

LET THE GARBAGE CRASH

Filipe Varjão
Developer @ Erlang Solutions



@filipevarjao
filipe.varjao@erlang-solutions.com

Erlang
SOLUTIONS

Why **Garbage** Collection?

- Facilitates software development.
- Increases robustness.
- Efficiency at runtime.

Memory management

- Mark-Sweep
- Reference counting
- Copy

Mark-Sweep

- J. McCarthy (1960)
- Lisp



Mark-Sweep

- Allocate

```
1 Function New() =  
2   | if free_pool is empty then  
3   |   | mark_sweep()  
4   |   | newcell = allocate()  
5   |   | return newcell
```



Mark-Sweep

- Collect

```
1 Function mark_sweep() =  
2   for R in Roots do  
3     | mark(R)  
4   sweep()  
5   if free_pool is empty then  
6     | abort "Memory Exhausted"
```



Mark-Sweep

- Mark

```
1 Function mark(N) =  
2   | if mark_bit(N) == unmarked then  
3   |   | mark_bit(N) = marked  
4   |   | for M in Children(N) do  
5   |   |   | mark(*M)
```



Mark-Sweep

- Sweep

```
1 Function sweep() =  
2   N = Heap_bottom  
3   while N < Heap_top do  
4     if mark_bit(N) == unmarked then  
5       free(N)  
6     else  
7       mark_bit(N) = unmarked  
8     N = N + size(N)
```



Reference Counting

- Collins (1960)



Reference Counting

- Allocate

```
1 Function allocate() =  
2   |   newcell = free_list  
3   |   free_list = next(free_list)  
4   |   return newcell  
5 Function new() =  
6   |   if free_list == nil then  
7   |   |   abort "Memory Exhausted"  
8   |   newcell = allocate()  
9   |   RC(newcell) = 1  
10  |   return newcell
```



Reference Counting

- Counting

```
1 Function free(N) =
2   |   next = free_list
3   |   free_list = N
4 Function delete(T) =
5   |   RC(T) = RC(T) - 1
6   |   if RC(T) == 0 then
7     |   for U in Children(T) do
8       |   |   delete(*U)
9       |   |   free(T)
10 Function Update(R, S) =
11   |   RC(S) = RC(S) + 1
12   |   delete(*R)
13   |   *R = S
```



Copy

- First version - Marvin Minsky (1963)
- Fenichel and Yochelson (1969)
- Cheney (1979)



Copy

- **Allocate**

```
1 Function init() =
2   |   To_space = Heap_bottom
3   |   space_size = Heap_size / 2
4   |   top_of_space = To_space + space_size
5   |   From_space = top_of_space + 1
6   |   free = To_space
7 Function New(n) =
8   |   if free + n > top_of_space then
9   |   |   flip()
10  |   if free + n > top_of_space then
11  |   |   abort ("Memory Exhausted")
12  |   newcell = free
13  |   free = free + n
14  |   return newcell
```



Copy

- Copying

```
1 Function flip() =  
2   From_spave, To_space = To_space, From_space  
3   top_of_space = to_space + space_size  
4   free = To_space  
5   for R in Roots do  
6     | R = copy(R)
```



Copy

- Collecting

```
1 Function copy(P) =  
2   if atomic(P) or P == nil then  
3     return P  
4   if not forwarded(P) then  
5     n = size(P)  
6     P' = free  
7     free = free + n  
8     temp = P[0]  
9     forwarding_address(P) = P'  
10    P'[0] = copy(temp)  
11    for i = 1 to n - 1 do  
12      return P'[i] = copy(P[i])  
13  return forwarding_address(P)
```



Golang


- Robert Griesemer, Rob Pike and Ken Thompson (2007)
- Google
- Hoare - Communicating Sequential Processes (1978)

Erlang


- Joe Armstrong, Robert Virding and Mike Williams
(1986)
- Ericsson

Golang memory management


- Mark-Sweep
- Tri-color algorithm (Dijkstra, 1978)




1



2



3



4

Erlang memory management

- Copying semi-space garbage collector
- Generational GC

Memory management

Erlang

- Process heap
- Atom (table)
- ETS (table)
- Binary

Golang

- Objects might live in heap or stack (escape analyses)
- Compiling analyses

Memory management

Golang

- The Go's compiler will first try to put it on stack but, often it cannot.
- The object is reachable from heap, e.g. assigned to a global
- Object's lifetime is beyond the current function, e.g. x is returned to caller

Memory management

Golang

- An object is sized, e.g. the backing store of slice/string
- When an object is passed to another function which makes x escape, e.g. call f(x) where f does "global = x"
- There is not enough information to decide if it doesn't escape, e.g. passing as arg to function pointer calls

Memory management

Erlang

- Atoms are never deleted.
- The EVM creates a fixed size to all atoms, system crashes when it gets full.

