

Fibers Are the Right Solution

Samuel Williams

samuel@ruby-lang.org www.codeotaku.com

@ioquatix



Fibers Are the Right Solution

Samuel Williams


samuel@ruby-lang.org www.codeotaku.com


@ioquatix



What is Scalability?

スケーラビリティとは何か？





Proportional Improvement

比例的なパフォーマンス改善



Why is scalability important?

なぜスケールビリティは大切なのか？



Photo from NASA.

Is Ruby scalable?

Rubyはスケーラブルか?




Over a million web
sites globally.


Rubyが使われた世界中のWebサイトについて



Is this scalable?



Processing 500 requests/min



Processing 500 requests/min

S3

Postgres

WebSocket

Redis

SMTP

Blocking

MySQL

ブロッキング

DNS

Disk

HTTP

How do we maximise hardware utilisation?

どのようにしてハードウェア最大限活用できるか?

Late 1950s

1950年代後半








1
SELECT
READY
LOAD REWIND
START
CHANGE BASKET
UNLOAD
STOP

IBM 729 II
MAGNETIC TAPE UNIT

2000 IBM Punch Cards
2000 IBM Punch Cards

FIRE EXTINGUISHER

IBM
EZ Head II
30 101
140041 Air Filter



**CPU Instruction
1 nanosecond**

CPU Instruction
1 nanosecond

CPU Cache
10 nanosecond

CPU Instruction
1 nanosecond

CPU Cache
10 nanosecond

Main Memory
100 nanosecond


Solid State Disk
1 millisecond

Solid State Disk
1 millisecond

Network Packet
10 millisecond


**Can we avoid being
idle?**

アイドル状態を避けることは可能か?




Time Sharing



Time Slicing

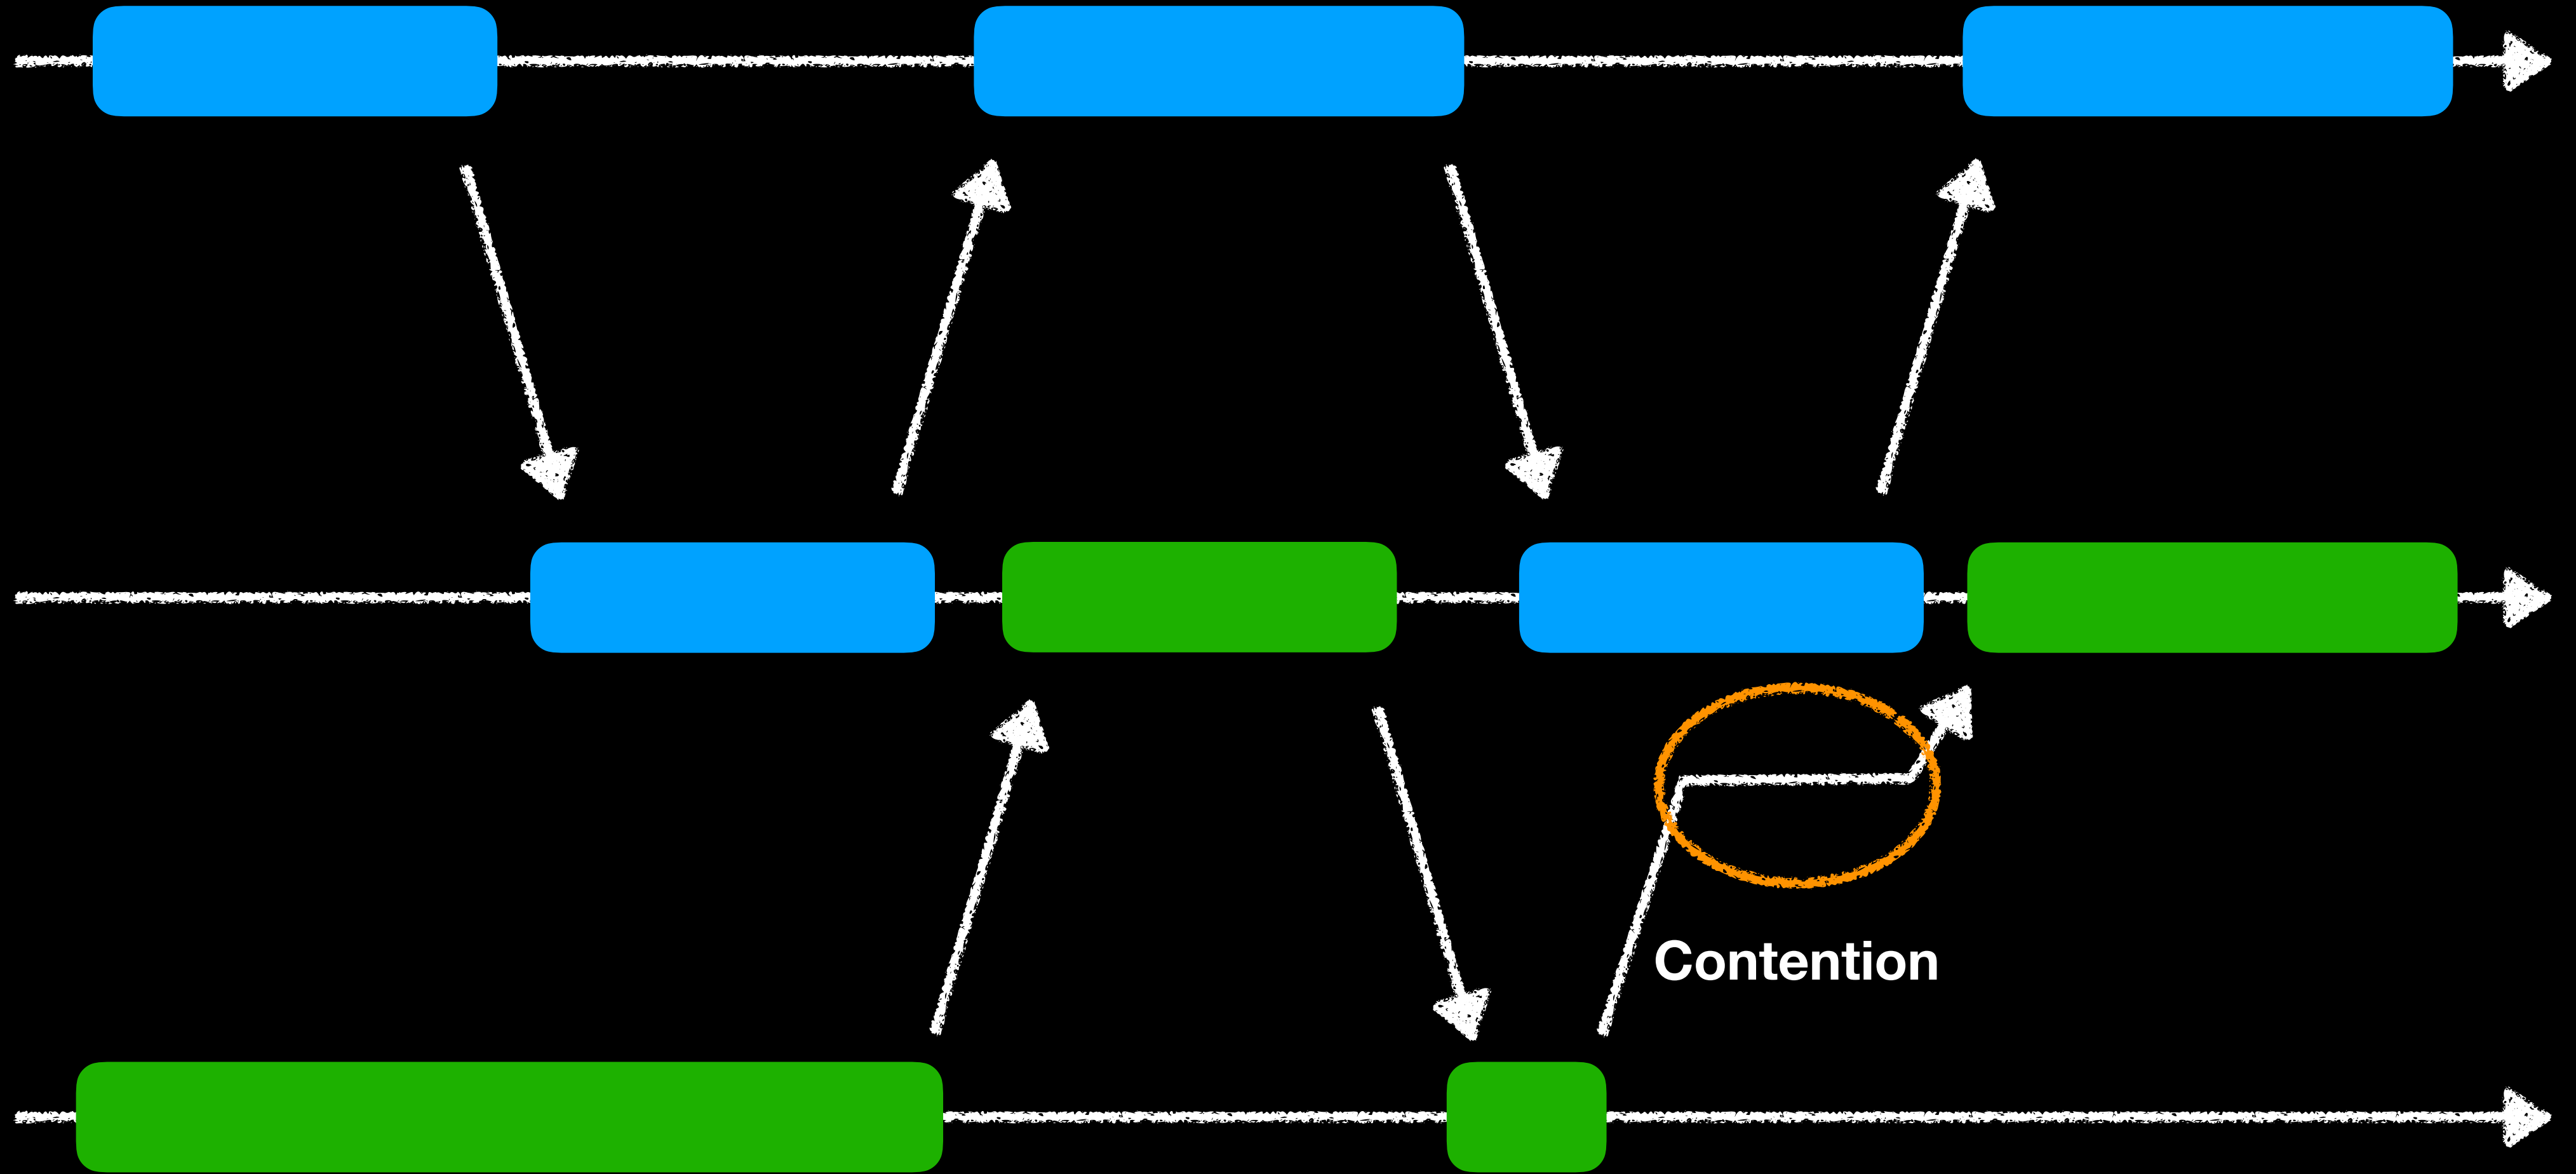



並行性

Concurrency








並列性

Parallelism



**How does this apply
to Ruby?**


どのようにしてRubyに應用するか?

