

[Buy in print and eBook](#)

## Table of Contents

Prologue
I. Language Concepts
II. Tools and Techniques
III. The Runtime System
19. Foreign Function Interface
20. Memory Representation of Values
<a href="#">21. Understanding the Garbage Collector</a>
22. The Compiler Frontend: Parsing and Type Checking
23. The Compiler Backend: Bytecode and Native code
Index

[Login with GitHub to view and add comments](#)

# Chapter 21. Understanding the Garbage Collector

We've described the runtime format of individual OCaml variables earlier, in [Chapter 20, Memory Representation of Values](#). When you execute your program, OCaml manages the lifecycle of these variables by regularly scanning allocated values and freeing them when they're no longer needed. This in turn means that your applications don't need to manually implement memory management, and it greatly reduces the likelihood of memory leaks creeping into your code.

The OCaml runtime is a C library that provides routines that can be called from running OCaml programs. The runtime manages a *heap*, which is a collection of memory regions that it obtains from the operating system. The runtime uses this memory to hold *heap blocks* that it fills up with OCaml values in response to allocation requests by the OCaml program.

## MARK AND SWEEP GARBAGE COLLECTION

When there isn't enough memory available to satisfy an allocation request from the pool of allocated heap blocks, the runtime system invokes the garbage collector (GC). An OCaml program can't explicitly free a value when it is done with it. Instead, the GC regularly determines which values are *live* and which values are *dead*, i.e., no longer in use. Dead values are collected and their memory made available for reuse by the application.

The GC doesn't keep constant track of values as they are allocated and used. Instead, it regularly scans them by starting from a set of *root* values that the application always has access to (such as the stack). The GC maintains a directed graph in which heap blocks are nodes, and there is an edge from heap block *b1* to heap block *b2* if some field of *b1* is a pointer to *b2*.

All blocks reachable from the roots by following edges in the graph must be retained, and unreachable blocks can be reused by the application. The algorithm used by OCaml to perform this heap traversal is commonly known as *mark and sweep* garbage collection, and we'll explain it further now.

## GENERATIONAL GARBAGE COLLECTION

The usual OCaml programming style involves allocating many small variables that are used for a short period of time and then never accessed again. OCaml takes advantage of this fact to improve performance by using a *generational* GC.

A generational GC maintains separate memory regions to hold blocks based on how long the blocks have been live. OCaml's heap is split into two such regions:

- A small, fixed-size *minor heap* where most blocks are initially allocated
- A larger, variable-size *major heap* for blocks that have been live longer

A typical functional programming style means that young blocks tend to die young and old blocks tend to stay around for longer than young ones. This is often referred to as the *generational hypothesis*.

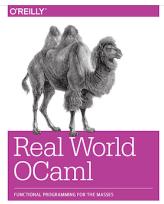
OCaml uses different memory layouts and garbage-collection algorithms for the major and minor heaps to account for this generational difference. We'll explain how they differ in more detail next.

## The Gc Module and OCAMLRUNPARAM

OCaml provides several mechanisms to query and alter the behavior of the runtime system. The `Gc` module provides this functionality from within OCaml code, and we'll frequently refer to it in the rest of the chapter. As with several other standard library modules, Core alters the `Gc` interface from the standard OCaml library. We'll assume that you've opened `Core.Std` in our explanations.

You can also control the behavior of OCaml programs by setting the `OCAMLRUNPARAM` environment variable before launching your application. This lets you set GC parameters without recompiling, for example to benchmark the effects of different settings. The format of `OCAMLRUNPARAM` is documented in the [OCaml manual](#).

## THE FAST MINOR HEAP



[Buy in print and eBook](#)

### Table of Contents

Prologue
I. Language Concepts
II. Tools and Techniques
III. The Runtime System
19. Foreign Function Interface
20. Memory Representation of Values
<a href="#">21. Understanding the Garbage Collector</a>
22. The Compiler Frontend: Parsing and Type Checking
23. The Compiler Backend: Bytecode and Native code
Index

[Login with GitHub to view and add comments](#)

The minor heap is where most of your short-lived values are held. It consists of one contiguous chunk of virtual memory containing a sequence of OCaml blocks. If there is space, allocating a new block is a fast, constant-time operation that requires just a couple of CPU instructions.

