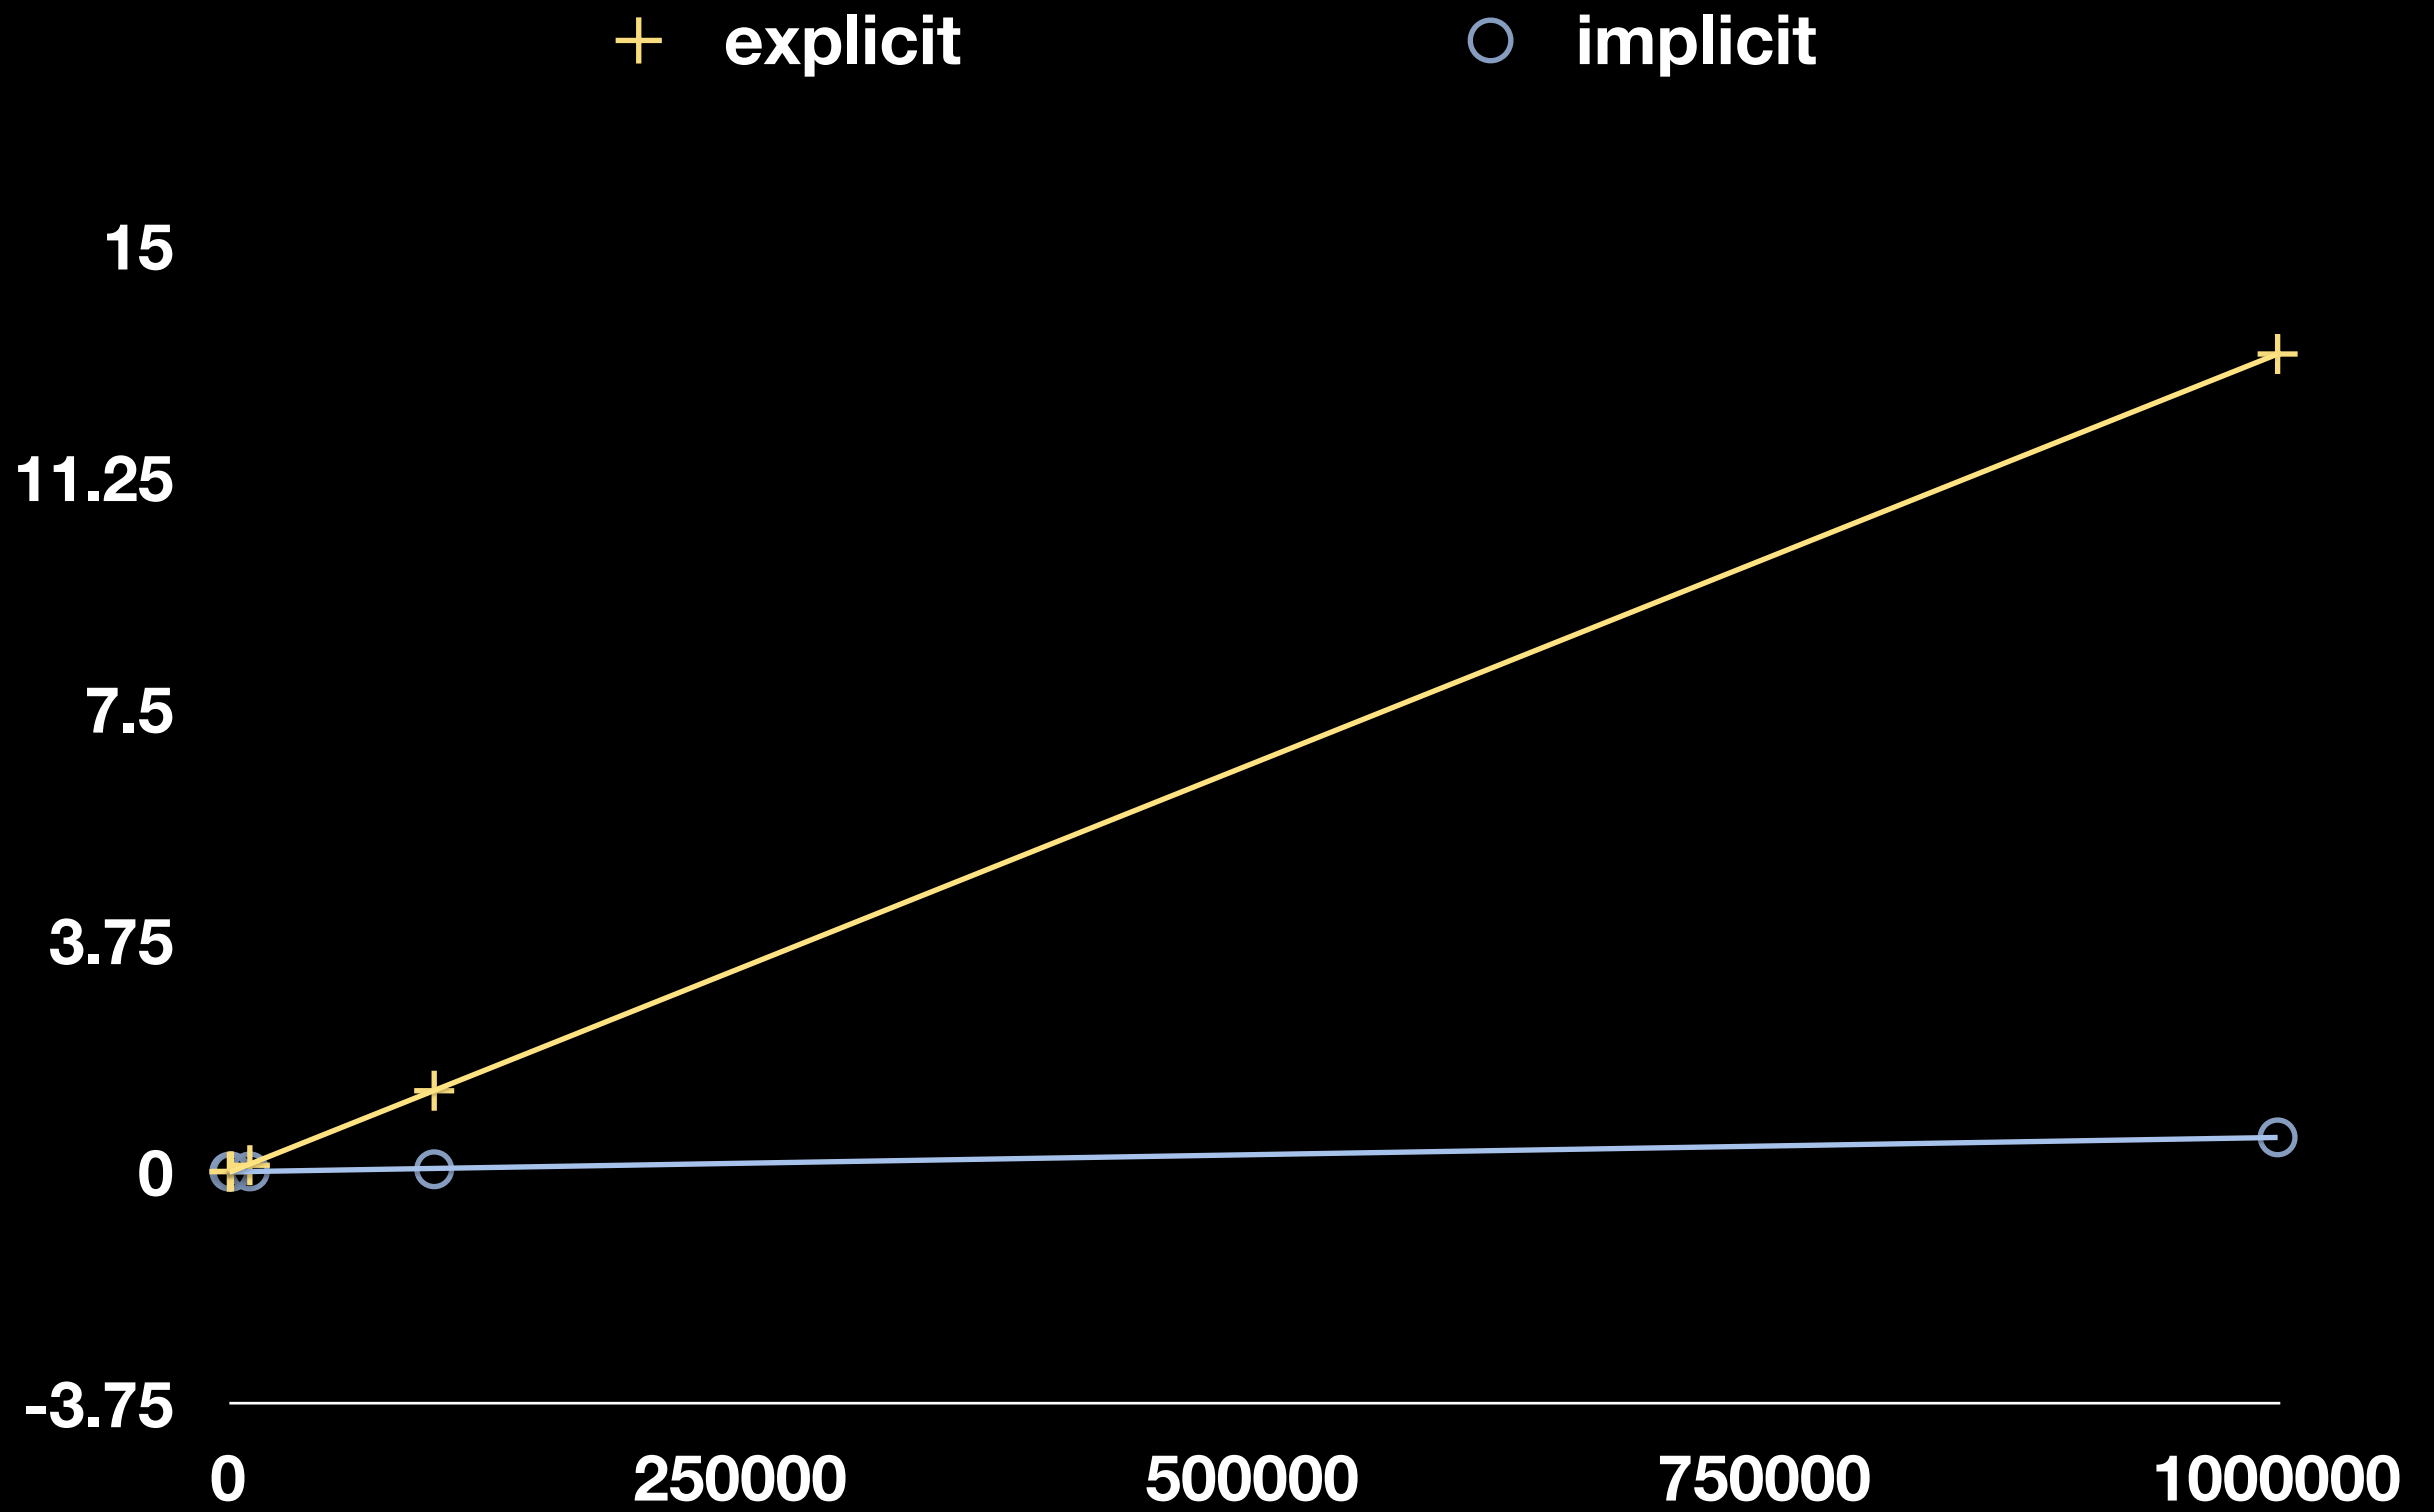
```
class Foo
  def foo; end
  define_method :bar do; end
  class_eval %{ def baz; end }
end
```



Explicit Block Parameters

```
class Foo
  def explicit &block
    yield
  end

  def implicit
    yield
  end
end
```



```
def sometimes_block
  if block_given?
    Proc.new.call
  end
end
```

