

○ ● ● dRuby and Rinda.

その実装と応用を札幌で

○●● 重要なことを先に

● 先に

¥ 3360

- 初刷まだ買えます



Who is translating it?

- 英語版はまだ買えません



<http://dx.doi.org/10.1007/s10766-008-0086-1>

\$32.00

- International Journal of PARALLEL PROGRAMING
- コレクターズアイテム
- 余剰資金のある研究室は Must buy!

Int J Parallel Prog
DOI 10.1007/s10766-008-0086-1

dRuby and Rinda: Implementation and Application of Distributed Ruby and its Parallel Coordination Mechanism

Masatoshi Seki

Received: 20 March 2008 / Accepted: 13 August 2008
© Springer Science+Business Media, LLC 2008

Abstract The object-oriented scripting language Ruby is admired by many programmers for being easy to write in, and for its flexible, dynamic nature. In the last few years, the Ruby on Rails web application framework, popular for its productivity benefits, has brought about a renewed attention to Ruby for enterprise use. As the focus of Ruby has broadened from small tools and scripts, to large applications, the demands on Ruby's distributed object environment have also increased, as has the need for information about its usage, performance and examples of common practices. dRuby and Rinda were developed by the author as the distributed object environment and shared tuplespace implementation for the Ruby language, and are included as part of Ruby's standard library. dRuby extends method calls across the network while retaining the benefits of Ruby. Rinda builds on dRuby to bring the functionality of Linda, the glue language for distributed co-ordination systems, to Ruby. This article discusses the design policy and implementation points of these two systems, and demonstrates their simplicity with sample code and examples of their usage in actual applications. In addition to dRuby and Rinda's appropriateness for prototyping distributed systems, this article will also demonstrate that dRuby and Rinda are building a reputation for being suitable infrastructure components for real-world applications.

Keywords Ruby · Linda · dRuby · Rinda · TupleSpace · RMI

M. Seki (✉)
Tochigi, Japan
e-mail: m_seki@mva.biglobe.ne.jp
URL: www.druby.org

 Springer

 Springer

Print ISSN 1076-6531
Online ISSN 1076-6549
© Springer Science+Business Media, LLC
2008

<http://dx.doi.org/10.1007/s10766-008-0086-1>

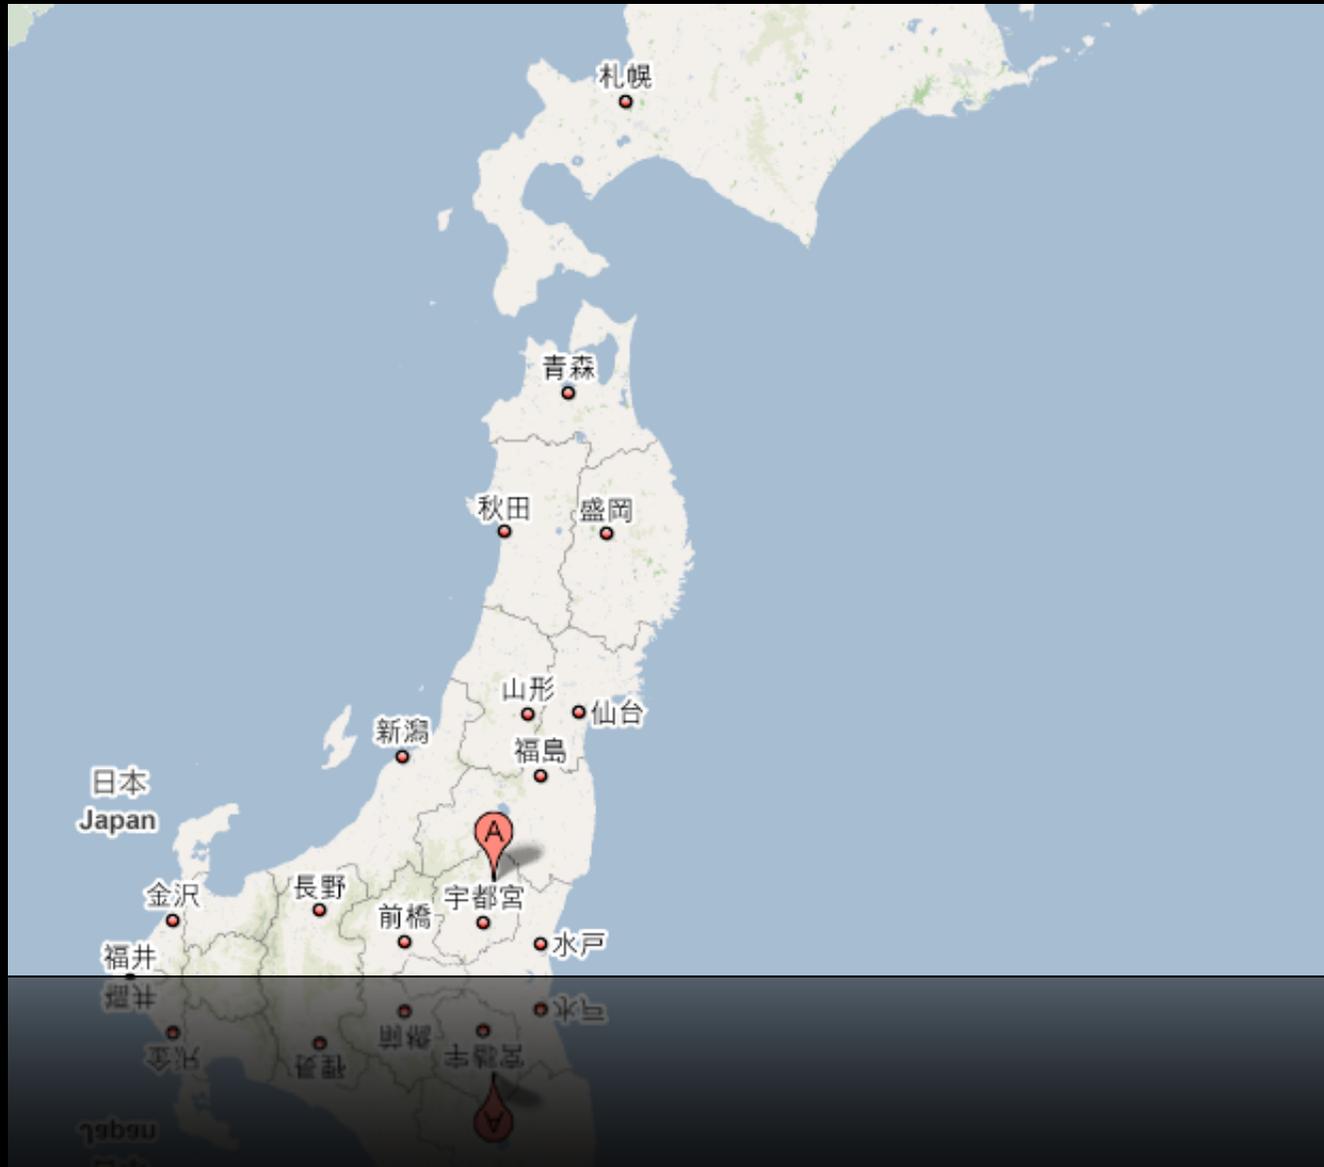
私について

- toRubyからきました
 - 池澤さんで一躍有名に
 - 毎月第一水曜日開催

会場



会場



○ ● ● Agenda

- 重要なことを先に
- 紹介
- 入門
- 応用
- 重要なことをもう一度

○ ● ● Rubyは奇妙だ

- Rubyっぽく書いているとOOPな自分になれたように勘違いさせてくれる超かっこいい言語。

○ ● ● Rubyの特徴の一部

- かっこいいクラスライブラリ
- 型のない変数
- 豪華なリフレクション
- ユーザレベルスレッド

DRbも奇妙だ

- The Good:
 - Stupid Easy
 - Reasonably Fast
- The Bad:
 - Kinda Flaky
 - Zero Redundancy
 - Tightly Coupled