To garbage-collect the minor heap, OCaml uses *copying collection* to move all live blocks in the minor heap to the major heap. This takes work proportional to the number of live blocks in the minor heap, which is typically small according to the generational hypothesis. The minor collection *stops the world* (that is, it halts the application) while it runs, which is why it's so important that it complete quickly to let the application resume running with minimal interruption.

### Allocating on the Minor Heap

The minor heap is a contiguous chunk of virtual memory that is usually a few megabytes in size so that it can be scanned quickly.

```

      <---- size ---->
base --- start ----- end
          limit      ptr <----->
                           blocks

```

Diagram \* gc/minor\_heap.ascii \* all code

The runtime stores the boundaries of the minor heap in two pointers that delimit the start and end of the heap region (`caml_young_start` and `caml_young_end`, but we will drop the `caml_` prefix for brevity). The `base` is the memory address returned by the system `malloc`, and `start` is aligned against the next nearest word boundary from `base` to make it easier to store OCaml values.

In a fresh minor heap, the `limit` equals the `start`, and the current `ptr` will equal the `end`. `ptr` decreases as blocks are allocated until it reaches `limit`, at which point a minor garbage collection is triggered.

Allocating a block in the minor heap just requires `ptr` to be decremented by the size of the block (including the header) and a check that it's not less than `limit`. If there isn't enough space left for the block without decrementing past `limit`, a minor garbage collection is triggered. This is a very fast check (with no branching) on most CPU architectures.

You may wonder why `limit` is required at all, since it always seems to equal `start`. It's because the easiest way for the runtime to schedule a minor heap collection is by setting `limit` to equal `end`. The next allocation will never have enough space after this is done and will always trigger a garbage collection. There are various internal reasons for such early collections, such as handling pending UNIX signals, and they don't ordinarily matter for application code.

### Setting the Size of the Minor Heap

The default minor heap size in OCaml is normally 2 MB on 64-bit platforms, but this is increased to 8 MB if you use Core (which generally prefers default settings that improve performance, but at the cost of a bigger memory profile). This setting can be overridden via the `s=<words>` argument to `OCAMLRUNPARAM`. You can change it after the program has started by calling the `Gc.set` function:

```

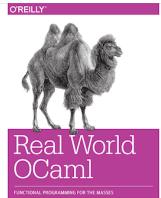
# let c = Gc.get () ;;
val c : Gc.control =
  {Core.Std.Gc.Control.minor_heap_size = 1000000;
   major_heap_increment = 1000448; space_overhead = 100; verbose = 0;
   max_overhead = 500; stack_limit = 1048576; allocation_policy = 0}
# Gc.tune ~minor_heap_size:(262144 * 2) () ;;
- : unit = ()

```

OCaml Utop \* gc/tune.topscript \* all code

Changing the GC size dynamically will trigger an immediate minor heap collection. Note that Core increases the default minor heap size from the standard OCaml installation quite significantly, and you'll want to reduce this if running in very memory-constrained environments.

## THE LONG-LIVED MAJOR HEAP



[Buy in print and eBook](#)

## Table of Contents

### Prologue

I. Language Concepts
II. Tools and Techniques
III. The Runtime System
19. Foreign Function Interface
20. Memory Representation of Values
<a href="#">21. Understanding the Garbage Collector</a>
22. The Compiler Frontend: Parsing and Type Checking
23. The Compiler Backend: Bytecode and Native code
Index

[Login with GitHub to view and add comments](#)

The major heap is where the bulk of the longer-lived and larger values in your program are stored. It consists of any number of noncontiguous chunks of virtual memory, each containing live blocks interspersed with regions of free memory. The runtime system maintains a free-list data structure that indexes all the free memory that it has allocated, and uses it to satisfy allocation requests for OCaml blocks.

