#### o dRuby and Rinda.

その実装と応用を(再び)関西で

#### 念のため

先月の札幌Ruby会議01を見ちゃった 人いますか?

### ○● 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.

 $\textbf{Keywords} \quad Ruby \cdot Linda \cdot dRuby \cdot Rinda \cdot TupleSpace \cdot RMI$ 

M. Seki (⊠) Tochigi, Japan e-mail: m\_seki@mva.biglobe.ne.jp URL: www.druby.org





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

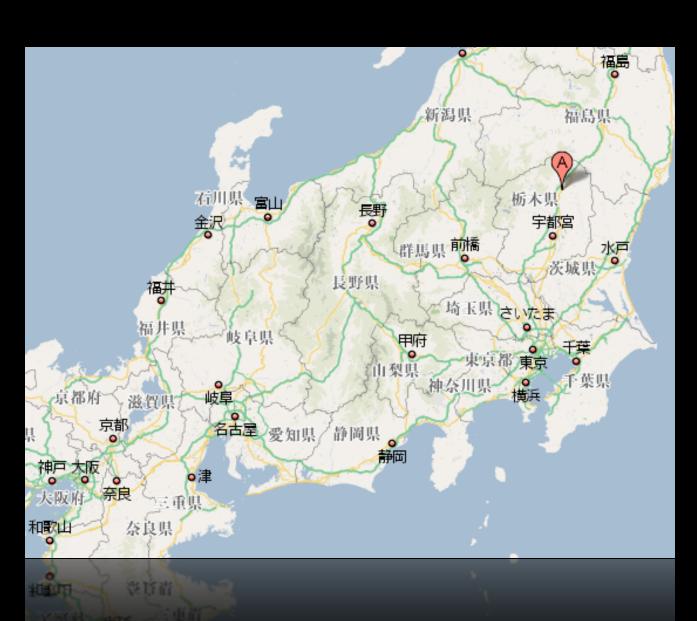
#### 私について

- toRubyからきました
  - 池澤さんのLTで一躍有名に

#### 会場



#### 会場



# 開催日

● 毎月第一水曜

# 。。dRubyによる分散オ ブジェクト環境

Linux Conference 2000 fall Perl/Ruby Conference 2000.11.29-12/1 国立京都国際会館

# o dRubyお披露目

- この京都のイベントがはじまり
  - 以来、進歩のないネタをずっと・・・
- いつかRuby関西へ!

#### Hello, Again.

- ついに関西Ruby会議01
  - 実は京都でやるもんだと勘違いしてた..