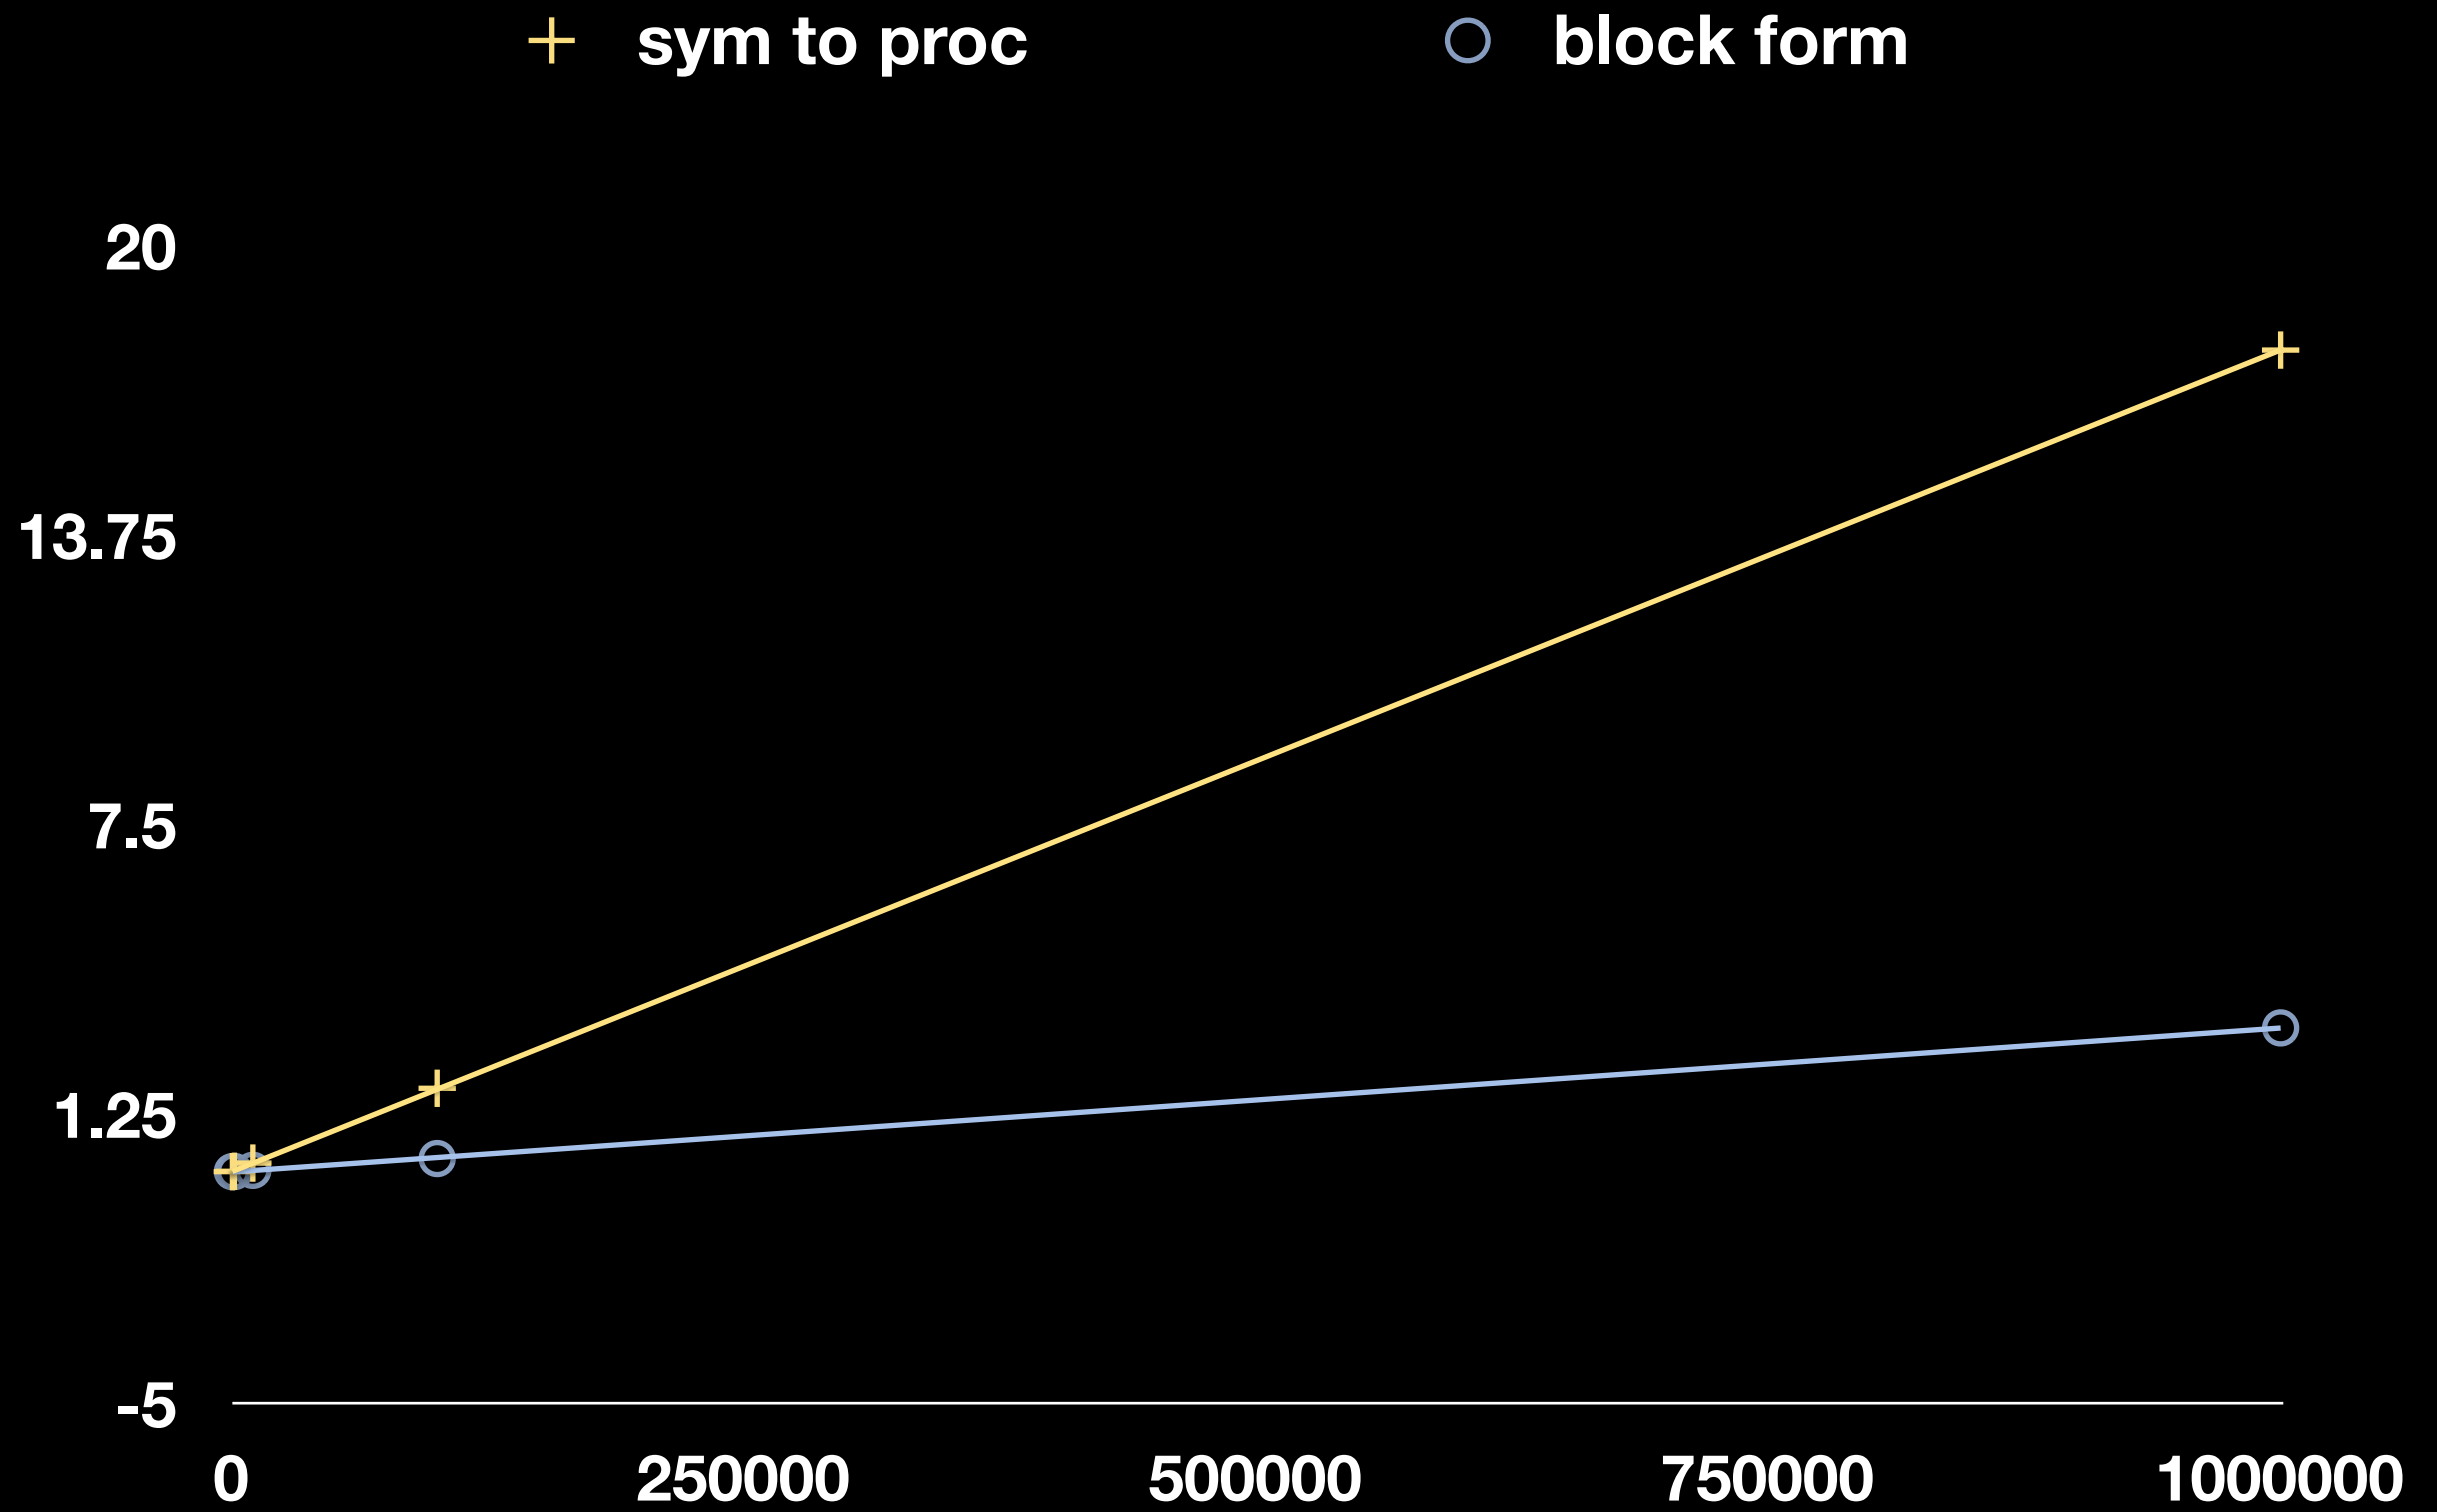
```
sometimes_block { puts "hi" }
sometimes_block
```


Symbol to Proc

```
@list.map(&:to_i)
```

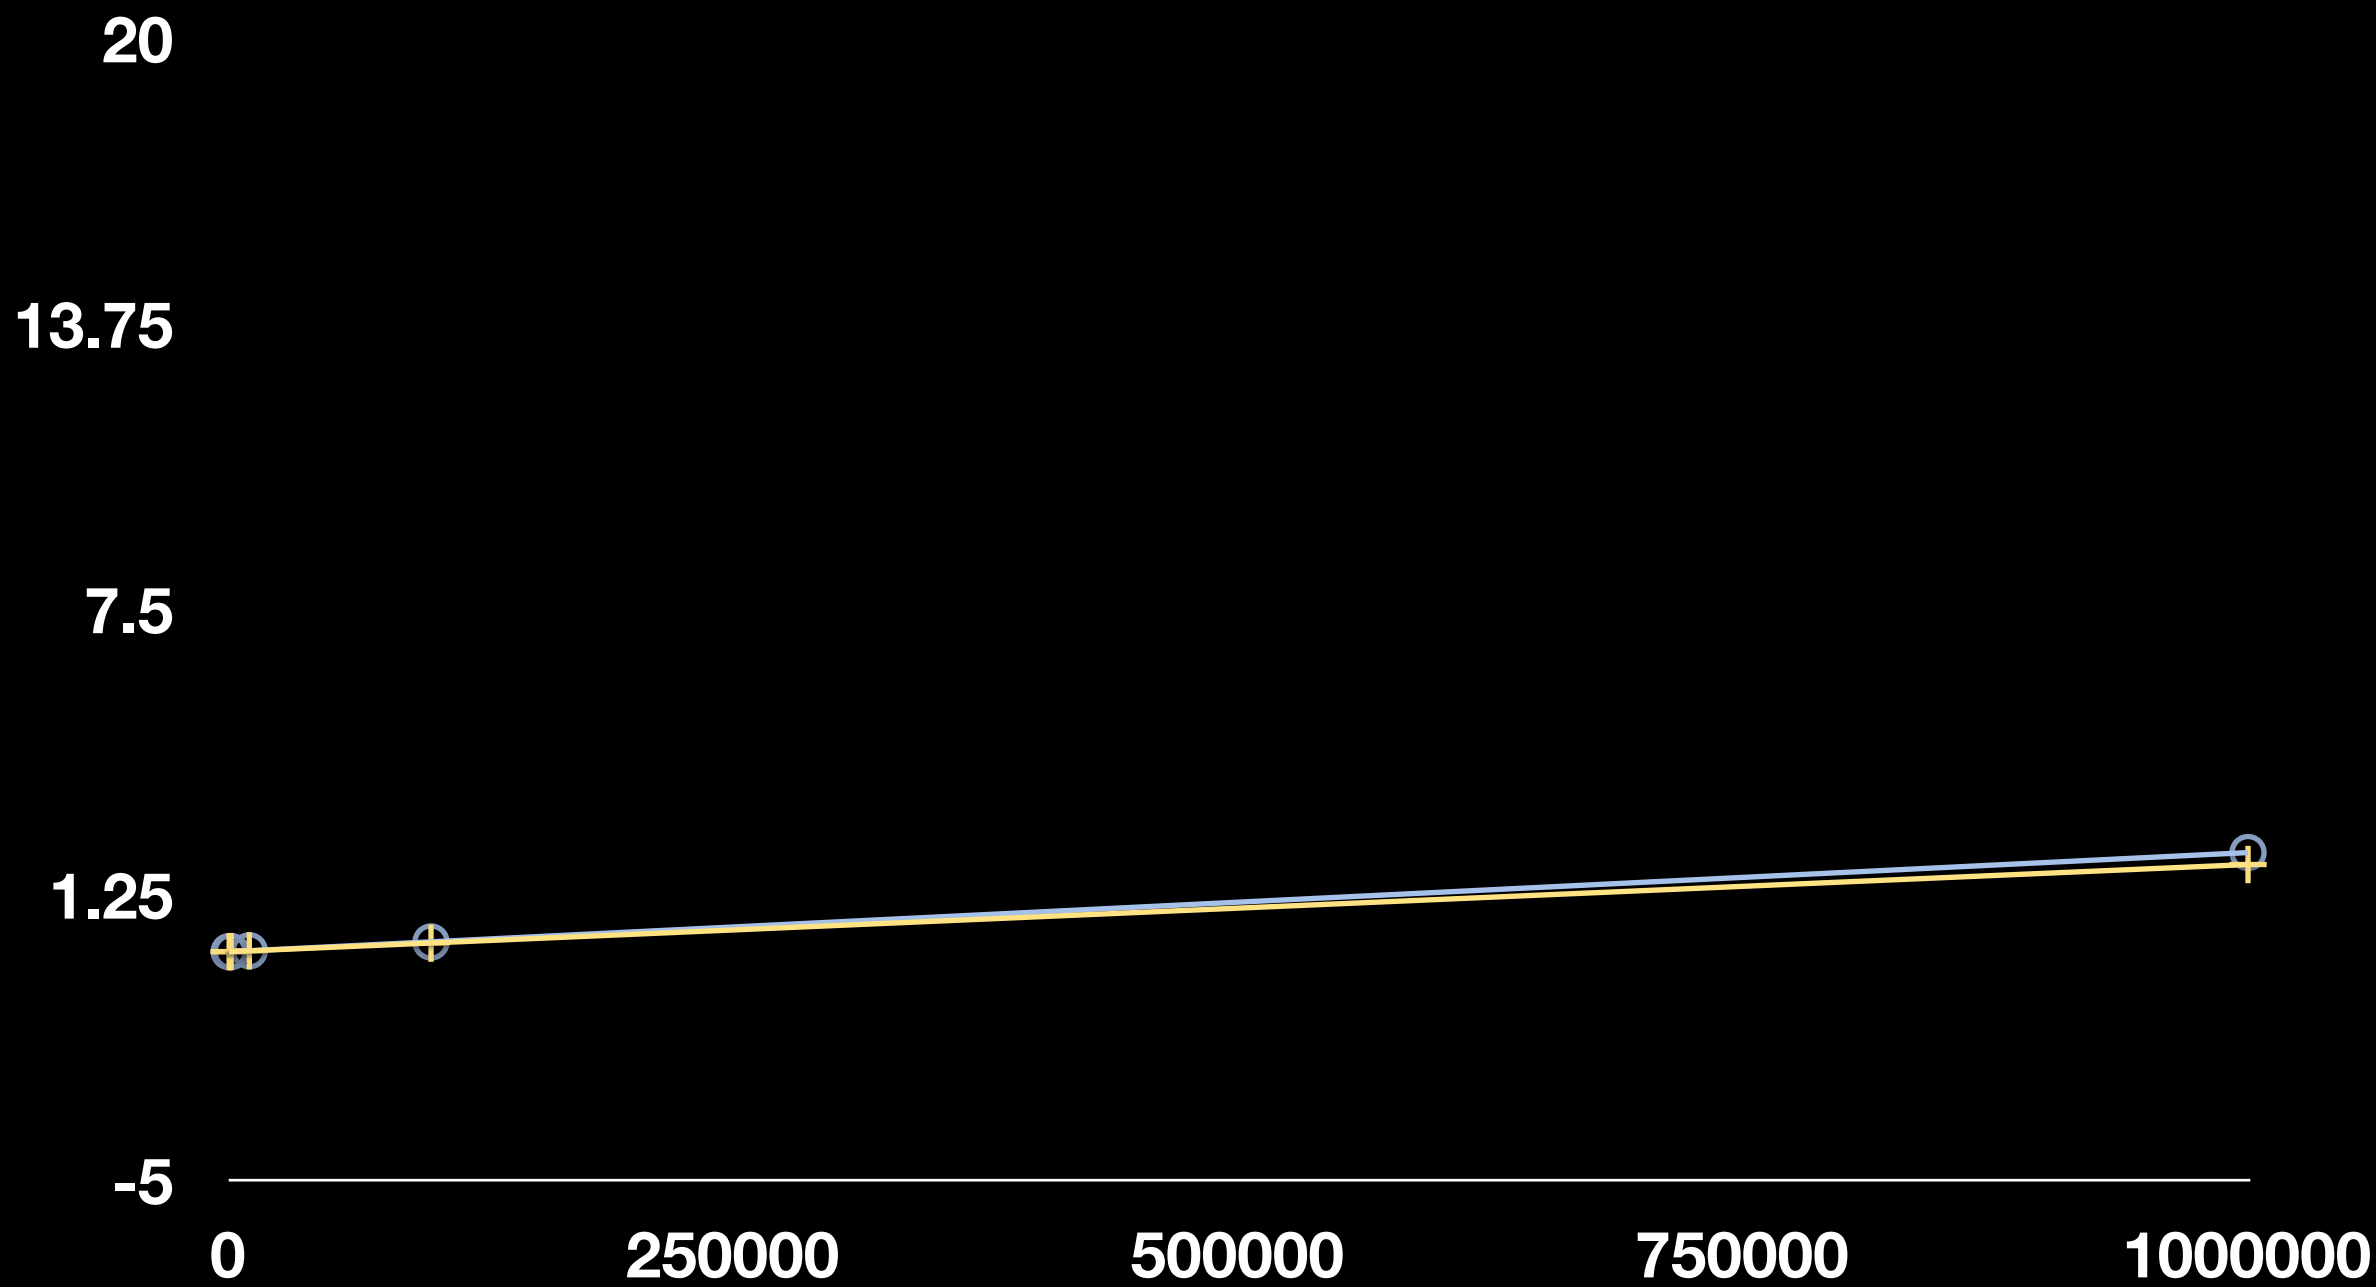
```
# vs
```

```
@list.map { |x| x.to_i }
```



+ symbol to proc

○ block method



Know Your Audience

return value caching

```
def some_method
  @some_method ||= some_expensive_op
end
```

How many times?

**Can the caller
cache?**


**Made our
improvements**

Feeling better!



:-D

+ before (3.0 beta) ○ after (3.0 beta) □ Rails 2.3



What do we do?