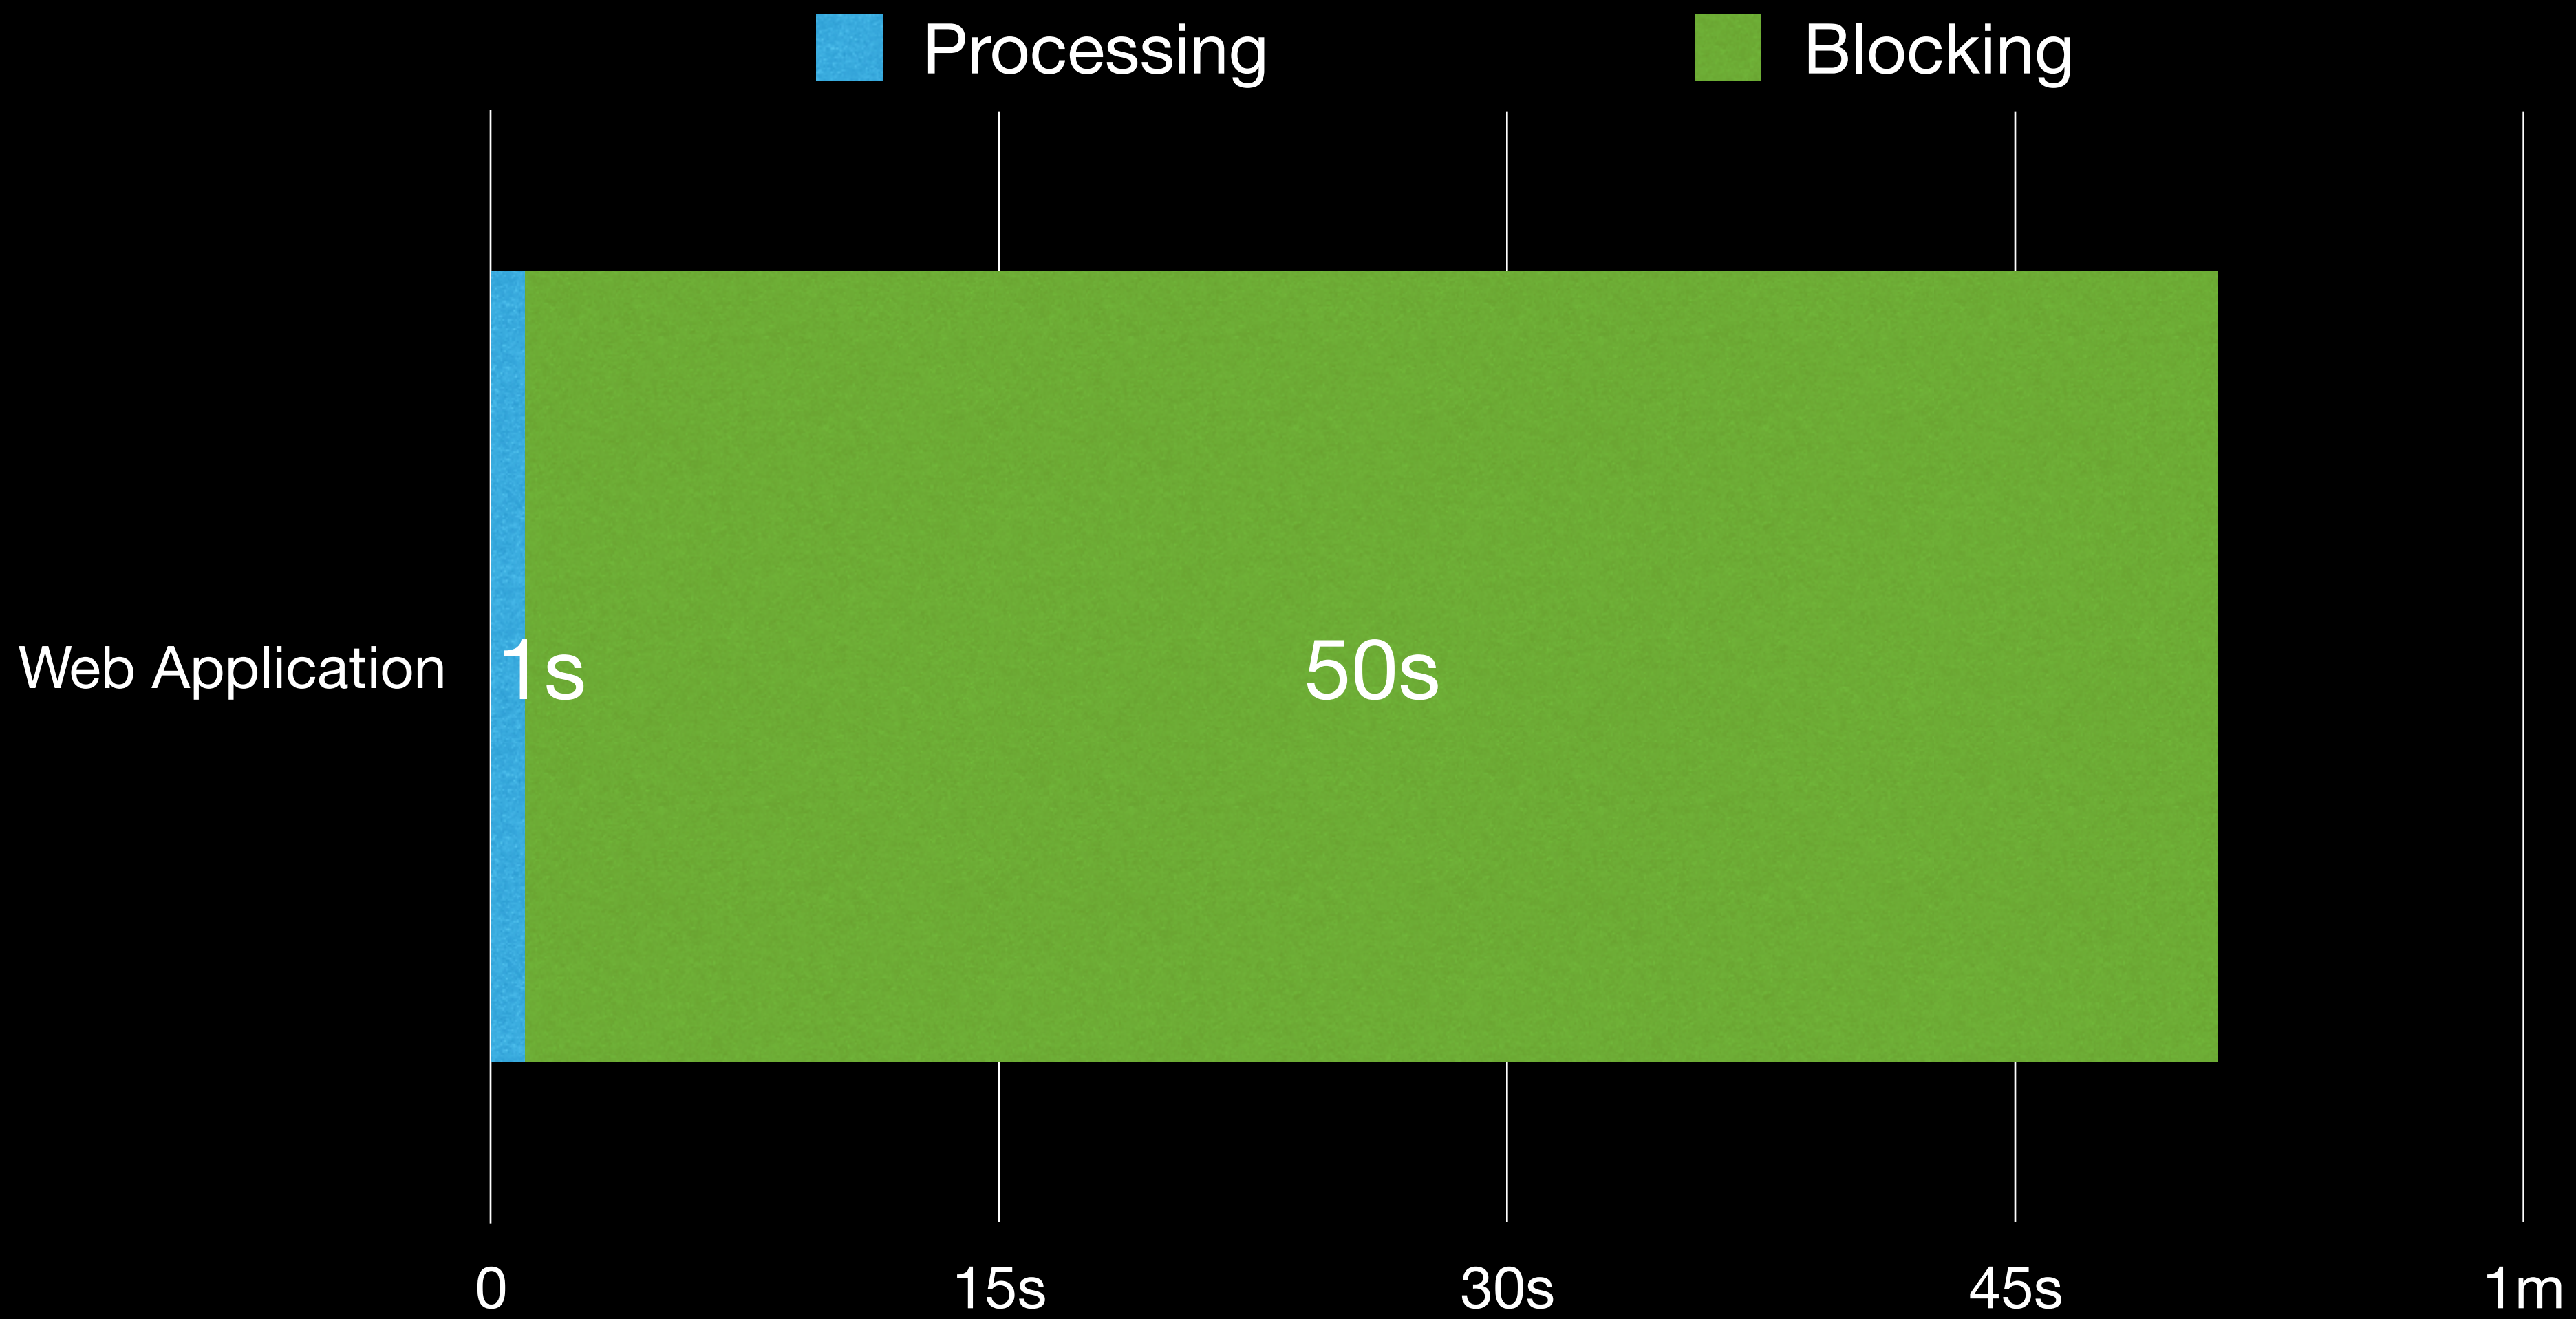
Is Ruby fast enough?

Rubyは十分早いのか?