Remote Method Invocation

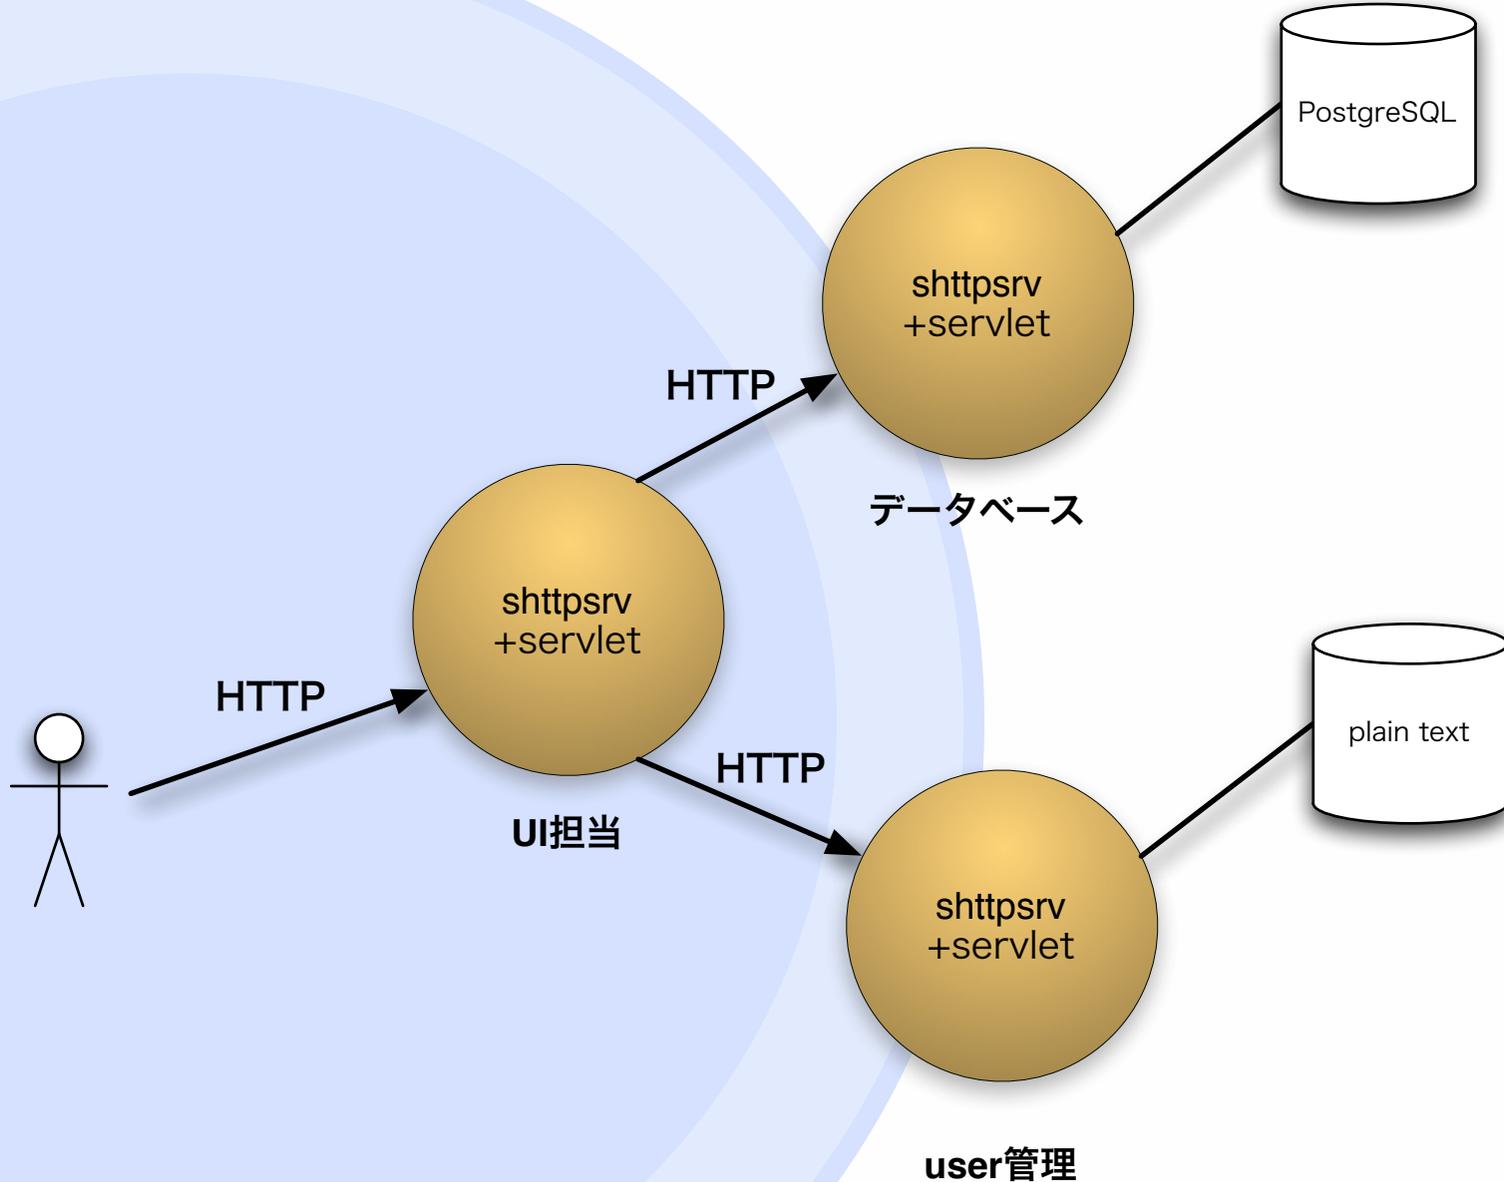
- 別のプロセス／マシンのオブジェクトにメッセージを送り、メソッドを起動する仕組み
- Remote Procedure Callと紙一重

例えばWeb Service

- まあ似てなくもない
 - XML-RPCとかいうぐらいだし
 - そういえば..



原型はこんなだった





発見

- あらゆるプロセスはサーバでありクライアントである
- まるでOOPのようだ

dRubyの特徴

- たくさんあるRMIライブラリの一つ
- でも既存のRMI, RPCのブリッジではなく
他のシステムからの移植版でもない

dRubyの特徴

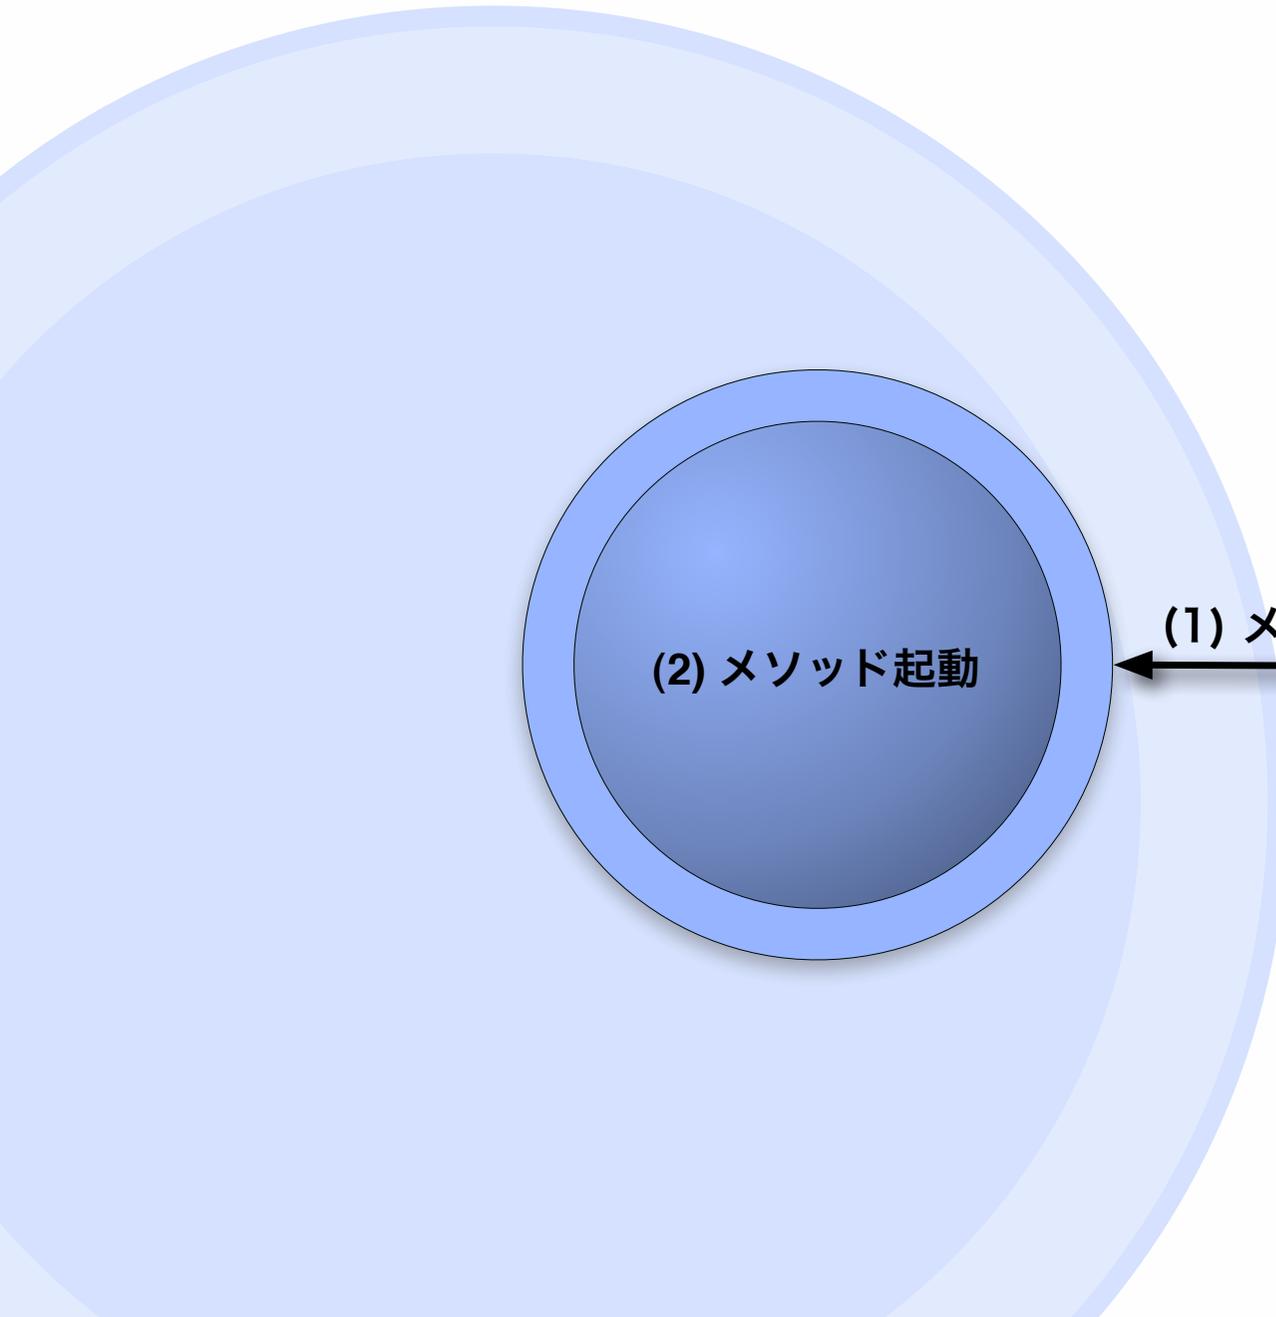
- Rubyのメソッド呼び出しを拡張
 - インタプリタを越えて
 - プロセスを
 - マシンを
- そこだけを守備範囲にしている