We have to go deeper

```
$ git grep 'include Relation' | wc -l  
6
```

```
$ git grep 'def bind' | wc -l  
12
```

```
$
```


**Everything is_a
Relation**

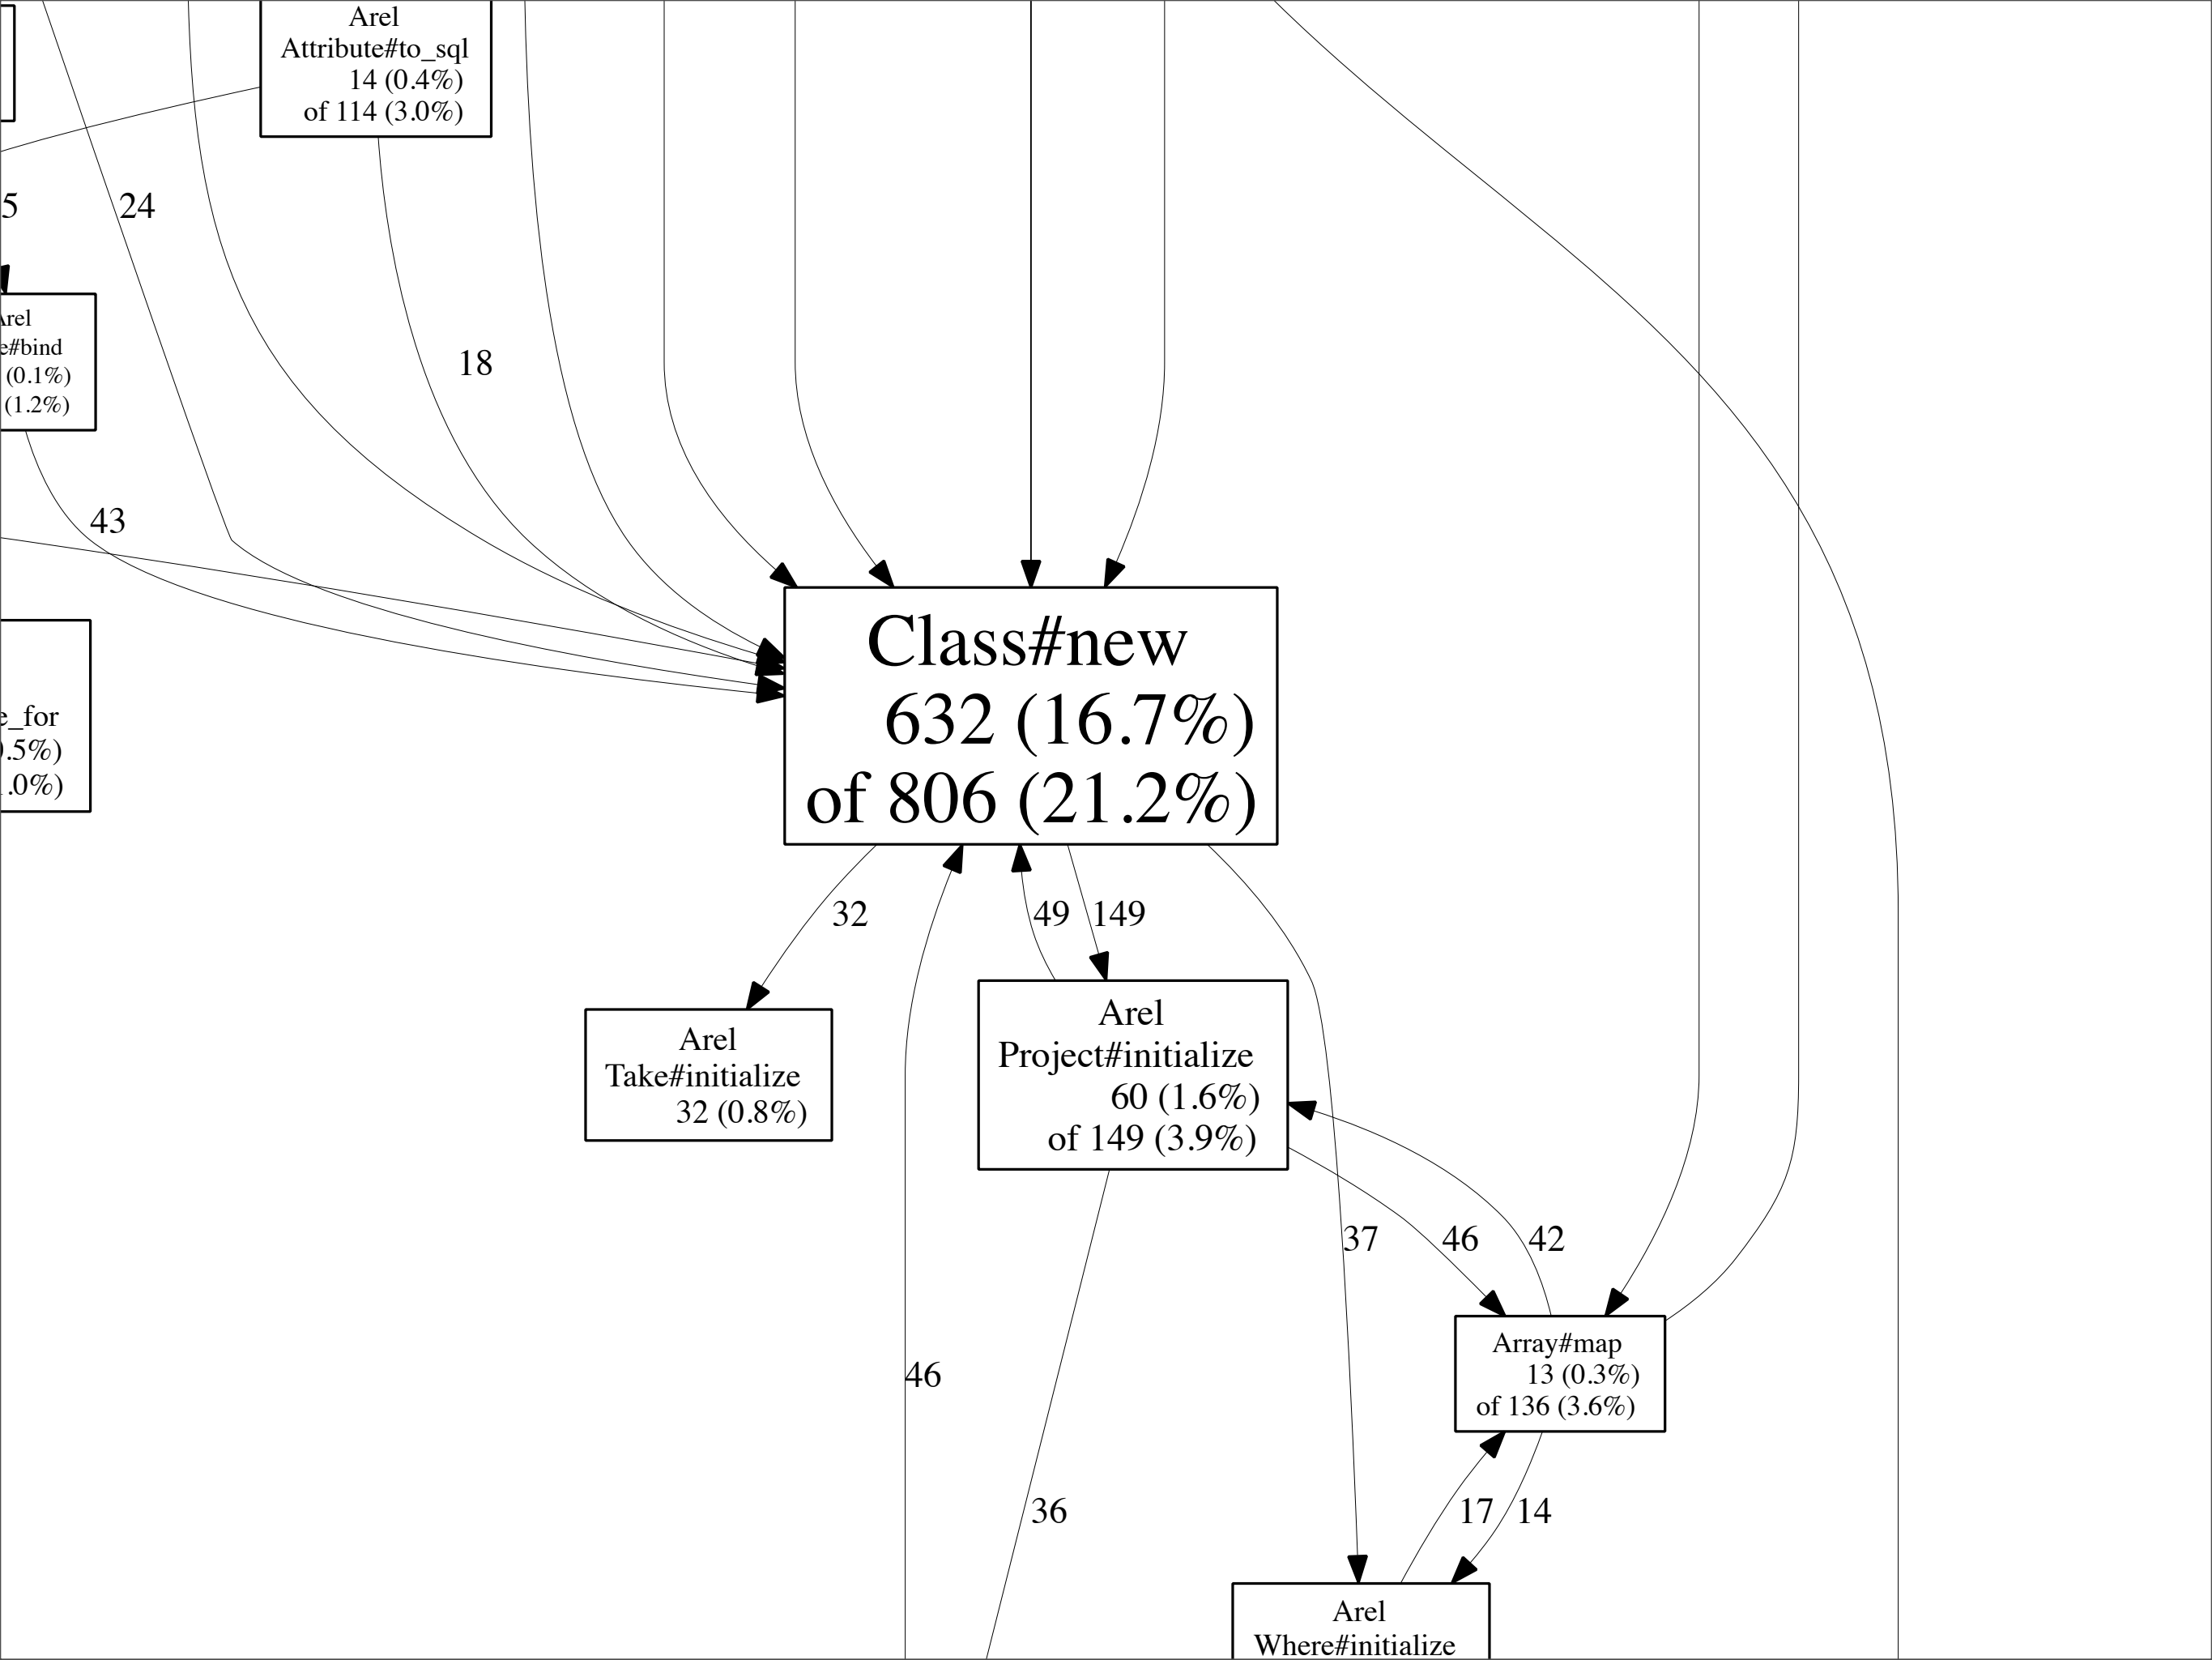
**Everything
Responds to "bind"**

**Everything has a
"relation"**

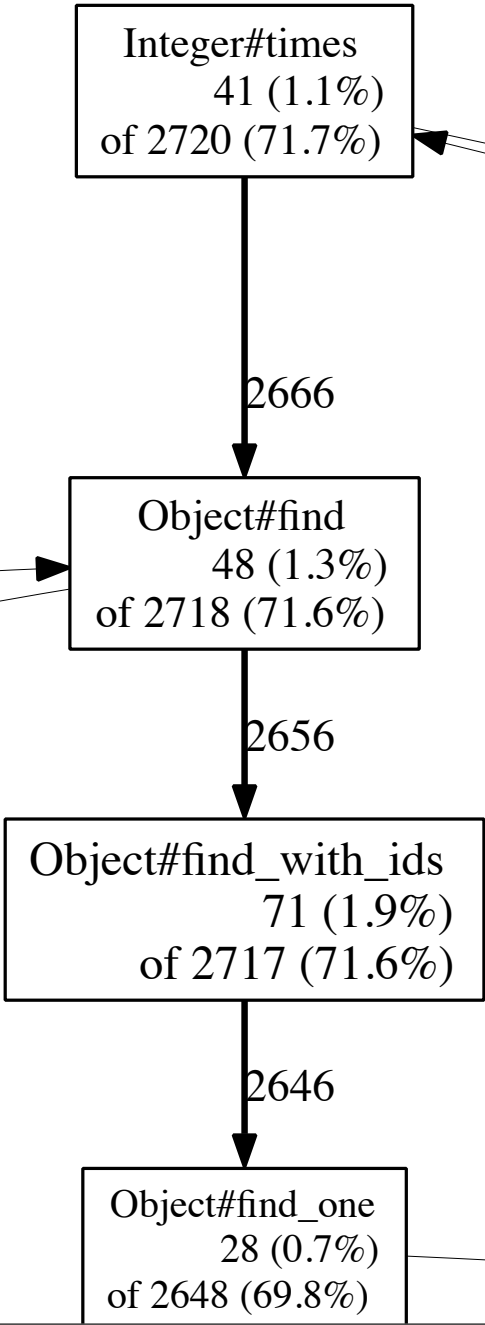
**bind() is recursively
called on relation**



How does it work?



garbage_collector
840 (22.1%)




Data Structure Analysis

Graphviz

graphviz.org

```
digraph "foo" {  
    node [width=0.375,height=0.25];  
    N1 [label="hello"];  
    N2 [label="world"];  
    N1 -> N2;  
}
```



Visitor Pattern

```
class Visitor
  def accept(object)
    method = object.class.name.split('::').join('_')
    send("visit_#{method}", object)
  end
end
```