**Let's make Ruby 10x
faster.**

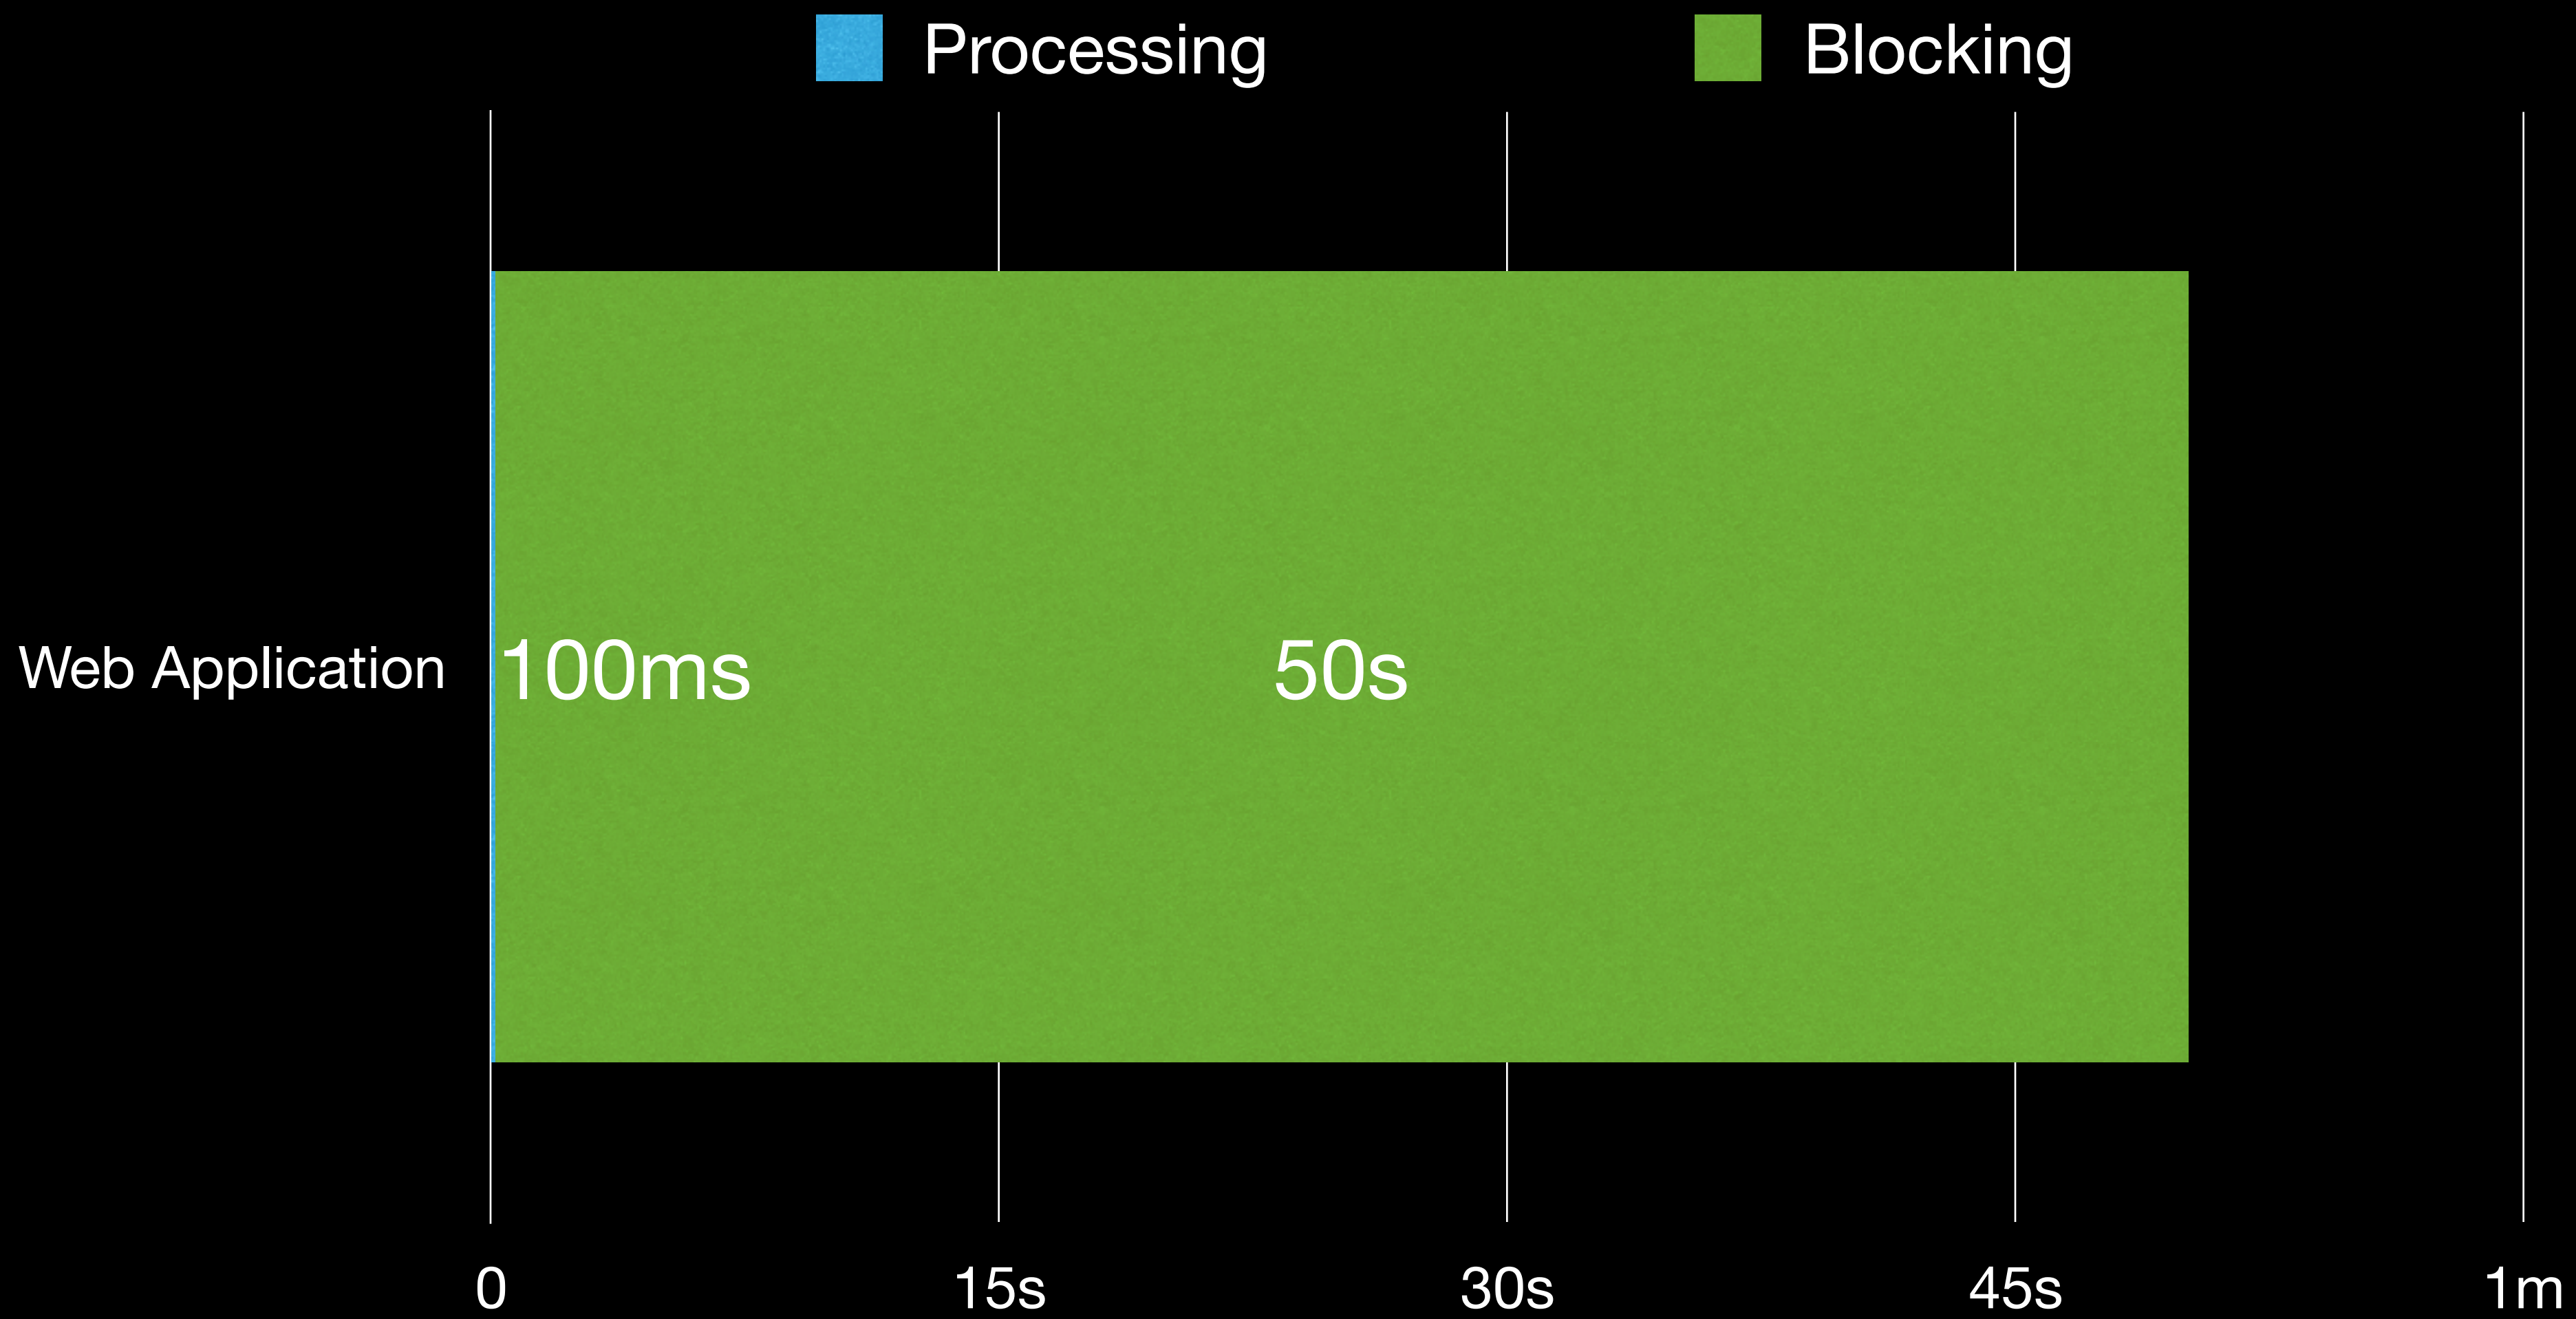
Rubyを10倍早くしよう





**Let's make Ruby 100x
faster.**

Rubyを100倍早くしよう



**How do we handle
more requests?**

どのようにしてもより多くのリクエストを処理するか?

Can we use multiple
processes?

複数のプロセスを利用することはできるか?

