

An Efficient Non-Moving Garbage Collector for Functional Languages

Katsuhiko Ueno
Atsushi Ohori
Toshiaki Otomo

RIEC , Tohoku University

ICFP 2011,
20 May 2011

Copying Collectors

Mark-and-Sweep Collectors

Copying Collectors

Mark-and-Sweep Collectors

Copying Collectors

Good!

fast allocation
by "bump pointer"

Mark-and-Sweep Collectors

Copying Collectors

Good!

fast allocation
by "bump pointer"

Good!

$O(\text{numLives})$
collection cost

Mark-and-Sweep Collectors

Copying Collectors

Good!

fast allocation
by "bump pointer"

Good!

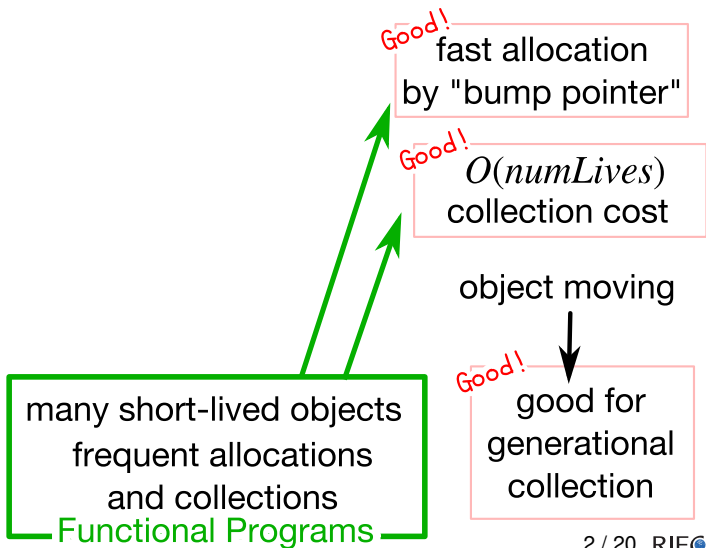
$O(\text{numLives})$
collection cost

many short-lived objects
frequent allocations
and collections

Functional Programs

Mark-and-Sweep Collectors

Copying Collectors



Mark-and-Sweep Collectors

Bad...
fragmentation
slow allocation

Copying Collectors

Good!
fast allocation
by "bump pointer"

Good!
 $O(\text{numLives})$
collection cost

object moving

Good!
good for
generational
collection

many short-lived objects
frequent allocations
and collections

Functional Programs

Mark-and-Sweep Collectors

Bad...
fragmentation
slow allocation

Bad...
 $O(\text{heapSize})$
collection cost

many short-lived objects
frequent allocations
and collections

Functional Programs

Copying Collectors

Good!
fast allocation
by "bump pointer"

Good!
 $O(\text{numLives})$
collection cost

object moving

Good!
good for
generational
collection

Mark-and-Sweep Collectors

Bad...
fragmentation
slow allocation

Bad...
 $O(\text{heapSize})$
collection cost

non-moving

Bad...
no way of
generational
collection

many short-lived objects
frequent allocations
and collections

Functional Programs

Copying Collectors

Good!
fast allocation
by "bump pointer"

Good!
 $O(\text{numLives})$
collection cost

object moving

Good!
good for
generational
collection

Mark-and-Sweep Collectors

avoid in practice

~~fragmentation~~
~~slow allocation~~

$O(\text{heapSize})$
collection cost

non-moving

no way of
generational
collection

many short-lived objects
frequent allocations
and collections

Functional Programs

Copying Collectors

fast allocation
by "bump pointer"

$O(\text{numLives})$
collection cost

object moving

good for
generational
collection

Mark-and-Sweep Collectors

avoid in practice

~~fragmentation~~
~~slow allocation~~

~~$O(\text{heapSize})$~~ $O(\text{numLives})$
collection cost

non-moving

no way of
generational
collection

many short-lived objects
frequent allocations
and collections

Functional Programs

Copying Collectors

fast allocation
by "bump pointer"

