

Building the Next-Generation Garbage Collector in Ruby

Peter Zhu

**Ruby Core Committer
Staff Software Engineer, Figma**

blog.peterzhu.ca/assets/rubykaigi_2026_slides.pdf



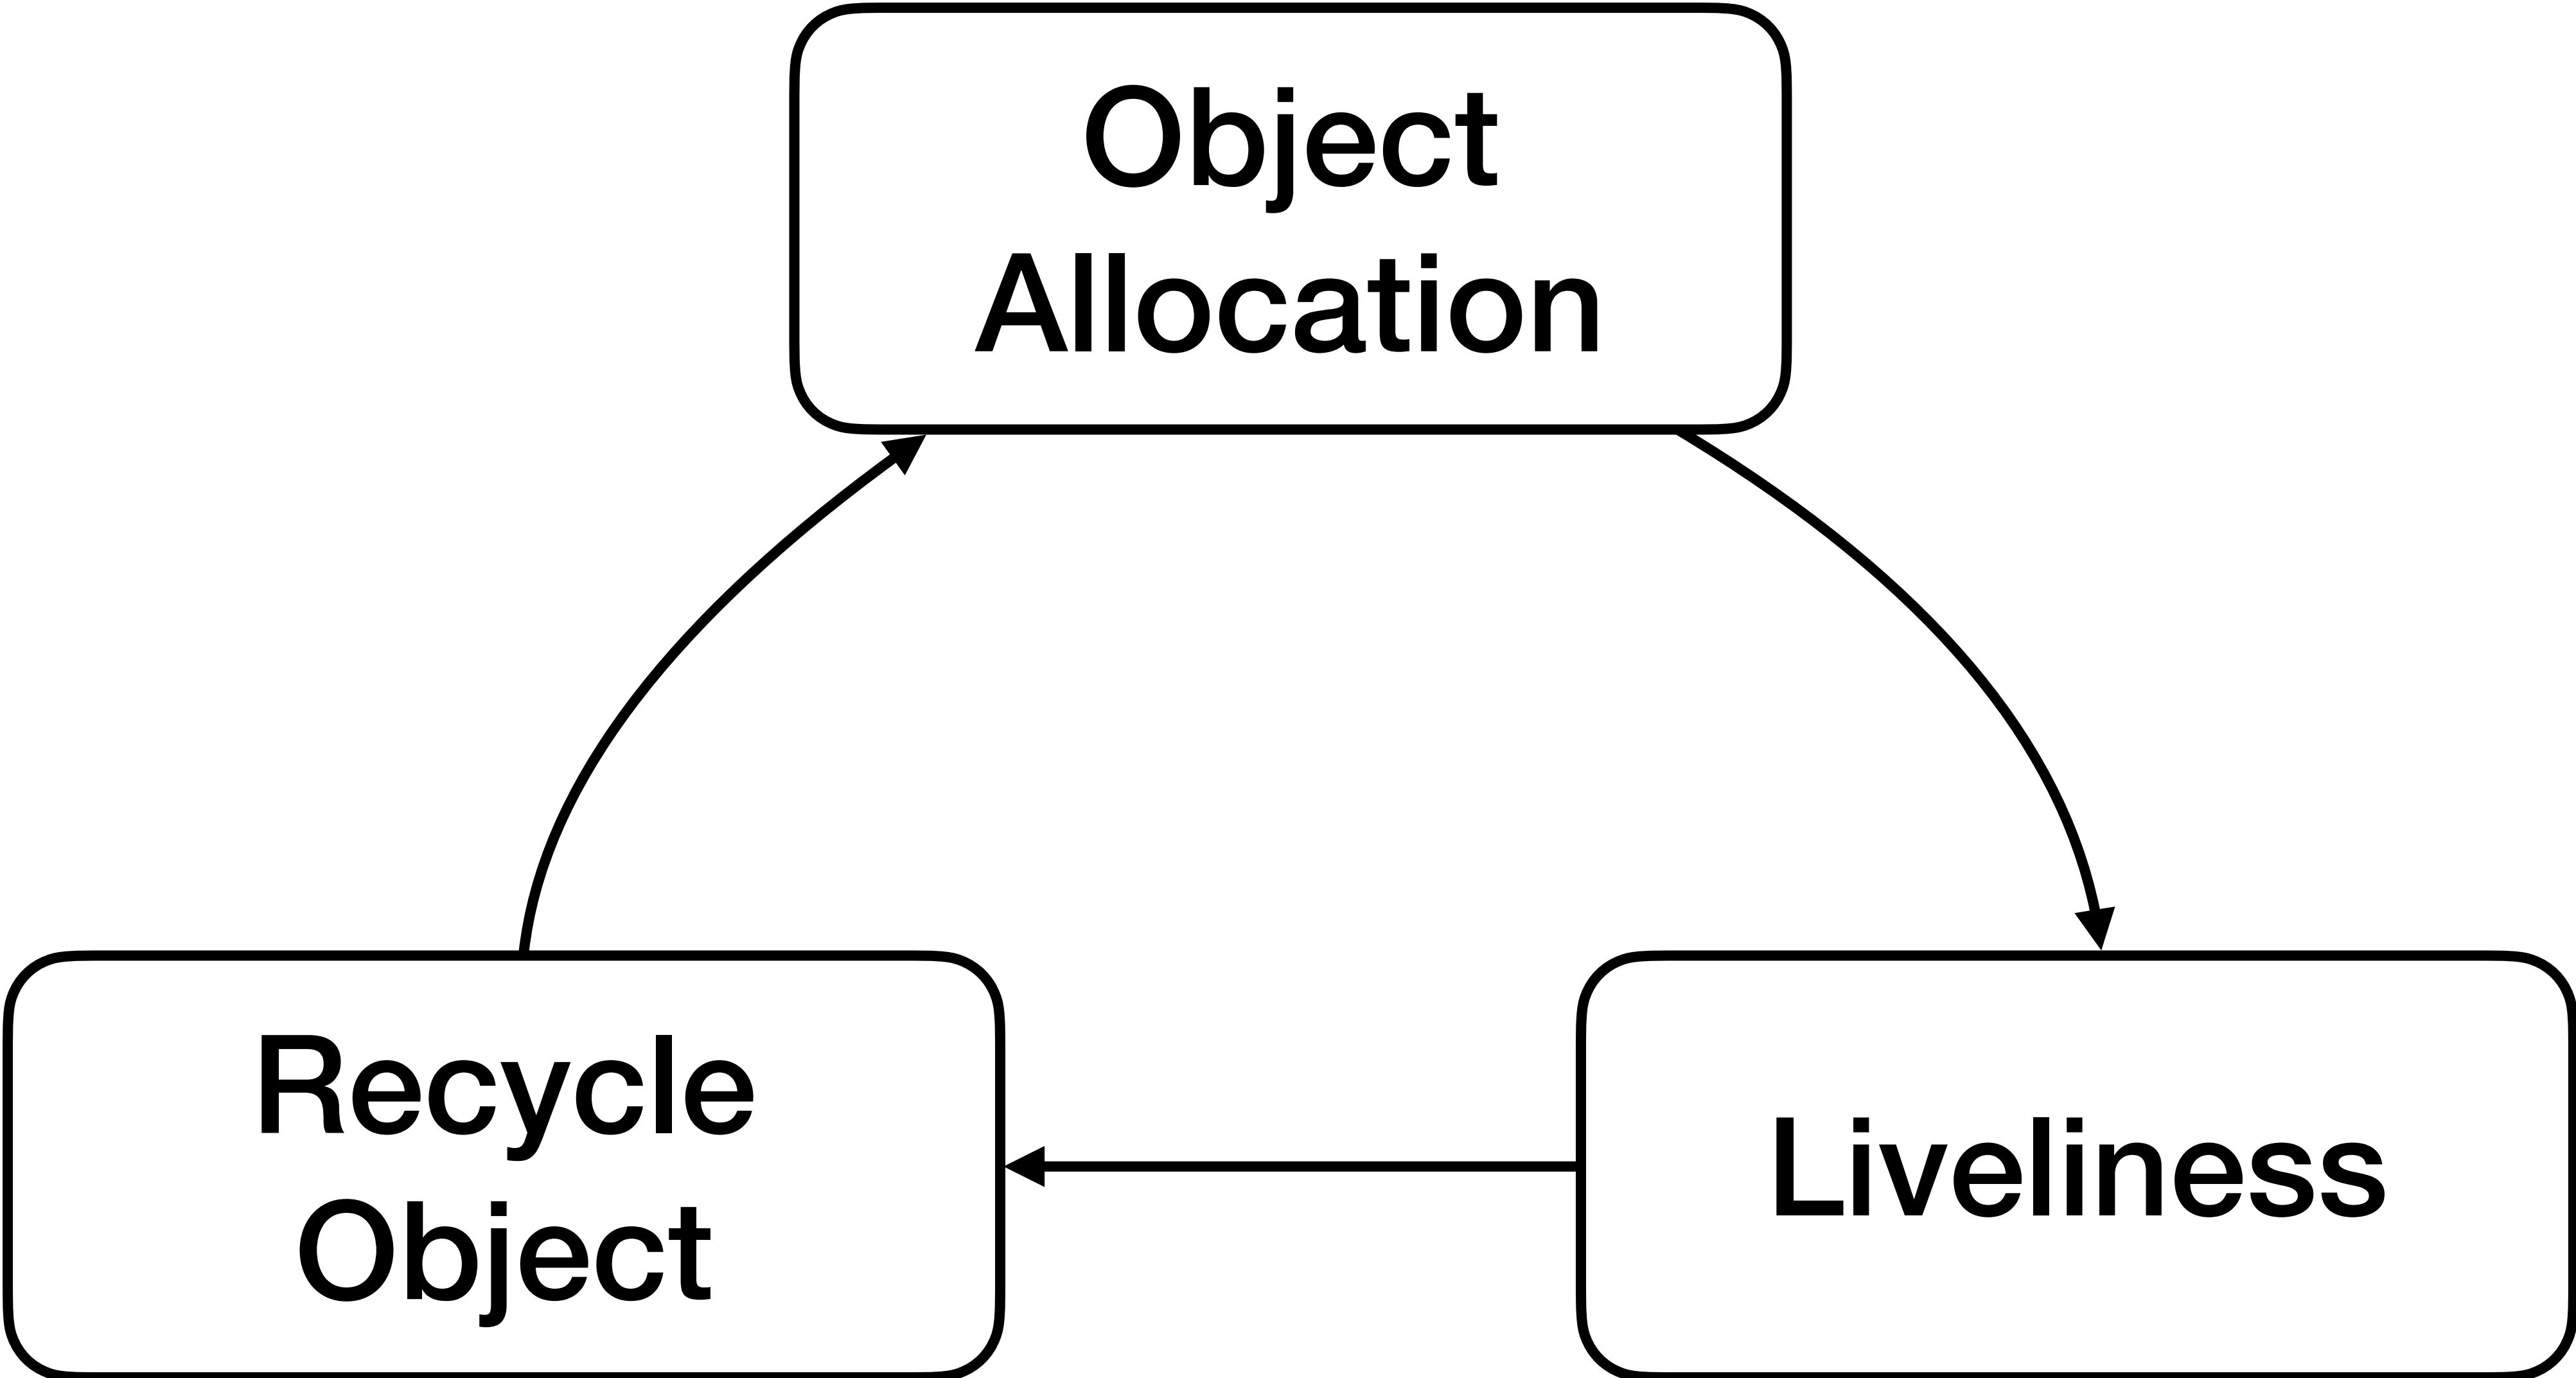
Peter Zhu

- Based in Toronto, Canada
- Ruby Core Committer
- Staff Software Engineer at Figma
- Co-author of Variable Width Allocation, RUBY_FREE_AT_EXIT, and Modular Garbage Collectors in Ruby
- Author of `ruby_memcheck` and `autotuner`



- Introduce Ruby's Garbage Collector
- Modular Garbage Collectors
- Memory Management Toolkit
 - Improvements since RubyKaigi 2025
 - Moving Immix
 - Bump Pointer Allocator Fast Path
 - Ruby Heap
 - Future Roadmap

What is a Garbage Collector?