# o • Agenda

- Rubyの紹介
- dRubyの紹介
- 入門
- 応用

# o 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

### ○ **ブは分散のブ**

- D is for Distributed
- dRubyは分散なRubyである

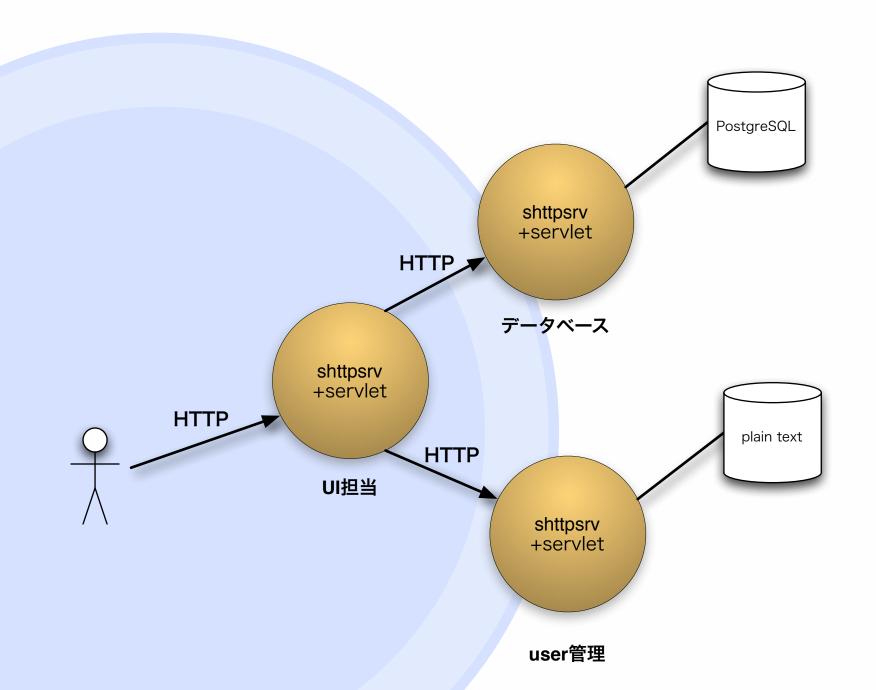
# Remote Method Invocation

- 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..

#### ○●原型はこんなのだった



#### 2 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

- dRubyはRubyのRMIである
  - 数多あるRMIのRuby版ではない

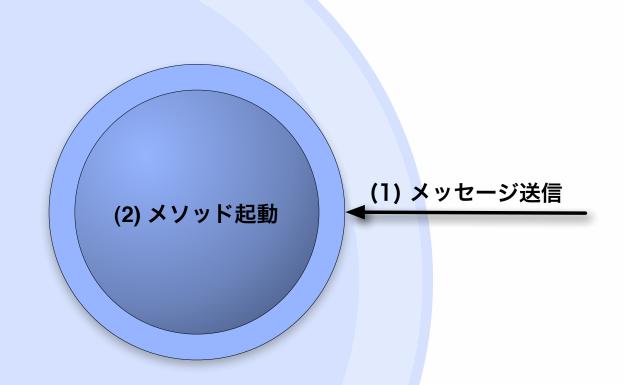
# dRubyの特徴

- 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

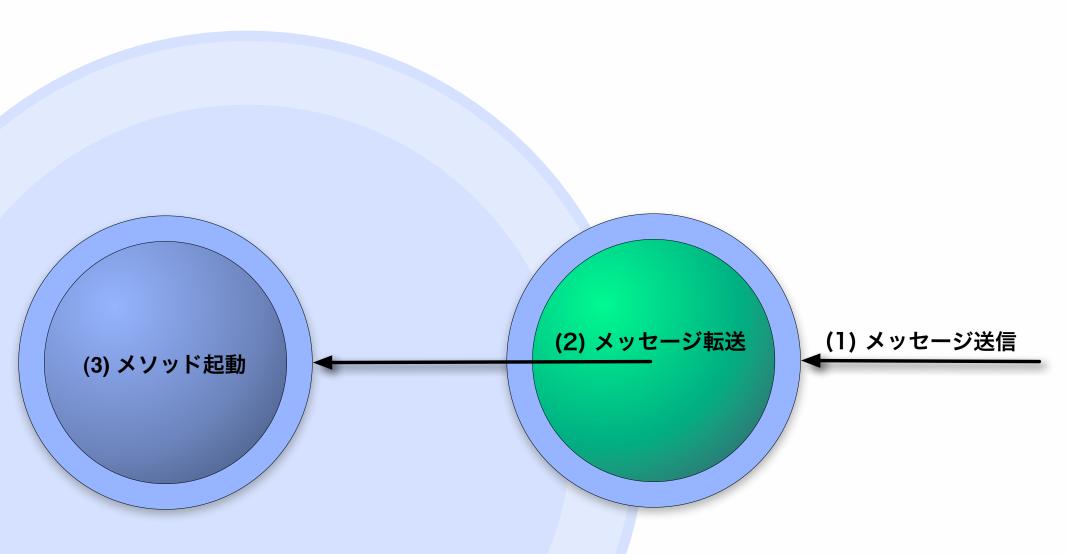
#### **メッセージング**

- 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

# o Ruby



# o dRuby



### O • 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
  - 一引数、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
```

#### 。<br/> 受け取ったものは何か

- 意識したくない
  - O 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・・

### 。<br/>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

# ○・値を送るべきか、それ

# とも参照か

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

# こんなの

M	a	rs	h	al
/ Y \	u	1 3		ЧI

dumpable

Number, String, Boolean...

can't dump

Proc, Thread, File ...

pass by value

pass by reference

#### ○<br /> 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べ る
- **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**

# 一。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

```
OSXの
% irb --simple-prompt --noreadlin
てadlineでスレッドがス
イッチしない人向け
```

% irb --simple-prompt --noreadline

## 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)
```

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", ....]
```

```
>> ro = DRbObject.new_wi

=> #<DRb::DRbObject:0x.