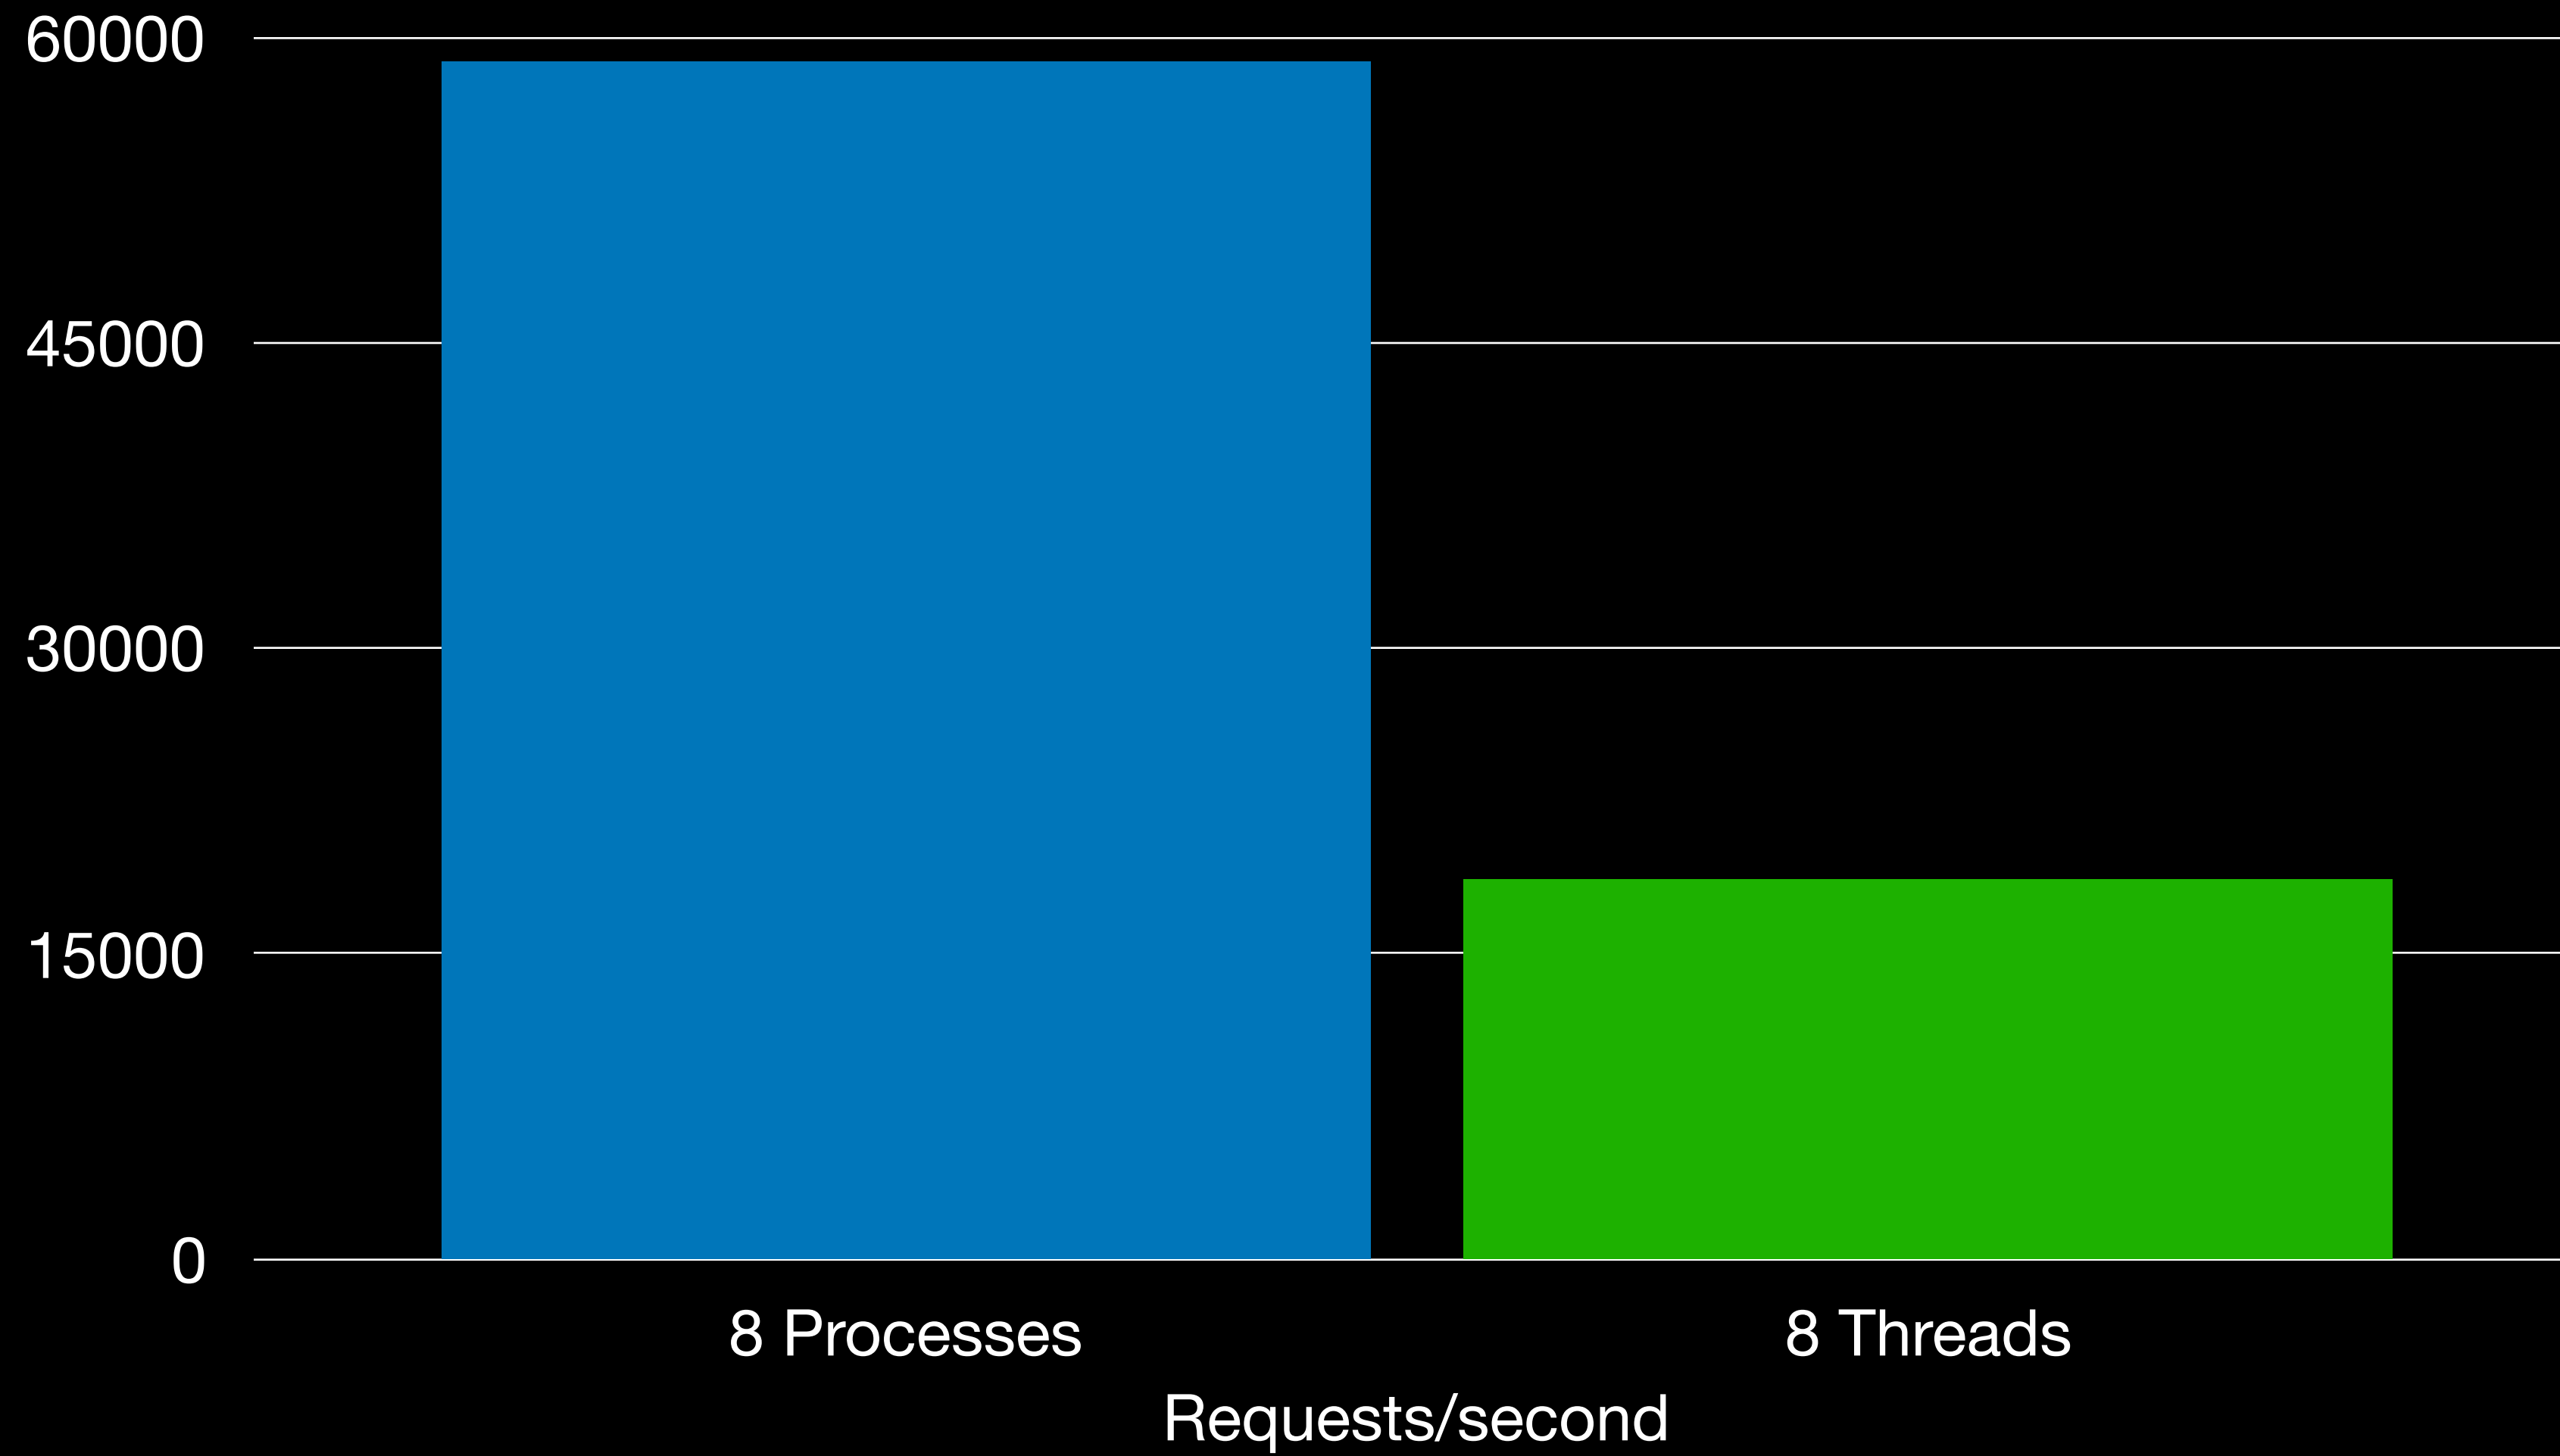
What about threads?

スレッドは?

**How bad is the global
interpreter lock?**

グローバルインタプリタロックはどれほど悪いのか

Falcon "Hello World" Web Server



**Are processes and
threads sufficient?**

プロセスとスレッドは十分か?

**How many processes
can we create?**

いくつのプロセスを作ることができるか?

**How many threads
can we create?**

いくつのスレッドを作ることができるか

100?

1,000?

10,000?

**What about long
running connections?**

ロングランニングコネクションはどうするか?

**What about 100,000
connected WebSockets?**

100,000のWebSocketsは?



We need to go deeper

Event driven non-blocking I/O

イベントドリブンのノンブロッキングI/O



```
while message = connection.read  
  handle(message)  
end
```



```
while message = connection.read  
  handle(message)  
end
```



```
while message = connection.read  
  handle(message)  
end
```



```
while true
  ready = IO.select(connections)
  ready.each{|connection| handle(connection.read)}
end
```

How do we handle user logic?

ユーザーロジックはどのようにして処理するか

```
def remote_size(host, port)
  peer = TCPSocket.new(host, port)
  count = 0

  while buffer = peer.read(1024)
    count += buffer.bytesize
  end

  return count
ensure
  peer&.close
end

puts remote_size(HOST, PORT)
```

Sequential is easy.