$O(\text{numLives})$
collection cost

object moving

good for
generational
collection

Mark-and-Sweep Collectors

avoid in practice

~~fragmentation~~
~~slow allocation~~

~~$O(\text{heapSize})$~~ $O(\text{numLives})$
collection cost

non-moving

~~non-moving~~
~~generational~~
~~collection~~

any short-lived objects
frequent allocations
and collections

Functional Programs

Copying Collectors

Good!
fast allocation
by "bump pointer"

Good!
 $O(\text{numLives})$
collection cost

object moving

Good!
good for
generational
collection

Mark-and-Sweep Collectors

avoid in practice

~~fragmentation~~
~~slow allocation~~

~~$O(\text{heapSize})$~~ $O(\text{numLives})$
collection cost

non-moving

~~non-moving~~
~~generational~~
~~collection~~

any short-lived objects
frequent allocations
and collections

Functional Programs

Copying Collectors

Good!
fast allocation
by "bump pointer"

Good!
 $O(\text{numLives})$
collection cost

object moving

Good!
good for
generational
collection

Mark-and-Sweep Collectors

avoid in practice

~~fragmentation~~
~~slow allocation~~

~~$O(\text{heapSize})$~~ $O(\text{numLives})$
collection cost

non-moving

~~non-moving~~
~~generational~~
~~collection~~

any short-lived objects
frequent allocations
and collections

Functional Programs

Copying Collectors

fast allocation
by "bump pointer"

fragmentation-free
by definition

collection cost

object moving

good for
generational
collection

Mark-and-Sweep Collectors

avoid in practice

~~fragmentation~~
~~slow allocation~~

~~$O(heapSize)$~~ $O(numLives)$
collection cost

non-moving

~~non-moving~~
~~generational~~
~~collection~~

any short-lived objects
frequent allocations
and collections

Functional Programs

Copying Collectors

fast allocation
by "bump pointer"

fragmentation-free
by definition

keep good locality

object moving

good for
generational
collection

Mark-and-Sweep Collectors

avoid in practice

~~fragmentation~~
~~slow allocation~~

~~$O(heapSize)$~~ $O(numLives)$
collection cost

non-moving

~~non-moving~~
~~generational~~
~~collection~~

any short-lived objects
frequent allocations
and collections

Functional Programs

Copying Collectors

fast allocation
by "bump pointer"

fragmentation-free
by definition

keep good locality

object moving

good for
generational
collection

Mark-and-Sweep Collectors

avoid in practice

non-moving itself is also a good property!

collection cost

non-moving

non-moving generational collection

any short-lived objects frequent allocations and collections

Functional Programs

Copying Collectors

fast allocation by "bump pointer"

fragmentation-free by definition

keep good locality

object moving

good for generational collection

Viva! Non-moving

Very easy to share heap objects.

- No need to do *pinning* when interacting with C and other languages.

Very easy to maximize concurrency.

- No need to stop threads for managing shared objects.

These are *free* when you choose non-moving!

Mark-and-Sweep Collectors

avoid in practice

~~fragmentation~~
~~slow allocation~~

~~$O(\text{heapSize})$~~ $O(\text{numLives})$
collection cost

non-moving

~~non-moving~~
~~generational~~
~~collection~~

any short-lived objects
frequent allocations
and collections

Functional Programs

Copying Collectors

Good!
fast allocation
by "bump pointer"

Good!
 $O(\text{numLives})$
collection cost

object moving

Good!
good for
generational
collection

Mark-and-Sweep Collectors

avoid in practice

~~fragmentation~~
~~slow allocation~~

~~$O(\text{heapSize})$~~ $O(\text{numLives})$
collection cost

non-moving

~~non-moving~~
~~generational~~
~~collection~~

any short-lived objects
frequent allocations
and collections

Functional Programs

Copying Collectors

Good!
fast allocation
by "bump pointer"

Good!
 $O(\text{numLives})$
collection cost

object moving

Good!
good for
generational
collection

Mark-and-Sweep Collectors

avoid in practice

~~fragmentation~~
~~slow allocation~~

~~$O(\text{heapSize})$~~ $O(\text{numLives})$
collection cost

non-moving

non-moving
generational
collection

any short-lived objects
frequent allocations
and collections

Functional Programs

Copying Collectors

fast allocation
by "bump pointer"

$O(\text{numLives})$
collection cost

object moving

good for
generational
collection

choice

The topic of this talk

We propose
a *non-moving* GC
which is as efficient as
Cheney's copying GC.