>> ro.keys

=> []

>> ro.class

=> DRb::DRbObject

>> ENV.each {|k, v| ro[k] = v}

=> {....}

>> ro.keys

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

```
>> 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", ....]
```

### ほんとう?

こっちのHashに入っ てる!

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

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.
```

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

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

** Hello, World.
```

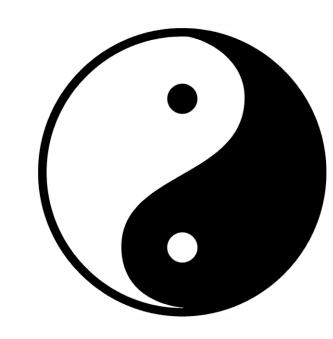
```
>> front['stdout'] = $stdout
=> #<IO:0x130d80>
>> ro['stdout']
=> #<DRb::DRbObject
>> ro['stdout'].p' 拍手ね。
>> Hello, World.
```

### 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

### 湿池

かなり強引だけどこんな感じ



### ・次はスクリプトで

- 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のスクリプトはあんなもん**

# 実装

- 初版を読もう
  - ruby-list:15406
  - **200**lines
  - いろいろバグってるが細かいことは 気にしない

URIとobject\_idで

```
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
```

- 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
```

```
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
```

```
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
          obi = 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

```
def proc
 ns = @soc.accept
  Thread.start do
    begin
     s = ns
     begin
                                               オブジェクトのメソッ
       ro, msg, argv = recv_request(s)
       if ro and ro.ref
                                                    ドを起動する
         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
```

```
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
                                       結果を返信
        result = $!
        succ = false
      end
      send_reply(s, succ, result)
    ensure
      close s if s
    end
  end
```

end

```
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, su
                           result)
    ensure
      close s if
    end
  end
end
```

# 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

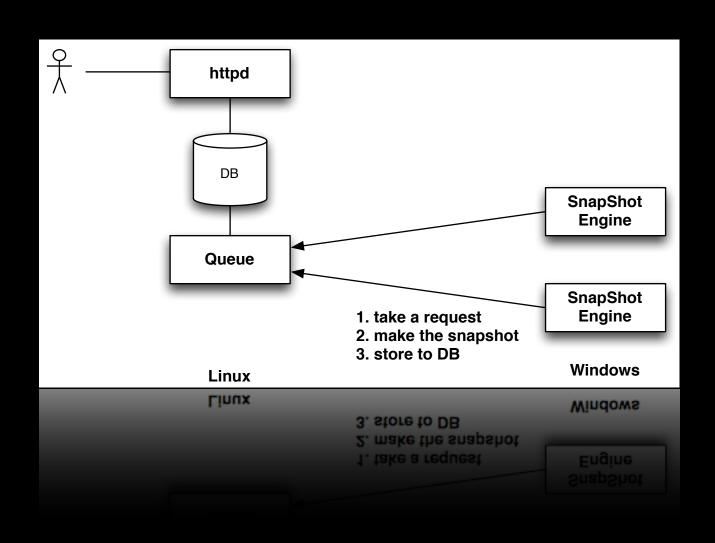
# ○ ■ dRuby作成入門でした

**■ dRubyを書けるようになりましたね** 

# 00015月

- 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なサーブレットと同じように書 けるCGI
- 環境に依って書き分ける必要がない
  - 抽象化いらないじゃん

#### 同じ値しか返さないカウンタ

```
require 'webrick/cgi'
require 'thread'
class SimpleCountCGI < WEBrick::CGI</pre>
  def initialize
    super
    @count = Queue.new
    @count.push(1)
  end
  def count
    value = @count.pop
  ensure
    @count.push(value + 1)
  end
  def do GET(req, res)
    res['content-type'] = 'text/plain'
    res.body = count.to s
  end
end
SimpleCountCGI.new.start
```