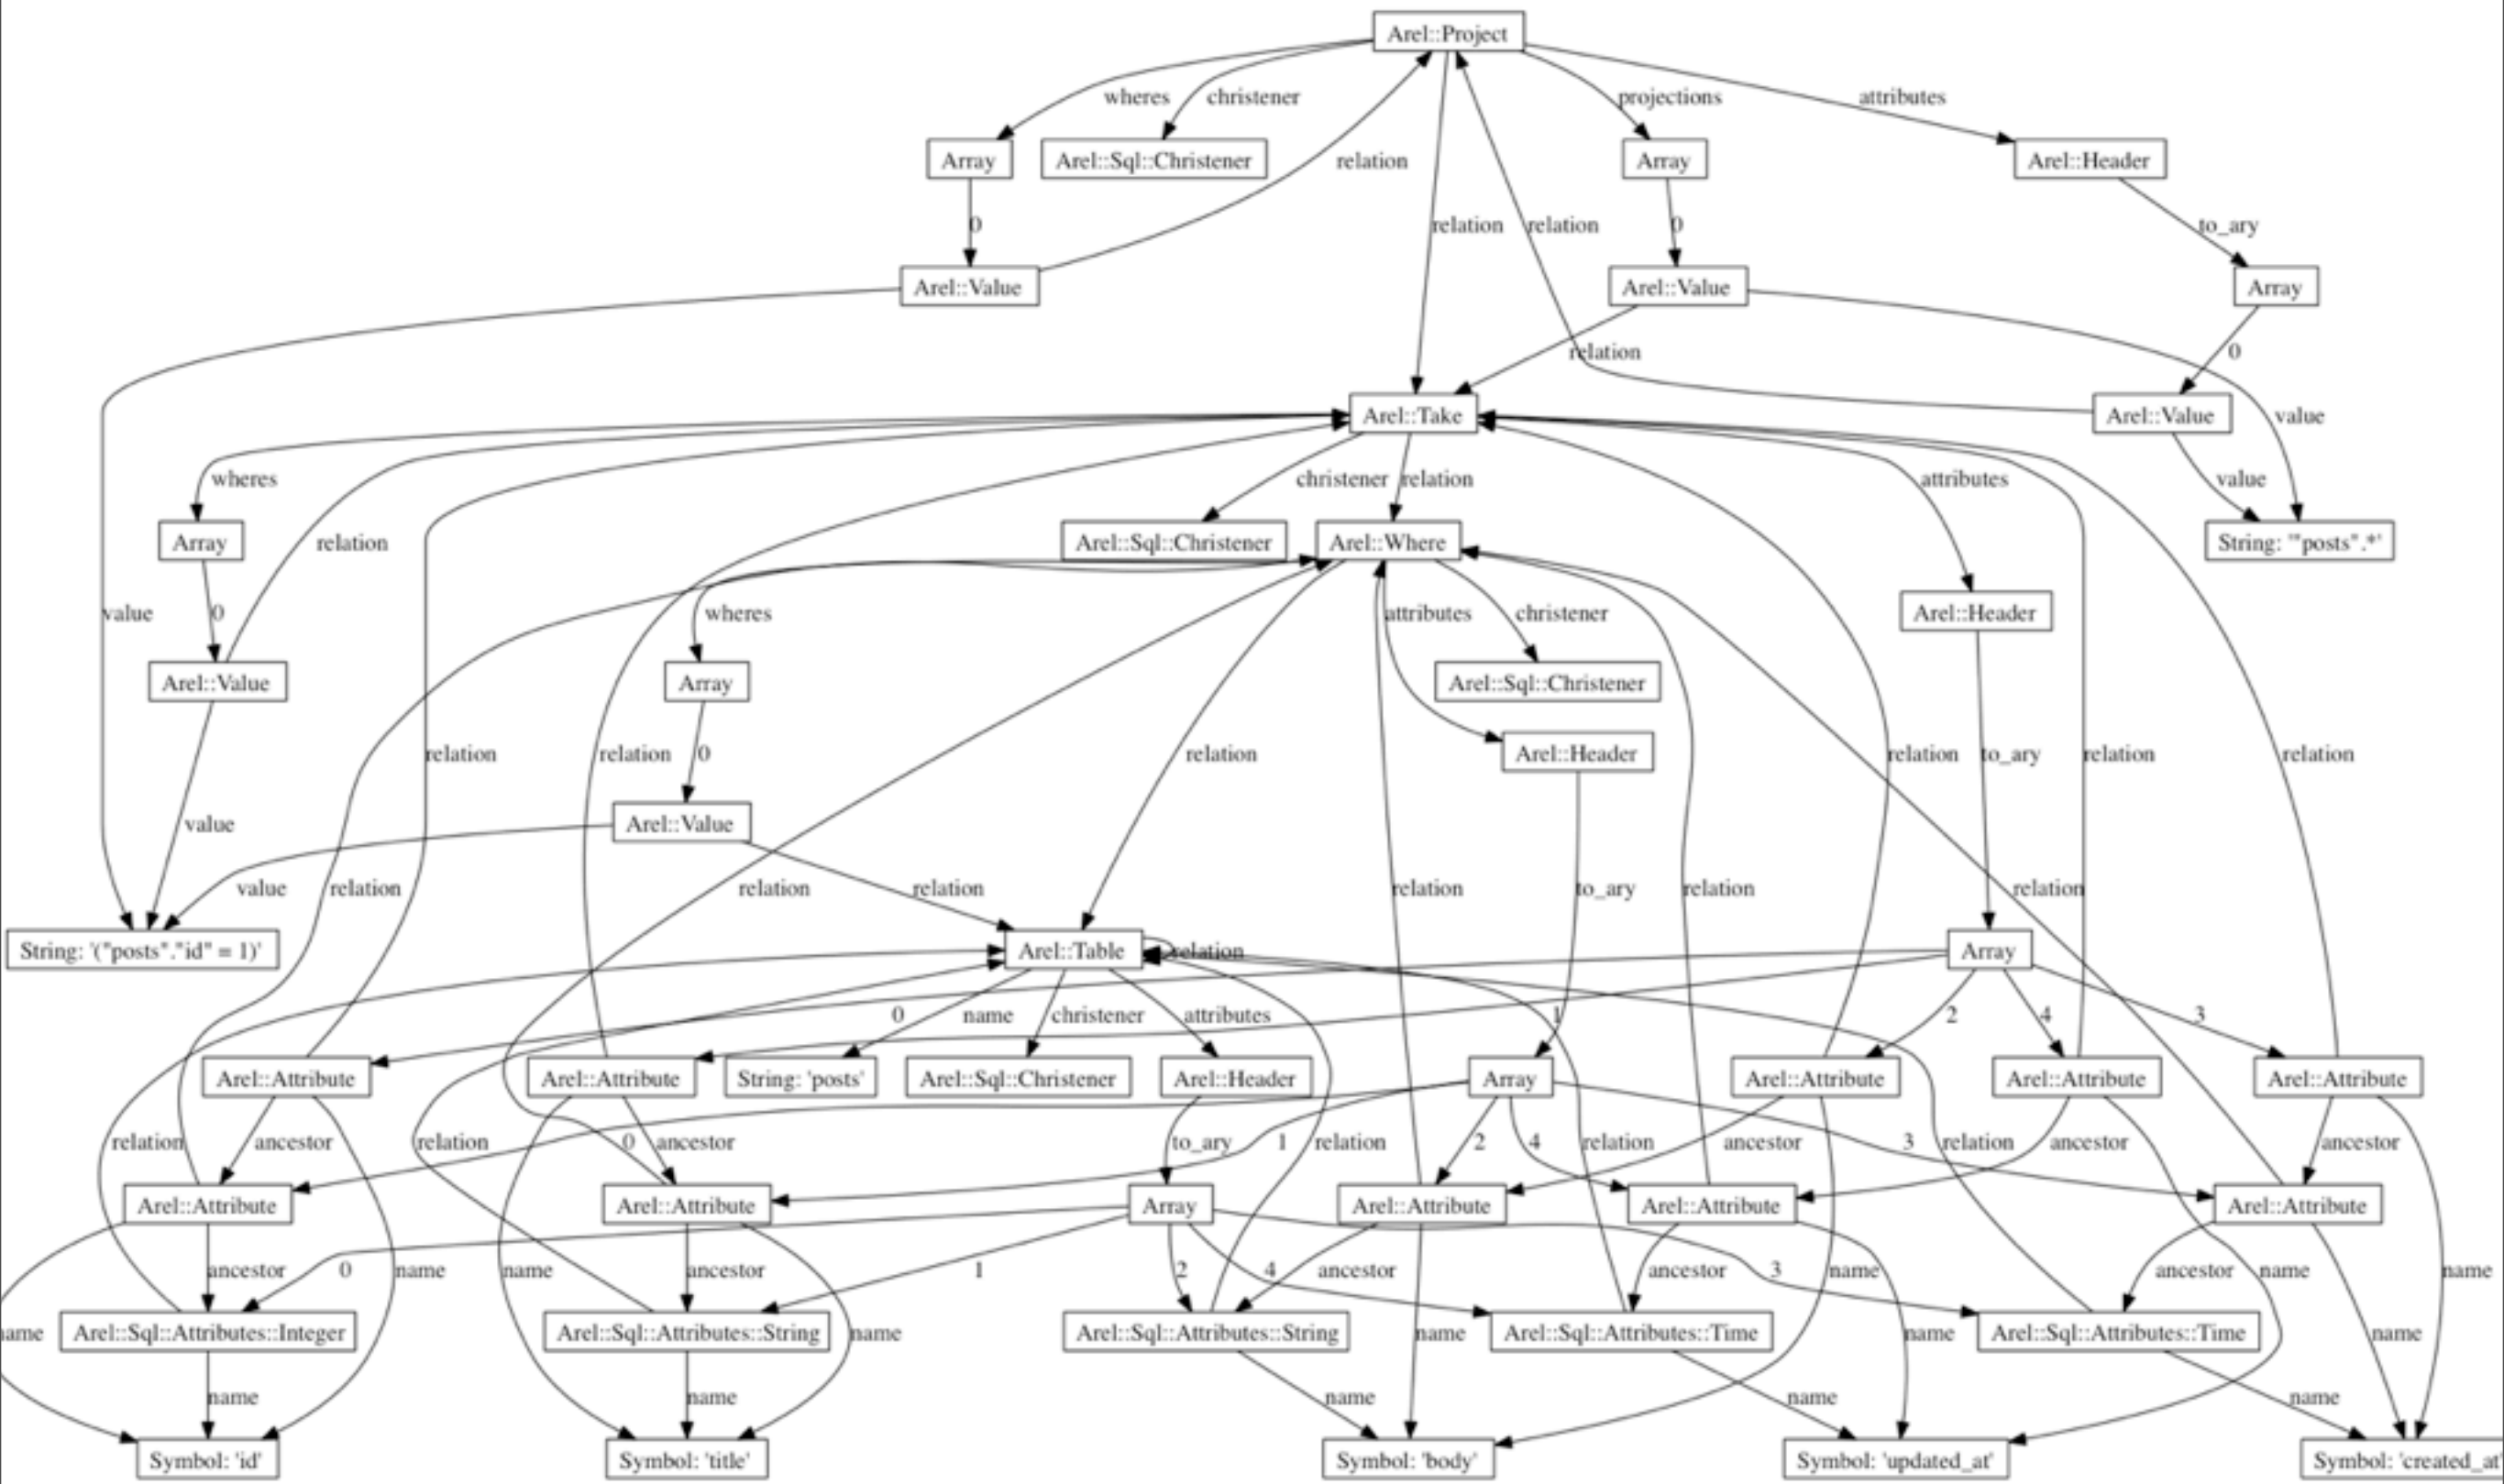
```
class Visitor
  def accept(object)
    method = object.class.name.split('::').join('_')
    send("visit_#{method}", object)
  end

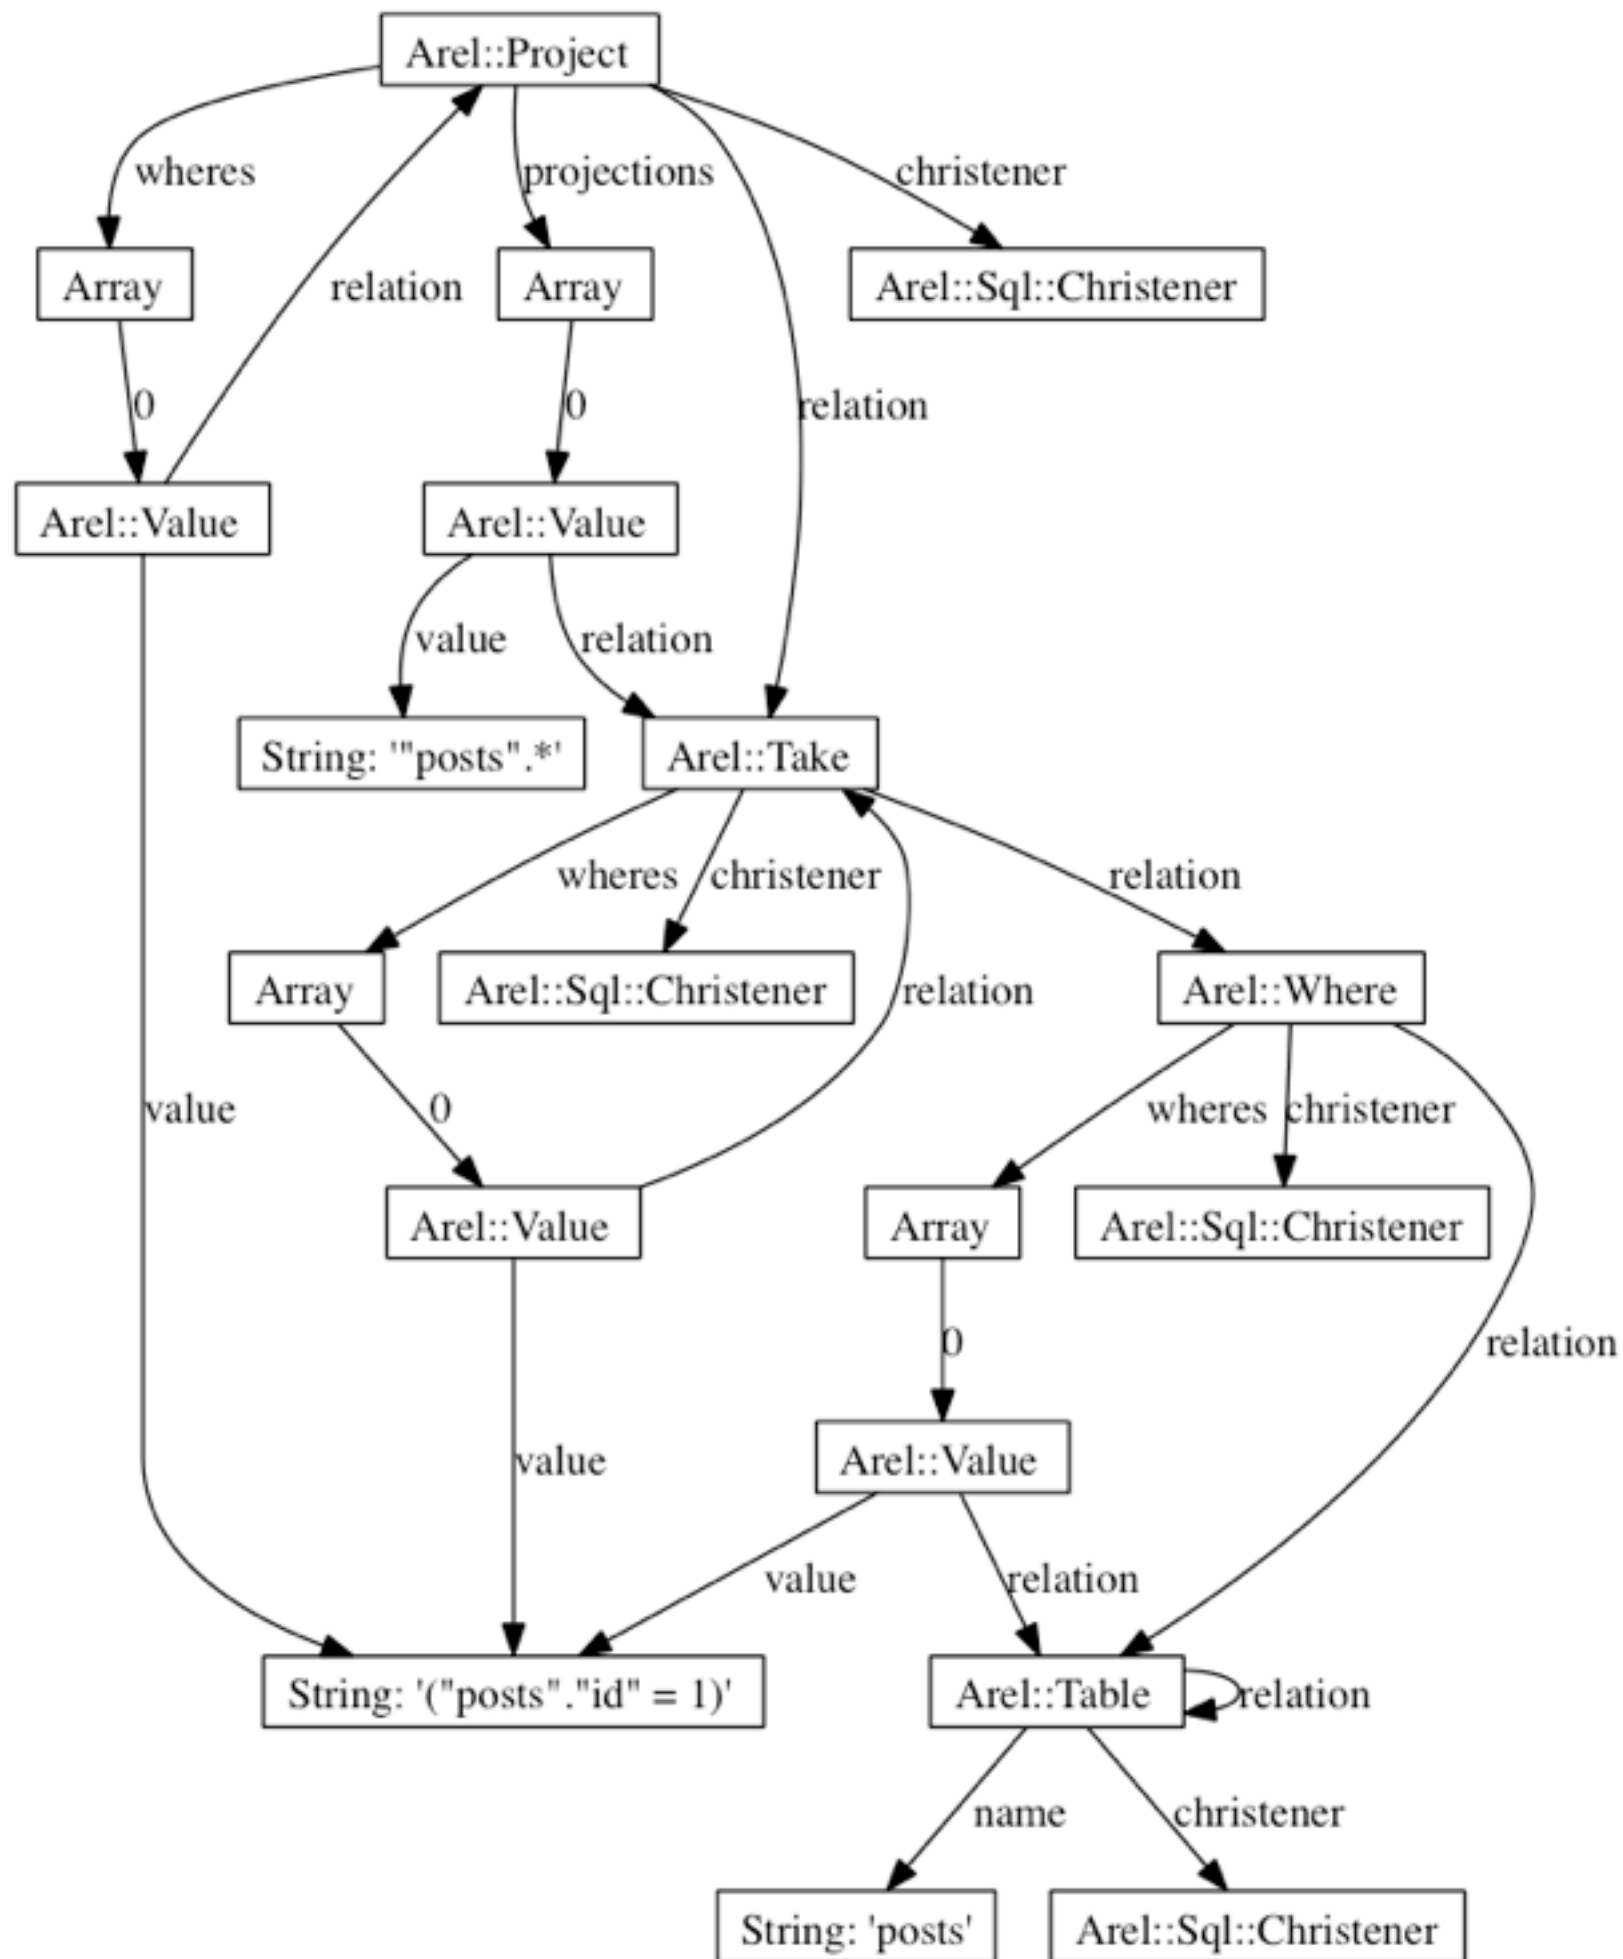
  def visit_Arel_Alias(node)
    # keep track of the node called
    accept(node.attribute)
  end
end
```

