

The Actor Model, Part Two

CSCI 5828: Foundations of Software Engineering
Lecture 14 — 10/06/2016

Goals

- Cover a more advanced example of using processes in Elixir
- Review the material on process linking
 - and show how we can then introduce the notion of process supervision

Fibonacci Calculator (I)

- Let's jump back into Elixir and the Actor model
 - We'll take a look at using Actors to calculate Fibonacci numbers
 - 0, 1, 1, 2, 3, 5, 8, 13, ...
- Our example will calculate a set of Fibonacci numbers using a different number of actors
 - starting with one actor and proceeding up to ten actors running at once

Elixir Function Composition

- In order to understand the source code of the example, we must review Elixir's function composition operator, also known as the "pipe operator"
- If you had a series of statement like this
 - $a = f(x); b = g(a); c = h(b)$
- You could also write it like this
 - $c = h(g(f(x)))$
- In Elixir, you would write it like this
 - $c = x \mid > f \mid > g \mid > h$
 - x is piped into f , the result is piped into g , the result is piped into h
- The functions on the right hand side can have parameters
 - $x \mid > f(y, z)$ is equivalent to calling $f(x, y, z)$ —the *value* being piped becomes the *first argument of the function on the right hand side*

Fibonacci Calculator (II)

- To start our Fibonacci example, we first design two actors
 - A **solver**: is able to calculate the *n*th Fibonacci number
 - A **scheduler**: distributes calculation requests to a set of 1 or more solvers
- A solver will sit in loop and do the following
 - It sends `{:ready, pid}` to the scheduler
 - It will then receive a `:fib` message asking it to calculate a number
 - When it is done, it will send an `:answer` message to the scheduler
- The solver will perform these actions until it receives a `:shutdown` message
- The scheduler will receive an array of integers that represent the Fibonacci numbers to calculate
 - it will send out `:fib` messages to solvers until all requests are complete

Fibonacci Calculator (III)

- The solver

```
1 defmodule FibSolver do
2
3   def fib(scheduler) do
4     send(scheduler, {:ready, self})
5     receive do
6       {:fib, n, client} ->
7         send(client, {:answer, n, fib_calc(n), self})
8         fib(scheduler)
9       {:shutdown} -> exit(:normal)
10    end
11  end
12
13  defp fib_calc(0) do 0 end
14  defp fib_calc(1) do 1 end
15  defp fib_calc(n) do fib_calc(n-1) + fib_calc(n-2) end
16 end
```

Fibonacci Calculator (IV): The Scheduler

```
1 defmodule Scheduler do
2
3   def run(num_processes, module, func, to_calculate) do
4     (1..num_processes)
5     |> Enum.map(fn(_) -> spawn(module, func, [self]) end)
6     |> schedule_processes(to_calculate, [])
7   end
8
9   defp schedule_processes(processes, queue, results) do
10    receive do
11      {:ready, pid} when length(queue) > 0 ->
12        [ next | tail ] = queue
13        send(pid, {:fib, next, self})
14        schedule_processes(processes, tail, results)
15
16      {:ready, pid} ->
17        send(pid, {:shutdown})
18        if length(processes) > 1 do
19          schedule_processes(List.delete(processes, pid), queue, results)
20        else
21          Enum.sort(results, fn ({n1, _}, {n2, _}) -> n1 <= n2 end)
22        end
23
24      {:answer, number, result, _pid} ->
25        schedule_processes(processes, queue, [ {number, result} | results])
26    end
27  end
28 end
```

Fibonacci Calculator (V): Main Program

```
47 to_process = [ 37, 37, 37, 37, 37, 37 ]
48
49 Enum.each(1..10, fn (num_processes) ->
50   {time, result} =
51     :timer.tc(Scheduler, :run,
52       [num_processes, FibSolver, :fib, to_process])
53
54   if num_processes == 1 do
55     IO.puts inspect result
56     IO.puts "\n # time (s)"
57   end
58   :io.format "~2B ~.2f~n", [num_processes, time/1000000.0]
59 end)
```