### WEBrick::CGI

- WEBrick::CGIをfrontにしたサービス
- FastCGIもどき
  - お手軽

# 長生きCGI

```
require 'webrick/cgi'
require 'drb/drb'
require 'thread'
class SimpleCountCGI < WEBrick::CGI</pre>
  def initialize
    super
    @count = Queue.new
    @count.push(1)
  end
  def count
    value = @count.pop
  ensure
    @count.push(value + 1)
  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()の実装は

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

## 4 lines CGI

```
#!/usr/local/bin/ruby
require 'drb/'drb'

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 & 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#I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

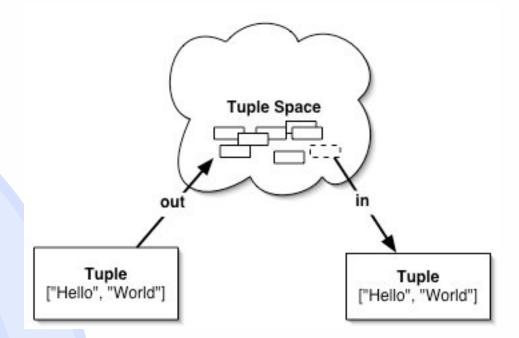
M. Seki (⊠) Tochigi, Japan e-mail: m\_seki@mva.biglobe.ne.jp URL: www.druby.org



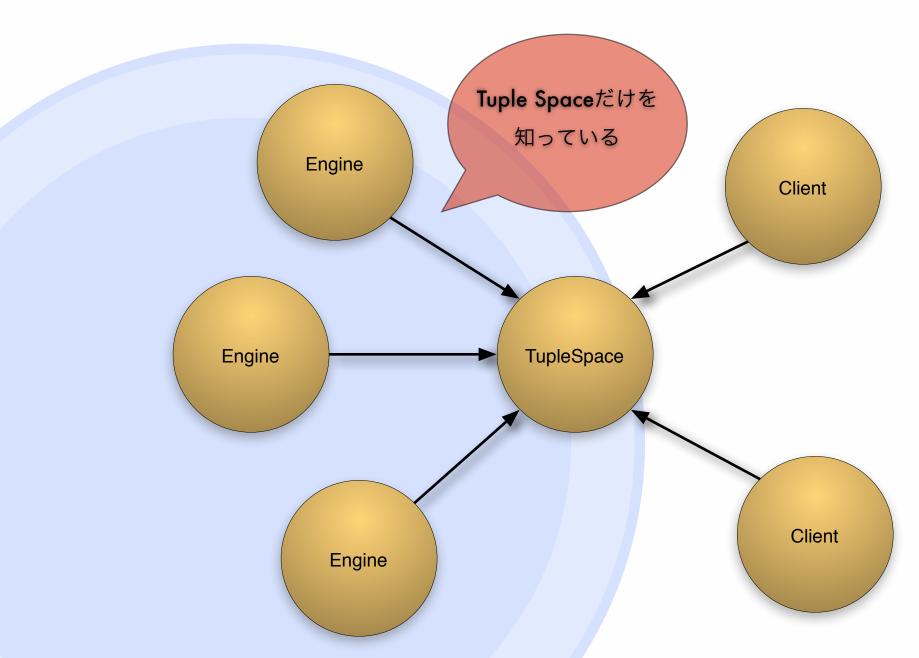


# o Lindaの協調

- 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



# o 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]]
```

# o 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:...>

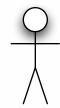
>>
```