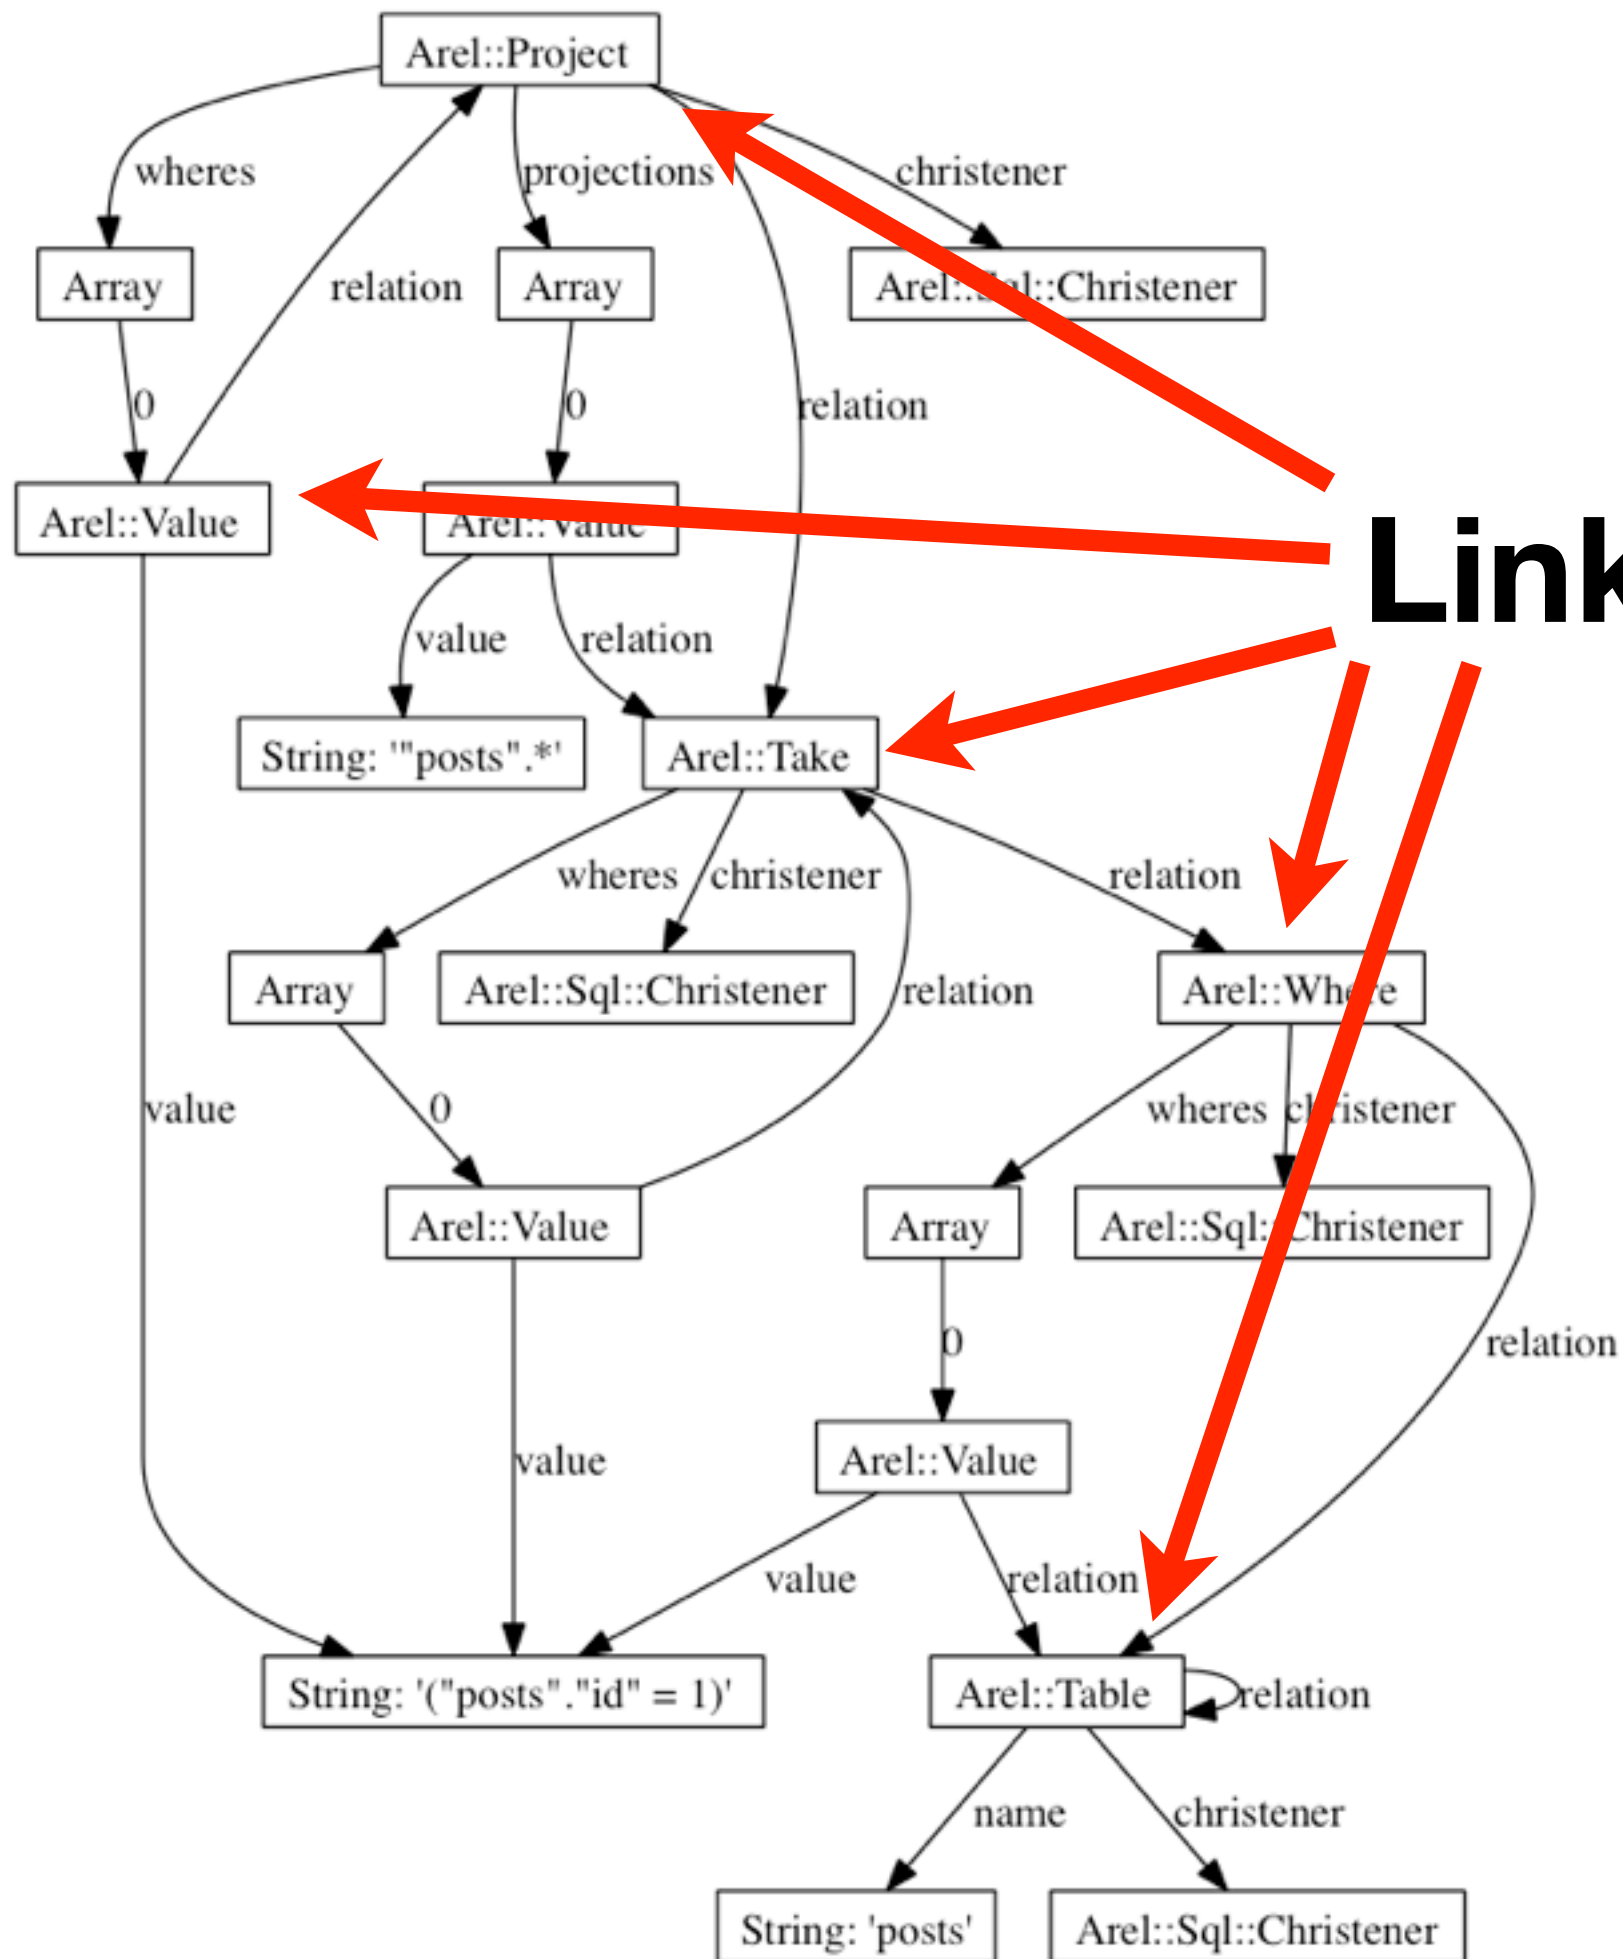
```
class Visitor
  def accept(object)
    method = object.class.name.split('::').join('_')
    send("visit_#{method}", object)
  end

  def visit_Arel_Alias(node)
    # keep track of the node called
    accept(node.attribute)
  end


  def visit_Arel_Table(node)
    # keep track of the node called
    accept(node.name)
    node.columns.each { |c| accept(c) }
  end
end
```