メッセージング

- オブジェクトにメッセージを送信
- メソッドを起動
- ちょっと嘘かも



Ruby

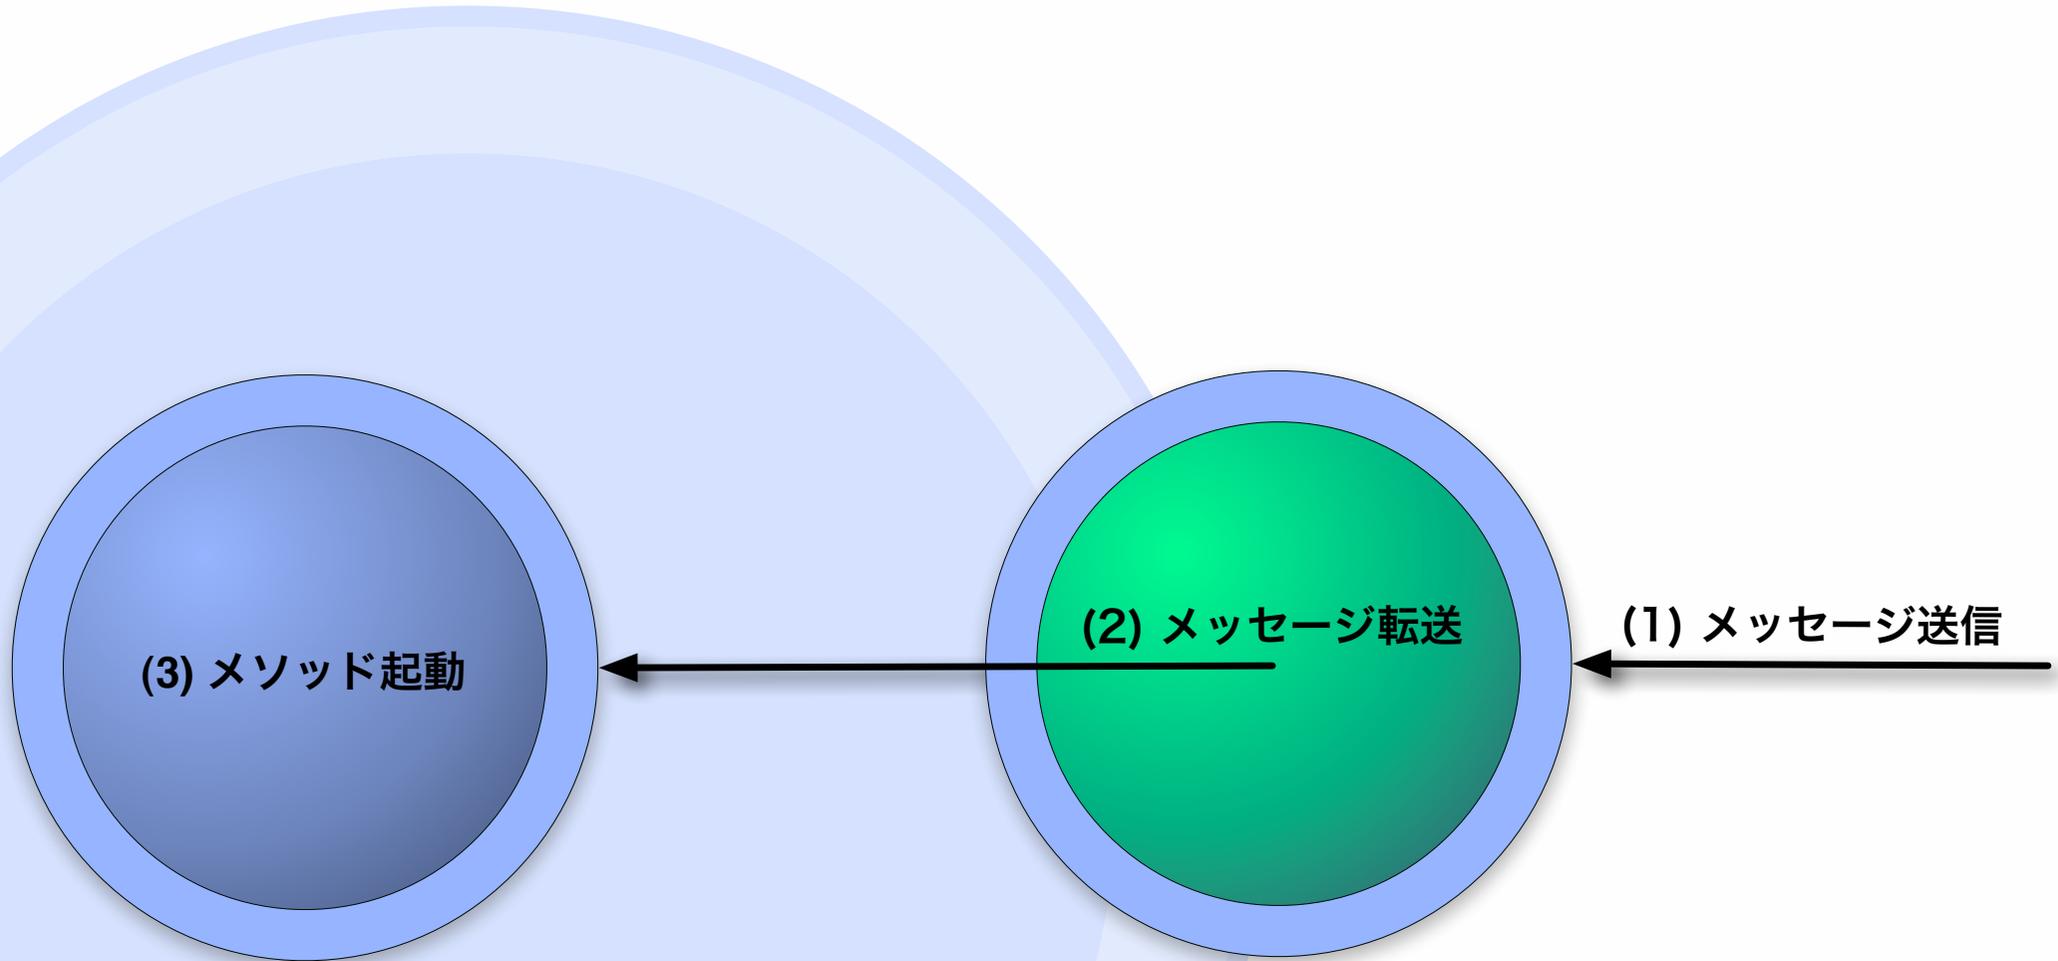


(2) メソッド起動

(1) メッセージ送信



● ● ● dRuby



○ ● ● Ruby風

- メソッド呼出しに特化したおかげで徹底したRuby風な振る舞いができたよ
- インターフェイスの定義が不要
- イテレータ
- スレッド同期メカニズム

●●● オブジェクトの交換

- それでもRuby風になりきれない部分
- 交換するのはコピーか参照か
 - 引数、戻り値、例外
- Rubyでは参照のやりとりしかないよ

オブジェクトの交換

```
def hoge(a1, a2, a3, &block)
  ...
  raise FooError unless some_cond
  ...
  return value
end
```

● ● ● 受け取ったものは何か

- 意識したくない
 - Rubyだし
- リモートのオブジェクトならdRubyがRMIしてくれる
- 複製注意
 - でもOOPなら複製しないよな..

○ ● ● 送るものは何か

- 意識したくない
 - Rubyだし
- すべてを参照にすべきか？
 - Rubyだからね
- でもいつか値が必要になる
- 値を送るか、参照を送るか

○ ● ● dRubyの戦略

- オブジェクトの交換方法を勝手に選択
 - Marshal可能かどうか条件



こんなの

Marshal

dumpable

Number, String, Boolean...

pass by value

can't dump

Proc, Thread, File ...

pass by reference

本当にこの選択でよいのか

- わかんないけどわりと実用的
- 気に入らなければ明示的に参照を選ぶ
- 手間だけど値渡しを選ぶこともできる
- カスタマイズしたdumpが必要



残念!

- セキュリティ

- Rubyの可視性を模倣している程度

- 悪意は防げない

- 分散GC

- dRubyはあくまでRubyのメソッド呼出しである



実験

- memcachedが流行ってるらしい
- Hashを一つ公開する実験
- 二つの端末でirbを動かしてみるよ

●●● はい準備です

● 今すぐTerminal.appを開いて!

