

Mastering Async with Ratpack

Daniel Hyun

2017 June 1

Table of Contents

1. Details	1
2. Need for Async	1
3. Scalability	1
4. State of async in Java land	3
5. Async is hard	4
6. Libraries to the rescue	4
7. Ratpack	4
8. Ratpack Thread Model	4
9. Ratpack Execution	5
10. Ratpack Async Primitives	5
11. Promises/Operations	5
12. Testability	5
13. ExecHarness	6
14. ExecHarness#yield	6
15. ExecHarness#run	7
16. Advanced Async	8
17. Examples	8
17.1. Anatomy of a Promise	12
18. Best practices	23

1. Details

Code: <https://github.com/danhyun/mastering-async-ratpack/>

Notes: <https://danhyun.github.io/mastering-async-ratpack/>

PDF: <https://danhyun.github.io/mastering-async-ratpack/notes.pdf>

2. Need for Async

We want to do more with less.

Resources are expensive: You pay for memory/compute/network usage

cost	\$\$\$	\$\$	\$
-----	-----	-----	-----
method	Multiple Processes	threads	event loop
example	apache	servlet	netty
-----	-----	-----	-----

<http://www.kegel.com/c10k.html>

3. Scalability



DHH ✓
@dhh

Følger

Remember: BUT DOES IT SCALE? Basecamp is peaking at 2,000 requests/second against the Rails application instances these days. 🙌

🌐 Oversæt fra engelsk

RETWEETS
141

LIKES
184



12.31 - 30. sep. 2015

🔄 22 🔄 141 ❤️ 184 ✉️



Tweet dit svar



Flavio W. Brasil @flaviowbrasil · 30. sep. 2015

Svarer @dhh

@dhh Is it 2k rps per instance?

🌐 Oversæt fra engelsk

🔄 1 🔄 ❤️ ✉️



DHH ✓ @dhh · 30. sep. 2015

@flaviowbrasil No, total.

🔄 1 🔄 ❤️ ✉️



DHH ✓ @dhh · 30. sep. 2015

@fbranczyk We run our own servers. 30 app servers for the current version of Basecamp. Don't know what AWS equivalent is.

🌐 Oversæt fra engelsk

🔄 3 🔄 ❤️ 1 ✉️

Source: <https://twitter.com/dhh/status/649260226521210880>

2000 peak rps for 30 servers sounds expensive...

Meanwhile at Apple...

Massive Scale

What does "Massive Scale" mean...

- Instances of Netty based Services in Production: 400,000+
- Data / Day: 10s of PetaBytes
- Requests / Second: 10s of Millions
- Versions: 3.x (migrating to 4.x), 4.x

This is not a contribution

<https://speakerdeck.com/normanmaurer/connectivity>

That's pretty big scale... Still expensive xD

Takeaways

- Need async to reduce footprint, especially important with each service that gets deployed
- Async is difficult, not everyone knows Netty
- Need to find balance between scale and usability

Some arguments

"Bottle neck is db calls"

This depends on the nature of your application. Not every app is a CRUD app, and even those that are serve of static/in memory content

4. State of async in Java land

Threads/Executors/Mutexes/AtomicReferences

We have the means, but...

- Testing
- Readability, ease of understanding
- Resource sharing (memory, system resources)

5. Async is hard

- Non-deterministic
- Callback hell
- Error handling/propagation

How many times have you forgotten to send a response to the user after making some async call in nodejs or Playframework? Susceptible to brain damage trying to track down what is happening.

6. Libraries to the rescue

It's no secret that concurrency is hard. Different concurrency models have found their way into the JVM ecosystem over the years:

- Futures (callbacks)
- Queues
- Fork/Join
- Actors
- Disruptor
- Continuations (Quasar/parallel universe)

Last few models try to limit or eliminate resource sharing between running threads.

<https://shipilev.net/blog/2014/jmm-pragmatics/>

<https://shipilev.net/blog/2016/close-encounters-of-jmm-kind/>

7. Ratpack

Ratpack is async and non-blocking (Netty Rocks!)