The major heap is typically much larger than the minor heap and can scale to gigabytes in size. It is cleaned via a mark-and-sweep garbage collection algorithm that operates in several phases:

- The *mark* phase scans the block graph and marks all live blocks by setting a bit in the tag of the block header (known as the *color* tag).
- The *sweep* phase sequentially scans the heap chunks and identifies dead blocks that weren't marked earlier.
- The *compact* phase relocates live blocks into a freshly allocated heap to eliminate gaps in the free list. This prevents the fragmentation of heap blocks in long-running programs and normally occurs much less frequently than the mark and sweep phases.

A major garbage collection must also stop the world to ensure that blocks can be moved around without this being observed by the live application. The mark-and-sweep phases run incrementally over slices of the heap to avoid pausing the application for long periods of time, and also precede each slice with a fast minor collection. Only the compaction phase touches all the memory in one go, and is a relatively rare operation.

### Allocating on the Major Heap

The major heap consists of a singly linked list of contiguous memory chunks sorted in increasing order of virtual address. Each chunk is a single memory region allocated via *malloc(3)* and consists of a header and data area which contains OCaml heap chunks. A heap chunk header contains:

- The *malloced* virtual address of the memory region containing the chunk
- The size in bytes of the data area
- An allocation size in bytes used during heap compaction to merge small blocks to defragment the heap
- A link to the next heap chunk in the list

Each chunk's data area starts on a page boundary, and its size is a multiple of the page size (4 KB). It contains a contiguous sequence of heap blocks that can be as small as one or two 4 KB pages, but are usually allocated in 1 MB chunks (or 512 KB on 32-bit architectures).

### Controlling Major Heap Growth

The `Gc` module uses the `major_heap_increment` value to control the major heap growth. This defines the number of words to add to the major heap per expansion and is the only memory allocation operation that the operating system observes from the OCaml runtime after initial startup (since the minor is fixed in size).

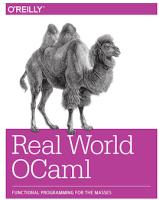
If you anticipate allocating some large OCaml values or many small values in one go, then setting the heap increment to a larger value will improve performance by reducing the amount of heap resizing required in order to satisfy the allocation requests. A small increment may result in lots of smaller heap chunks spread across different regions of virtual memory that require more housekeeping in the OCaml runtime to keep track of them:

```
# Gc.tune ~major_heap_increment:(1000448 * 4) () ;;
- : unit = ()
```

OCaml Utop \* gc/tune.topscript , continued (part 1) \* all code

Allocating an OCaml value on the major heap first checks the free list of blocks for a suitable region to place it. If there isn't enough room on the free list, the runtime expands the major heap by allocating a fresh heap chunk that will be large enough. That chunk is then added to the free list, and the free list is checked again (and this time will definitely succeed).

Remember that most allocations to the major heap will go via the minor heap and only be promoted if they are still used by the program after a minor collection. The one exception to this



[Buy in print and eBook.](#)

## Table of Contents

Prologue
I. Language Concepts
II. Tools and Techniques
III. The Runtime System
19. Foreign Function Interface
20. Memory Representation of Values
21. Understanding the Garbage Collector
22. The Compiler Frontend: Parsing and Type Checking
23. The Compiler Backend: Bytecode and Native code
Index

[Login with GitHub to view and add comments](#)

is for values larger than 256 words (that is, 2 KB on 64-bit platforms). These will be allocated directly on the major heap, since an allocation on the minor heap would likely trigger an immediate collection and copy it to the major heap anyway.

## Memory Allocation Strategies

The major heap does its best to manage memory allocation as efficiently as possible and relies on heap compaction to ensure that memory stays contiguous and unfragmented. The default allocation policy normally works fine for most applications, but it's worth bearing in mind that there are other options, too.