シーケンシャルな処理は簡単

Callbacks...

```
def remote_size(host, port)
  TCPSocket.new(host, port) do |peer, error|
    end
end

remote_size(HOST, PORT) do |size, error|
  puts size
end
```

```
def remote_size(host, port)
  TCPSocket.new(host, port) do |peer, error|
    if error
      yield nil, error
    else
      count = 0

      peer.read(1024) do |buffer, error|
        end
      end
    end
  end
end
```

```
    yield nil, error
else
  count = 0

  read_more = lambda do
    peer.read(1024) do |buffer, error|
      if error
        yield nil, error
      elsif buffer
        count += buffer.bytesize

        read_more.call
      else
        yield count, nil
      end
    end
  end
end

read_more.call
end
end
end
```

```
while buffer = peer.read(1024)
  count += buffer.bytesize
end

return count
```



```
    yield nil, error
else
  count = 0

  read_more = lambda do
    peer.read(1024) do |buffer, error|
      if error
        peer.close
        yield nil, error
      elsif buffer
        count += buffer.bytesize

        read_more.call
      else
        peer.close
        yield count, nil
      end
    end
  end
end

read_more.call
end
```

```
    yield nil, error
else
  count = 0

  read_more = lambda do
    peer.read(1024) do |buffer, error|
      if error
        peer.close
        yield nil, error
      elsif buffer
        count += buffer.bytesize

        read_more.call
      else
        peer.close
        yield count, nil
      end
    end
  end
end

read_more.call
end
```

Callback Hell

Async/Await...

```
async def remote_size(host, port)
  peer = await TCPSocket.new(host, port)
  count = 0

  while buffer = await peer.read(1024)
    count += buffer.bytesize
  end

  return count
ensure
  peer&.close
end

async lambda do
  puts await remote_size(HOST, PORT)
end.call
```

```
async lambda do
  puts(await remote_size(HOST, PORT))
end.call
```

```
async lambda do
  await puts(await remote_size(HOST, PORT))
end.call
```

```
async def remote_size(host, port)
  peer = await TCPSocket.new(host, port)
  count = 0

  while buffer = await peer.read(1024)
    count += buffer.bytesize
  end

  return count
ensure
  peer&.close
end
```

Async/Await Hell

```
async lambda do
  await puts(await remote_size(HOST, PORT))
end.call
```

Can we do better?

改善できるか

Should we rewrite
existing code?

既存のコードを書き直すべきか？

**What about using
fibers?**

ファイバーを使うのはどうか

What are fibers?

ファイバーとは何か

```
def add(sum)
  return sum
end
```

```
puts add(10) # => 10
puts add(20) # => 20
```

add(10)

call



sum = 10

puts(10)

return



return sum

add(20)

call



sum = 20

puts(20)