It provides its own concurrency model (execution model) for managing and handling web requests

8. Ratpack Thread Model

Netty's event loop used for compute bound code:

- Entry point to executing your Chain
- Handling NIO events (e.g. read/write)
- Scheduling/Coordinating executions



Never block the compute thread! Don't make any syscalls that block CPU until operation completes!

Thread pool for running blocking code

- For long running computations or code that blocks CPU until further notice

9. Ratpack Execution

An execution is a collection of units of work to be executed and managed by Ratpack

These units of work are called execution segment

Any http handling code is always done within an execution

Users schedule execution segments via Promise/Operation/Blocking facilities

Execution segments for a given execution are always executed on the same thread

<http://ldaley.com/post/97376696242/ratpack-execution-model-part-1>

<http://ldaley.com/post/102495950257/ratpacks-execution-model-in-practice>

10. Ratpack Async Primitives

- Promise
- Operation
- Blocking

All Ratpack primitives are implemented with the Execution api

If you can't find a method in Promise/Operation/Blocking you can always build it yourself!

11. Promises/Operations

- Creating Promises schedules execution segments
- Promises are executed in the order they were created (except for forked promises)
- They run on cpu or blocking threads, determined at time of creation.
- Easy to adapt with other async libraries (rx-mongo, thread pools)

12. Testability

Ratpack provides ExecHarness test fixture for easy testability

It allows you to run executions without starting a Ratpack server

13. ExecHarness

- java.lang.AutoCloseable
- Convenience methods that let ExecHarness manage start/close

Can use Java 7 try-with-resources using new Parrot project (Bridging the gap between Groovy and Java syntax)

14. ExecHarness#yield

Great for unit testing, executes a given promise and returns the value from the Promise in a blocking fashion

Automatically subscribes to the Promise, no need to call Promise#then

Comes in two varieties:

ExecHarness#yield

Keeps the ExecHarness running

```
given:  
ExecHarness execHarness = ExecHarness.harness() ①  
  
when:  
ExecResult result = execHarness.yield { ②  
    Promise.value('ratpack') ③  
}  
  
then:  
result.value == 'ratpack' ④  
  
cleanup:  
execHarness.close() ⑤
```

- ① Get an instance of ExecHarness
- ② Invoke yield on the instance
- ③ Return a Ratpack Promise from the Closure
- ④ Extract the value from the Promise
- ⑤ Remember to clean up after ourselves!

ExecHarness.yieldSingle

Creates and cleans up ExecHarness on each invocation


```

when:
ExecResult result = ExecHarness.yieldSingle { ①
    Promise.value('ratpack') ②
}

then:
result.value == 'ratpack' ③

```

- ① Invoke static method ExecHarness.yieldSingle
- ② Return Ratpack Promise from Closure to ExecHarness
- ③ Extract value from Promise

15. ExecHarness#run

Great for seeing Promises in action, closer to coding experience in Ratpack code

No return value

Promises are not automatically subscribed

Comes in two varieties:

ExecHarness#run

Starts and blocks execution until completed

```

given:
ExecHarness execHarness = ExecHarness.harness() ①

expect:
execHarness.run { ②
    Promise.value('ratpack') ③
        .then { String value -> ④
            assert value == 'ratpack' ⑤
        }
}

cleanup:
execHarness.close() ⑥

```

- ① Get an instance of ExecHarness
- ② Pass a closure to ExecHarness#run method
- ③ Create Ratpack Promise
- ④ Subscribe to the Promise
- ⑤ Use Groovy Power Assert to make assertion on value from Promise

ExecHarness.runSingle

Creates and cleans up ExecHarness on each invocation

```
expect:  
ExecHarness.runSingle { ①  
    Promise.value('ratpack') ②  
        .then { String value -> ③  
            assert value == 'ratpack' ④  
        }  
}
```

- ① Invoke static method ExecHarness.runSingle()
- ② Create Ratpack Promise
- ③ Subscribe to Ratpack Promise
- ④ Use Groovy Power assert to assert the value resolved from the Promise



When making assertions from within closures you need to make sure that you use Groovy Power Assert. Spock does not apply assertions to values from within the Closure

16. Advanced Async

- Promised
- SerialBatch/ParallelBatch

17. Examples

Promise.sync

```
boolean notExecuted = true  
Promise p = Promise.sync {  
    notExecuted = false  
}  
  
assert notExecuted
```

Promises don't execute when you create them.