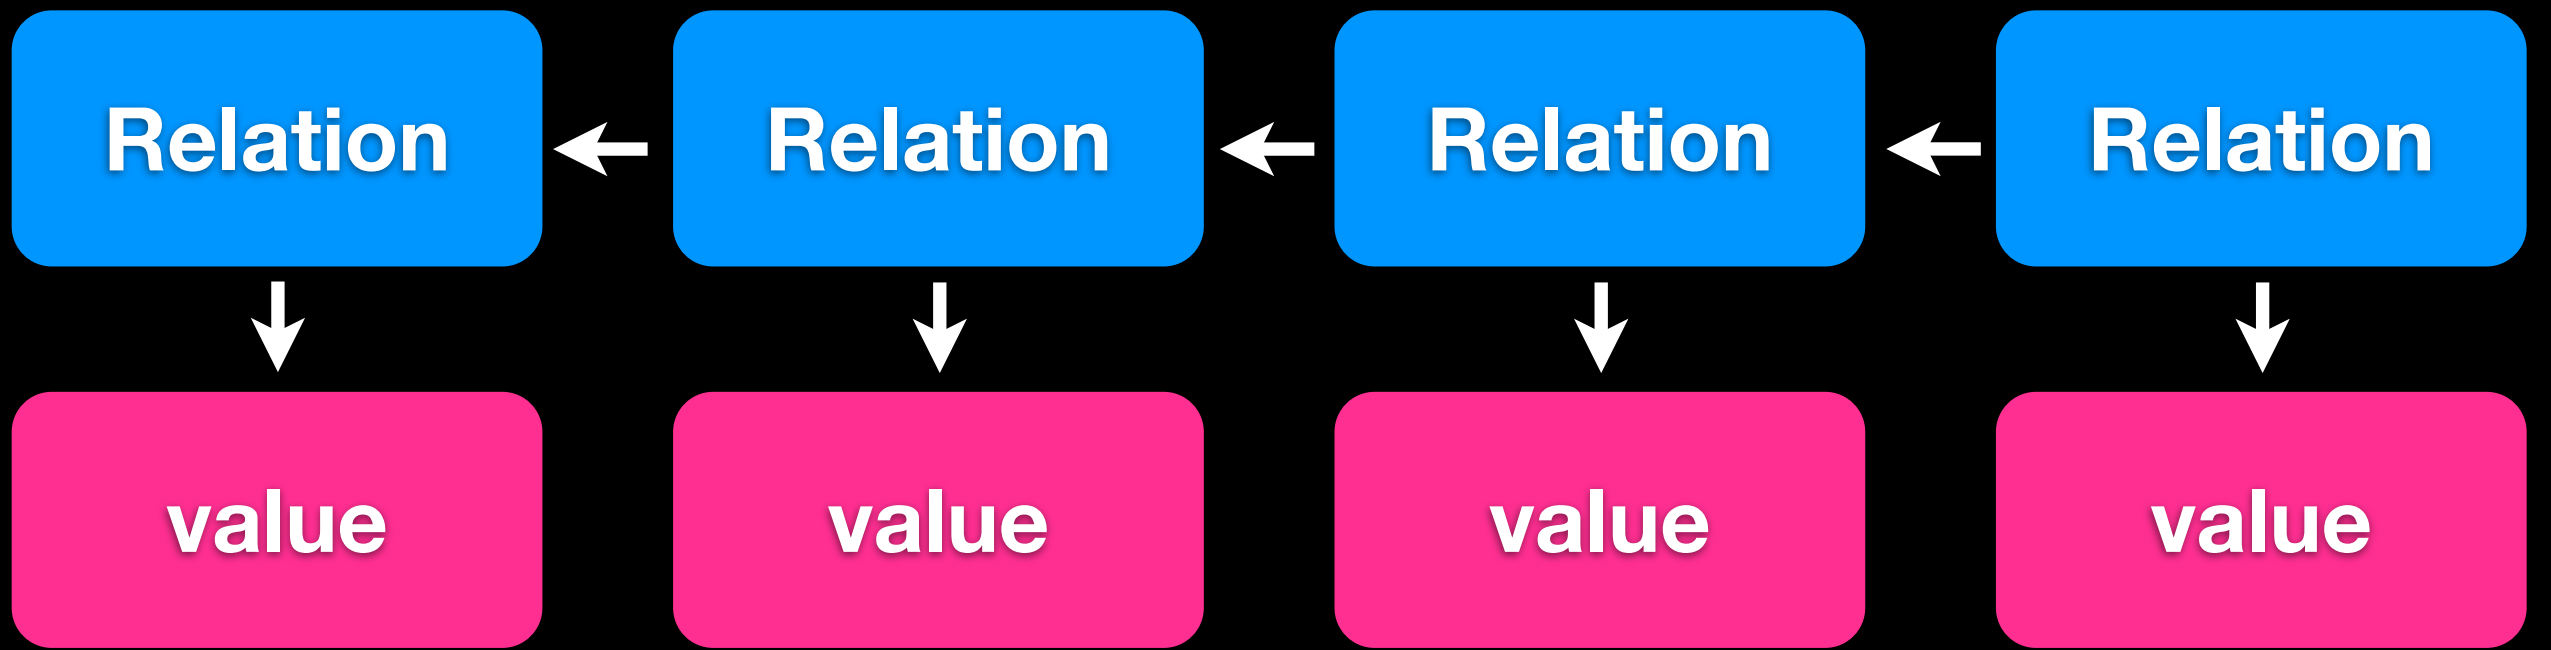






Linked List





Relation



**Bind() +
Dup**

Relation



**Bind() +
Dup**

Relation

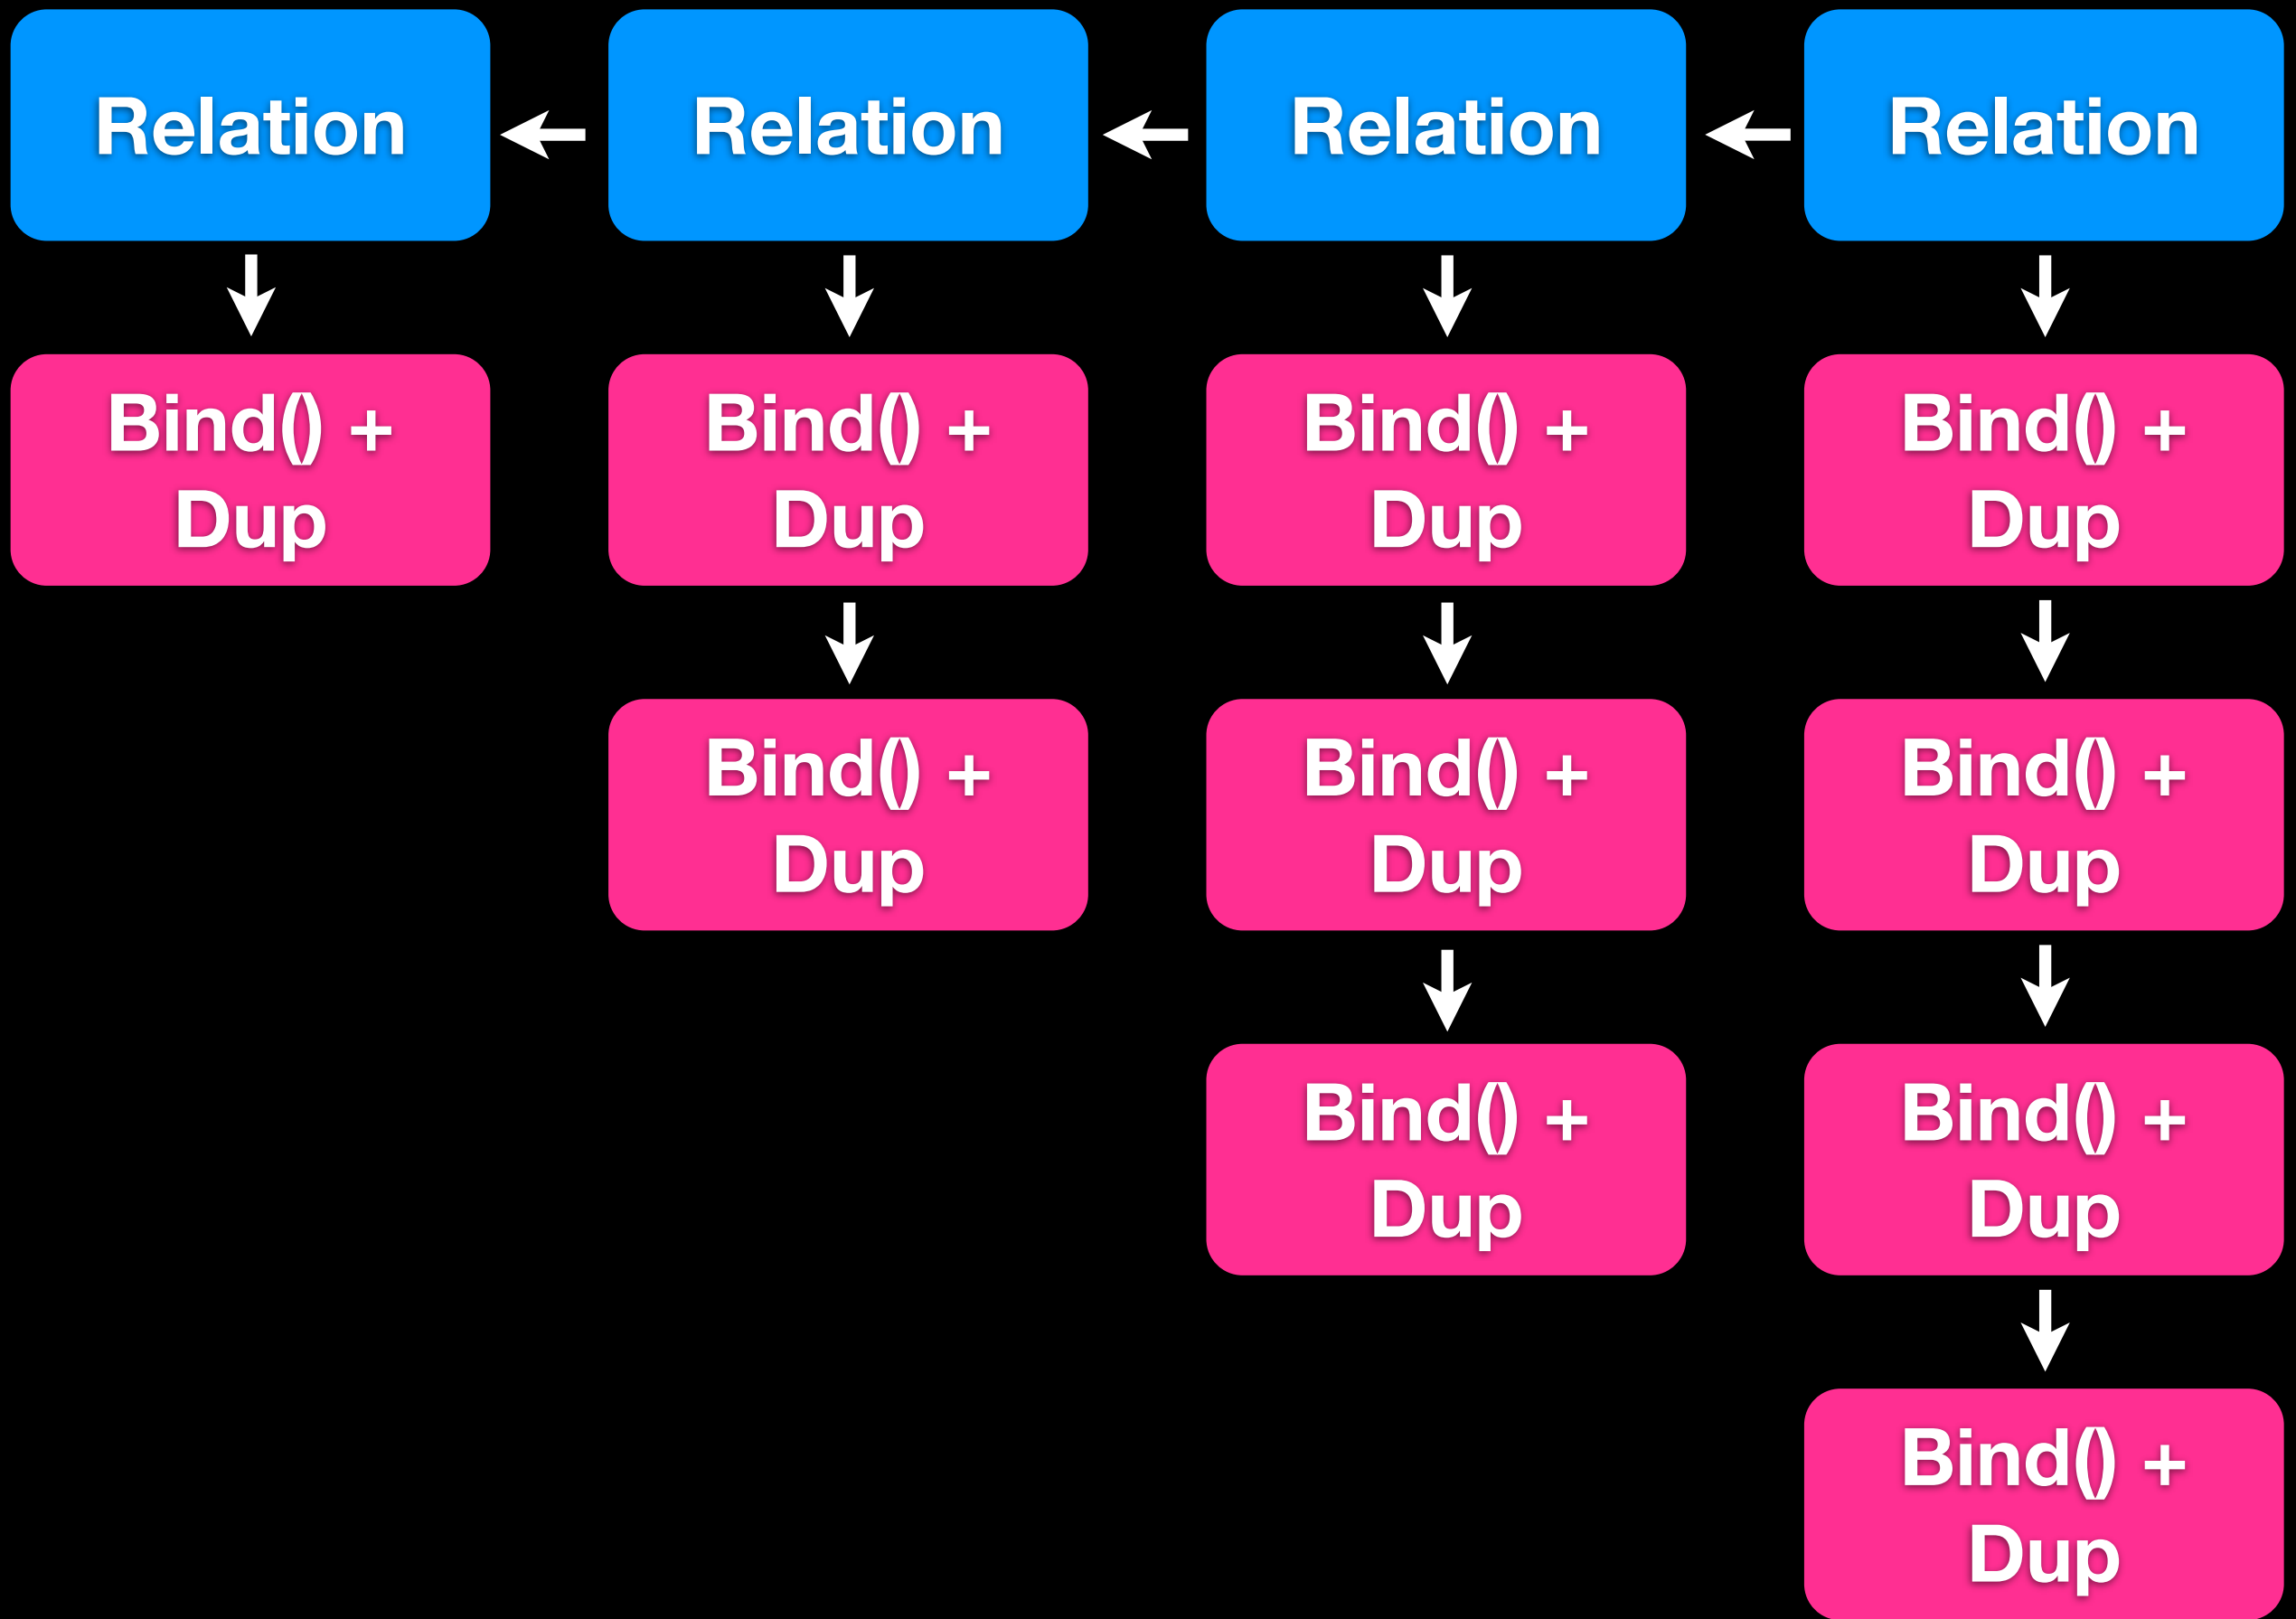


**Bind() +
Dup**