○ 早く早く!!

irbを二つ用意してね

```
% irb --simple-prompt --noredline
```

OSXの
readlineでスレッドがス
イッチしない人向け

```
% irb --simple-prompt --noredline
```

DRb.start_service

参加者はまず
DRb.start_service

```
>> require 'drb/drb'  
>> DRb.start_service  
=> .....
```

front

frontは入り口になる
オブジェクト

URIと結びつける

```
>> require 'drb/drb'  
>> front = {}  
>> DRb.start_service('druby://localhost:12345', front)
```

DRbObject

URIを与えて明示的に
参照を作る

```
>> ro = DRbObject.new_with_uri('druby://localhost:12345')
=> #<DRb::DRbObject:0x.....>
>> ro.keys
=> []
>> ro.class
=> DRb::DRbObject
>> ENV.each {|k, v| ro[k] = v}
=> {...}
>> ro.keys
=> ["MANPATH", "CVS_RSH", "USER", ...]
```

DRbObject

```
>> ro = DRbObject.new_with_data({'12345'})
=> #<DRb::DRbObject:0x...>
>> ro.keys
=> []
>> ro.class
=> DRb::DRbObject
>> ENV.each {|k, v| ro[k] = v}
=> {...}
>> ro.keys
=> ["MANPATH", "CVS_RSH", "USER", ...]
```

classはDRbObject

DRbObject

```
>> ro = DRbObject.new_with_uri('druby://localhost:12345')
=> #<DRb::DRbObject:0x.....>
>> ro.keys
=> []
>> ro.class
=> DRb::DRbObject
>> ENV.each {|k, v| ro[k] = v}
=> {...}
>> ro.keys
=> ["MANPATH", "CVS_RSH", "USER", ...]
```

でもHashっぽい

ほんとう？

こっちのHashに入っ
てる!

```
>> front.keys  
=> ["MANPATH", "CVS_RSH", "USER", ...]
```

暗黙的な参照

IOを入れてみるよ

```
>> front['stdout'] = $stdout  
=> #<IO:0x130d80>
```

```
>> ro['stdout']  
=> #<DRb::DRbObject:... @uri="druby://localhost:12345">  
>> ro['stdout'].puts('Hello, World.')
```

```
>> Hello, World.
```

暗黙的な参照

```
>> front['stdout'] = $stdout  
=> #<IO:0x130d80>
```

リモートからは
DRbObjectに見える

```
>> ro['stdout']  
=> #<DRb::DRbObject:... @uri="druby://localhost:12345">  
>> ro['stdout'].puts('Hello, World.')
```

```
>> Hello, World.
```

暗黙的な参照

```
>> front['stdout'] = $stdout  
=> #<IO:0x130d80>
```

せっかくなのでputsだ

```
>> ro['stdout']  
=> #<DRb::DRbObject:... @uri="druby://localhost:12345">  
>> ro['stdout'].puts('Hello, World.')
```

```
>> Hello, World.
```

暗黙的な参照

```
>> front['stdout'] = $stdout  
=> #<IO:0x130d80>
```

```
>> ro['stdout']  
=> #<DRb::DRbObject:0x130d80@host:12345">  
>> ro['stdout'].p
```

拍手ね。

```
>> Hello, World.
```



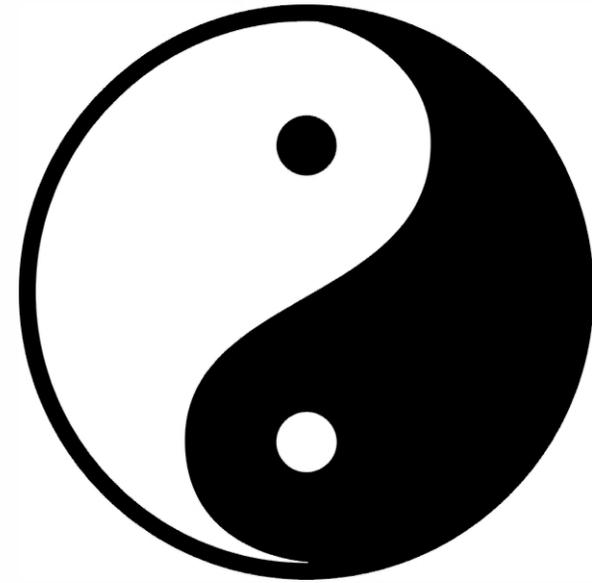
混沌

- オブジェクトがプロセスの境界を越え
混じり合う
- でもやっぱり別プロセス



混沌

- かなり強引だけどこんな感じ



● ● ● 次はスク립トで

- 典型的な生産者消費者問題を
SizedQueueで解く
- スレッド同期メカニズムもそのまま動く
- という例なので眺めてて

SizedQueue server

```
require 'thread'  
require 'drb/drb'  
  
queue = SizedQueue.new(10)  
DRb.start_service('druby://localhost:12345', queue)  
sleep
```

Producer

```
require 'drb/drb'  
  
DRb.start_service  
queue = DRbObject.new_with_uri('druby://localhost:12345')  
  
100.times do |n|  
  sleep(rand)  
  queue.push(n)  
end
```

Consumer

```
require 'drb/drb'  
  
DRb.start_service  
queue = DRbObject.new_with_uri('druby://localhost:12345')  
  
100.times do  
  sleep(rand)  
  puts queue.pop  
end
```

● ● ● 地味だ

- あんまりおもしろくないですね
- dRubyのスク립トはあんなもん



実装

- 初版を読もう

- `ruby-list:15406`

- 200lines

- いろいろバグってるが細かいことは気にしない

DRbObject

URIとobject_idで
できてる

```
class DRbObject
  def initialize(obj, uri=nil)
    @uri = uri || DRb.uri
    @ref = obj.id if obj
  end

  def method_missing(msg_id, *a)
    succ, result = DRbConn.new(@uri).send_message(self, msg_id, *a)
    raise result if ! succ
    result
  end

  attr :ref
end
```

DRbObject

- URIはインタプリタを特定し
- refはオブジェクトを特定する

method_missing

```
class DRbObject
  def initialize(obj, uri=nil)
    @uri = uri || DRb.uri
    @ref = obj.id if obj
  end

  def method_missing(msg_id, *a)
    succ, result = DRbConn.new(@uri).send_message(self, msg_id, *a)
    raise result if ! succ
    result
  end

  attr :ref
end
```

serverに

Marshalしたメッセージと
引数を転送する

DRbServer

```
def proc
  ns = @soc.accept
  Thread.start do
    begin
      s = ns
      begin
        ro, msg, argv = recv_request(s)
        if ro and ro.ref
          obj = ObjectSpace._id2ref(ro.ref)
        else
          obj = DRb.front
        end
        result = obj.__send__(msg.intern, *argv)
        succ = true
      rescue
        result = $!
        succ = false
      end
      send_reply(s, succ, result)
    ensure
      close s if s
    end
  end
end
```

acceptしたら処理は
別スレッドで。

```
def proc
  ns = @soc.accept
  Thread.start do
    begin
      s = ns
      begin
        ro, msg, argv = recv_request(s)
        if ro and ro.ref
          obj = ObjectSpace._id2ref(ro.ref)
        else
          obj = DRb.front
        end
        result = obj.__send__(msg.intern, *argv)
        succ = true
      rescue
        result = $!
        succ = false
      end
      send_reply(s, succ, result)
    ensure
      close s if s
    end
  end
end
```

DRbServer

```
def proc
  ns = @soc.accept
  Thread.start do
    begin
      s = ns
      begin
        ro, msg, argv = recv_request(s)
        if ro and ro.ref
          obj = ObjectSpace._id2ref(ro.ref)
        else
          obj = DRb.front
        end
        result = obj.__send__(msg.intern, *argv)
        succ = true
      rescue
        result = $!
        succ = false
      end
      send_reply(s, succ, result)
    ensure
      close s if s
    end
  end
end
```

参照からObjectへ変換

DRbServer

```
def proc
  ns = @soc.accept
  Thread.start do
    begin
      s = ns
      begin
        ro, msg, argv = recv_request(s)
        if ro and ro.ref
          obj = ObjectSpace._id2ref(ro.ref)
        else
          obj = DRb.front
        end
        result = obj.__send__(msg.intern, *argv)
        succ = true
      rescue
        result = $!
        succ = false
      end
      send_reply(s, succ, result)
    ensure
      close s if s
    end
  end
end
```

オブジェクトのメソッド
を起動する

DRbServer

```
def proc
  ns = @soc.accept
  Thread.start do
    begin
      s = ns
      begin
        ro, msg, argv = recv_request(s)
        if ro and ro.ref
          obj = ObjectSpace._id2ref(ro.ref)
        else
          obj = DRb.front
        end
        result = obj.__send__(msg.intern, *argv)
        succ = true
      rescue
        result = $!
        succ = false
      end
      send_reply(s, succ, result)
    ensure
      close s if s
    end
  end
end
```



結果を返信

DRbServer

```
def proc
  ns = @soc.accept
  Thread.start do
    begin
      s = ns
      begin
        ro, msg, argv = recv_request(s)
        if ro and ro.ref
          obj = ObjectSpace._id2ref(ro.ref)
        else
          obj = DRb.front
        end
        result = obj.__send__(msg.intern, *argv)
        succ = true
      rescue
        result = $!
        succ = false
      end
      send_reply(s, succ, result)
    ensure
      close s if s
    end
  end
end
```

あ。
間違い..