Promise#then

```
given:
Promise<String> p = Promise.sync {
  println 'sync'
  'sync'
}

when:
p.then { s ->
  assert s == 'sync'
}

then:
throw UnmanagedThreadException
```

Promises execute with you subscribe via `Promise#then`

Promise.sync yield

```
given:
Promise p = Promise.sync {
  println 'sync'
  'sync'
}

when:
ExecResult result = yield { p }

then:
result.value == 'sync'
```

Promises need to execute in Ratpack managed thread.

`ExecHarness` provides Ratpack managed threads as do `RatpackServers` of all varieties.

Promise.sync run

```
p.then { String s ->
  println 'then'
  assert s == 'sync'
}
```

Promise.value

```
given:
Promise p = Promise.value('value') ①

when:
ExecResult result = yield { p }

then:
result.value == 'value'
```

- ① `Promise.value` creates promise from an already available value, unlike `Promise.sync` which will wait until the promise is subscribed in order to generate the value.

Blocking.get()

```
given:
Promise p = Blocking.get {
  'from blocking'
}

when:
ExecResult result = yield { p }

then:
result.value == 'from blocking'
```

Blocking runs on different threadpool

```
given:
Closure getCurrentThreadName = {
  return Thread.currentThread().name.split('-')[1]
}

and:
Promise p = Blocking.get {
  Thread.sleep(1000)
  getCurrentThreadName()
} map { String nameFromBlocking ->
  String name = getCurrentThreadName()
  return [nameFromBlocking, name].join(' -> ')
}

when:
ExecResult result = yield { p }

then:
result.value == 'blocking -> compute'
```

Promise.async

```
Thread externalThirdPartyAsyncLibraryWithCallback(Closure callback) {
  Thread.start {
    println 'Thread started'
    (1..5).each { i ->
      println(i)
      sleep(1000)
    }

    callback('async computation complete')
    println 'Thread finished'
  }
}

given:
Promise p = Promise.async { Downstream downstream ->
  println 'async start'
  externalThirdPartyAsyncLibraryWithCallback(
    downstream.&success
  )
  println 'async end'
}

when:
ExecResult result = yield { p }

then:
result.value == 'async computation complete'
```

You can think of Operations as a `Promise<Void>`, they don't share a common type but there are ways to switch back and forth between Promises and Operations

Operation.of

```
Operation.of { ①
  println 'hello from operation'
}.then { ②
  println 'nothing returned from operation'
}
```

① Factory to queue up an Operation

② Note that Operations don't return anything so there is nothing to receive in the subscriber

Operation to Promise

```
Promise<Void> p = Operation.of {
  println 'hello from operation'
}.promise() ①

p.map { v ->
  assert v == null ②
  println "v is void $v"
  "promise"
}.then { String msg -> ③
  println "We transformed from operation to $msg"
}
```

- ① Invoke `Operation#promise` to create a `Promise<Void>`
- ② See that we get `null`
- ③ Note that we can still work with this transformed `Promise`

Promise to Operation

```
Promise.value("foo")
  .operation { String msg -> ①
    println "found value $msg"
  }
  .then {
    println "no value emitted from operation"
    assert it == null ②
  }
```

- ① We can turn a `Promise` into an `Operation`, however note that we still get the previous `Promise` value
- ② See that `Operation` doesn't return anything

17.1. Anatomy of a Promise

```
Promise.sync {
  return 'hello'
}.map { s ->
  s.toUpperCase()
}.then { s ->
  ctx.render(s)
}
```

As the methods `sync`, `map`, `then` are invoked, execution segments get queued.