Relation



**Bind() +
Dup**



Where



$$1 = 1$$

Where



$$2 = 2$$

Where

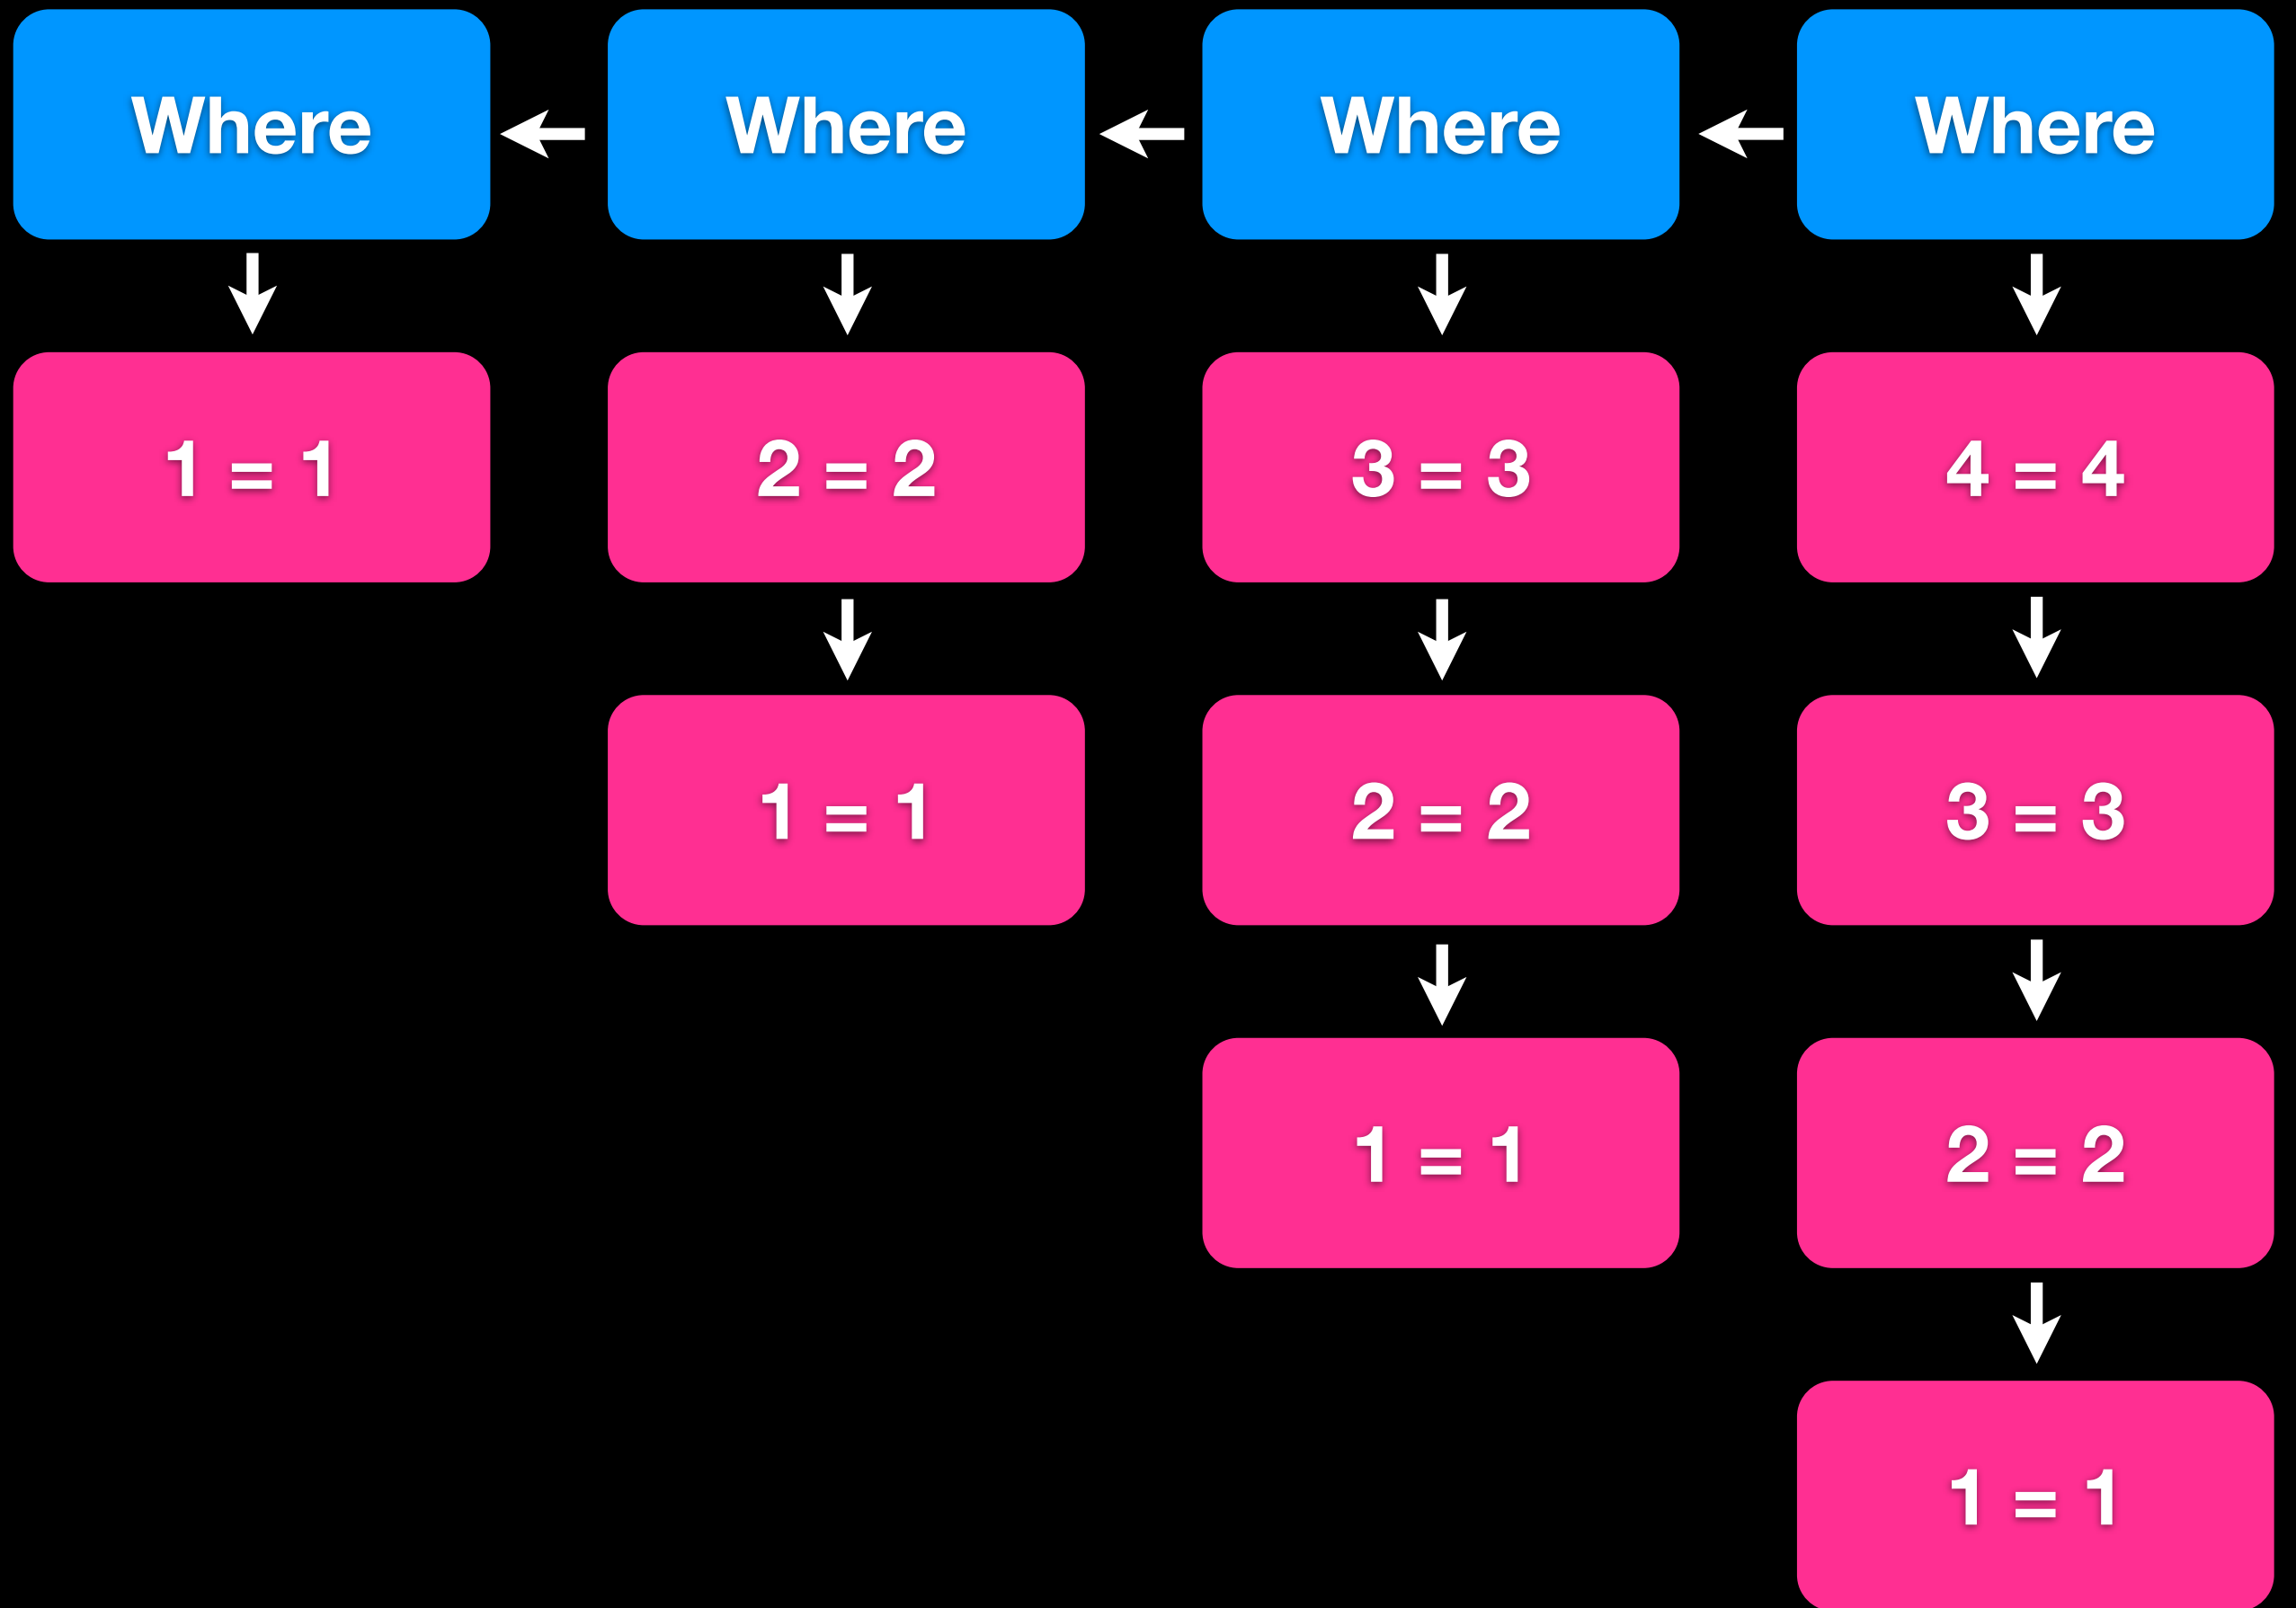


$$3 = 3$$

Where



$$4 = 4$$



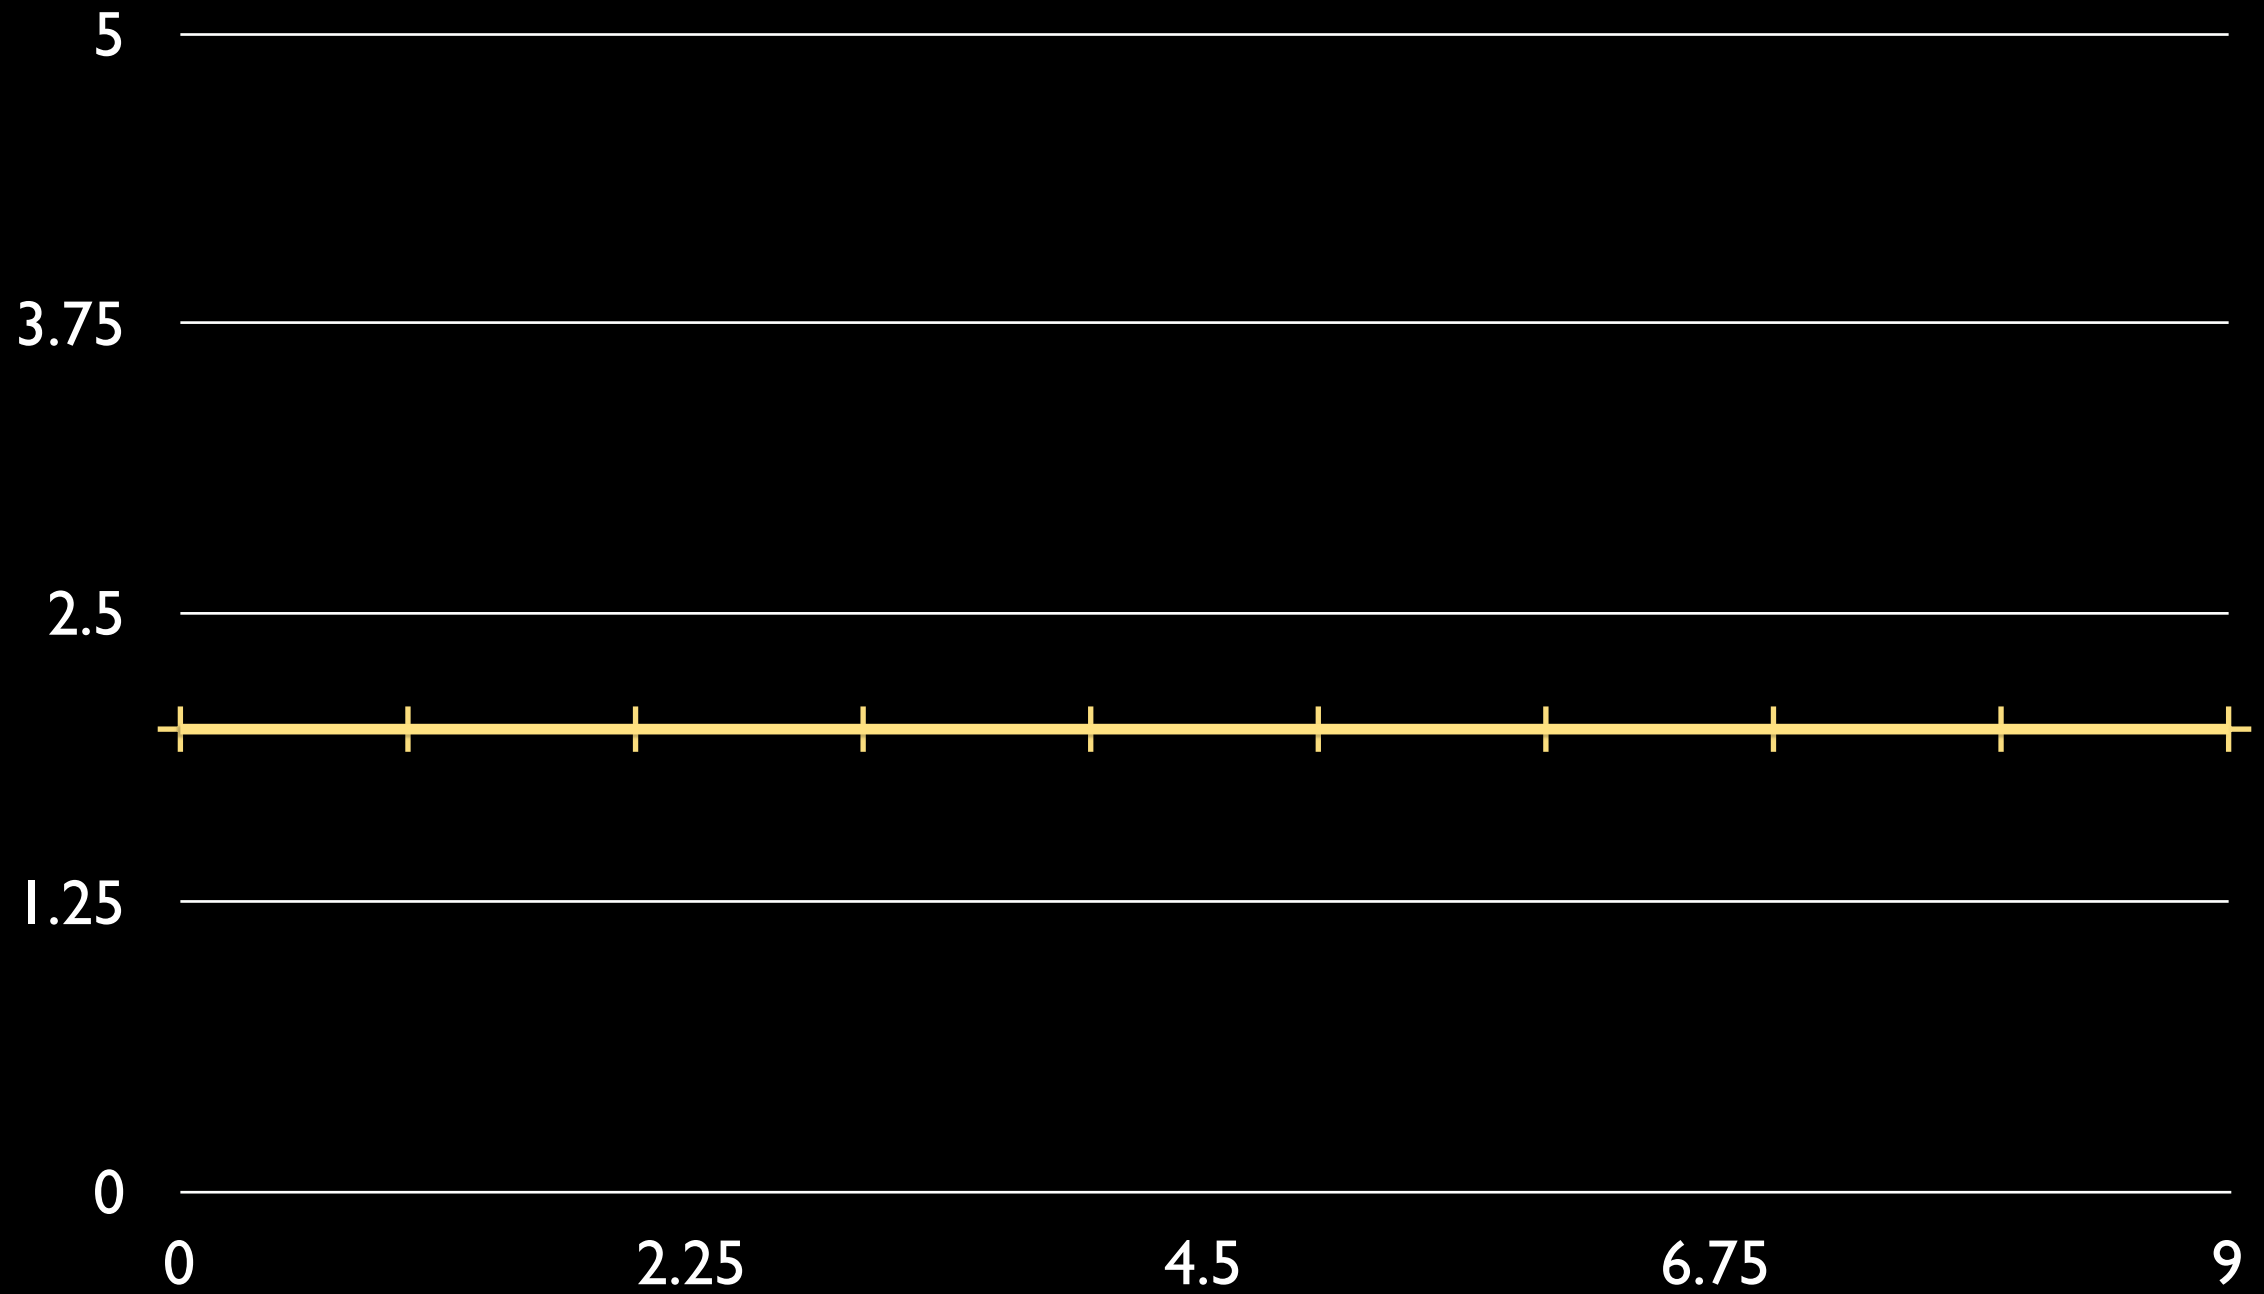
Big O!