Weakness of mark-and-sweep GC

- Fragmentation, and slow allocation
- High sweep cost ($O(\text{heapSize})$)
- No known method for extending it to non-moving generational GC

We choose a well-known idea of **bitmap marking** as our start point to overcome these weaknesses...

Strategy

But bitmap marking strategy alone does not yield an efficient GC.


We re-organize the bitmap marking with

- fragmentation avoiding heap organization
- tree structured bitmap
- automatic heap size adjustment
- non-moving generational extension

with a series of optimized bit manipulation algorithms.

Fragmentation

Varied size objects incur fragmentation.




If all objects were same size, heap could be a fixed size array without incurring fragmentation.




Avoid fragmentation

Separate the heap for each object size.




Avoid fragmentation


Separate the heap for each object size.




8 bytes/block



16 bytes/block



32 bytes/block




64 bytes/block

⋮


Tree-structured bitmaps

For fast allocation and collection



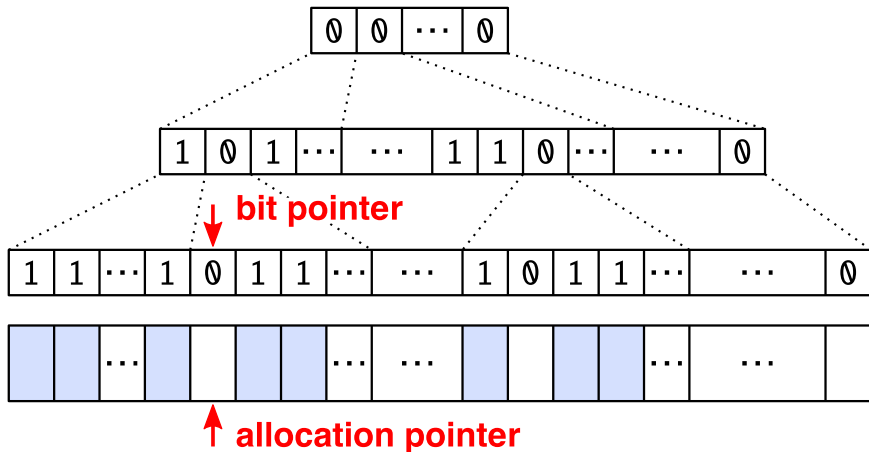
Tree-structured bitmaps

For fast allocation and collection




Tree-structured bitmaps

For fast allocation and collection




Tree-structured bitmaps

For fast allocation and collection



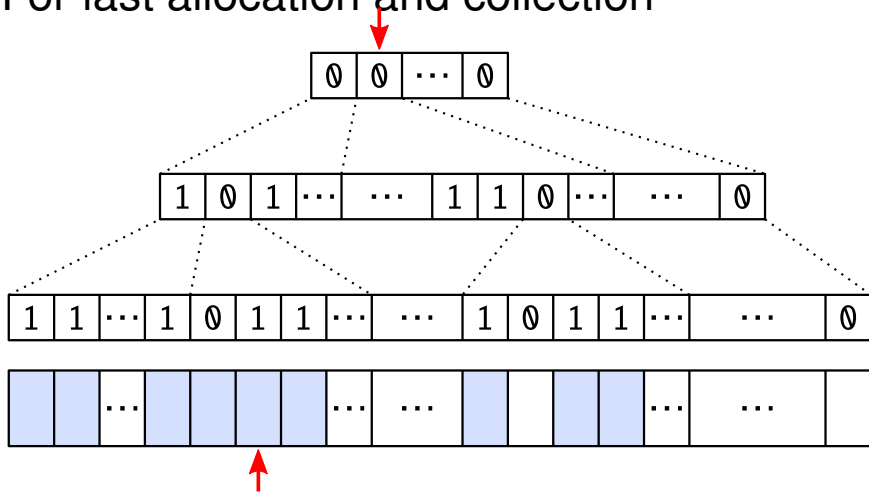
Tree-structured bitmaps

For fast allocation and collection




Tree-structured bitmaps

For fast allocation and collection




Tree-structured bitmaps

For fast allocation and collection




Tree-structured bitmaps

For fast allocation and collection




Tree-structured bitmaps

For fast allocation and collection




Tree-structured bitmaps

For fast allocation and collection




Tree-structured bitmaps