Fibonacci Calculator (VI): Results

- On my 8-core machine, the results are:
 - # time (s)
 - 1 6.22
 - 2 3.07
 - 3 2.10
 - 4 2.14
 - 5 2.43
 - 6 1.65 <== almost 4 times as fast
 - 7 1.72
 - 8 1.77
 - 9 1.78
 - 10 1.89 <== roughly 3.3 times as fast on average

Discussion

- Striking how simple the implementation of the FibSolver Actor is
 - small piece of code with a defined “message API”
 - program can then spin up as many of these actors as they want
- The scheduler is more complex BUT
 - it implemented scheduling in a very generic way
 - the function being calculated was completely abstracted away
 - the logic simply took care of providing work to all ready actors
 - and then shutting down actors when no more work was available
- With 11 active actors (10 solvers + 1 scheduler): Elixir has flexibility as to how those actors are distributed across the cores of the machine

Making Fibonacci More Efficient

- See page 204 of our textbook to understand why our Fibonacci solver takes a while to calculate the result of fib(37)
- We can make our solver way more efficient (and eliminate the need for our program above) using Elixir's Agent module.
 - Using agents, we can quickly specify the state of an actor and how that state can be updated
- For Fibonacci, the basic idea of making it more efficient is to remember all of our previous calculations; have code like this in a function called "do_fib"

```
{ n_1, cache } = do_fib(cache, n-1)
result         = n_1 + cache[n-2]
{ result, Map.put(cache, n, result) }
```

Error Handling and Resilience

- Actors provide the ability to write fault-tolerant code
 - We can assign a supervisor to a set of actors that detects when an actor has crashed and can do something about it
 - such as restart the actor
 - The way they do this is by linking the actors together (as we saw in Lecture 20)
 - First: `Process.flag(:trap_exit, true)`
 - Second: `pid = spawn_link(...)`
 - Third: `receive do {:EXIT, pid, reason}`
- We're going to build up an example that demonstrates these concepts

An Actor to Test Links: LinkTest

```
1 defmodule LinkTest do
2   def loop do
3     receive do
4       {:exit_because, reason} -> exit(reason)
5       {:link_to, pid} -> Process.link(pid)
6       {:EXIT, pid, reason} -> IO.puts("#{inspect(pid)} exited because #{reason}")
7     end
8   loop
9 end
10
11 def loop_system do
12   Process.flag(:trap_exit, true)
13   loop
14 end
15 end
```

An actor that can link to other actors via `:link_to`; otherwise it can be told to die by sending it a `:exit_because` message

If we want to receive `:EXIT` messages, we need to invoke this actor with the `loop_system` call. Otherwise, we can just call `loop` to see what happens when an actor exits for a non `:normal` reason

Example: Linked Actors; Non-Normal Exit

- Create two instances of the actor
 - `pid1 = spawn(LinkTest, :loop, [])`
 - `pid2 = spawn(LinkTest, :loop, [])`
- Link them (links are bidirectional)
 - `send(pid1, {:link_to, pid2})`
- Tell one to quit for a non-normal reason (it doesn't matter which actor)
 - `send(pid2, {:exit_because, :bad_thing})`
- The result?
 - BOTH actors die; no `:EXIT` message received
 - We can check this with `Process.info`: `Process.info(pid2, :status)`

Example: Linked Actors; Normal Exit

- Create two instances of the actor
 - `pid1 = spawn(&LinkTest.loop/0)`
 - `pid2 = spawn(&LinkTest.loop/0)`
- Link them (links are bidirectional)
 - `send(pid1, {:link_to, pid2})`
- Tell one to quit for a normal reason (it doesn't matter which actor)
 - `send(pid2, {:exit_because, :normal})`
- The result?
 - Actor 2 dies; Actor 1 lives; still no `:EXIT` message received

Example: Linked System Actors; Non-Normal Exit

- Create two instances of the actor
 - `pid1 = spawn(LinkTest, :loop_system, [])`
 - `pid2 = spawn(&LinkTest.loop/0)`
- Link them (links are bidirectional)
 - `send(pid1, {:link_to, pid2})`
- Tell one to quit for a normal reason (it doesn't matter which actor)
 - `send(pid2, {:exit_because, :bad_thing})`
- The result?
 - Actor 2 dies; Actor 1 lives; `:EXIT` message received and logged