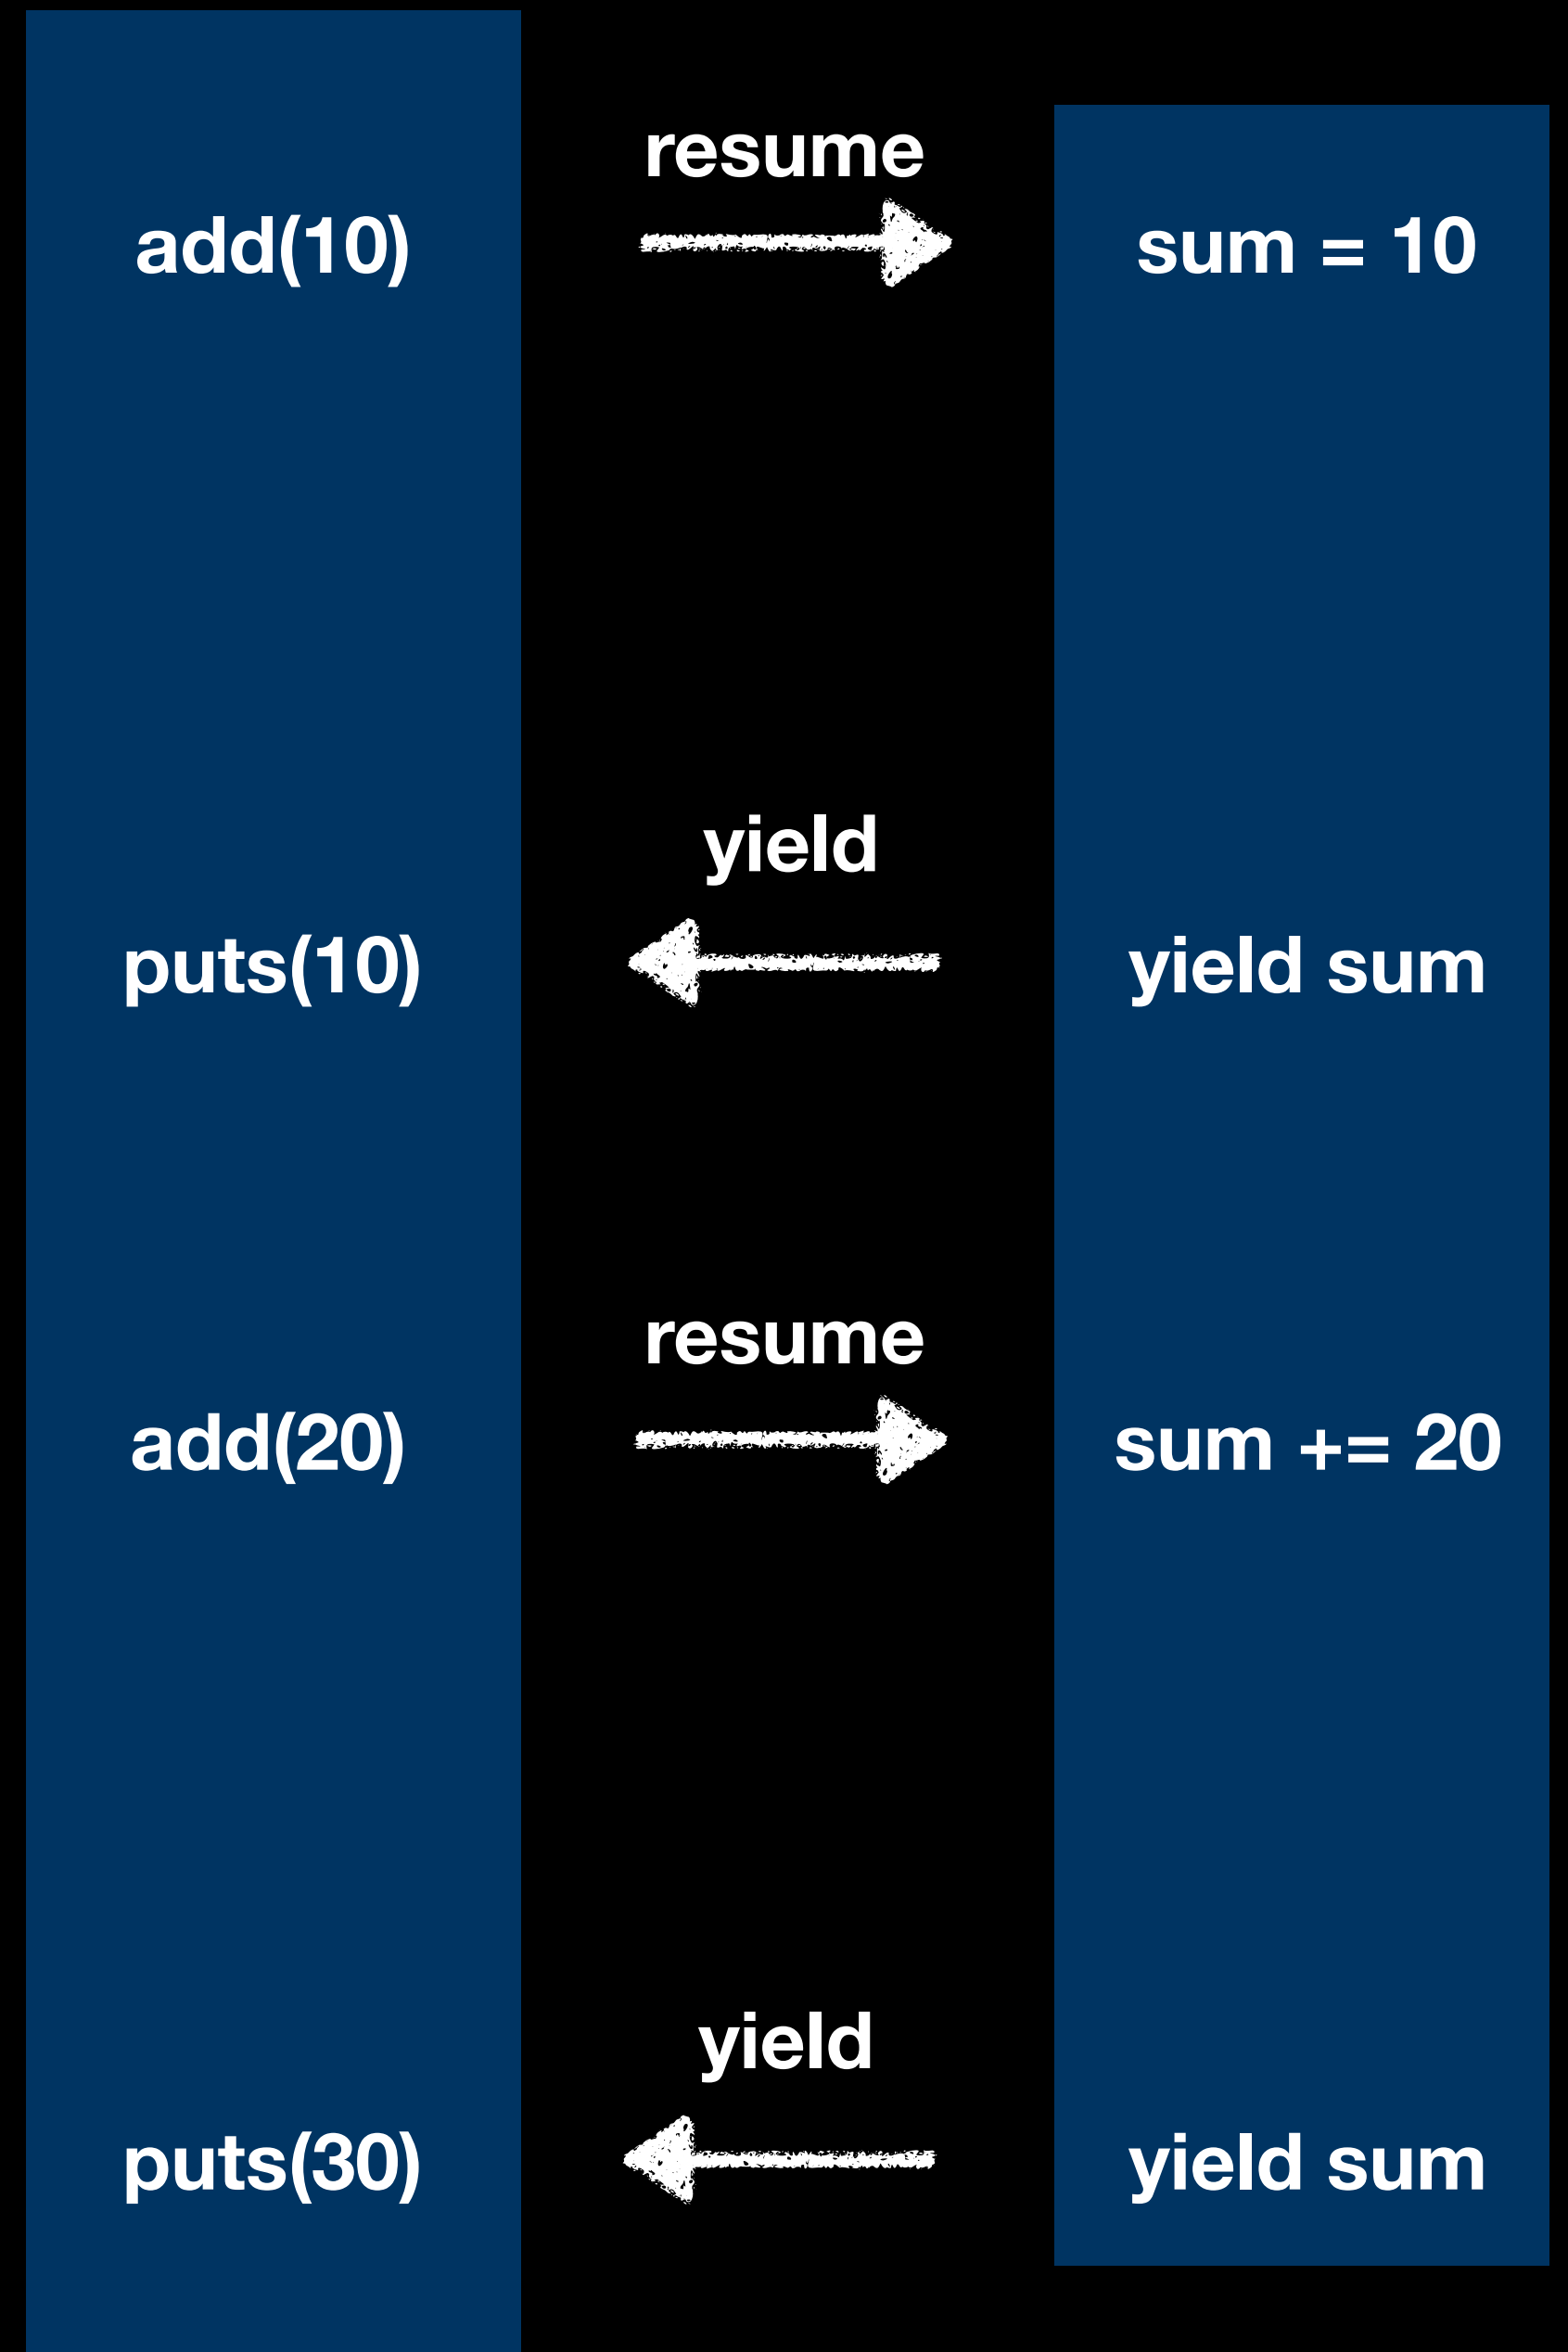
return



return sum

```
add = Fiber.new do |sum|
  while true
    sum += Fiber.yield(sum)
  end
end

puts add.resume(10) # => 10
puts add.resume(20) # => 30
```

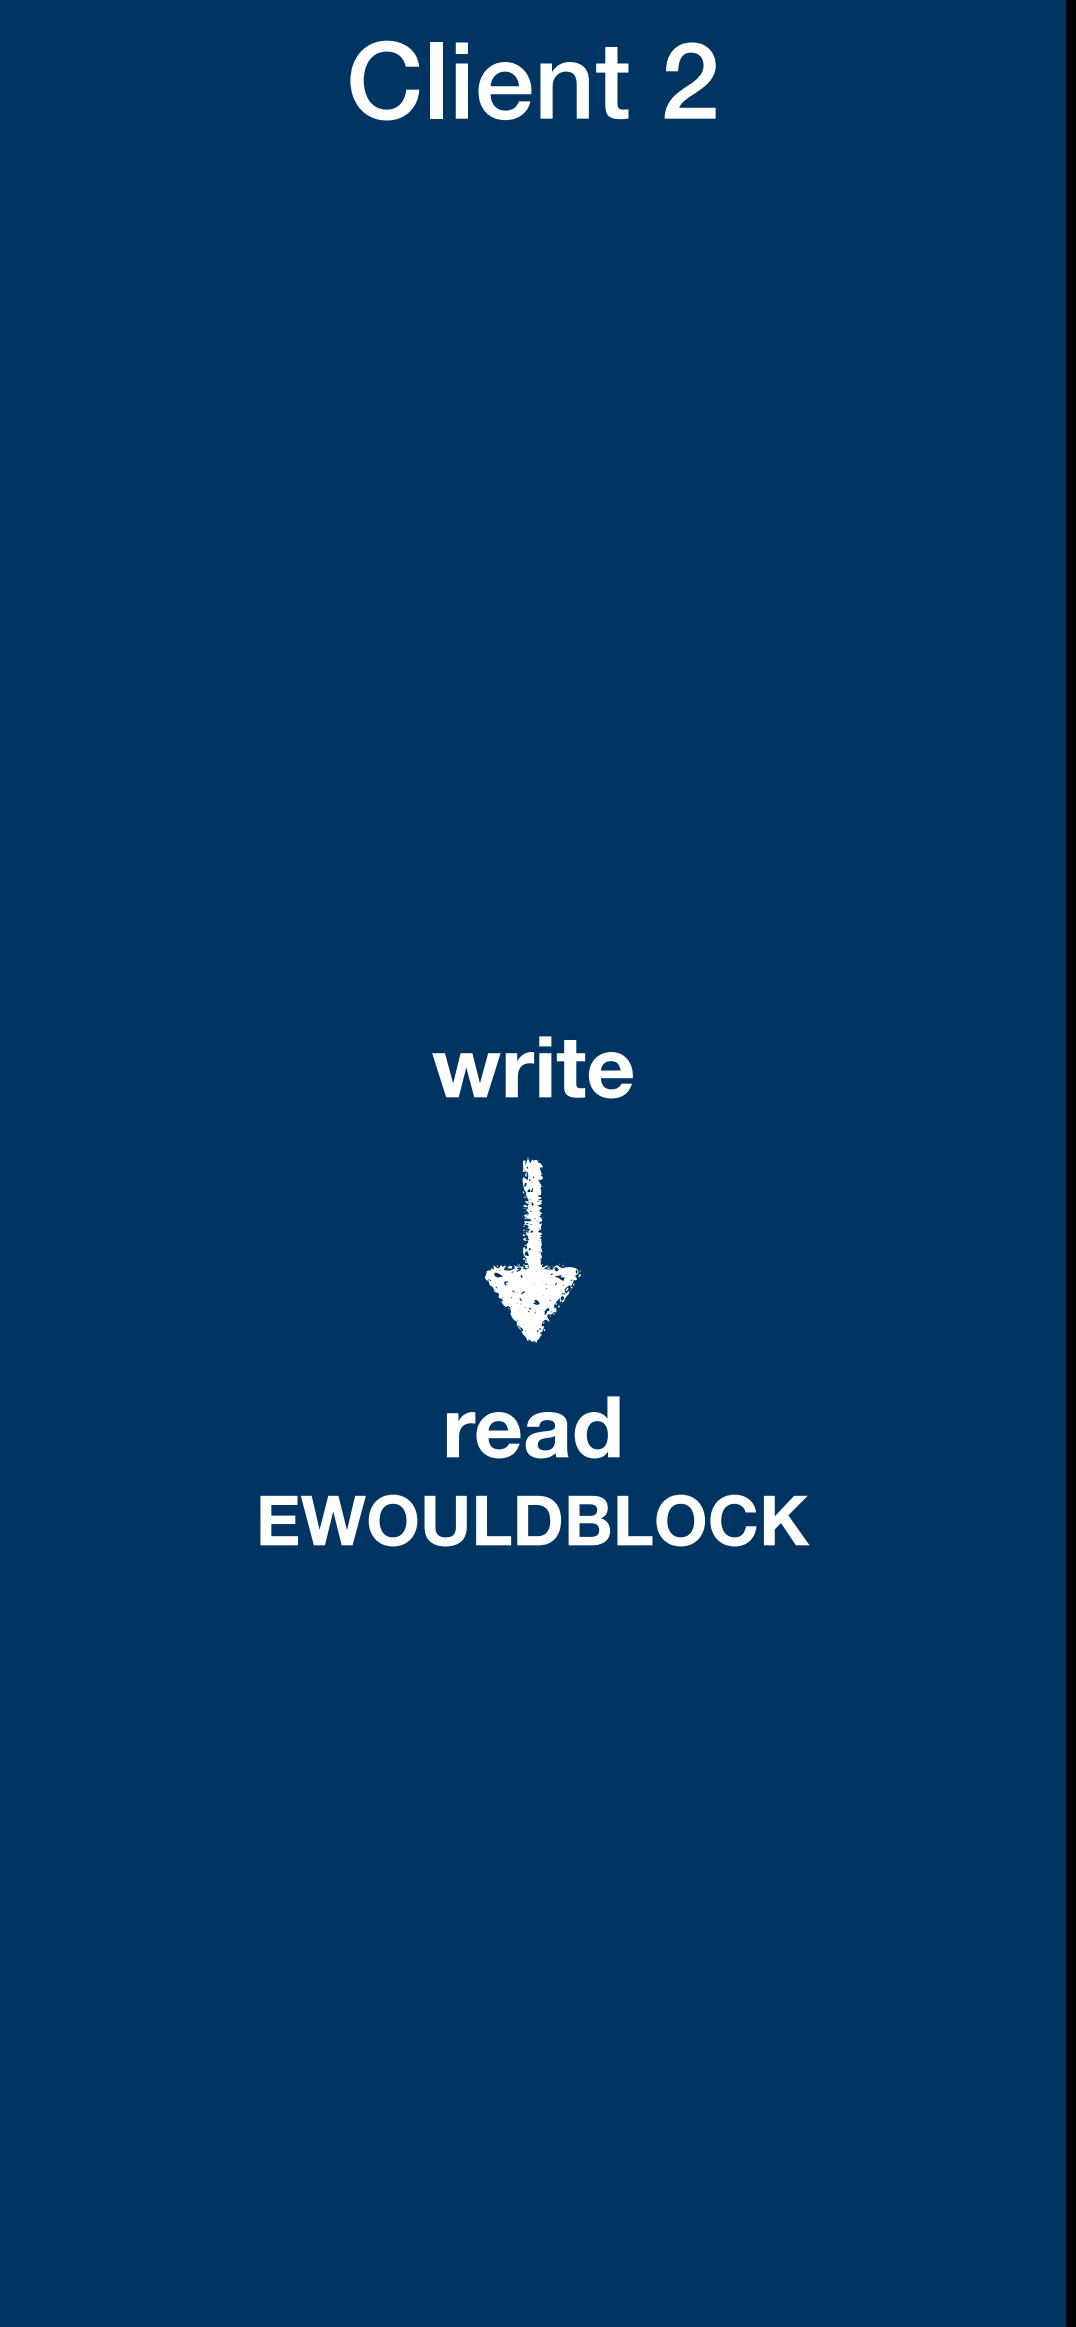
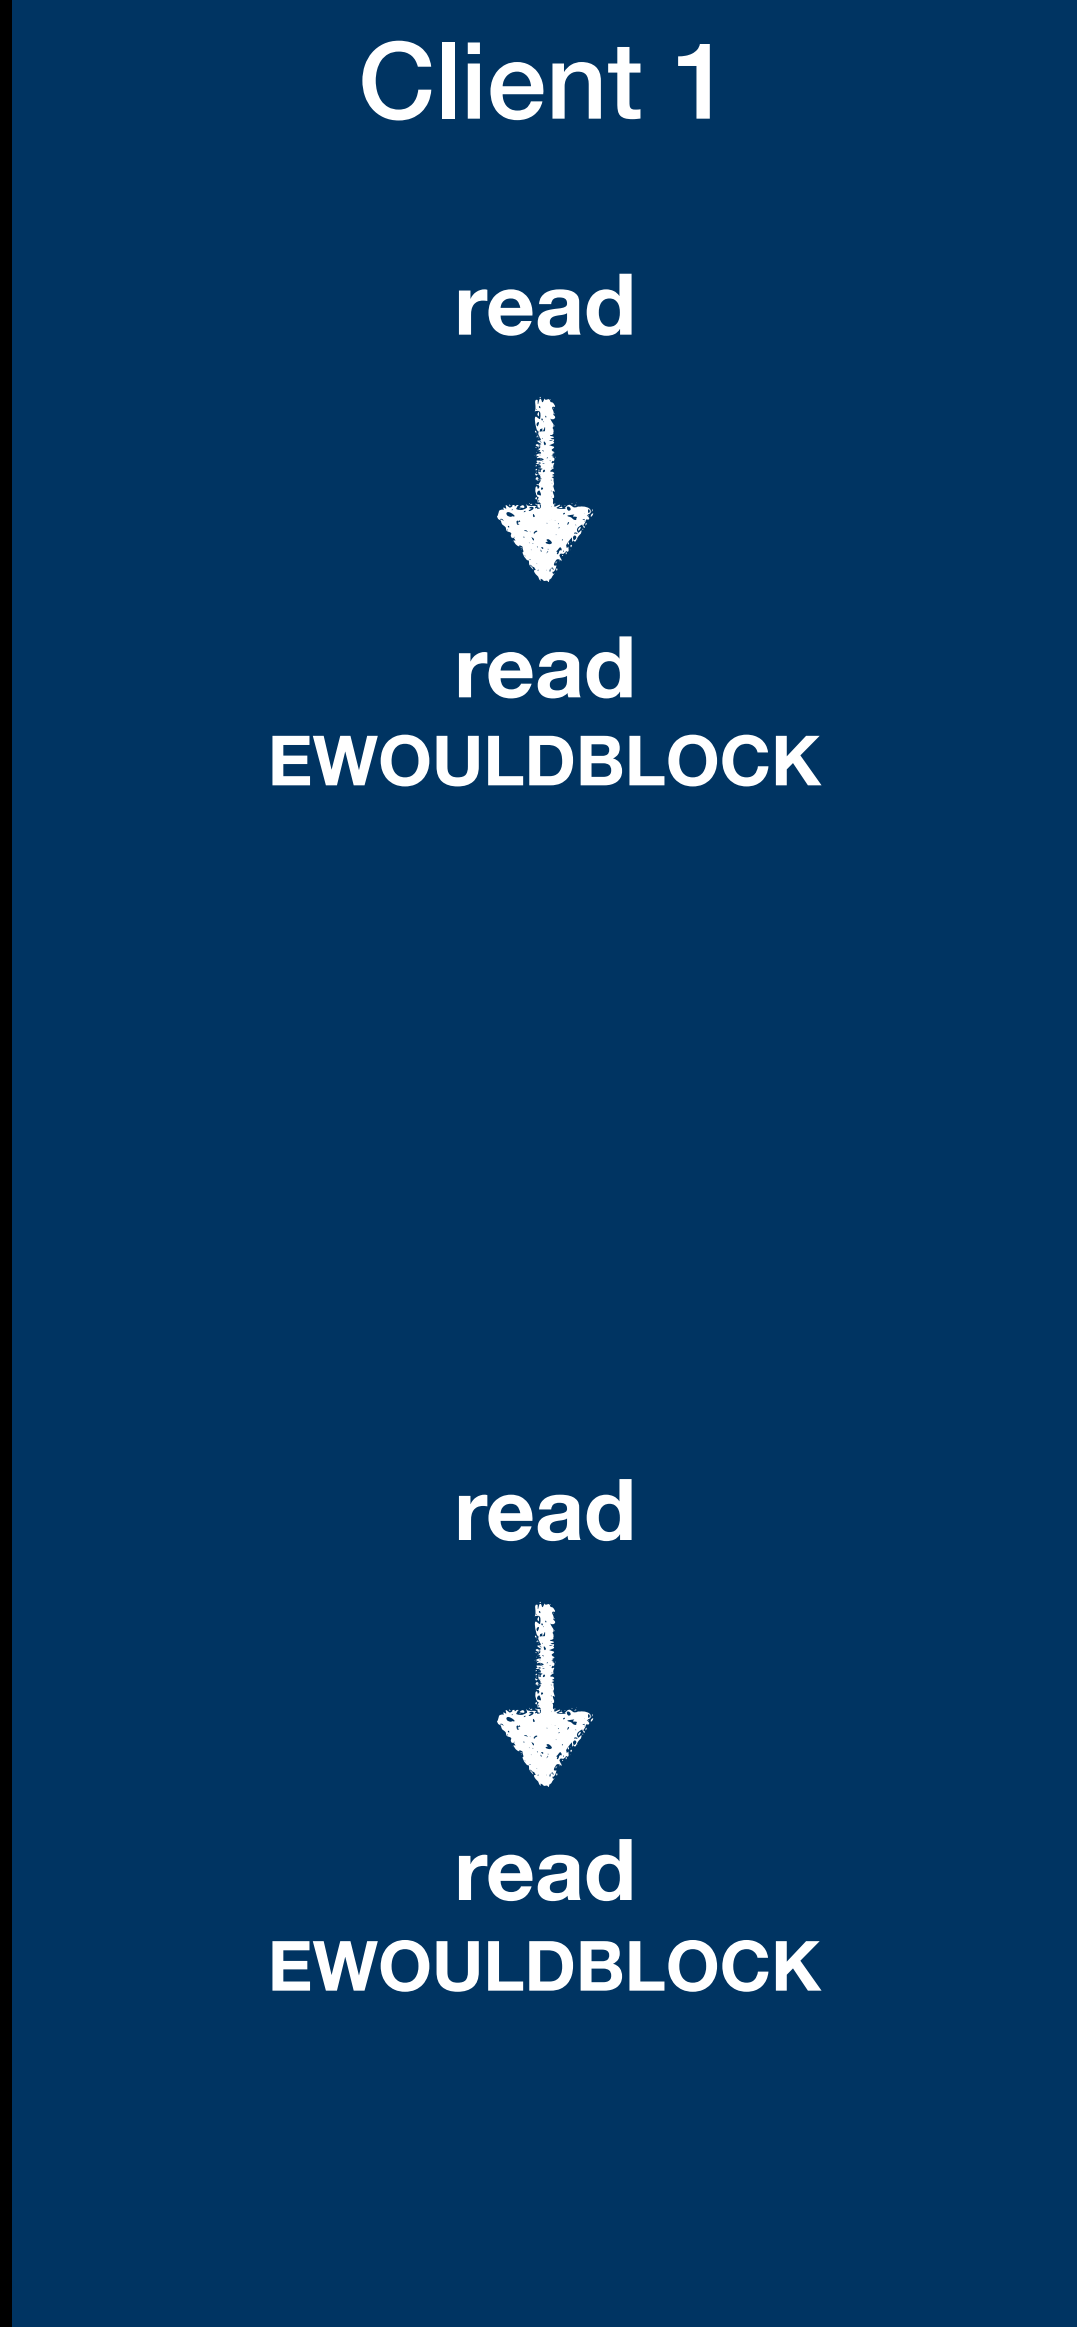