The free list of blocks is always checked first when allocating a new block in the major heap. The default free list search is called *next-fit allocation*, with an alternative *first-fit* algorithm also available.

### Next-fit allocation

Next-fit allocation keeps a pointer to the block in the free list that was most recently used to satisfy a request. When a new request comes in, the allocator searches from the next block to the end of the free list, and then from the beginning of the free list up to that block.

Next-fit allocation is the default allocation strategy. It's quite a cheap allocation mechanism, since the same heap chunk can be reused across allocation requests until it runs out. This in turn means that there is good memory locality to use CPU caches better.

### First-fit allocation

If your program allocates values of many varied sizes, you may sometimes find that your free list becomes fragmented. In this situation, the GC is forced to perform an expensive compaction despite there being free chunks, since none of the chunks alone are big enough to satisfy the request.

First-fit allocation focuses on reducing memory fragmentation (and hence the number of compactations), but at the expense of slower memory allocation. Every allocation scans the free list from the beginning for a suitable free chunk, instead of reusing the most recent heap chunk as the next-fit allocator does.

For some workloads that need more real-time behavior under load, the reduction in the frequency of heap compaction will outweigh the extra allocation cost.

### Controlling the Heap Allocation Policy

You can set the heap allocation policy via the `Gc.allocation_policy` field. A value of 0 (the default) sets it to next-fit, and 1 to the first-fit allocator.

The same behavior can be controlled at runtime by setting `a=0` or `a=1` in `OCAMLRUNPARAM`.

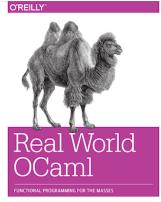
## Marking and Scanning the Heap

The marking process can take a long time to run over the complete major heap and has to pause the main application while it's active. It therefore runs incrementally by marking the heap in *slices*. Each value in the heap has a 2-bit *color* field in its header that is used to store information about whether the value has been marked so that the GC can resume easily between slices.

**Table 21.1. Tag color statuses**

Tag color	Block status
Blue	On the free list and not currently in use
White (during marking)	Not reached yet, but possibly reachable
White (during sweeping)	Unreachable and can be freed
Gray	Reachable, but its fields have not been scanned
Black	Reachable, and its fields have been scanned

The color tags in the value headers store most of the state of the marking process, allowing it to be paused and resumed later. The GC and application alternate between marking a slice of the major heap and actually getting on with executing the program logic. The OCaml runtime calculates a sensible value for the size of each major heap slice based on the rate of allocation and available memory.



[Buy in print and eBook.](#)

## Table of Contents

### Prologue

- I. Language Concepts
- II. Tools and Techniques
- III. The Runtime System
  - 19. Foreign Function Interface
  - 20. Memory Representation of Values
  - [21. Understanding the Garbage Collector](#)
  - 22. The Compiler Frontend: Parsing and Type Checking
  - 23. The Compiler Backend: Bytecode and Native code
- Index

[Login with GitHub to view and add comments](#)

The marking process starts with a set of *root* values that are always live (such as the application stack). All values on the heap are initially marked as white values that are possibly reachable but haven't been scanned yet. It recursively follows all the fields in the roots via a depth-first search, and pushes newly encountered white blocks onto an intermediate stack of *gray values* while it follows their fields. When a gray value's fields have all been followed, it is popped off the stack and colored black.

This process is repeated until the gray value stack is empty and there are no further values to mark. There's one important edge case in this process, though. The gray value stack can only grow to a certain size, after which the GC can no longer recurse into intermediate values since it has nowhere to store them while it follows their fields. If this happens, the heap is marked as *impure* and a more expensive check is initiated once the existing gray values have been processed.

To mark an impure heap, the GC first marks it as pure and walks through the entire heap block-by-block in increasing order of memory address. If it finds a gray block, it adds it to the gray list and recursively marks it using the usual strategy for a pure heap. Once the scan of the complete heap is finished, the mark phase checks again whether the heap has again become impure and repeats the scan until it is pure again. These full-heap scans will continue until a successful scan completes without overflowing the gray list.