Exception thrown from Promise

```
given:
Promise p = Promise.sync {
  throw new Exception("oh no")
}

when:
yield { p }.valueOrThrow

then:
thrown Exception
```

But we can handle it and short circuit

Promise#onError

```
when:
p = Promise.sync {
  throw new Exception("oh no")
}.onError { Exception e ->
  println e.message
}.map {
  'map'
}

and:
def value = yield { p }.valueOrThrow

then:
notThrown Exception
value != 'map'
value == null
```

Or we can handle and continue processing

Promise#mapError

```
given:
Promise p = Promise.sync {
  throw new Exception("oh no")
} mapError { Throwable t ->
  'default value'
} map({ String s -> s.toUpperCase()} as Function)

when:
String value = yield { p }.valueOrThrow

then:
notThrown Exception
value == 'DEFAULT VALUE'
```



Promises are immutable, methods like `Promise#map` always return new promises

Promises are immutable

```
given:
boolean exception = false

when:
Promise p = Promise.sync { throw new Exception() }

p.onError { ①
  exception = true
  println 'oops'
}
p.map { 'map' } ②

and:
yield { p }.valueOrThrow

then:
exception == false
thrown Exception
```

① `Promise#onError` returns a *new* Promise

② `Promise#map` returns a *new* Promise

Promise api allows you to chain promise manipulation

```
when:  
p = Promise.sync { throw new Exception() }  
  .onError { ①  
    exception = true  
    println 'oops'  
  }  
  .map { 'map' }  
  
and:  
yield { p }.valueOrThrow  
  
then:  
exception  
notThrown Exception
```

① Promise#onError returns a new Promise

Error mapping also available in flatmap flavor Promise#flatMapError.

Promise#map

```
Promise.value(3)  
  .map { int i -> 'A' * i } ①  
  .then { String s ->  
    println 'from map'  
    assert s == 'AAA'  
  }
```

① Promise#map runs on compute thread, don't block here

Promise#flatMap

```
Promise.value(3)  
  .flatMap { int i -> ①  
    Blocking.get { 'A' * i } ②  
  }.then { String s -> ③  
    println 'from flatMap blocking'  
    assert s == 'AAA'  
  }
```

① Promise#flatMap runs on compute thread, don't block here

② Since Blocking#get returns a Promise, need to use Promise#flatMap in order to "unpack" the nested promise and continue working with the value as in <3>

Promise#blockingMap

```
Promise.value(3)
  .blockingMap { int i -> ①
    'A' * i
  }.then { String s ->
    println 'from blockingMap'
    assert s == 'AAA'
  }
```

- ① A convenience method for executing blocking code inline without having to do `Promise#flatMap { Blocking.get {} }` as in the previous example

Promise#flatMap with async

```
Promise.value(3)
  .flatMap { int i ->
    Promise.async { Downstream d -> ①
      Thread.start {
        println 'starting thread'
        Thread.sleep(100)
        d.success('A' * i)
      }
    }
  }.then { String s ->
    println 'from flatMap async Thread'
    assert s == 'AAA'
  }
```

- ① Integrating with an externally managed threadpool/async library

Promise#left/right

```
given:
List<Map<String, Object>> users = [[name: 'danny', interests: ['dancing', 'cooking']]]
Map<String, List<String>> interests = [dancing: ['fun'], cooking: ['a great catch'],
kotlin: ['with it']]

expect:
run {
  Blocking.get {
    users.find { Map user -> user.name == 'danny' } ①
  }.right { Map user -> ②
    user.interests.collectMany { interest -> ③
      interests[interest]
    }
  }.map { Pair<Map, List<String>> pair -> ④
    "${pair.left.name} is ${pair.right.join(' and is ')}"
  }.then { String msg ->
    assert msg == 'danny is fun and is a great catch'
  }
}
```

- ① Imagine some blocking jdbc lookup by name
- ② Promise#right takes the result of the previous promise and creates a tuple like graph datastructure called Pair. The value returned from the closure/lambda is then pushed to the "right" position of the Pair
- ③ Imagine some in memory lookup for interests by user that we just looked up from <1>
- ④ The result from the previous call is of type `Pair<A, B>` where `A` is the result of the `Blocking.get{}` call and `B` is ther result of `Promise#right` call

Pairs are handy when working with small number of arguments.