**Stack and instruction
pointer is not lost.**

スタックとインストラクションポインタはなくなっていない

How do we use fibers for blocking I/O?

どのようにしてファイバーをブロッキングI/Oに使うか



yield(socket)



resume



yield(socket)



resume



yield(socket)



```
def remote_size(*address)
  Async do
    peer = Async::IO::TCPSocket.new(*address)
    count = 0

    while buffer = peer.read(1024)
      count += buffer.bytesize
    end

    return count
  ensure
    peer&.close
  end.wait
end


puts remote_size(HOST, PORT)
```

How do we make
existing code scalable?

どのようにして既存のコードをスケーラブルにするか

Proof of concept of light weight Thread selector implementation. #1870

 Open

ioquatix wants to merge 4 commits into `ruby:trunk` from `ioquatix:thread-selector` 

Transparently make all I/O non-blocking

コードを変更することなく全てのI/Oをノンブロッキングにする



⌘ ⌘ @@ -1114,6 +1114,12 @@ io_fflush(rb_io_t *fptr)

```
1114 1114     int
1115 1115     rb_io_wait_readable(int f)
1116 1116     {
1117 +     VALUE selector = rb_current_thread_selector();
1118 +     if (selector != Qnil) {
1119 +         VALUE result = rb_funcall(selector, rb_intern("wait_readable"), 1, INT2NUM(f));
1120 +         return RTEST(result);
1121 +     }
1122 +
1117 1123     io_fd_check_closed(f);
1118 1124     switch (errno) {
1119 1125     case EINTR:
```


⌘ ⌘ @@ -1138,6 +1144,12 @@ rb_io_wait_readable(int f)

```
1138 1144     int
1139 1145     rb_io_wait_writable(int f)
1140 1146     {
1147 +     VALUE selector = rb_current_thread_selector();
1148 +     if (selector != Qnil) {
1149 +         VALUE result = rb_funcall(selector, rb_intern("wait_writable"), 1, INT2NUM(f));
1150 +         return RTEST(result);
1151 +     }
1152 +
1141 1153     io_fd_check_closed(f);
1142 1154     switch (errno) {
1143 1155     case EINTR:
```

⌘ ⌘

```
12     thread = Thread.new do
13         selector = Selector.new
14         Thread.current.selector = selector
15
16         i, o = IO.pipe
17         i.nonblock = true
18         o.nonblock = true
19         e = i.to_enum(:each_char)
20
21         Fiber.new do
22             o.write("Hello World")
23             o.close
24         end.resume
25
26         Fiber.new do
27             while c = (e.next rescue nil)
28                 message << c
29             end
30         end.resume
31
32         selector.run
33     end
```

```
13 def run
14   while @readable.any? or @writable.any?
15     readable, writable = IO.select(@readable.keys, @writable.keys, [])
16
17     readable.each do |io|
18       @readable[io].transfer
19     end
20
21     writable.each do |io|
22       @writable[io].transfer
23     end
24   end
25 end
26
27 def wait_readable(fd)
28   io = IO.for_fd(fd)
29
30   @readable[io] = Fiber.current
31
32   @fiber.transfer
33
34   @readable.delete(io)
35
36   return true
37 end
```



<https://github.com/socketry/async>



:P

Bryan Powell

bryanp

Block or report user

“Async is the right model because web apps are almost always I/O bound. The Ruby web ecosystem is really lacking in scalability (e.g. WebSockets on Puma). Async unlocks the next tier of scalability in the most Ruby-like way possible.”

Bryan Powell, on migrating from Puma to Falcon.