Memory management

Erlang

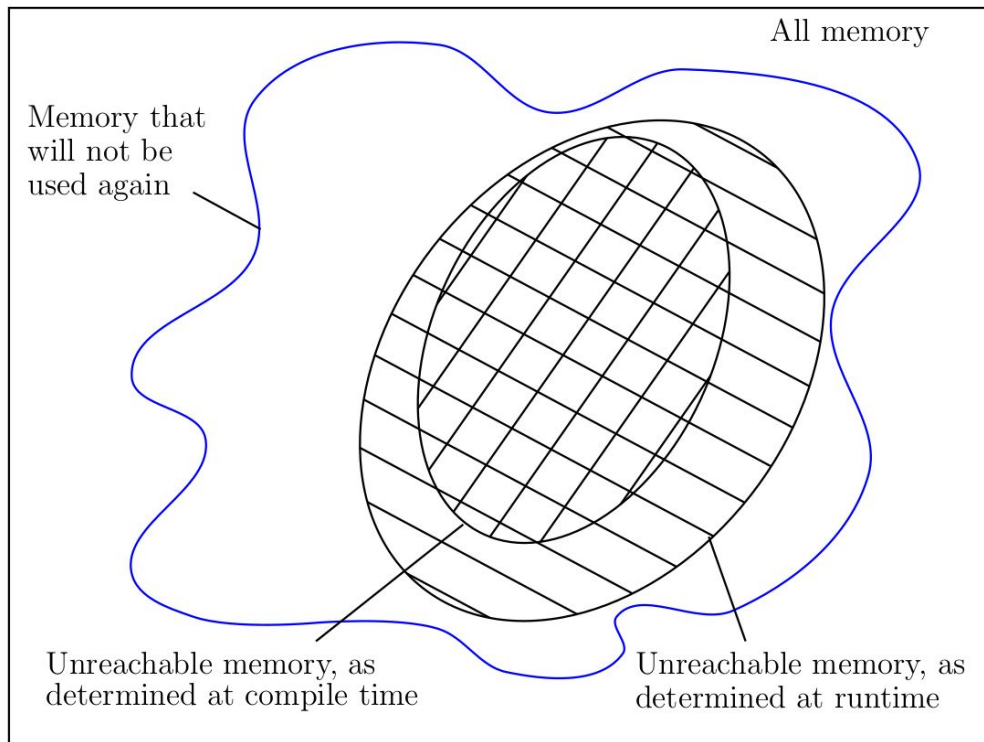
- Small binaries (< 64 bytes)
- Large binaries
- ETS tables are totally isolated where items are copied from/to process heaps

Memory management

Erlang

- Process lifetime is short

Region-Based **Memory** management



Davis (2015)

Memory management

- Region-Based Cyclic Reference Counting
- gitlab.com/filipevarjao/go-rbmm-refcount

Experiment

Setup

name	CPU	Memory
M1	AMD Phenom II X6 1090T 3,2Ghz	4GB
M2	Intel i7-3770, 3,4Ghz	8GB
M3	Intel i7-855oU 1,80Ghz	16GB

* *Ubuntu Server 16.04 LTS(Xenial Xerus) 64-bit*

Memory management

Benchmark Pause - M1

<i>collector</i>	<i>exec time</i>	<i>worst pause time</i>	<i>total memory</i>	<i>pauses</i>
ms	1,54s	18920us	1,02GB	11
rc	1,43s	630,45us	1,02GB	0
rbmm-rc	3,37s	68,03us	24,77KB	0

Memory management

Benchmark Pause - M2


<i>collector</i>	<i>exec time</i>	<i>worst pause time</i>	<i>total memory</i>	<i>pauses</i>
ms	668,48ms	6220us	1,02GB	11
rc	645,66ms	282,93us	1,15GB	0
rbmm-rc	295000ms	21,97us	24,76KB	0

Memory management

Benchmark Pause - M3

<i>collector</i>	<i>exec time</i>	<i>worst pause time</i>	<i>total memory</i>	<i>pauses</i>
ms	705,97ms	5220us	1,02GB	11
rc	702,63ms	297,09us	1,02GB	0
rbmm-rc	311000ms	106,93us	24,44KB	0

Profiling



Memory management

Conclusions

- A new concurrent Garbage Collection algorithm.
- Non-suspensive, a feature of paramount importance in real-time systems.
- There is no silver-bullet.

THANK YOU

Q&A

Filipe Varjão
Developer @ Erlang Solutions



@filipevarjao
filipe.varjao@erlang-solutions.com

Erlang
SOLUTIONS