### Controlling Major Heap Collections

You can trigger a single slice of the major GC via the `major_slice` call. This performs a minor collection first, and then a single slice. The size of the slice is normally automatically computed by the GC to an appropriate value and returns this value so that you can modify it in future calls if necessary:

```
# Gc.major_slice 0 ;;
- : int = 280015
# Gc.full_major () ;;
- : unit = ()
```

OCaml Utop \* gc/tune.topscript , continued (part 2) \* all code

The `space_overhead` setting controls how aggressive the GC is about setting the slice size to a large size. This represents the proportion of memory used for live data that will be "wasted" because the GC doesn't immediately collect unreachable blocks. Core defaults this to `100` to reflect a typical system that isn't overly memory-constrained. Set this even higher if you have lots of memory, or lower to cause the GC to work harder and collect blocks faster at the expense of using more CPU time.

### Heap Compaction

After a certain number of major GC cycles have completed, the heap may begin to be fragmented due to values being deallocated out of order from how they were allocated. This makes it harder for the GC to find a contiguous block of memory for fresh allocations, which in turn would require the heap to be grown unnecessarily.

The heap compaction cycle avoids this by relocating all the values in the major heap into a fresh heap that places them all contiguously in memory again. A naive implementation of the algorithm would require extra memory to store the new heap, but OCaml performs the compaction in place within the existing heap.

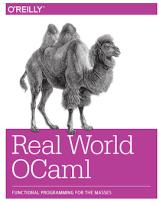
### Controlling Frequency of Compactions

The `max_overhead` setting in the `Gc` module defines the connection between free memory and allocated memory after which compaction is activated.

A value of `0` triggers a compaction after every major garbage collection cycle, whereas the maximum value of `1000000` disables heap compaction completely. The default settings should be fine unless you have unusual allocation patterns that are causing a higher-than-usual rate of compactions:

```
# Gc.tune ~max_overhead:0 () ;;
- : unit = ()
```

OCaml Utop \* gc/tune.topscript , continued (part 3) \* all code



[Buy in print and eBook.](#)

## Table of Contents

### Prologue

- I. Language Concepts
- II. Tools and Techniques
- III. The Runtime System
  - 19. Foreign Function Interface
  - 20. Memory Representation of Values
  - [21. Understanding the Garbage Collector](#)
  - 22. The Compiler Frontend: Parsing and Type Checking
  - 23. The Compiler Backend: Bytecode and Native code

### Index

[Login with GitHub to view and add comments](#)

## Intergenerational Pointers

One complexity of generational collection arises from the fact that minor heap sweeps are much more frequent than major heap collections. In order to know which blocks in the minor heap are live, the collector must track which minor-heap blocks are directly pointed to by major-heap blocks. Without this information, each minor collection would also require scanning the much larger major heap.

OCaml maintains a set of such *intergenerational pointers* to avoid this dependency between a major and minor heap collection. The compiler introduces a write barrier to update this so-called *remembered set* whenever a major-heap block is modified to point at a minor-heap block.

### The mutable write barrier

The write barrier can have profound implications for the structure of your code. It's one of the reasons using immutable data structures and allocating a fresh copy with changes can sometimes be faster than mutating a record in place.

The OCaml compiler keeps track of any mutable types and adds a call to the runtime `caml_modify` function before making the change. This checks the location of the target write and the value it's being changed to, and ensures that the remembered set is consistent. Although the write barrier is reasonably efficient, it can sometimes be slower than simply allocating a fresh value on the fast minor heap and doing some extra minor collections.