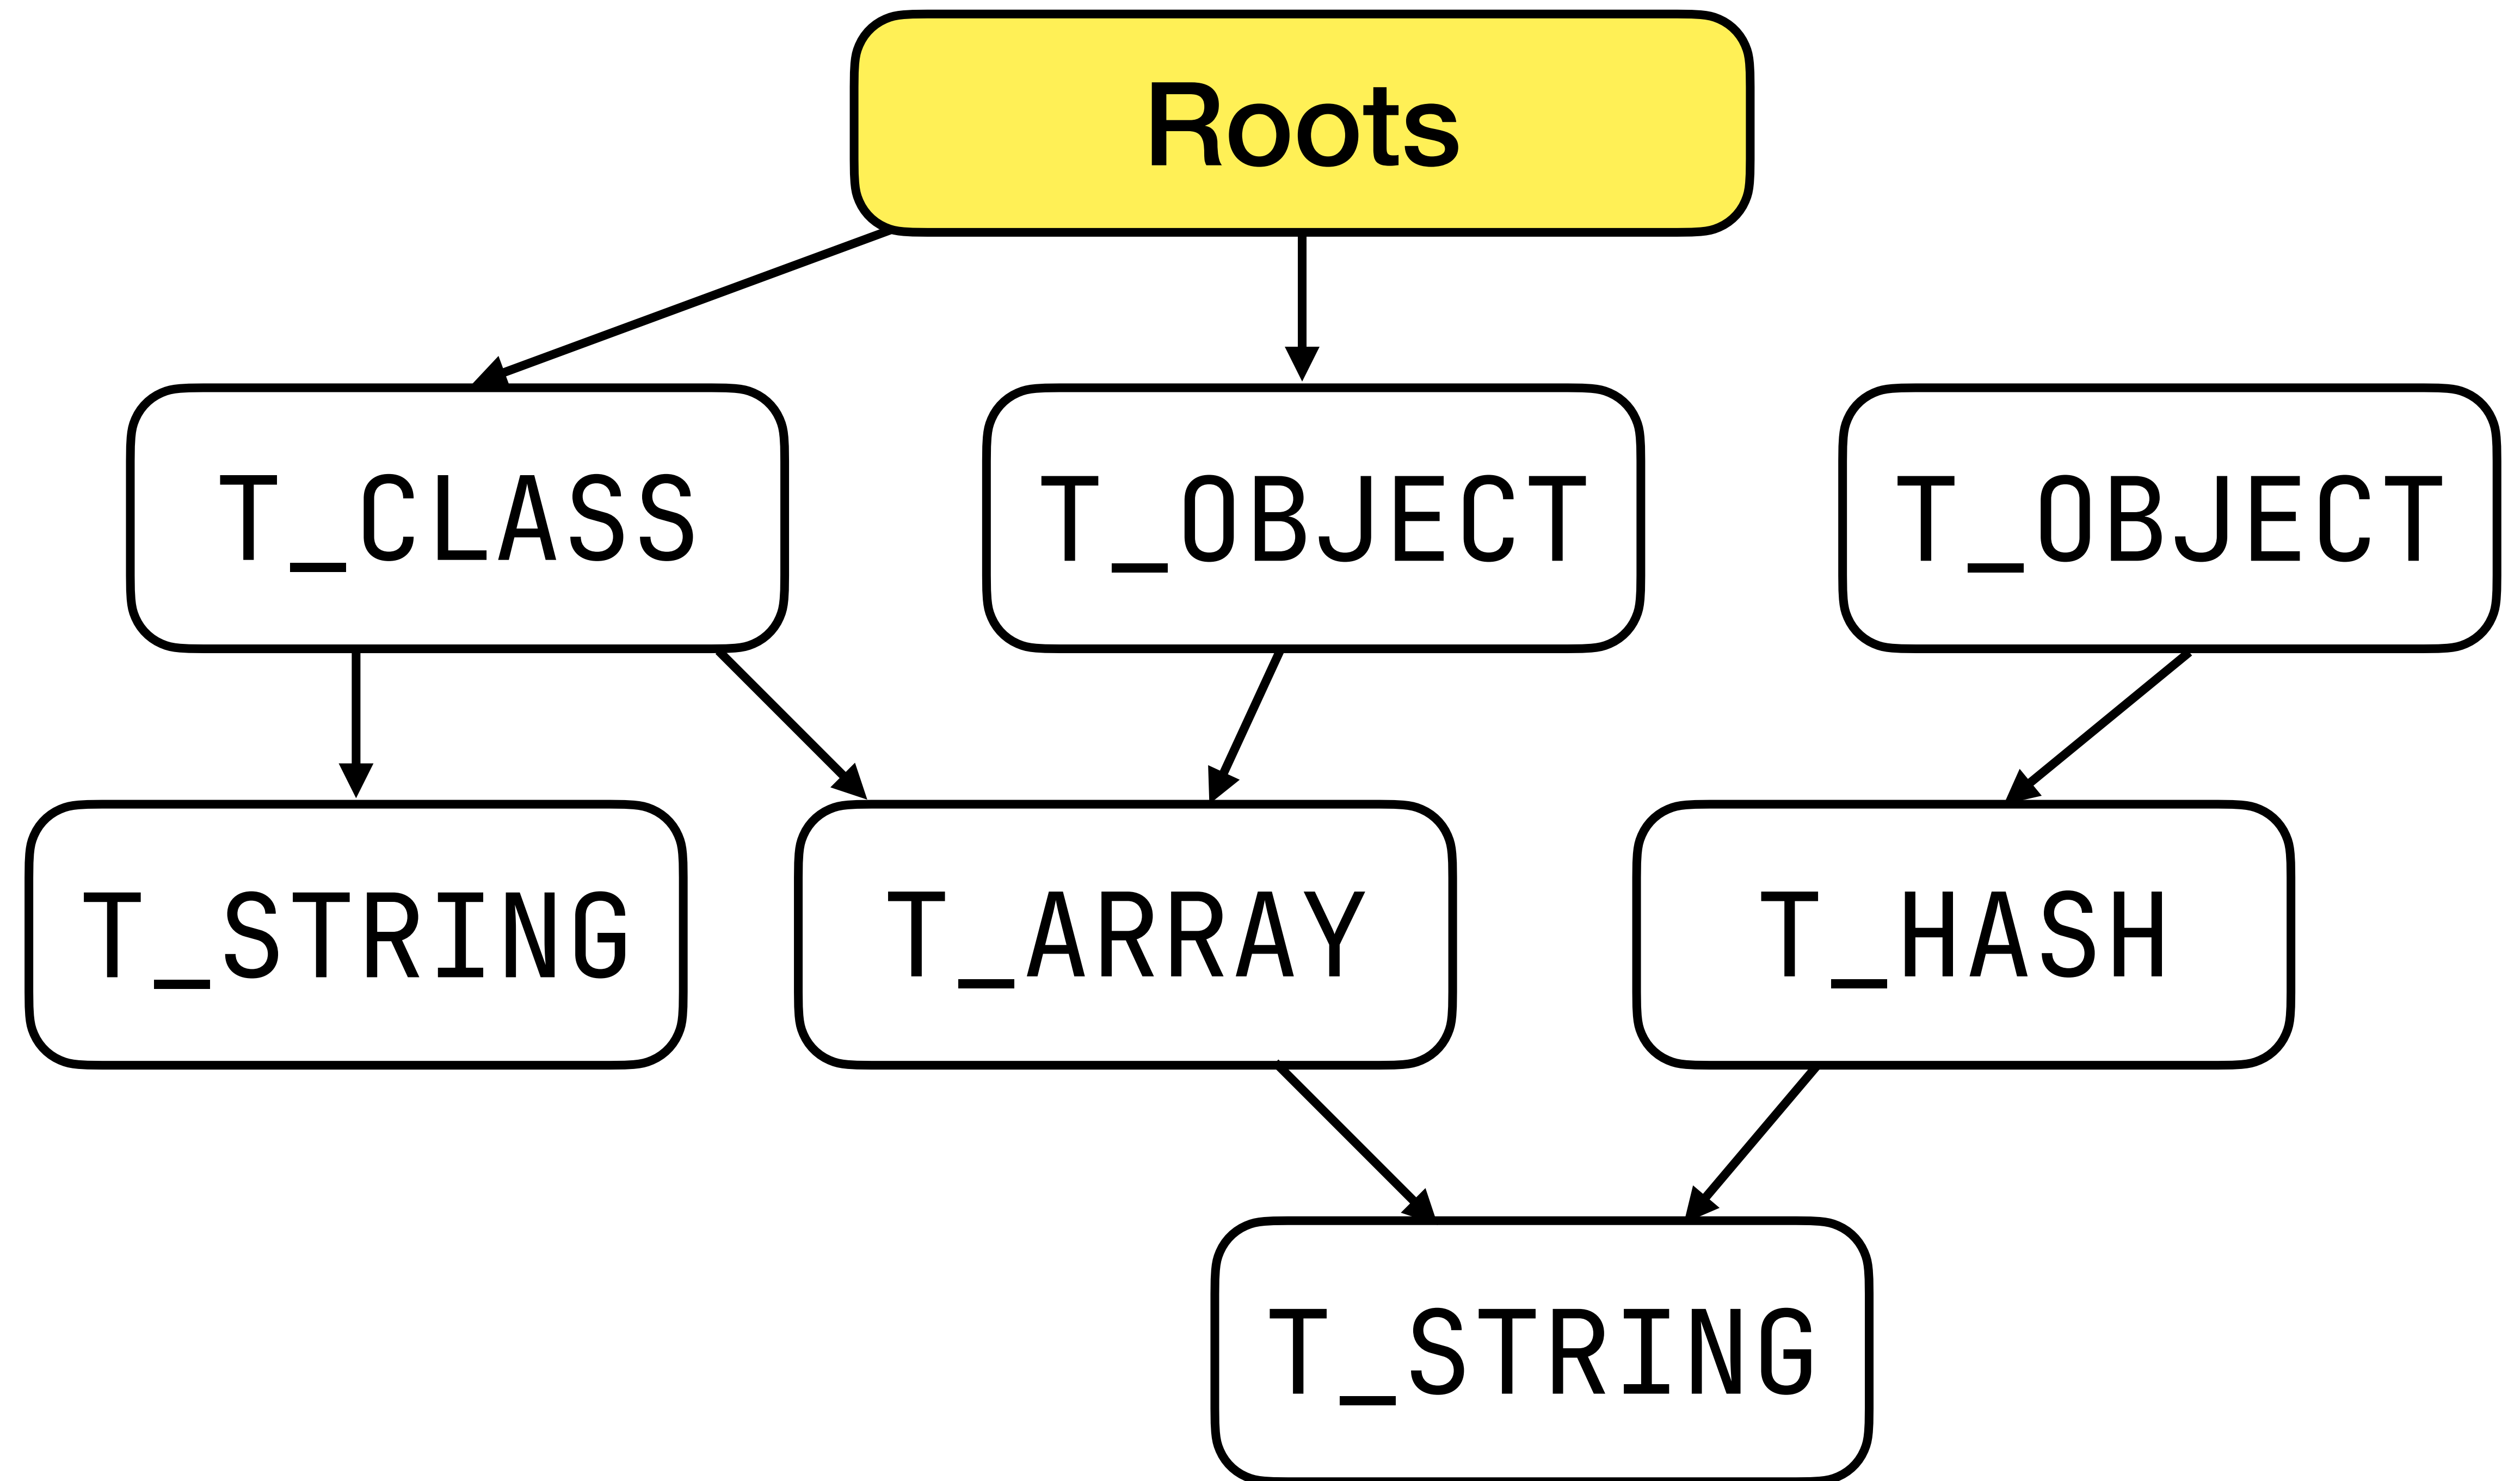