Creating a Supervisor

- We now have enough knowledge to create an actor and its supervisor
 - The idea is that we can implement a process that monitors the state of other processes and, if they crash, attempts to restart them
- We will create an actor that will "cache" values for us
- The cache will be able to
 - receive a request to store something in the cache
 - receive a request to retrieve something in the cache
 - receive a request to return the size of the cache (in bytes)
- The supervisor will create a cache actor and monitor its status
 - If it goes down, it will restart the cache

Cache

We can cause this actor to crash by sending nil for page in a :put message

```
1 defmodule Cache do
2   def loop(pages, size) do
3     receive do
4       {:put, url, page} ->
5         new_pages = Dict.put(pages, url, page)
6         new_size = size + byte_size(page)
7         loop(new_pages, new_size)
8       {:get, sender, ref, url} ->
9         send(sender, {:ok, ref, pages[url]})
10        loop(pages, size)
11       {:size, sender, ref} ->
12         send(sender, {:ok, ref, size})
13        loop(pages, size)
14       {:terminate} -> # Terminate request - don't recurse
15     end
16   end
17 end
```

Cache Helper Routines

```
18  def start_link do
19    pid = spawn_link(__MODULE__, :loop, [HashDict.new, 0])
20    Process.register(pid, :cache)
21    pid
22  end
23
24  def put(url, page) do
25    send(:cache, {:put, url, page})
26  end
27
28  def get(url) do
29    ref = make_ref()
30    send(:cache, {:get, self(), ref, url})
31    receive do
32      {:ok, ^ref, page} -> page
33    end
34  end
35
36  def size do
37    ref = make_ref()
38    send(:cache, {:size, self(), ref})
39    receive do
40      {:ok, ^ref, s} -> s
41    end
42  end
43
44  def terminate do
45    send(:cache, {:terminate})
46  end
```

These functions provide an “API” to the Cache. We can call them and not worry about starting actors and sending messages.

Cache Supervisor

DEMO

```
1 defmodule CacheSupervisor do
2
3   def start do
4     spawn(__MODULE__, :loop_system, [])
5   end
6
7   def loop do
8     pid = Cache.start_link
9     receive do
10      {:EXIT, ^pid, :normal} ->
11        IO.puts("Cache exited normally")
12        :ok
13      {:EXIT, ^pid, reason} ->
14        IO.puts("Cache failed with reason #{inspect reason} - restarting it")
15      loop
16    end
17  end
18
19  def loop_system do
20    Process.flag(:trap_exit, true)
21    loop
22  end
23 end
```

Start up a Cache. If it crashes, restart it; otherwise quit

Make sure we call :trap_exit to receive :EXIT messages

Using the Cache

- In iex, compile both modules
 - `c("cache.ex")`
 - `c("cache_supervisor.ex")`
- Start by creating the supervisor (which creates the Cache, its worker)
 - `CacheSupervisor.start_link`
- Then just use the Cache
 - `Cache.size => 0`
 - `Cache.put "foo", "bar" => :ok`
 - `Cache.size => 3`
 - `Cache.put "ohnoes", nil => error message; auto restart`
 - `Cache.size => 0`
- To cleanly kill both processes, just type `Cache.terminate`

Discussion

- This example illustrates a generic approach to concurrent actor systems
 - Keep the supervisors as small and as simple as possible
 - So simple that they are easy to debug and get correct
 - Have the actors that they supervise crash when things go wrong
 - Let the supervisors detect those crashes and decide what to do
- This approach maximizes simplicity
 - rather than adding lots of error checking code in the workers
 - implement the success case and let all error cases cause a crash that gets handled by the supervisor => a nice separation of concerns

Wrapping Up

- We saw a more advanced example of processes via the Fibonacci example
 - The scheduler demonstrated how we make use of immutable data structures to maintain state
 - and how to transition to a new state on well-defined boundaries
- We saw that our implementation of the Fibonacci calculation was inefficient and that Elixir's agents module can be used to implement caching
- We then returned to the notion of linking processes and saw how it forms the basis of process supervision

- Up next: distributing processes over multiple nodes