OMG


ZOMG

Mathematical Representation


$$y = 2$$



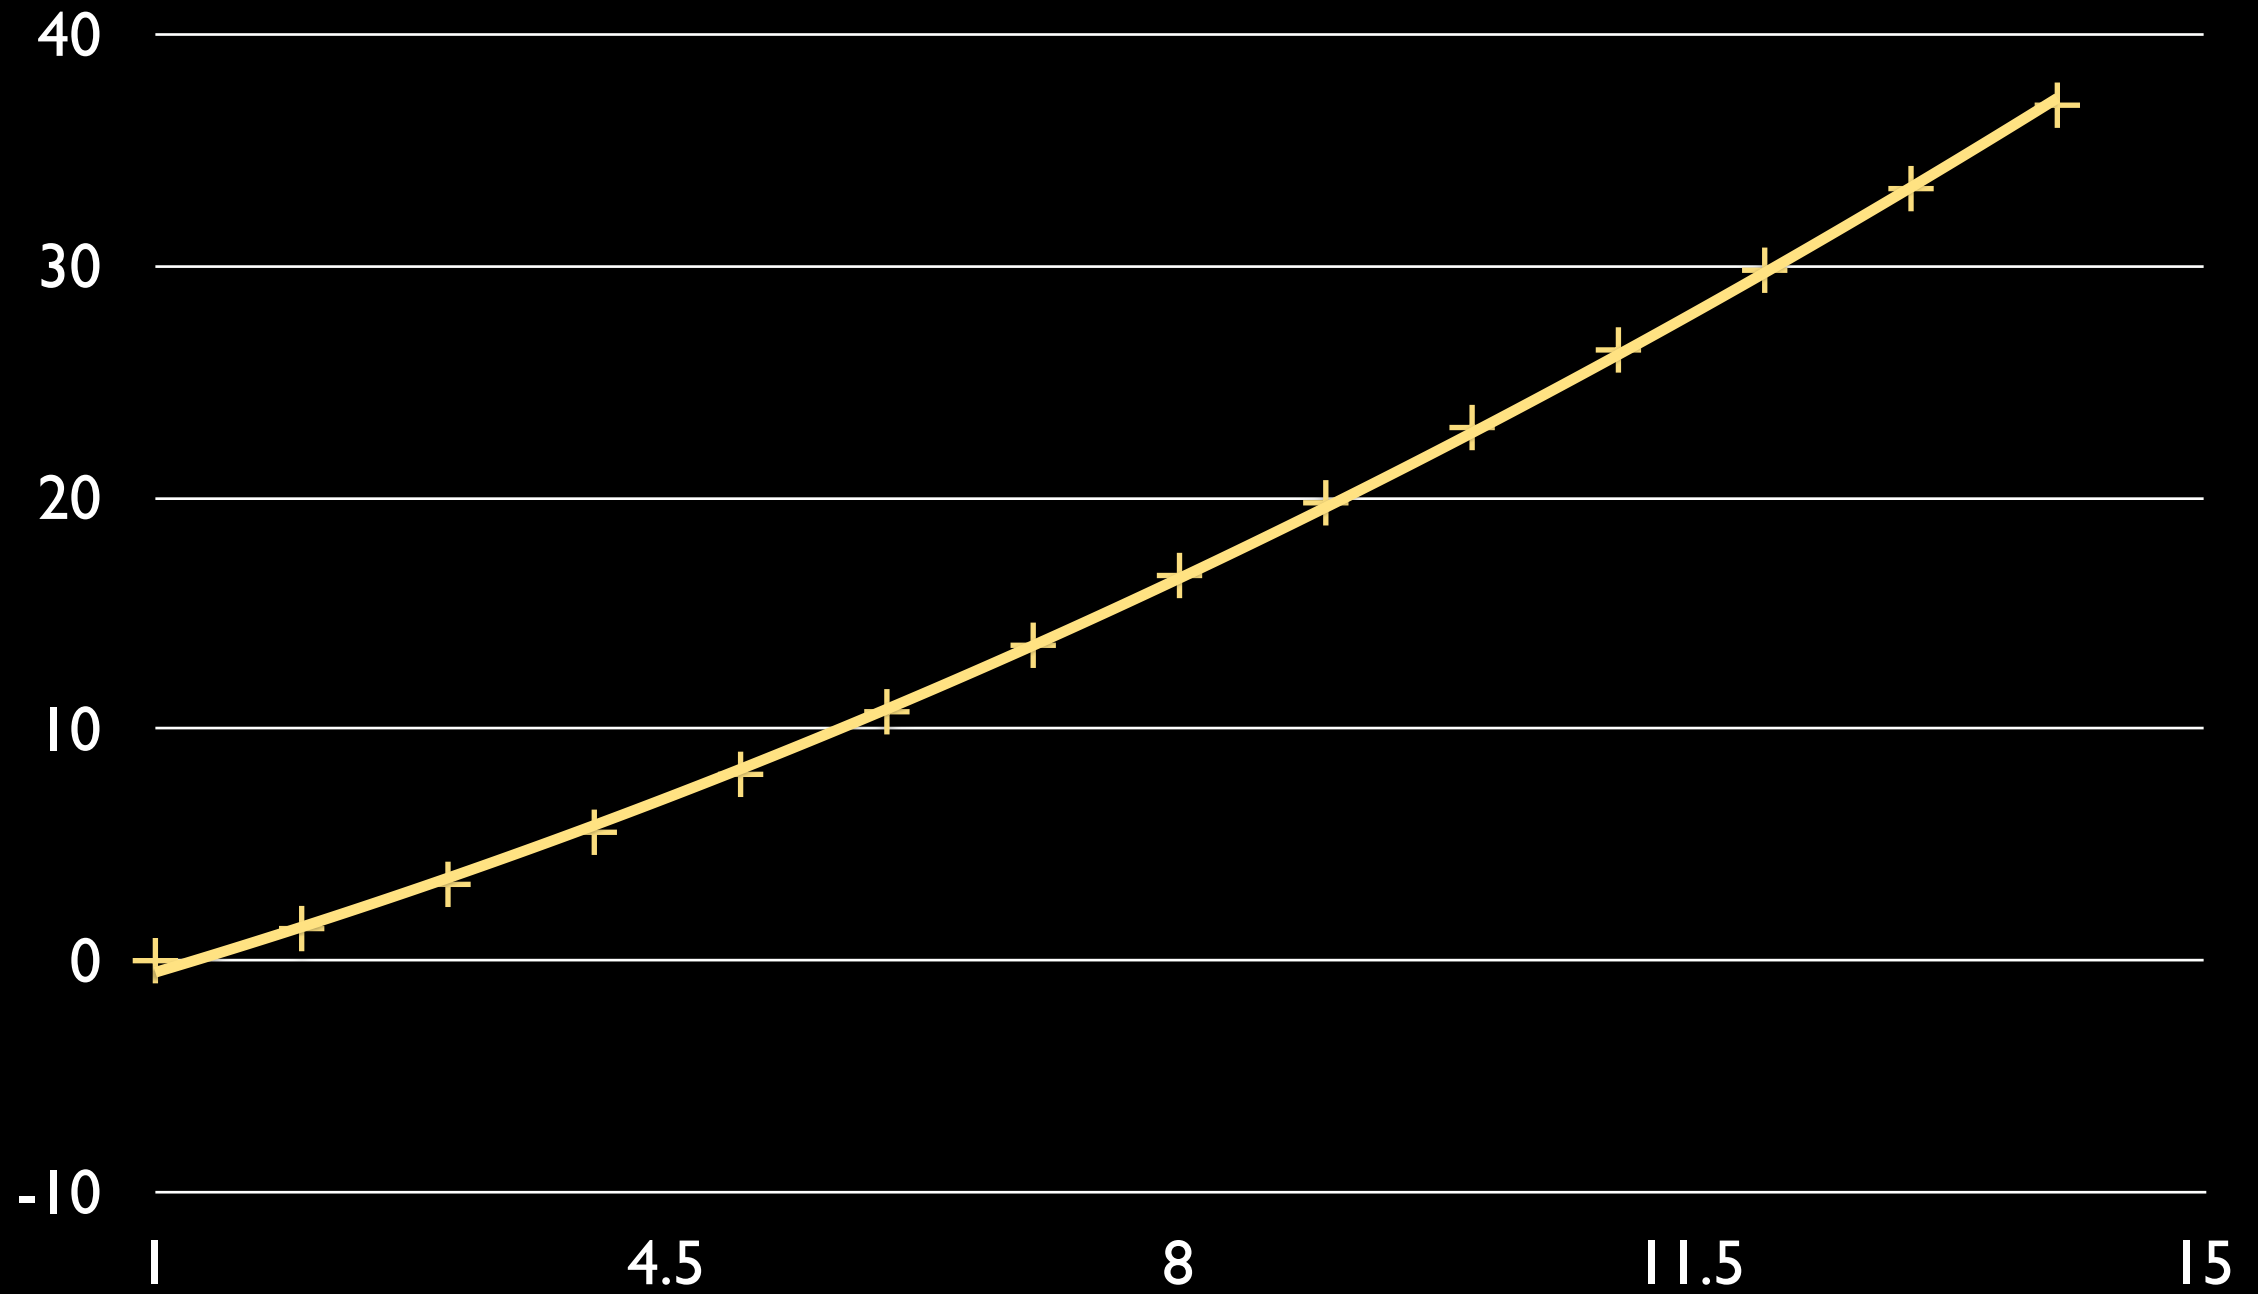
$y = \log(n)$




$$y = x$$



$$y = n \log(n)$$



$$y = n^2$$




Finding Big O

- **Give input**
- **Measure output**
- **Plot Results**


ARel's Big O

+ objects and funcalls




ARel's Big O

+ objects and funcalls




ARel's Big O

+ objects and funcalls




ARel's Big O

+ objects and funcalls



ARel's Big O



$$y = \frac{1}{2} (n^2 + n)$$


$O(n^2)$

**ActiveRecord/Arel
takes over 2
minutes to generate
a pseudo-complex
SQL query.**

<http://bit.ly/omgslow2>

Deep Improvements

System Impact



AST + Visitor

$O(n)$

Should I rewrite?

- **Clear solution**
- **Tests are numerous (Rails)**
- **Public API is limited**

YES

6 Weeks Later...

ARel Today

Data Sheet

- **$O(n)$**
- **6 Weeks to Rewrite**
- **2x faster (for simple queries)**
- **Adapter Specific code is DRY**

flog (before)

2533.1: flog total

6.8: flog/method average

116.6: OracleCompiler#select_sql

78.0: PostgreSQLCompiler#select_sql

64.9: main#none

59.4: GenericCompiler#insert_sql

52.4: Join#joins

flog (after)

```
1864.6: flog total
  6.5: flog/method average

81.4: main#none
75.9: Dot#none
59.1: Oracle# lib/arel/visitors/oracle.rb:6
54.0: ToSql#lib/arel/visitors/to_sql.rb:31
51.6: ToSql#none
```

flay (before)

Total score (lower is better) = 684

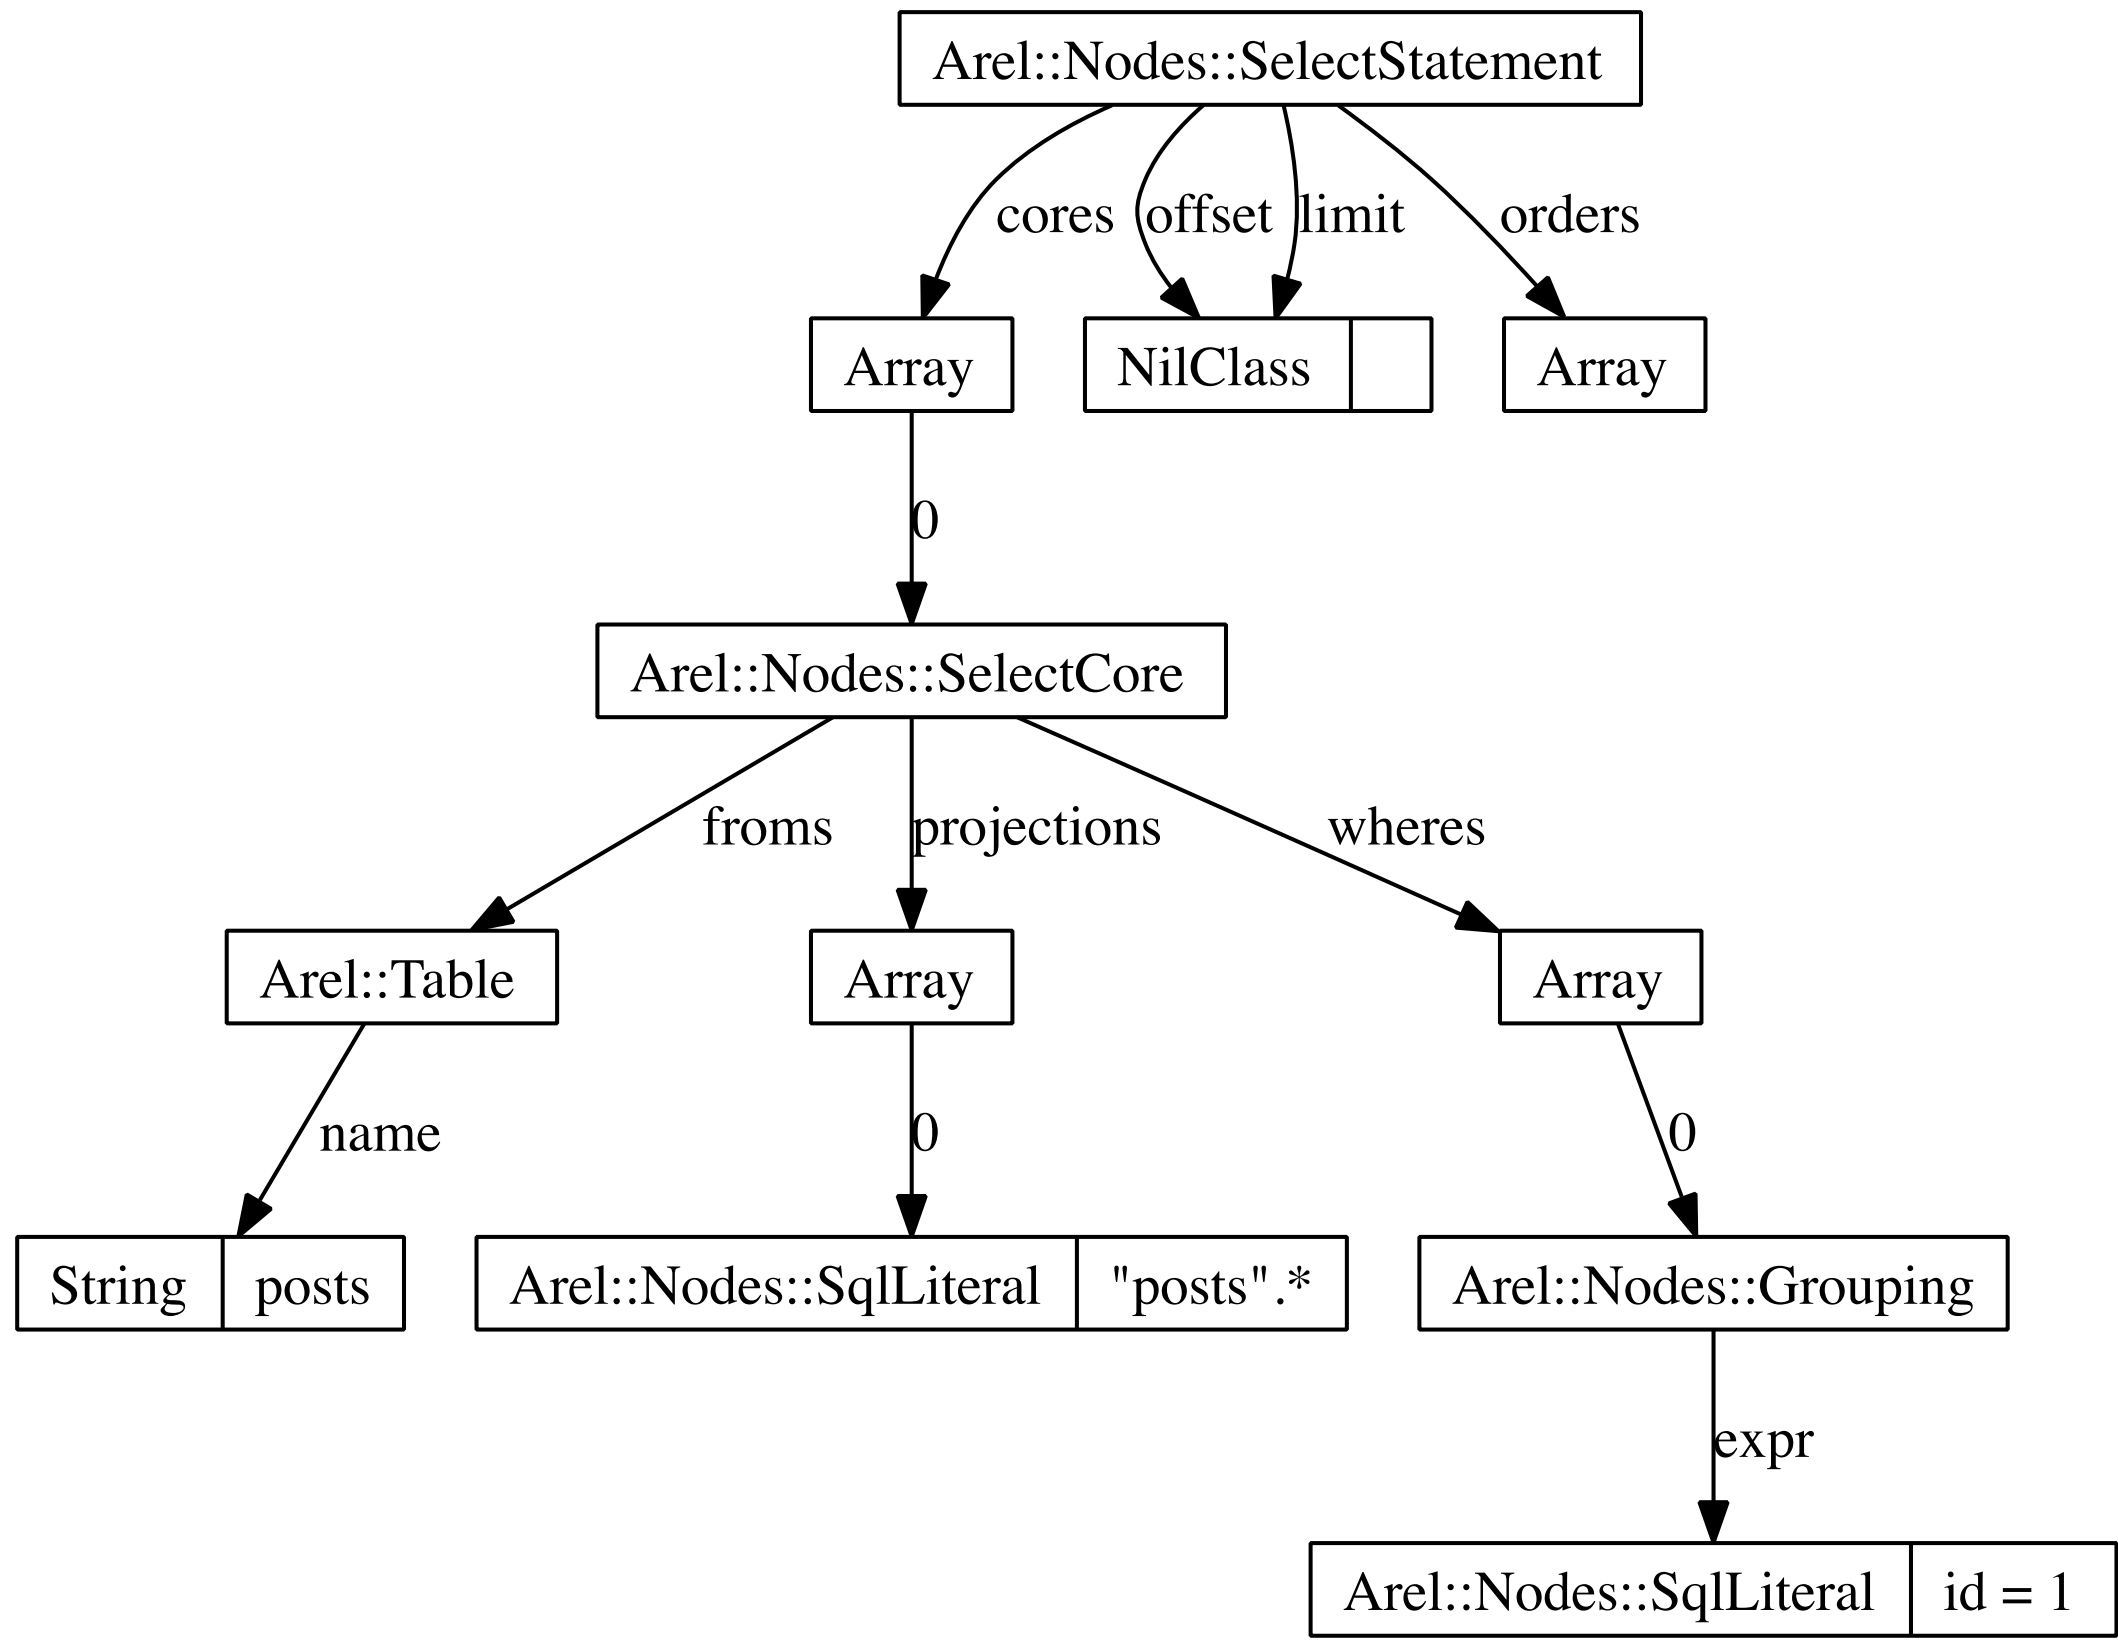
12 complaints

flay (after)

Total score (lower is better) = 420

7 complaints

Post.where("id = 1").to_dot




AReI Tomorrow

Conclusion

AKA:

Things I've Learned

System Impact



When Should I Rewrite?

Rewrite Timeline

Rewrite Timeline



We emphasize
the art
of Code

We should not forget
the Science

Learn
The Specific

But Embrace
The Generic

Photo Credits

- **DHH: <http://www.flickr.com/photos/pdcawley/250813158/>**
- **Matz: <http://www.flickr.com/photos/kakutani/4127354831/>**
- **Chad Fowler: <http://www.flickr.com/photos/fraserspeirs/3386558579/>**

Thanks @ebiltwin

Thank you!


One More Thing...

It's **Ruby**Conf 10!

Give Ryan a **Kiss!**

<3





Questions?