Let's see this for ourselves with a simple test program. You'll need to install the Core benchmarking suite via `opam install core_bench` before you compile this code:

```
open Core.Std
open Core_bench.Std

type t1 = { mutable iters1: int; mutable count1: float }
type t2 = { iters2: int; count2: float }

let rec test Mutable t1 =
  match t1.iters1 with
  | 0 -> ()
  | _ ->
    t1.iters1 <- t1.iters1 - 1;
    t1.count1 <- t1.count1 +. 1.0;
    test Mutable t1

let rec test Immutable t2 =
  match t2.iters2 with
  | 0 -> ()
  | n ->
    let iters2 = n - 1 in
    let count2 = t2.count2 +. 1.0 in
    test Immutable { iters2; count2 }

let () =
  let iters = 1000000 in
  let tests = [
    Bench.Test.create ~name:"mutable"
      (fun () -> test Mutable { iters1=iters; count1=0.0 });
    Bench.Test.create ~name:"immutable"
      (fun () -> test Immutable { iters2=iters; count2=0.0 })
  ] in
  Bench.make_command tests |> Command.run

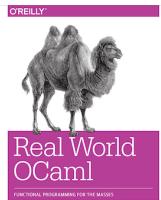
OCaml * gc/barrier_bench.ml * all code
```

This program defines a type `t1` that is mutable and `t2` that is immutable. The benchmark loop iterates over both fields and increments a counter. Compile and execute this with some extra options to show the amount of garbage collection occurring:

```
$ corebuild -pkg core_bench barrier_bench.native
$ ./barrier_bench.native -ascii name alloc
Estimated testing time 20s (change using -quota SECS).
```

Name	Time/Run	Minor	Major	Promoted	% of max
mutable	4_862_333	2_000_005	9.58	9.58	100.00
immutable	4_558_309	5_000_006	0.56	0.56	93.75

Terminal \* gc/run\_barrier\_bench.out \* all code

Buy in [print](#) and [eBook](#).

## Table of Contents

Prologue
I. Language Concepts
II. Tools and Techniques
III. The Runtime System
19. Foreign Function Interface
20. Memory Representation of Values
<a href="#">21. Understanding the Garbage Collector</a>
22. The Compiler Frontend: Parsing and Type Checking
23. The Compiler Backend: Bytecode and Native code
Index

Login with GitHub to view and add comments

There is a stark space/time trade-off here. The mutable version takes significantly longer to complete than the immutable one but allocates many fewer minor-heap words than the immutable version. Minor allocation in OCaml is very fast, and so it is often better to use immutable data structures in preference to the more conventional mutable versions. On the other hand, if you only rarely mutate a value, it can be faster to take the write-barrier hit and not allocate at all.

The only way to know for sure is to benchmark your program under real-world scenarios using `Core_bench` and experiment with the trade-offs. The command-line benchmark binaries have a number of useful options that affect garbage collection behavior:

```
$ ./barrier_bench.native -help
Benchmark for mutable, immutable

barrier_bench.native [COLUMN ...]

Columns that can be specified are:
  name      - Name of the test.
  cycles    - Number of CPU cycles (RDTSC) taken.
  time      - Number of nano secs taken.
  confidence - 95% confidence interval and R^2 error for predictors.
  alloc     - Allocation of major, minor and promoted words.
  gc        - Show major and minor collections per 1000 runs.
  percentage - Relative execution time as a percentage.
  speedup   - Relative execution cost as a speedup.
  samples   - Number of samples collected for profiling.
```

*R<sup>2</sup> is the fraction of the variance of the responder (such as runtime) that is accounted for by the predictors (such as number of runs). More informally, it describes how good a fit we're getting, with R<sup>2</sup> = 1 indicating a perfect fit and R<sup>2</sup> = 0 indicating a horrible fit. Because we expect runtime to be very highly correlated with our predictors, values very close to 1 are typical; a value less than 0.99 should cause some suspicion, and a value less than 0.9 probably indicates either a shortage of data or that the data is erroneous or peculiar in some way.*

*Also see: [http://en.wikipedia.org/wiki/Coefficient\\_of\\_determination](http://en.wikipedia.org/wiki/Coefficient_of_determination)*