You can also nest Pairs, you can have something like `Pair<Pair<A, B>, C>` but this quickly becomes hard to track.

If only Java had tuple support :(

This is also available in flatMap flavor `Promise#flatLeft/flatRight`

- `ParallelBatch`, `SerialBatch` You'll often want to take a list of promises and transform them into a list of resolved values. `ParallelBatch` and `SerialBatch` help achieve this.

```

given:
Random random = new Random(1)
Closure<Promise> getPrice = { int id -> ①
  Blocking.get {
    [id: id, price: random.nextInt(10)]
  }
}

expect:
run {
  Promise.value(50)
  .map { int i -> (1..i) }
  .flatMap { ids ->
    List<Promise> promises = ids.collect { int id -> getPrice(id) } ②
    batch(promises).yield() ③
  }.map { List<Map<String, Object>> results -> ④
    results.price.sum()
  }.map { int price ->
    "Total is \${price}"
  }.then { String msg ->
    assert msg == 'Total is $238'
  }
}

where:

type      | batch ⑤
'serial'  | SerialBatch.&of
'parallel'| ParallelBatch.&of

```

- ① Simulate price lookup in a blocking manner
- ② Use a Groovy range to generate a list of promises to lookup the price for the given id
- ③ Invoke either `SerialBatch.of` or `ParallelBatch.of` and yield a `Promise<List<Map<String, Object>>>`
- ④ Use the handy value as a single list
- ⑤ Spock datatable for making `batch` pluggable

Batch is great for when you have `List<Promise<A>>` but want to work with `Promise<List<A>>` in subsequent calculations.

Forking execution

Promise#fork

```
given:
List list = []
Closure addToList = { ①
  println it
  list << it
  it
}

expect:
run {
  Blocking.get {
    addToList('foo') ②
  }.right(
    Promise.async { Downstream d ->
      d.success(addToList('bar')) ③
    }.fork() ④
  ).map { Pair<String, String> pair ->
    pair.left + pair.right
  }.then { String msg ->
    assert list == ['bar', 'foo'] ⑤
    assert msg == 'foobar'
  }
}
```

- ① Convenience closure for printing and adding to list
- ② Add foo in a blocking manner
- ③ Add bar in async manner
- ④ Fork the bar async promise
- ⑤ Assert that bar was entered before foo



When forking Promises, they execute immediately!

"Flow control"

Promise#mapIf

```
given:
Closure fizzbuzz = { candidate -> ①
  Promise.sync {
    println "${LocalDateTime.now()} EXECUTING FIZZBUZZ for $candidate"
    candidate
  }
  .mapIf({i -> i instanceof Number && i % 3 == 0 && i % 5 == 0}, { 'fizzbuzz' }) ②
  .mapIf({i -> i instanceof Number && i % 3 == 0}, { 'fizz' }) ②
  .mapIf({i -> i instanceof Number && i % 5 == 0}, { 'buzz' }) ②
}

expect:
run {
  Promise.value(15)
  .map { 1..it }
  .flatMap { range ->
    ParallelBatch.of(range.collect { fizzbuzz(it) }).yield() ③
  }.then { list -> ④
    assert list == [
      1, 2, 'fizz', 4, 'buzz',
      'fizz', 7, 8, 'fizz', 'buzz',
      11, 'fizz', 13, 14, 'fizzbuzz'
    ]
  }
}
```

- ① Convenience closure
- ② Example of using `Promise#mapIf(Predicate, Function)`, only maps if predicate passes
- ③ Execute these promises in parallel
- ④ Assert that our list is fizzbuzzed correctly

Also available in flatMap variety `Promise#flatMapIf`

Other useful methods:

- `Promise#onNull`
- `Promise#route`



`Promise#route` is a terminating call, the rest of the promise chain is no longer executed!!