For fast allocation and collection



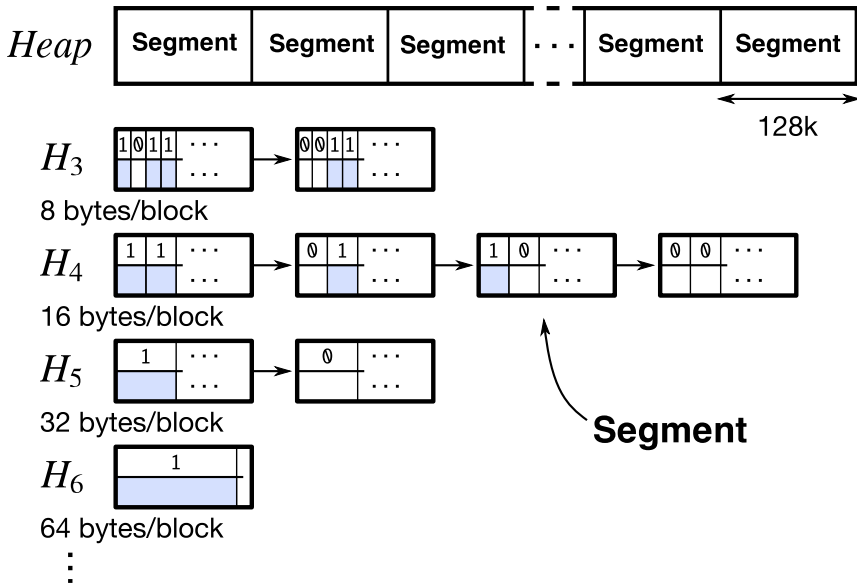
Sub-heap size adjustment



Which layout is appropriate?
... cannot determine it in advance.



Sub-heap size adjustment



Generational extension


Refining the idea of partial GC (Demers et al 1990)

- generations = disjoint sets of objects
- sets of objects = bitmaps

$$\mathcal{G}_3 = \{ \quad O_3, \quad O_4, \quad \dots \}$$

$$\mathcal{G}_2 = \{ \quad O_2, \quad \quad \quad O_6, \quad \quad \quad \dots \}$$

$$\mathcal{G}_1 = \{ O_1, \quad \quad \quad O_5, \quad \quad \quad O_7, O_8, \dots \}$$



Generational extension

Refining the idea of partial GC (Demers et al 1990)

- generations = disjoint sets of objects
- sets of objects = bitmaps

$$\mathcal{G}_3 = \begin{array}{|c|c|c|c|c|c|c|c|c|c|c|c|c|c|} \hline 0 & 0 & \dots & 1 & 0 & 1 & 0 & \dots & \dots & 0 & 0 & 0 & 0 & \dots \\ \hline \end{array}$$

$$\mathcal{G}_3 \cup \mathcal{G}_2 = \begin{array}{|c|c|c|c|c|c|c|c|c|c|c|c|c|c|} \hline 0 & 1 & \dots & 1 & 0 & 1 & 0 & \dots & \dots & 1 & 0 & 0 & 0 & \dots \\ \hline \end{array}$$

$$\mathcal{G}_3 \cup \mathcal{G}_2 \cup \mathcal{G}_1 = \begin{array}{|c|c|c|c|c|c|c|c|c|c|c|c|c|c|} \hline 1 & 1 & \dots & 1 & 0 & 1 & 1 & \dots & \dots & 1 & 0 & 1 & 1 & \dots \\ \hline \end{array}$$

$$\begin{array}{l} \text{block} \\ \text{array} \end{array} \begin{array}{|c|c|c|c|c|c|c|c|c|c|c|c|c|c|} \hline O_1 & O_2 & \dots & O_3 & & O_4 & O_5 & \dots & \dots & O_6 & & O_7 & O_8 & \dots \\ \hline \end{array}$$

Performance evaluation

We compared

- our method,
- Cheney's copying collector,
- our method with 2 generations, and
- generational copying (based on Reppy 1994)




by extensive benchmarks on our SML# compiler.