# • write & take

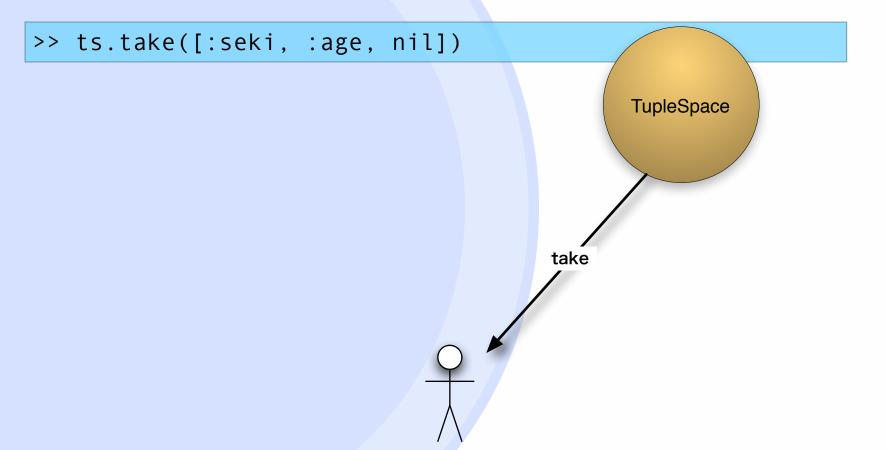
```
>> ts.write([:seki, :age, 20])
=> #<Rinda::TupleEntry:...>
>>
                                                 TupleSpace
                                                        :seki
                                                              :age
                                             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

# • write & take

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

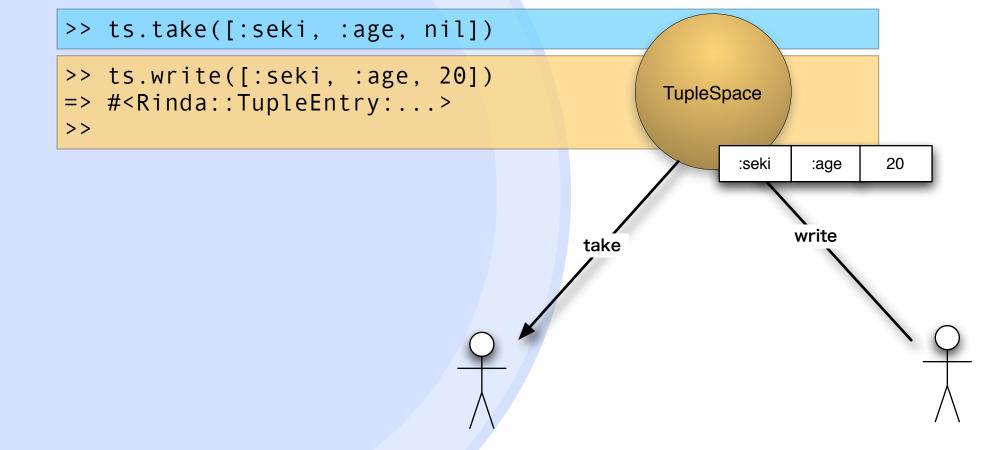


# 待ち合わせ



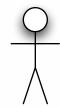


# 待ち合わせ



# 。<br/>待ち合わせ

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



## 。。Lindaはかっこいい

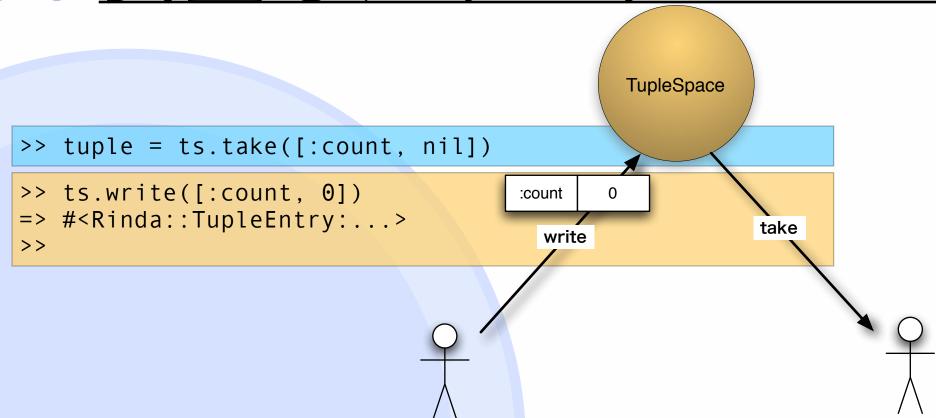
- 限られた操作で並列処理のさまざまな 協調を記述できる
  - 一般的な同期メカニズムを記述することもできる
- コンセプトのスケッチにぴったり

# ・ 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

**TupleSpace** >> tuple = ts.take([:count, nil]) take ・安全なカウンタ



・安全なカウンタ

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

・安全なカウンタ

```
TupleSpace
>> tuple = ts.take([:count, nil])
>> ts.write([:count, 0])
                                      :count
=> #<Rinda::TupleEntry:...>
                                                        write
                                        take
>>
=> [:count, 0]
>>
>> tuple = ts.take([:count, n1])
>> 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!
  - そんなにたくさんの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

# - 間違いの例

MapReduceのストレージに使う!