*GC stats indicate how many collections or compactions happen per 1000 runs of the benchmarked function.*

*The following columns will be displayed by default:*

```
name time +percentage
```

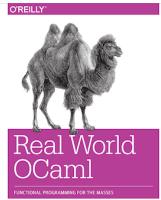
*By default, columns that have no values are suppressed. To force displaying empty columns, prefix the column name with a '+'.*

*Experimental feature: Internally, the library does a linear regression between the time taken as the predicted value and the number of runs as the predictor. This can be changed to include one or more of the additional predictors below, using the flag called "-predictors":*

```
m : the number of minor collections
c : the number of compactions
```

**==== flags ===**

<code>[-ascii]</code>	<i>Display data in simple ascii based tables.</i>
<code>[-ci-absolute]</code>	<i>Display 95% confidence interval in absolute numbers</i>
<code>[-clear-columns]</code>	<i>Don't display default columns. Only show user specified ones.</i>
<code>[-display STYLE]</code>	<i>Table style (short, tall, line, blank or column). Default short.</i>
<code>[-fork]</code>	<i>Fork and run each benchmark in separate child-process</i>
<code>[-geometric SCALE]</code>	<i>Use geometric sampling. (default 1.01)</i>
<code>[-linear INCREMENT]</code>	<i>Use linear sampling to explore number of runs, example 1.</i>
<code>[-Load FILE]</code>	<i>Analyze previously saved data files and don't run tests. <code>[-Load]</code> can be specified multiple times.</i>



[Buy in print and eBook](#)

## Table of Contents

Prologue
I. Language Concepts
II. Tools and Techniques
III. The Runtime System
19. Foreign Function Interface
20. Memory Representation of Values
<a href="#">21. Understanding the Garbage Collector</a>
22. The Compiler Frontend: Parsing and Type Checking
23. The Compiler Backend: Bytecode and Native code
Index

[Login with GitHub to view and add comments](#)

<code>[-no-compactions]</code>	<i>Disable GC compactions.</i>
<code>[-predictors m,c]</code>	<i>Include additional predictors in regression (see help).</i>
<code>[-quota SECS]</code>	<i>Time quota allowed per test (default 10s).</i>
<code>[-save]</code>	<i>Save benchmark data to &lt;test name&gt;.txt files.</i>
<code>[-stabilize-gc]</code>	<i>Stabilize GC between each sample capture.</i>
<code>[-v]</code>	<i>High verbosity level.</i>
<code>[-width WIDTH]</code>	<i>width limit on column display (default 170).</i>
<code>[-build-info]</code>	<i>print info about this build and exit</i>
<code>[-version]</code>	<i>print the version of this build and exit</i>
<code>[-help]</code>	<i>print this help text and exit</i>
	<i>(alias: -?)</i>

Terminal \* gc/show\_barrier\_bench\_help.out \* all code

The `-no-compactions` and `-stabilize-gc` options can help force a situation where your application has fragmented memory. This can simulate the behavior of a long-running application without you having to actually wait that long to re-create the behavior in a performance unit test.

## ATTACHING FINALIZER FUNCTIONS TO VALUES

OCaml's automatic memory management guarantees that a value will eventually be freed when it's no longer in use, either via the GC sweeping it or the program terminating. It's sometimes useful to run extra code just before a value is freed by the GC, for example, to check that a file descriptor has been closed, or that a log message is recorded.

### What Values Can Be Finalized?

Various values cannot have finalizers attached since they aren't heap-allocated. Some examples of values that are not heap-allocated are integers, constant constructors, Booleans, the empty array, the empty list, and the unit value. The exact list of what is heap-allocated or not is implementation-dependent, which is why Core provides the `Heap_block` module to explicitly check before attaching the finalizer.

Some constant values can be heap-allocated but never deallocated during the lifetime of the program, for example, a list of integer constants. `Heap_block` explicitly checks to see if the value is in the major or minor heap, and rejects most constant values. Compiler optimizations may also duplicate some immutable values such as floating-point values in arrays. These may be finalized while another duplicate copy is being used by the program.