オブジェクトの交換

```
def dump(obj, soc)
  begin
    str = Marshal::dump(obj)
  rescue
    ro = DRbObject.new(obj)
    str = Marshal::dump(ro)
  end
  soc.write(str) if soc
  return str
end
```

dumpに失敗したら
DRbObjectをdumpする

○ ● ● dRuby作成入門でした

- dRubyを書けるようになりましたね

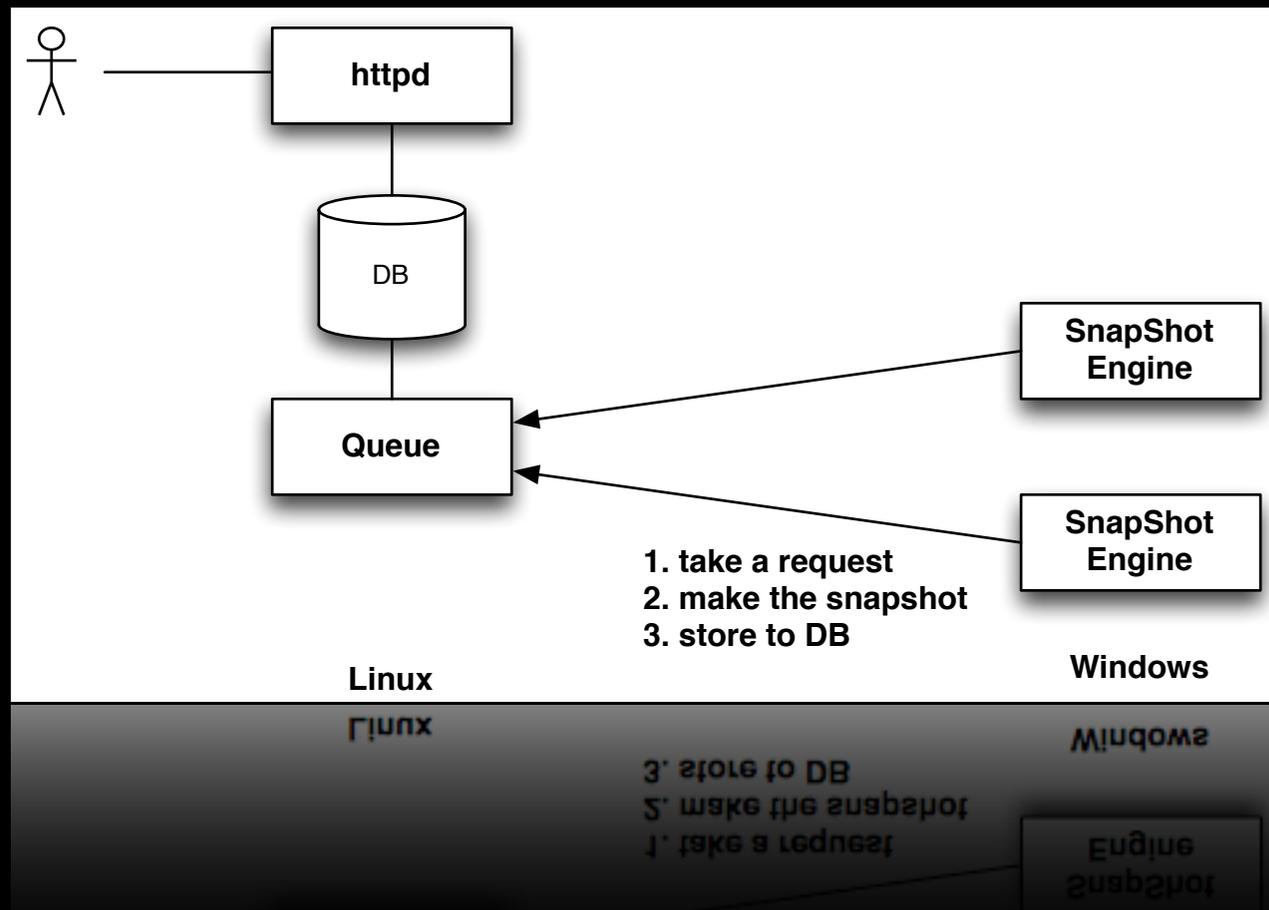


応用

- たぶんたくさんありそう
- IJPPの論文で引用したものから

hatena screenshot

id:secondlife @ RK06



RWiki

- In-Memory Database with logging
 - すべてメモリの中にあるWiki
 - Objectが潜んでいる
- 25000ページくらいのWikiは実績あるよ
- dRubyやERBなどのサンプル

WEBrick::CGI

- みんな使ってるよね
- WEBrick::CGIをfrontにしたサービス

長生きCGI

```
require 'webrick/cgi'
require 'drb/drb'
require 'thread'

class SimpleCountCGI < WEBrick::CGI
  def initialize
    super
    @count = Queue.new
    @count.push(1)
  end

  def count
    value = @count.pop
    ensure
      @count.push(value + 1)
    end
  end

  def do_GET(req, res)
    res['content-type'] = 'text/plain'
    res.body = count.to_s
  end
end

DRb.start_service('druby://localhost:12321', SimpleCountCGI.new)
sleep
```

cgi.start()

```
module WEBrick
  class CGI
    ....
    def start(env=ENV, stdin=$stdin, stdout=$stdout)
    ....
  end
end
```

4 lines CGI

```
#!/usr/local/bin/ruby  
  
require 'drb/drbr'  
  
DRb.start_service('druby://localhost:0')  
ro = DRbObject.new_with_uri('druby://localhost:12321')  
ro.start(ENV.to_hash, $stdin, $stdout)
```

さよなら遺言モデル

- 状態の永続化とか排他制御とか
- 遺言とか相続問題とか
- 死ななきゃ良いのにね

○ ● ● RindaとLinda

- Lindaは並列処理糊言語
- タプルとタプルスペース
- パターンマッチング
- Lindaを作った人は

author of Linda

- Reviewer#1の人
- 鋭い指摘でずいぶん詳しいと思ったら作者だった
- 余剰資金のある研究室は
Must buy!

Int J Parallel Prog
DOI 10.1007/s10766-008-0086-1

dRuby and Rinda: Implementation and Application of Distributed Ruby and its Parallel Coordination Mechanism

Masatoshi Seki

Received: 20 March 2008 / Accepted: 13 August 2008
© Springer Science+Business Media, LLC 2008

Abstract The object-oriented scripting language Ruby is admired by many programmers for being easy to write in, and for its flexible, dynamic nature. In the last few years, the Ruby on Rails web application framework, popular for its productivity benefits, has brought about a renewed attention to Ruby for enterprise use. As the focus of Ruby has broadened from small tools and scripts, to large applications, the demands on Ruby's distributed object environment have also increased, as has the need for information about its usage, performance and examples of common practices. dRuby and Rinda were developed by the author as the distributed object environment and shared tuplespace implementation for the Ruby language, and are included as part of Ruby's standard library. dRuby extends method calls across the network while retaining the benefits of Ruby. Rinda builds on dRuby to bring the functionality of Linda, the glue language for distributed co-ordination systems, to Ruby. This article discusses the design policy and implementation points of these two systems, and demonstrates their simplicity with sample code and examples of their usage in actual applications. In addition to dRuby and Rinda's appropriateness for prototyping distributed systems, this article will also demonstrate that dRuby and Rinda are building a reputation for being suitable infrastructure components for real-world applications.

Keywords Ruby · Linda · dRuby · Rinda · TupleSpace · RMI

M. Seki (✉)
Tochigi, Japan
e-mail: m_seki@mva.biglobe.ne.jp
URL: www.druby.org

 Springer

 Springer

Int J Parallel Prog
DOI 10.1007/s10766-008-0086-1
© Springer Science+Business Media, LLC 2008