Promise#route(Predicate, Action)

```
Promise.sync(LocalDate.now)
  .route({ldt -> (ldt.get(ChronoField.MILLI_OF_SECOND) & 1) == 0}, { ldt -> ①
    println "TERMINATING: $ldt is even"
  }).map { ldt ->
    println "MAPPING LocalDateTime to String"
    "$ldt is odd"
  }.then { String msg ->
    println "SUCCESS: $msg"
  }
}
```

① Note that we terminate the promise chain here if predicate passes

Throttling Promises

Throttle acts as a semaphore only allowing n number of promises to run in parallel.

Promise#throttled

```
given:
Throttle throttle = Throttle.ofSize(3) ①
Closure fizzbuzzer = { candidate ->
  Blocking.get {
    Thread.sleep(2000)
    println "${LocalDateTime.now()} EXECUTING FIZZBUZZ for $candidate"
    candidate
  }
  .mapIf({i -> i instanceof Number && i % 3 == 0 && i % 5 == 0}, { 'fizzbuzz' })
  .mapIf({i -> i instanceof Number && i % 3 == 0}, { 'fizz' })
  .mapIf({i -> i instanceof Number && i % 5 == 0}, { 'buzz' })
  .throttled(throttle) ②
}

expect:
run {
  Promise.value(15)
  .map { 1..it }
  .flatMap { range ->
    ParallelBatch.of(range.collect { fizzbuzzer(it) }).yield()
  }
  .then { list ->
    assert list == [
      1, 2, 'fizz', 4, 'buzz',
      'fizz', 7, 8, 'fizz', 'buzz',
      11, 'fizz', 13, 14, 'fizzbuzz'
    ]
  }
}
}
```

① Declare a throttle of size 3

- ② Make sure that the promise we submit to the ParallelBatch is throttled

Sample output

```
2017-05-30T07:42:48.294 EXECUTING FIZZBUZZ for 4
2017-05-30T07:42:48.308 EXECUTING FIZZBUZZ for 5
2017-05-30T07:42:48.294 EXECUTING FIZZBUZZ for 1
2017-05-30T07:42:50.367 EXECUTING FIZZBUZZ for 10
2017-05-30T07:42:50.367 EXECUTING FIZZBUZZ for 12
2017-05-30T07:42:50.367 EXECUTING FIZZBUZZ for 2
2017-05-30T07:42:52.371 EXECUTING FIZZBUZZ for 13
2017-05-30T07:42:52.371 EXECUTING FIZZBUZZ for 9
2017-05-30T07:42:52.371 EXECUTING FIZZBUZZ for 3
2017-05-30T07:42:54.374 EXECUTING FIZZBUZZ for 7
2017-05-30T07:42:54.374 EXECUTING FIZZBUZZ for 6
2017-05-30T07:42:54.374 EXECUTING FIZZBUZZ for 11
2017-05-30T07:42:56.377 EXECUTING FIZZBUZZ for 8
2017-05-30T07:42:56.377 EXECUTING FIZZBUZZ for 15
2017-05-30T07:42:56.377 EXECUTING FIZZBUZZ for 14
```

You can see there is a tight grouping of 3 per time period

Spying

If you wish to observe items as they get processed in the promise chain, you can make use of `Promise#wiretap`

Promise#wiretap

```
Promise.sync(LocalDate.&now)
  .map { ldt -> ldt.getDayOfWeek() }
  .wiretap { Result<DayOfWeek> r -> ①
    println "Found: ${r.value.getDisplayName(TextStyle.FULL_STANDALONE, Locale.KOREA)}" ②
  }.then { DayOfWeek dow -> ③
    println "Today is $dow"
  }
```

- ① Invoke wiretap method
- ② Note that we have to invoke `Result#getValue` in order to get the value produced from the previous Promise
- ③ Note that the next processor gets the same result as the wiretap

18. Best practices

- Avoid multiple then blocks
- Try to linearize/flatten data flow