For this reason, attach finalizers only to values that you are explicitly sure are heap-allocated and aren't immutable. A common use is to attach them to file descriptors to ensure they are closed. However, the finalizer normally shouldn't be the primary way of closing the file descriptor, since it depends on the GC running in order to collect the value. For a busy system, you can easily run out of a scarce resource such as file descriptors before the GC catches up.

Core provides a `Heap_block` module that dynamically checks if a given value is suitable for finalizing. This block is then passed to `Async's Gc.add_finalizer` function that schedules the finalizer safely with respect to all the other concurrent program threads.

Let's explore this with a small example that finalizes values of different types, some of which are heap-allocated and others which are compile-time constants:

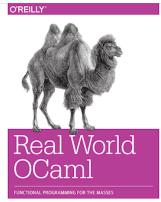
```

open Core.Std
open Async.Std

let attach_finalizer n v =
  match Heap_block.create v with
  | None -> printf "%20s: FAIL\n%" n
  | Some hb ->
    let final _ = printf "%20s: OK\n%" n in
    Gc.add_finalizer hb final

type t = { foo: bool }

let main () =
  let alloced_float = Unix.gettimeofday () in
  let alloced_bool = alloced_float > 0.0 in
  
```



[Buy in print and eBook](#)

## Table of Contents

Prologue
I. Language Concepts
II. Tools and Techniques
III. The Runtime System
19. Foreign Function Interface
20. Memory Representation of Values
<a href="#">21. Understanding the Garbage Collector</a>
22. The Compiler Frontend: Parsing and Type Checking
23. The Compiler Backend: Bytecode and Native code
Index

[Login with GitHub to view and add comments](#)

```
let allocoed_string = String.create 4 in
  attach_finalizer "immediate int" 1;
  attach_finalizer "immediate float" 1.0;
  attach_finalizer "immediate variant" ( `Foo "hello");
  attach_finalizer "immediate string" "hello world";
  attach_finalizer "immediate record" { foo=false };
  attach_finalizer "allocated float" allocoed_float;
  attach_finalizer "allocated bool" allocoed_bool;
  attach_finalizer "allocated variant" ( `Foo allocoed_bool);
  attach_finalizer "allocated string" allocoed_string;
  attach_finalizer "allocated record" { foo=allocoed_bool };
  Gc.compact ();
  return ()
```

```
let () =
  Command.async_basic ~summary:"Testing finalizers"
  Command.Spec.empty main
  |> Command.run
```

OCaml \* gc/finalizer.ml \* all code

Building and running this should show the following output:

```
$ corebuild -pkg async finalizer.native
$ ./finalizer.native
  immediate int: FAIL
  immediate float: FAIL
  immediate variant: FAIL
  immediate string: FAIL
  immediate record: FAIL
  allocated bool: FAIL
  allocated record: OK
  allocated string: OK
  allocated variant: OK
  allocated float: OK
```

Terminal \* gc/run\_finalizer.out \* all code

The GC calls the finalization functions in the order of the deallocation. If several values become unreachable during the same GC cycle, the finalization functions will be called in the reverse order of the corresponding calls to `add_finalizer`. Each call to `add_finalizer` adds to the set of functions, which are run when the value becomes unreachable. You can have many finalizers all pointing to the same heap block if you wish.

After a garbage collection determines that a heap block `b` is unreachable, it removes from the set of finalizers all the functions associated with `b`, and serially applies each of those functions to `b`. Thus, every finalizer function attached to `b` will run at most once. However, program termination will not cause all the finalizers to be run before the runtime exits.

The finalizer can use all features of OCaml, including assignments that make the value reachable again and thus prevent it from being garbage-collected. It can also loop forever, which will cause other finalizers to be interleaved with it.

[< Previous](#)

[Next >](#)