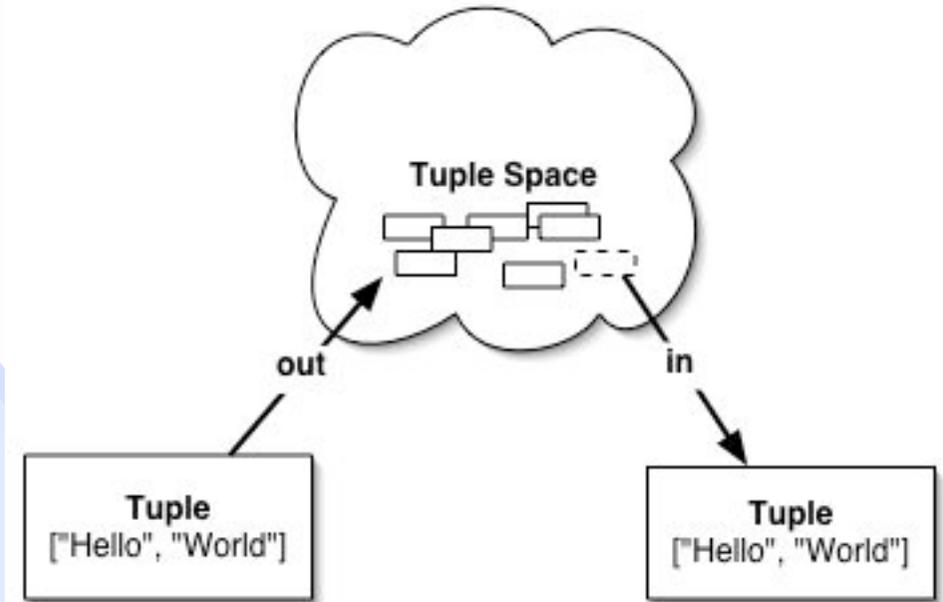
Lindaの協調

● out

● inとrd

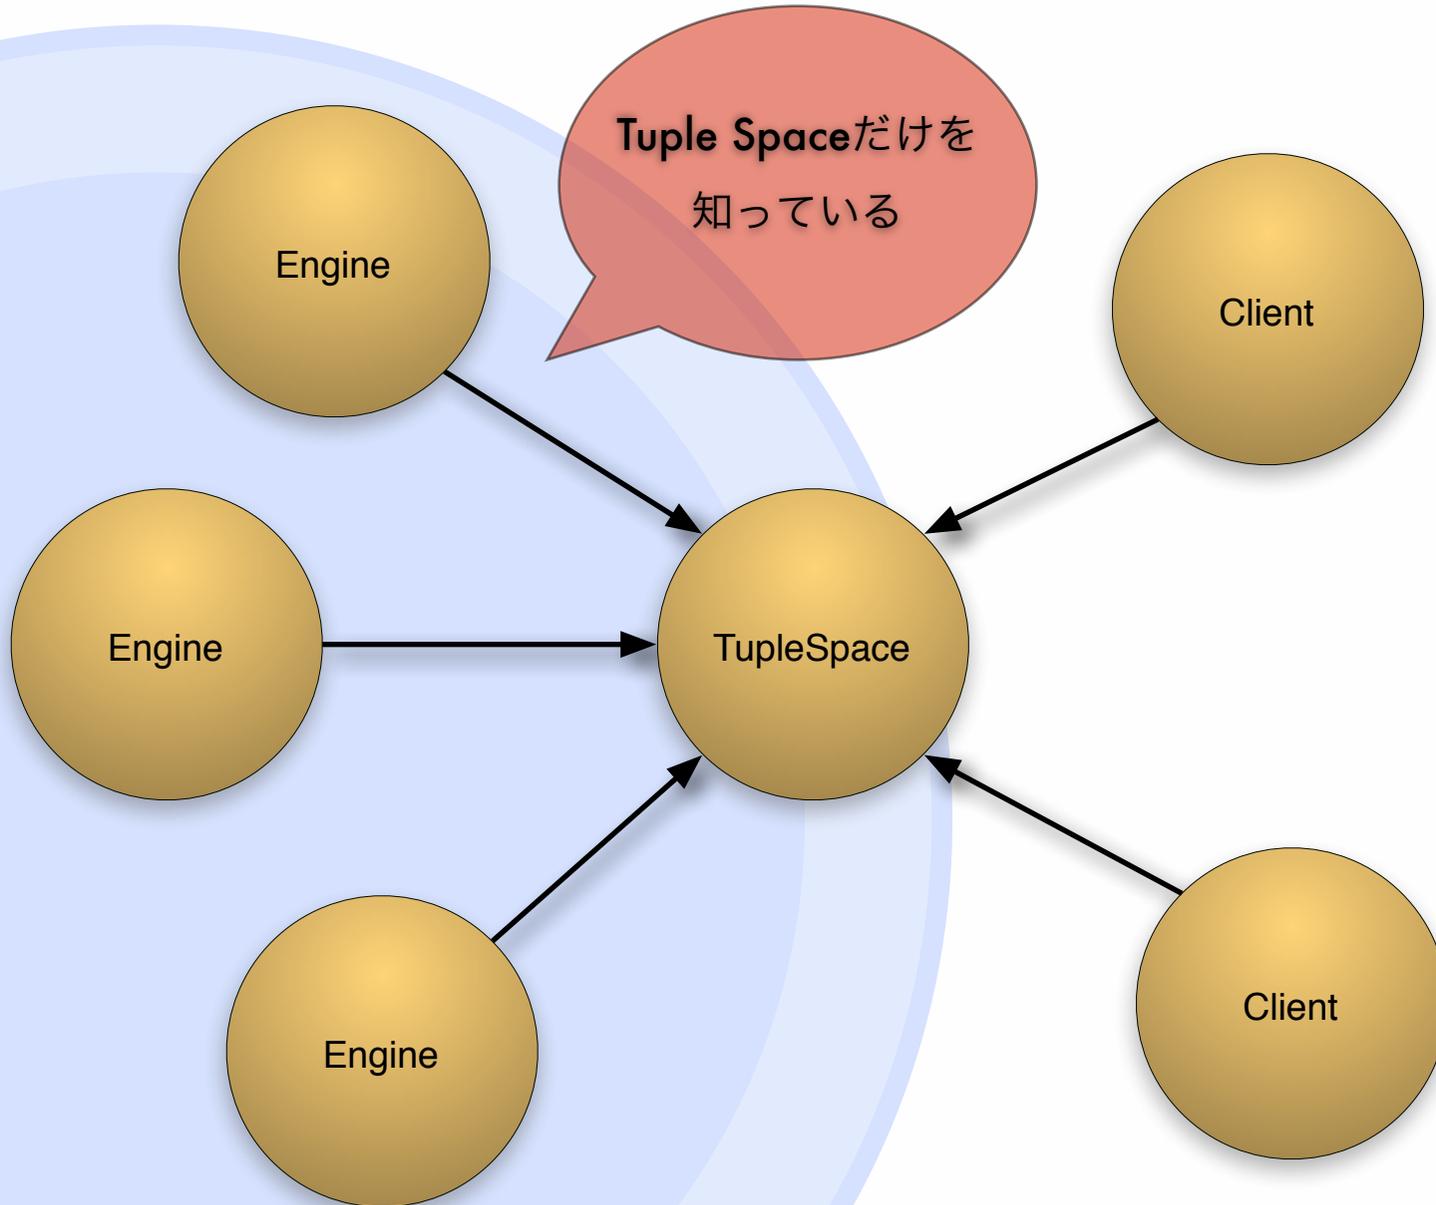
○ パターンで指定

○ ブロックできる





かっこいい並列処理



○ ● ● Rinda

- RubyによるLindaの実装
- dRubyを意識してくれる
- TupleはArray | Hashで

Tuple

- Arrayで表現する
 - Hashはやめたほうがいい
- 要素ごとにMarshalする
 - Marshalできない要素も混ぜられる
 - dRubyを意識

Tuple

```
[:chopstick, 2]  
[:room_ticket]  
['abc', 2, 5]  
[:matrix, 1.6, 3.14]  
['family', 'is-sister', 'Carolyn', 'Elinor']
```

Pattern

- Tupleの要素と===で比較する
 - case equalsなのでパターンっぽく動く
 - Regexp, Range, Classとか
- nilはワイルドカード

Pattern

```
[/^A/, nil, nil]  
[:matrix, Numeric, Numeric]  
['family', 'is-sister', 'Carolyn', nil]  
[nil, 'age', (0..18)]
```

['seki', 'age', 20]

```
>> require 'rinda/tuplespace'
>> ts = Rinda::TupleSpace.new
>> ts.write(['seki', 'age', 20])
>> ts.write(['sougo', 'age', 18])
>> ts.write(['leonard', 'age', 18])
>> ts.read_all([nil, 'age', 0..19])
=> [{"sougo", "age", 18}, {"leonard", "age", 18}]
>> ts.read_all([/^s/, 'age', Numeric])
=> [{"seki", "age", 20}, {"sougo", "age", 18}]
```

○ ● ● TupleSpace

```
>> ts = Rinda::TupleSpace.new
=> #<Rinda::TupleSpace:...>
>> DRb.front['ts'] = ts
=> #<Rinda::TupleSpace:...>
>>
```

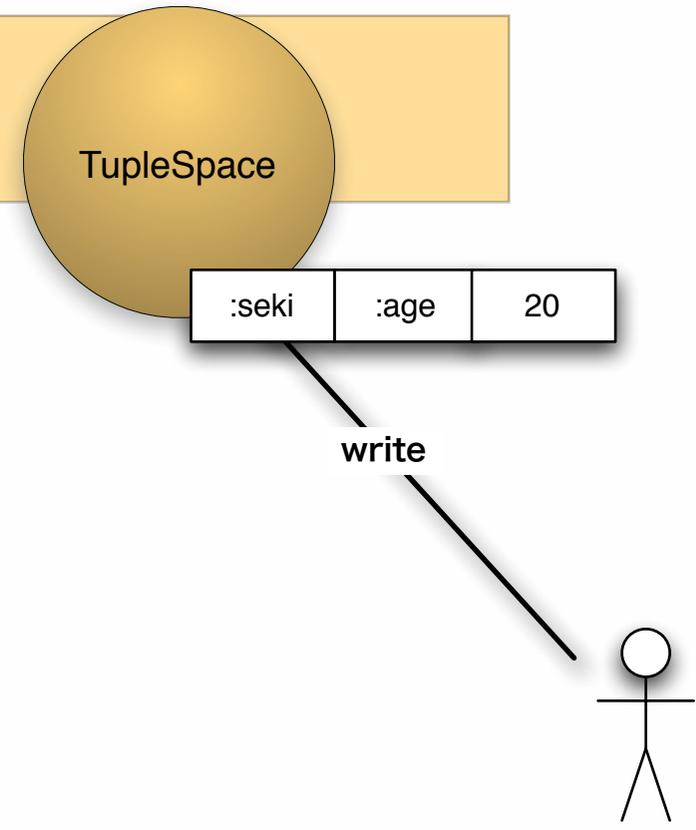
```
>> it = DRbObject.new_with_uri('druby://localhost:12345?
ts')
=> #<DRb::DRbObject:...>
>> ts = Rinda::TupleSpaceProxy.new(it)
=> #<Rinda::TupleSpaceProxy:...>
>>
```