<https://github.com/socketry/falcon>

Multi-process

Multi-thread

Multi-fiber

HTTP/1
HTTP/2 & TLS

WebSockets

**Event
Loop**

yield



resume



**Synchronous
Rack Middleware**

**Asynchronous
Faraday**

**Synchronous
Rack Middleware**

yield



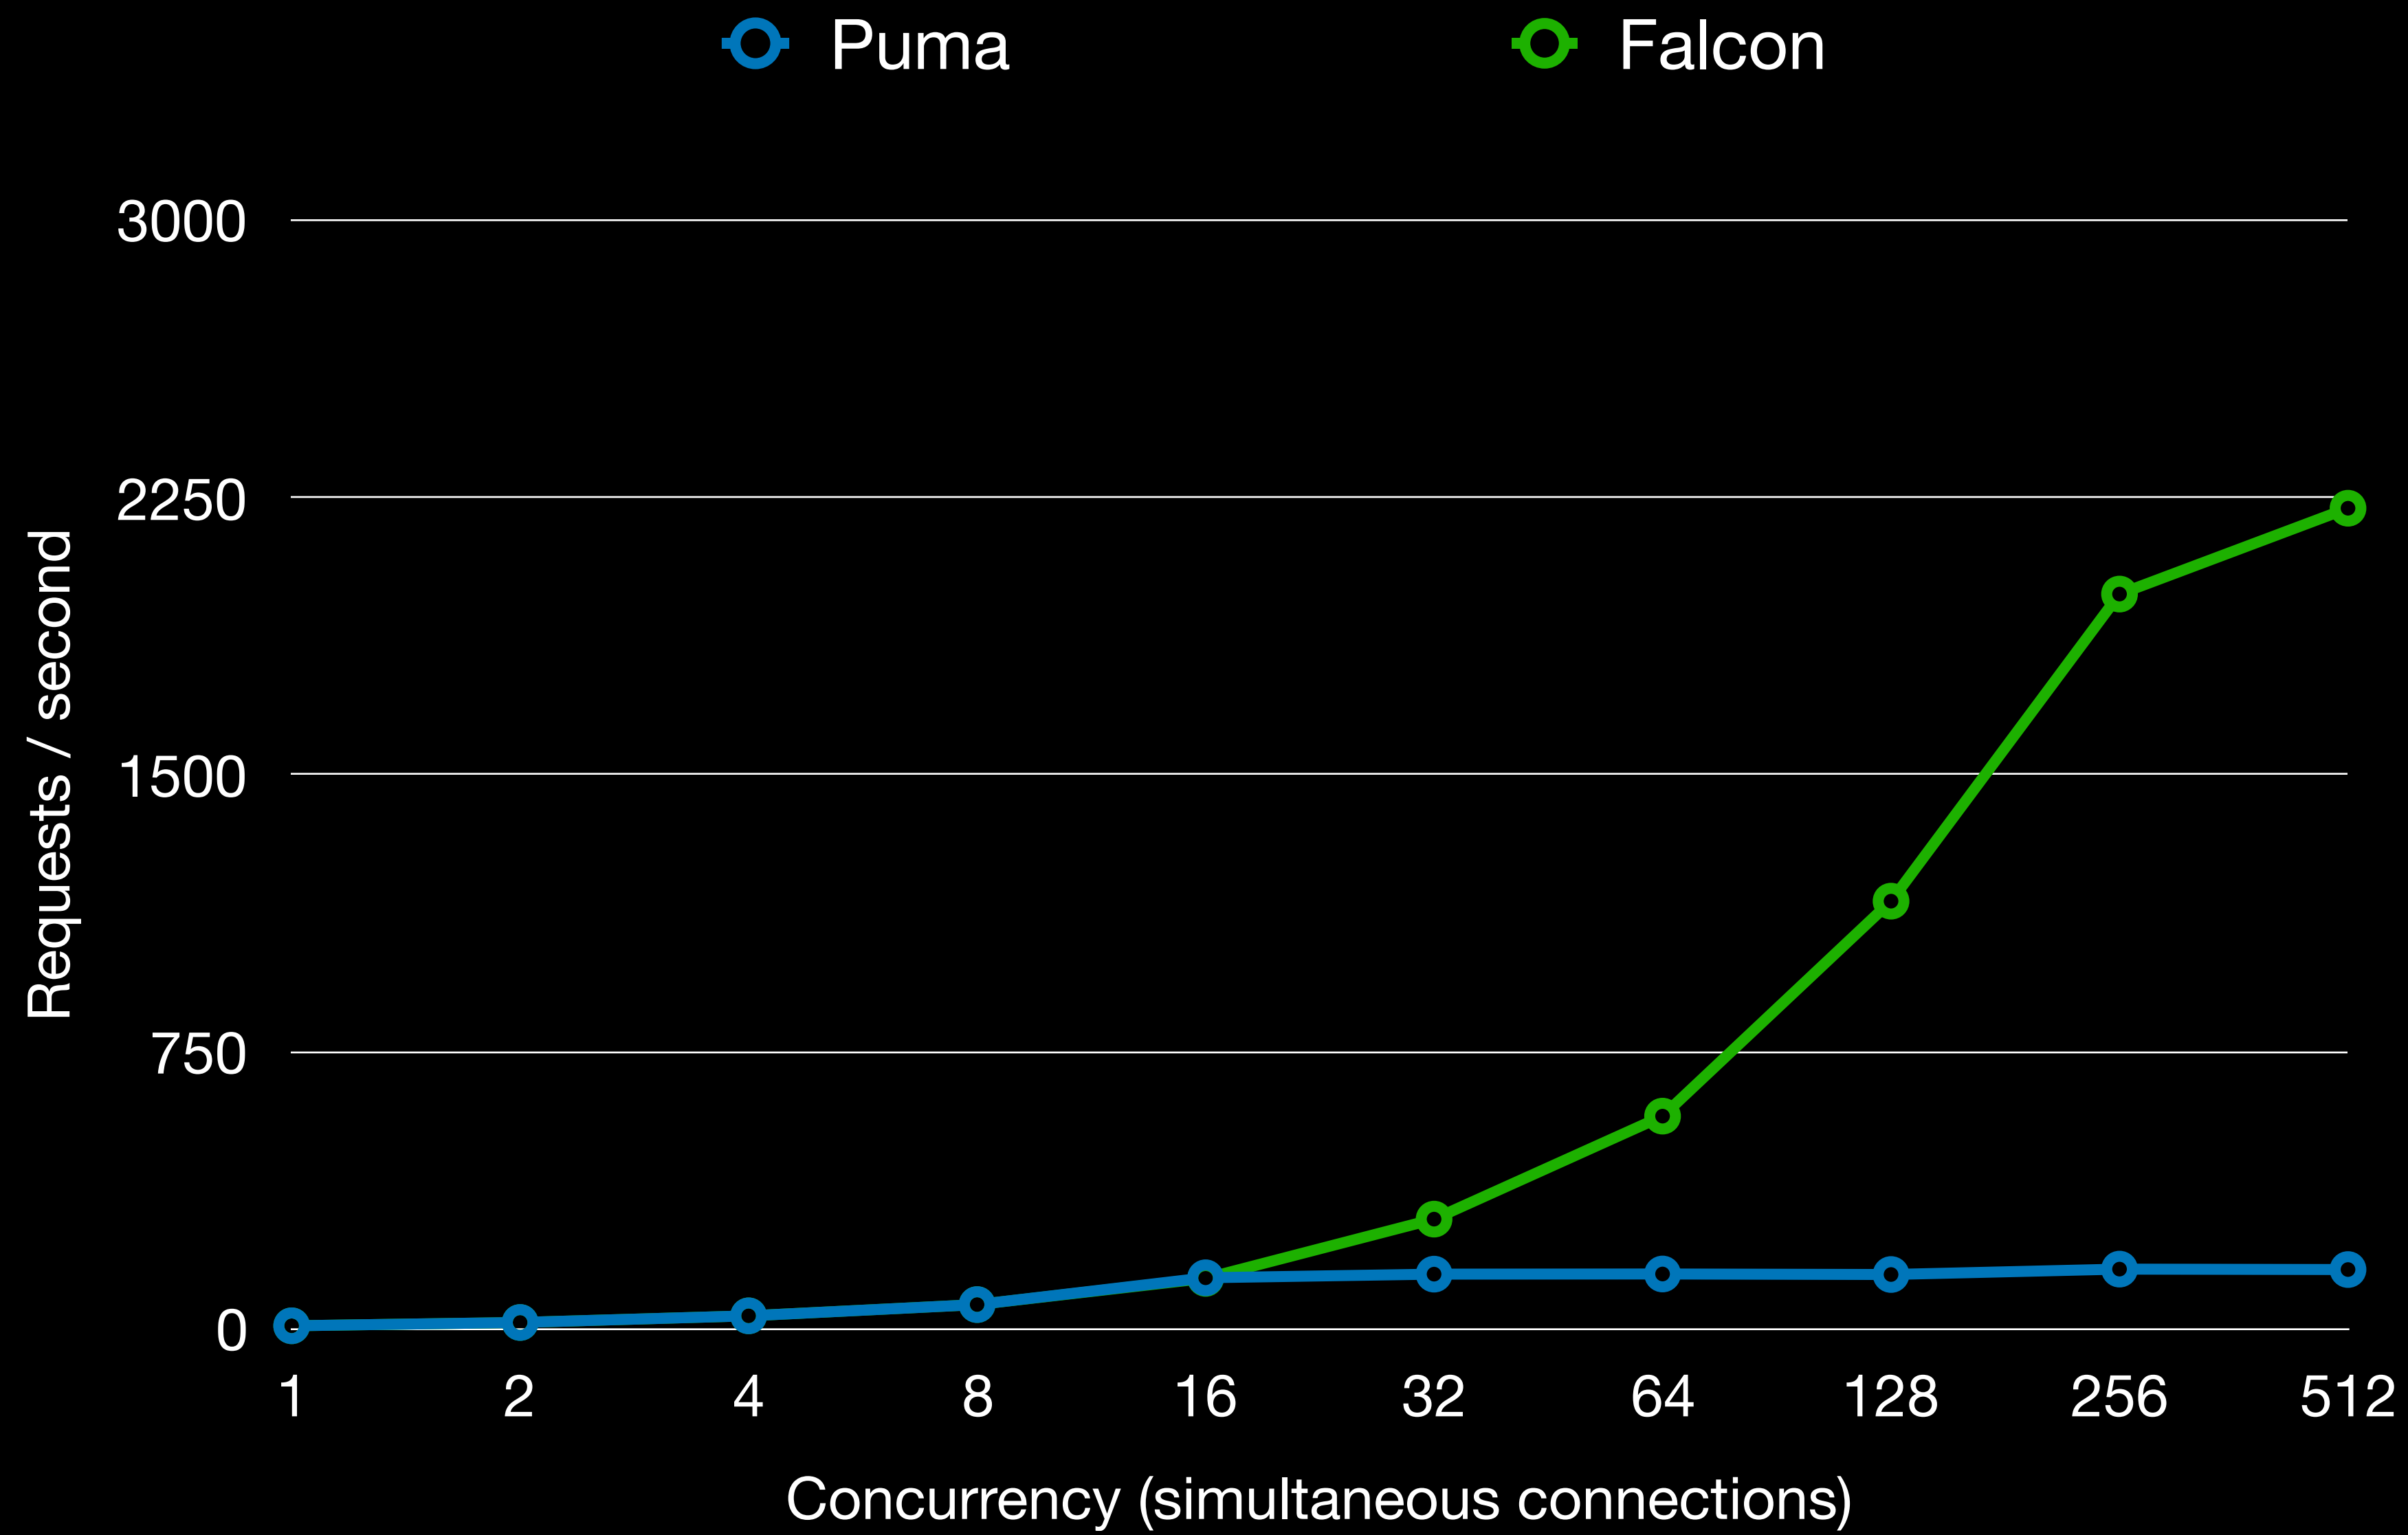
resume



**Asynchronous
Postgres**

How does it perform?

パフォーマンスはどうか



**Are fibers the right
solution?**

ファイバーが正しい選択か

**Fibers scale better
than threads.**

スレッドよりもファイバーの方がスケールする

**Fibers are easier than
threads.**

スレッドよりもファイバーの方が簡単

**Fibers improve the
scalability of existing code.**

ファイバーは既存コードのスケールビリティを改善する

S3

Postgres

WebSocket

Redis

Fibers Are the Right Solution

MySQL

SMTP

ファイバーが正しい選択

HTTP

DNS

Disk

Fibers Are the Right Solution