Memory usage

benchmark	size (MB)	our method		copying	
		live	occ.	live	occ.
count_graphs	2	6.02	55.4	6.02	48.7
cpstak	2	5.56	51.6	5.55	48.9
knuth_bendix	12	10.45	61.1	10.17	46.1
ratio_regions	20	11.97	62.3	11.95	47.4
gcbench	65	10.61	65.9	10.25	42.9
perm9	190	22.36	57.6	16.91	41.6

- live : ratio of survivals against GC
- occ. : ratio of memory amount filled with data

Benchmark : GC time






knuth_bendix

count_graphs

gcbench

Benchmark : mutator time



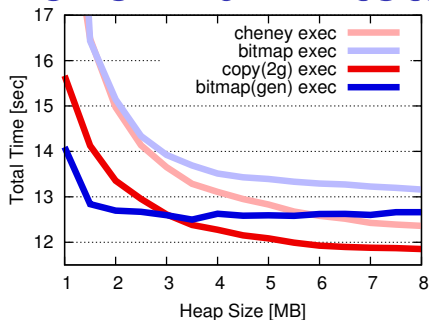


knuth_bendix

count_graphs

gcbench

Benchmark : total exec time



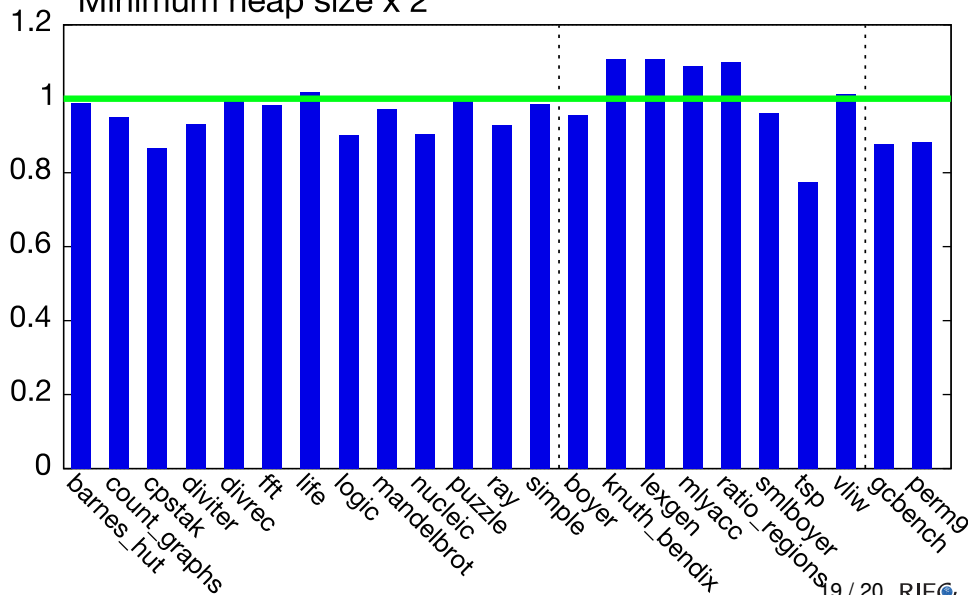
knuth_bendix

count_graphs

gcbench


Ratio of total time (ours / copying)

Minimum heap size x 2




Ratio of total time (ours / copying)

Minimum heap size x 3




Ratio of total time (ours / copying)

Minimum heap size x 4




Ratio of total time (ours / copying)

Minimum heap size x 5




Ratio of total time (ours / copying)

Minimum heap size x 6



Ratio of total time (ours / copying)



Conclusion

We have developed an efficient non-moving GC.

- avoid fragmentation by separating the heap.
- fast allocation and GC through bitmap trees.
- generational GC through multiple bitmaps.

A viable alternative to copying GC for functional languages.

Further Development

non-moving concurrent GC