TupleSpaceProxyは
dRuby下で安全に動作
させるためのもの



write & take

```
>> ts.write([:seki, :age, 20])  
=> #<Rinda::TupleEntry:...>  
>>
```

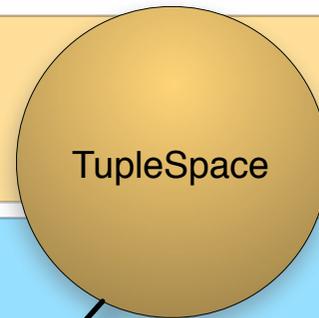




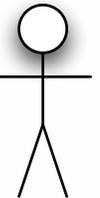
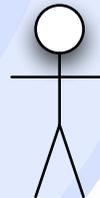
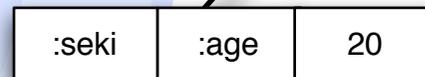
write & take

```
>> ts.write([:seki, :age, 20])  
=> #<Rinda::TupleEntry:...>  
>>
```

```
>> ts.take([:seki, :age, nil])  
=> [:seki, :age, 20]  
>>
```



take

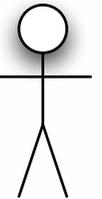
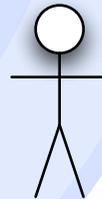


● ● ● 待ち合わせ

```
>> ts.take([:seki, :age, nil])
```

TupleSpace

take





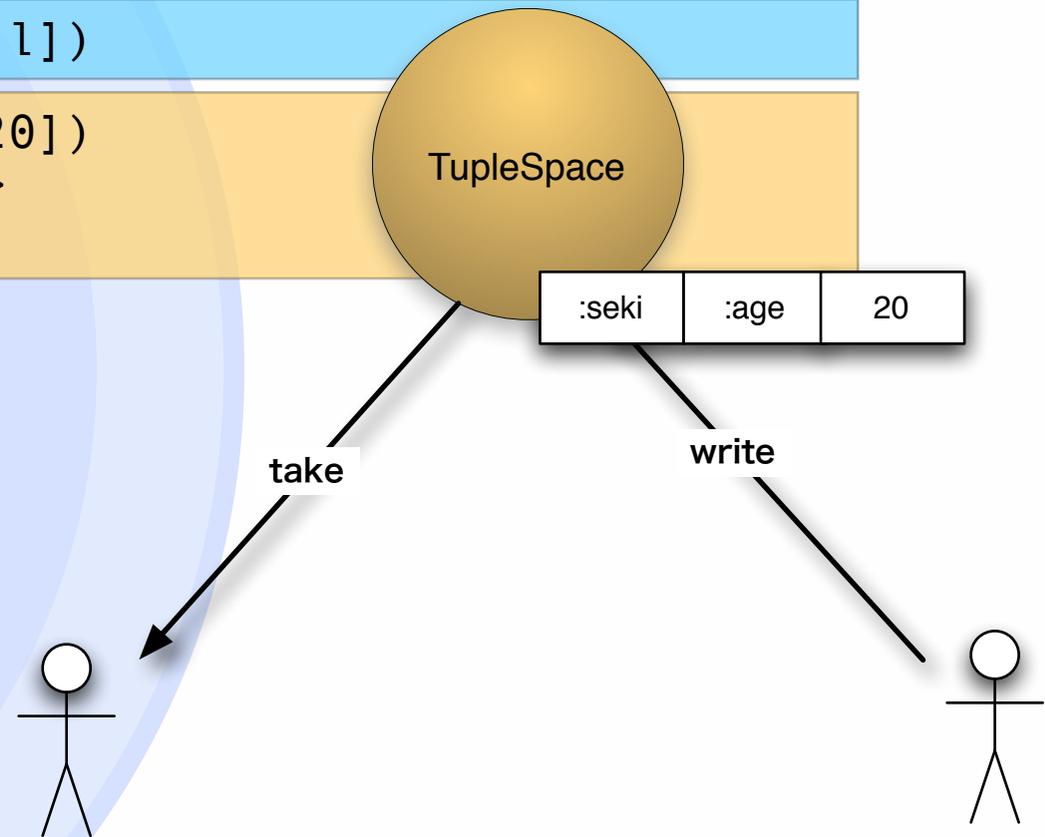
待ち合わせ

```
>> ts.take([:seki, :age, nil])
```

```
>> ts.write([:seki, :age, 20])
```

```
=> #<Rinda::TupleEntry:...>
```

```
>>
```



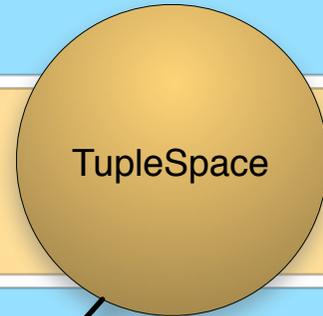


待ち合わせ

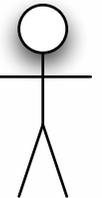
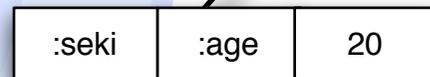
```
>> ts.take([:seki, :age, nil])
```

```
>> ts.write([:seki, :age, 20])  
=> #<Rinda::TupleEntry:...>  
>>
```

```
=> [:seki, :age, 20]  
>>
```



take



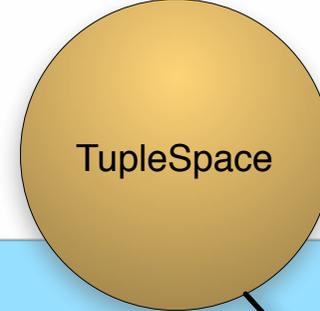
● ● ● 安全なカウンタ

- 値のタプルを資源と考える
- 同時にただ一つのプロセスだけが値を更新できる

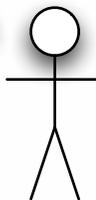
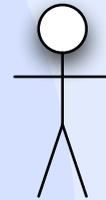


安全なカウンタ

```
>> tuple = ts.take([:count, nil])
```



take

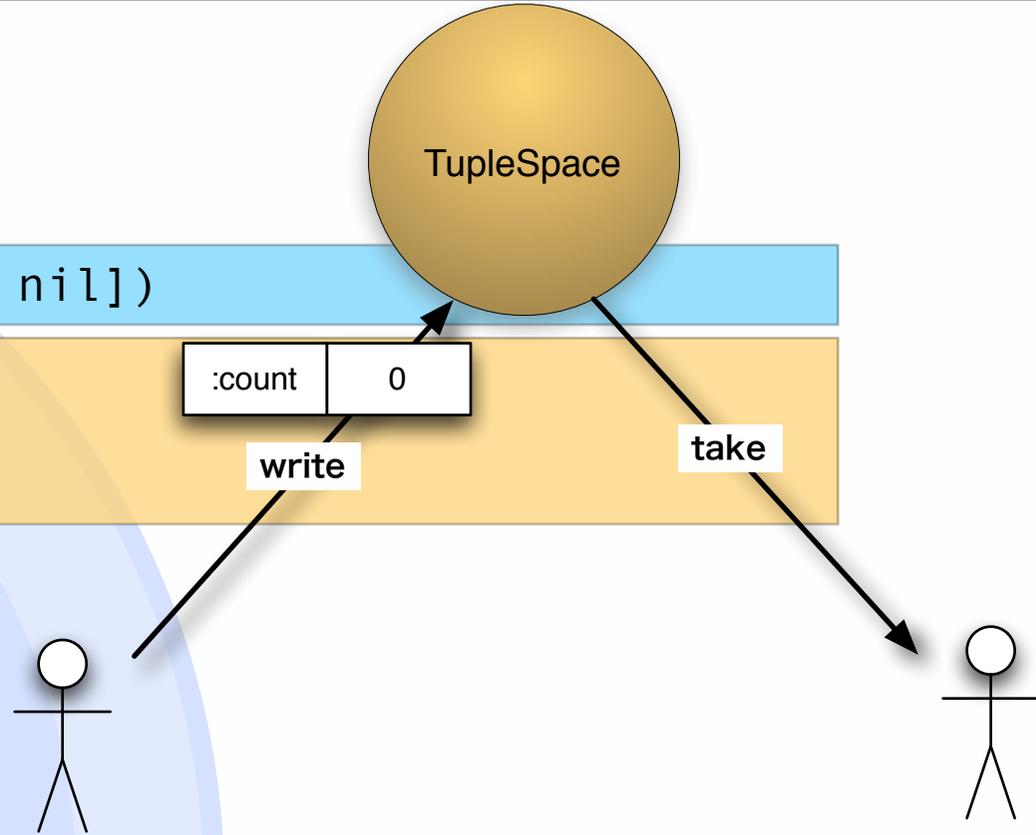




安全なカウンタ

```
>> tuple = ts.take([:count, nil])
```

```
>> ts.write([:count, 0])  
=> #<Rinda::TupleEntry:...>  
>>
```