# MapReduce

- 【名前, 値]のタプルを扱う小さなプログラムで大量の情報を処理するかっこいい作戦
- あとは自分で調べてね
  - **WEB+DB #47に良い記事があったよ**

# ○ · 元ネタを発生させる

```
bin = []
while line = gets
  line.scan(/\w+/) do |w|
  bin << [w, 1]
  end
end</pre>
```

# - <u>キー順に並べる</u>

bin.sort\_by {|pair| pair[0]}

# - キーごとにまとめる

```
word = nil
sum = 0
bin.sort_by {|pair| pair[0]}.each do |key, value|
  if word && word != key
    p [word, sum]
    sum = 0
  end
  word = key
  sum += value
end
p [word, sum]
```

# O Hashでさぼることも

小さい空間に向く

```
bin = Hash.new {|h, k| h[k] = []}
while line = gets
  line.scan(/\w+/) do |w|
  bin[w] << 1
  end
end

bin.keys.sort.each do |k|
  p [k, bin[k].reduce(:+)]
end</pre>
```

### • TupleSpace & Hash?

- TupleSpaceはHashみたいなもん
- 知ってるキーへのアクセスが得意
  - **つまり**
- 知らないキーが苦手
- 知らないかもしれないキーも苦手

## TupleSpaceの憂鬱

- 知らないキーへのアクセスが面倒
- 新規のキーを特別扱いするのも辛い

# Tiny MapReduce by TupleSpace

```
bin = Rinda::TupleSpace.new
while line = gets
  line.scan(/\w+/) do |w|
    bin.write([w, 1])
  end
end
```

#### 小さいキーから順番に欲しい

#### 全部読むの!?

```
while true
  key ,= bin.read_all([String, nil]).min_by {|k, v| k}
  break unless key
  sum = 0
  while true
    begin
       key, value = bin.take([key, nil], 0)
       sum += value
    rescue
       break
    end
  end
  p [key, sum]
end
```

#### そのキーのデータを全部くれ

```
while true
  key ,= bin.read_all([String, nil]).min
  break unless key
  sum = 0
  while true
    begin
      key, value = bin.take([key, nil], 0)
      sum += value
    rescue
      break
  end
  end
  p [key, sum]
end
```

takeが失敗 するまで またtake...

### **そういえば**

- よく考えたらMapReduceはデータの やりとりのところは待ち合わせしなく ていいじゃん
- タスク(フェーズ)遷移の待ち合わせに 向いてる

# TokyoCabinet 一族かっこいい

- いつも並んでいる集合をくれる、QDBM とかTokyoCabinetとかBigTableとか
- 広大な空間をシーケンシャルに扱う
- どんなデータ構造もたどるときは線形だ

#### OODB Koya

- ここからさらに脱線しますよ
- 小さなOODBの実装
- 任意の時刻の状態に戻ることができる

# Koyaの変遷

- PRb PostgreSQLをメモリ空間としたOODB
- 世界はオブジェクトIDの並びである
- めちゃくちゃ遅い

# Koya with SQLite2

- SQLiteは文字列だけを扱うawkみたいなSQL風ストレージでUNIXの匂いがする
- 型の呪縛からの解放
- インスタンス変数を知るには検索になる

# Koya with QDBM

- Hashではない
- いつも並んでいる集合を提供
- あるデータのとなりのデータを触れる
- キーを工夫するとオブジェクトと属性を まとめて扱える

#### 世界はHashだと思ってた

- 実はBigTableだったことに気付いた
  - ちょっと嘘
- 世界はシーケンシャルにアクセス

### ○● 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





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