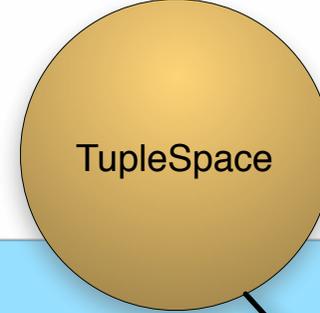


安全なカウンタ

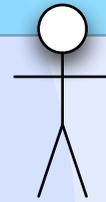
```
>> tuple = ts.take([:count, nil])
```

```
>> ts.write([:count, 0])  
=> #<Rinda::TupleEntry:...>  
>>
```

```
=> [:count, 0]  
>>
```



take



安全なカウンタ

```
>> tuple = ts.take([:count, nil])
```

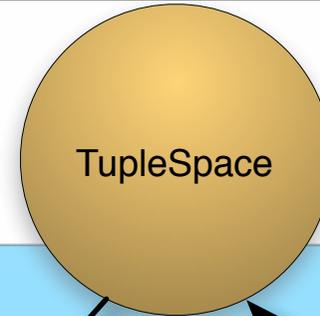
```
>> ts.write([:count, 0])  
=> #<Rinda::TupleEntry:...>  
>>
```

```
=> [:count, 0]  
>>
```

```
>> tuple = ts.take([:count, nil])
```

```
>> tuple[1] += 1  
=> 1  
>> ts.write(tuple)  
=> #<DRb::DRbObject:...>  
>>
```

```
=> [:count, 1]  
>>
```



不公平な最適化

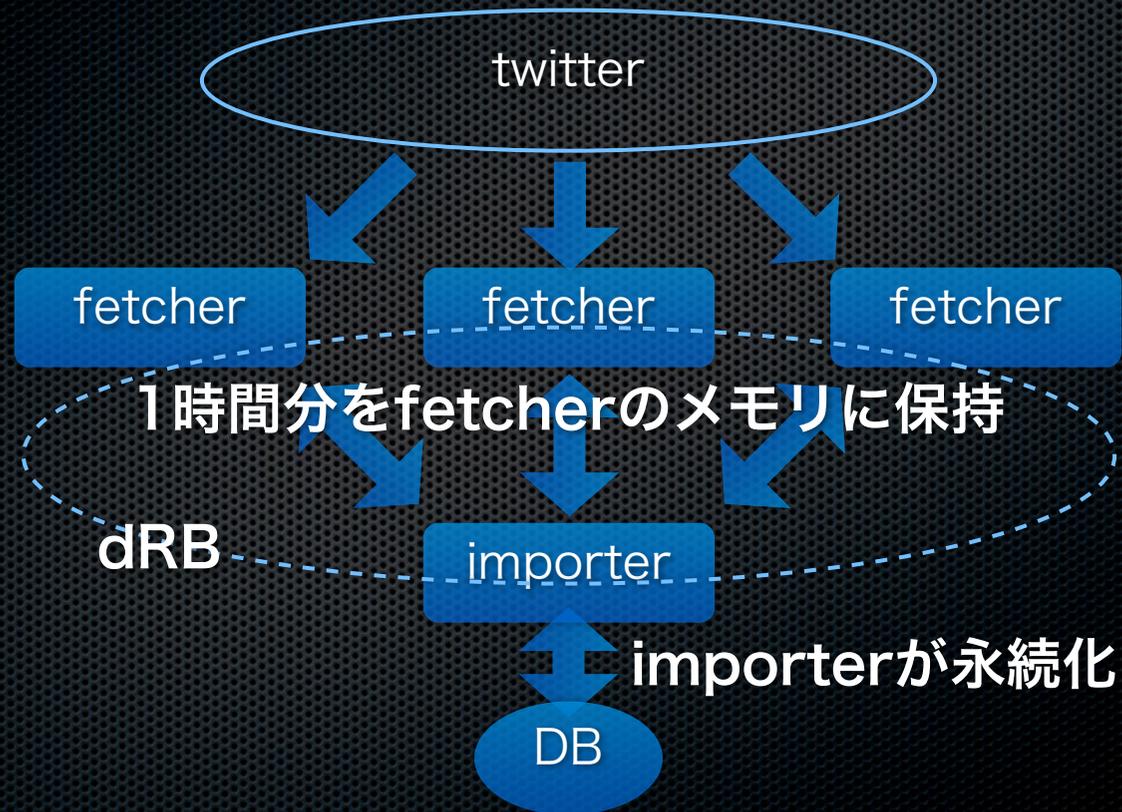
- 先頭の要素がSymbolである場合に検索が速く動作するようになってるよ!
- そんなにたくさんTupleを投入してるとしたらなにか間違ってると思うけどね

Application of Rinda

- Sapporo といえは

Buzztter !

分散クローラ w/dRB, Rinda



Buzztter !

コメントを書く

コメントを削除する

  m_seki

2007/12/20 01:36

http://www.druby.org/imaco_doc/ijpp_text.htmlの最後の方にBuzztterの紹介を書き始めました。うそ、まちがい、これは書いてくれないと困る！などのことがありましたら教えてください。

 darashi

2007/12/20 21:28

紹介していただいてどうもありがとうございます。論文にbuzztterのようなちょっと遊びっぽい応用例が載って良いのかすこし不安になってしまいますが。。うそ、まちがいは無いと思います。ほかに何か必要な情報がありましたらお知らせください。ちなみに、buzztterではダブルスペースの冗長性を確保するために、ダブルスペースをfetcher側に持たせています。importerがRingを使って適時ネットワーク上にある複数のダブルスペース(fetcher)を探し、そこからデータを集めてくるようにしています。dRubyとってもステキです。ありがとうございます！

  m_seki

2008/01/11 08:58

資料のタイトルの英語表記はありませんか？

  m_seki

2008/01/11 09:13

とりあえずInside Buzztterとしました。

 darashi

2008/01/11 12:13

特にないので、Inside Buzztterが良いと思います。僕が真っ先に思いついたのもそれでした。

特にないので、Inside Buzztterが良いと思います。僕が真っ先に思いついたのもそれでした。

 darashi

2008/01/11 15:13

特にないので、Inside Buzztterが良いと思います。

Buzztter !

6.4 Application of Rinda

Buzztter is a Web service that processes Twitter sentences. Twitter is a SNS⁶ specialized for short sentences. Buzztter collects sentences posted into Twitter, and figures out which words are used more often than usual. By doing so, Buzztter understands the overall trend of words usage at a given moment in Twitter.

Buzztter is composed of several subsystems. One of these is a distributed crawler sub-system that uses Twitter's API (HTTP based) to collect sentences. The crawler subsystem has multiple fetchers that fetch information from Twitter and importers to make it persistent. Rinda and dRuby are used as mediators between the fetchers and the importers. For reference, the following indicates the data rates handled by Buzztter (as of Nov 3, 2007)

- 72MB per day
- 125000 transaction per day

- 125000 transaction per day
- 72MB per day

data rates handled by Buzztter (as of Nov 3, 2007)

between the fetchers and the importers. For reference, the following indicates the

ここで発表者が

id:darashiと交代します

ここで発表者が
id:darashiと交代します

- 嘘です

○●● 重要なことをもう一度

● もう一度

¥ 3360

- 初刷まだ買えます



Who is translating it?

- 英語版はまだ買えません



\$32.00

- International Journal of PARALLEL PROGRAMING
- コレクターズアイテム
- 余剰資金のある研究室は Must buy!

Int J Parallel Prog
DOI 10.1007/s10766-008-0086-1

dRuby and Rinda: Implementation and Application of Distributed Ruby and its Parallel Coordination Mechanism

Masatoshi Seki

Received: 20 March 2008 / Accepted: 13 August 2008
© Springer Science+Business Media, LLC 2008

Abstract The object-oriented scripting language Ruby is admired by many programmers for being easy to write in, and for its flexible, dynamic nature. In the last few years, the Ruby on Rails web application framework, popular for its productivity benefits, has brought about a renewed attention to Ruby for enterprise use. As the focus of Ruby has broadened from small tools and scripts, to large applications, the demands on Ruby's distributed object environment have also increased, as has the need for information about its usage, performance and examples of common practices. dRuby and Rinda were developed by the author as the distributed object environment and shared tuplespace implementation for the Ruby language, and are included as part of Ruby's standard library. dRuby extends method calls across the network while retaining the benefits of Ruby. Rinda builds on dRuby to bring the functionality of Linda, the glue language for distributed co-ordination systems, to Ruby. This article discusses the design policy and implementation points of these two systems, and demonstrates their simplicity with sample code and examples of their usage in actual applications. In addition to dRuby and Rinda's appropriateness for prototyping distributed systems, this article will also demonstrate that dRuby and Rinda are building a reputation for being suitable infrastructure components for real-world applications.

Keywords Ruby · Linda · dRuby · Rinda · TupleSpace · RMI

M. Seki (✉)
Tochigi, Japan
e-mail: m_seki@mva.biglobe.ne.jp
URL: www.druby.org

 Springer

 Springer

Print ISSN 1076-6545
Online ISSN 1076-6553
Copyright © 2008 Springer
08/08/08

<http://dx.doi.org/10.1007/s10766-008-0086-1>

まとめ

- まだ初刷買えます。
- 余剰資金のある研究室はMust buy!
- dRubyの作り方入門
- そしてBuzztter.