**Object
Allocation**

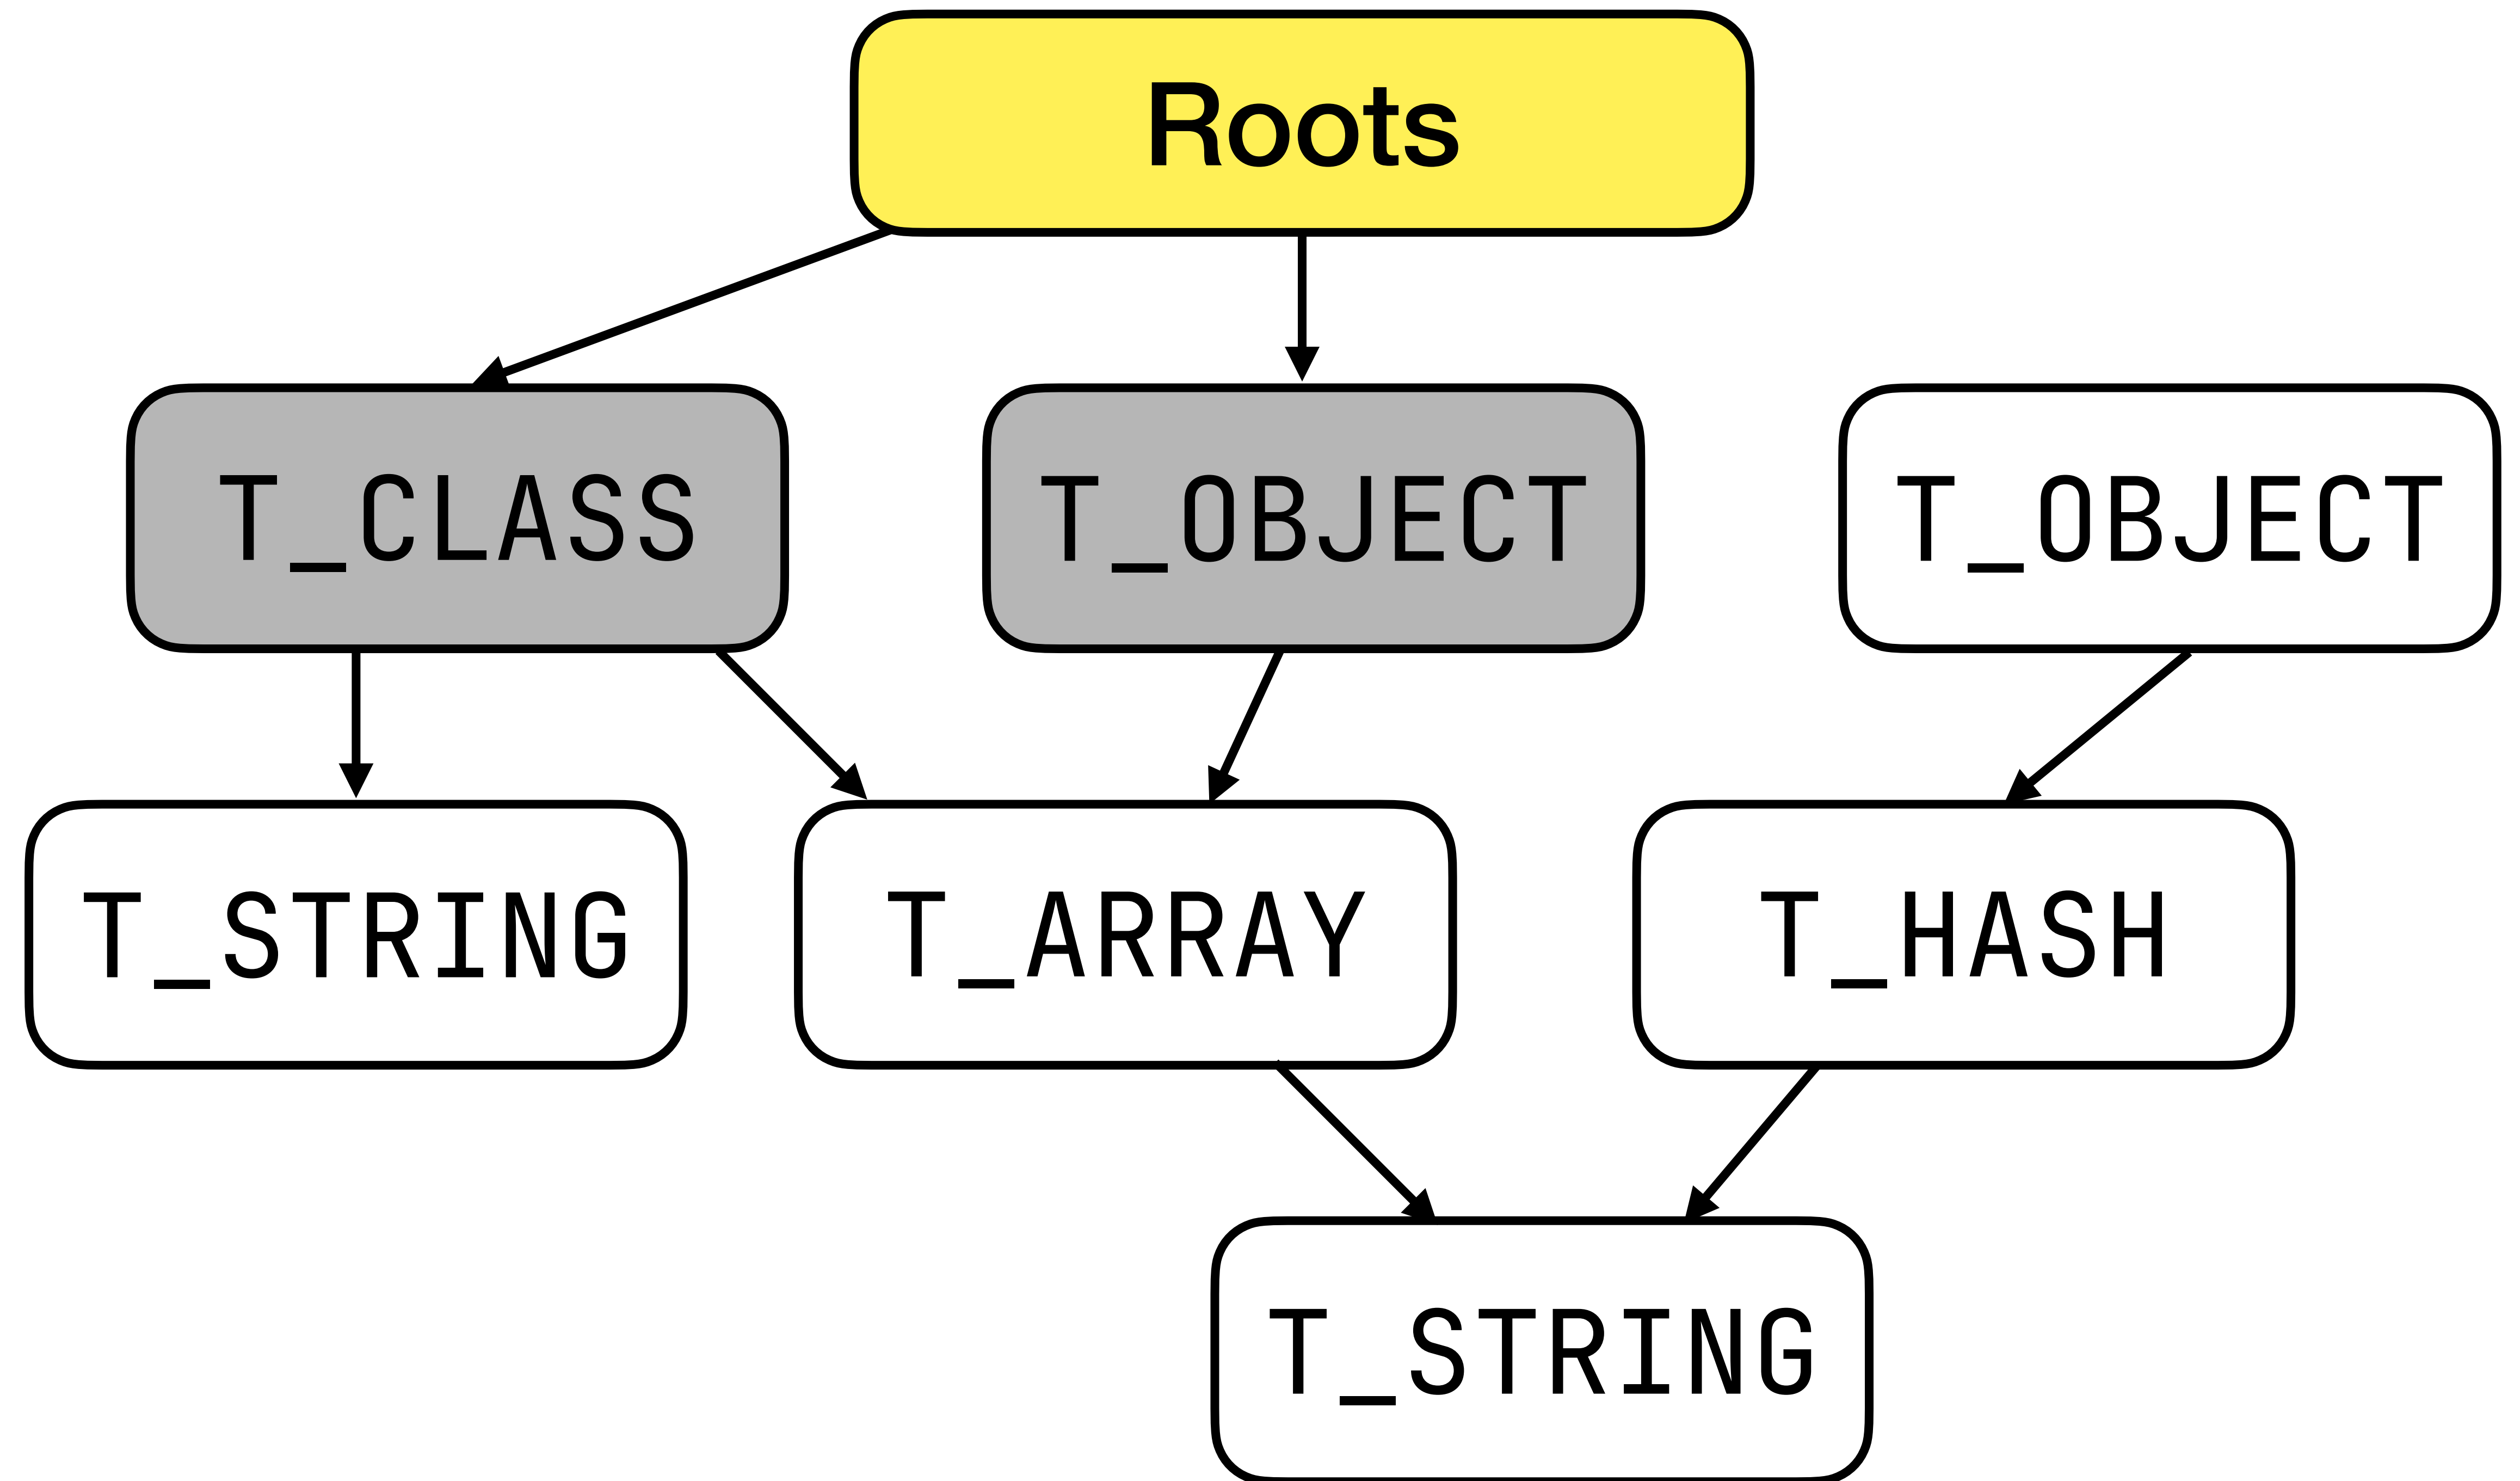
**Recycle
Object**

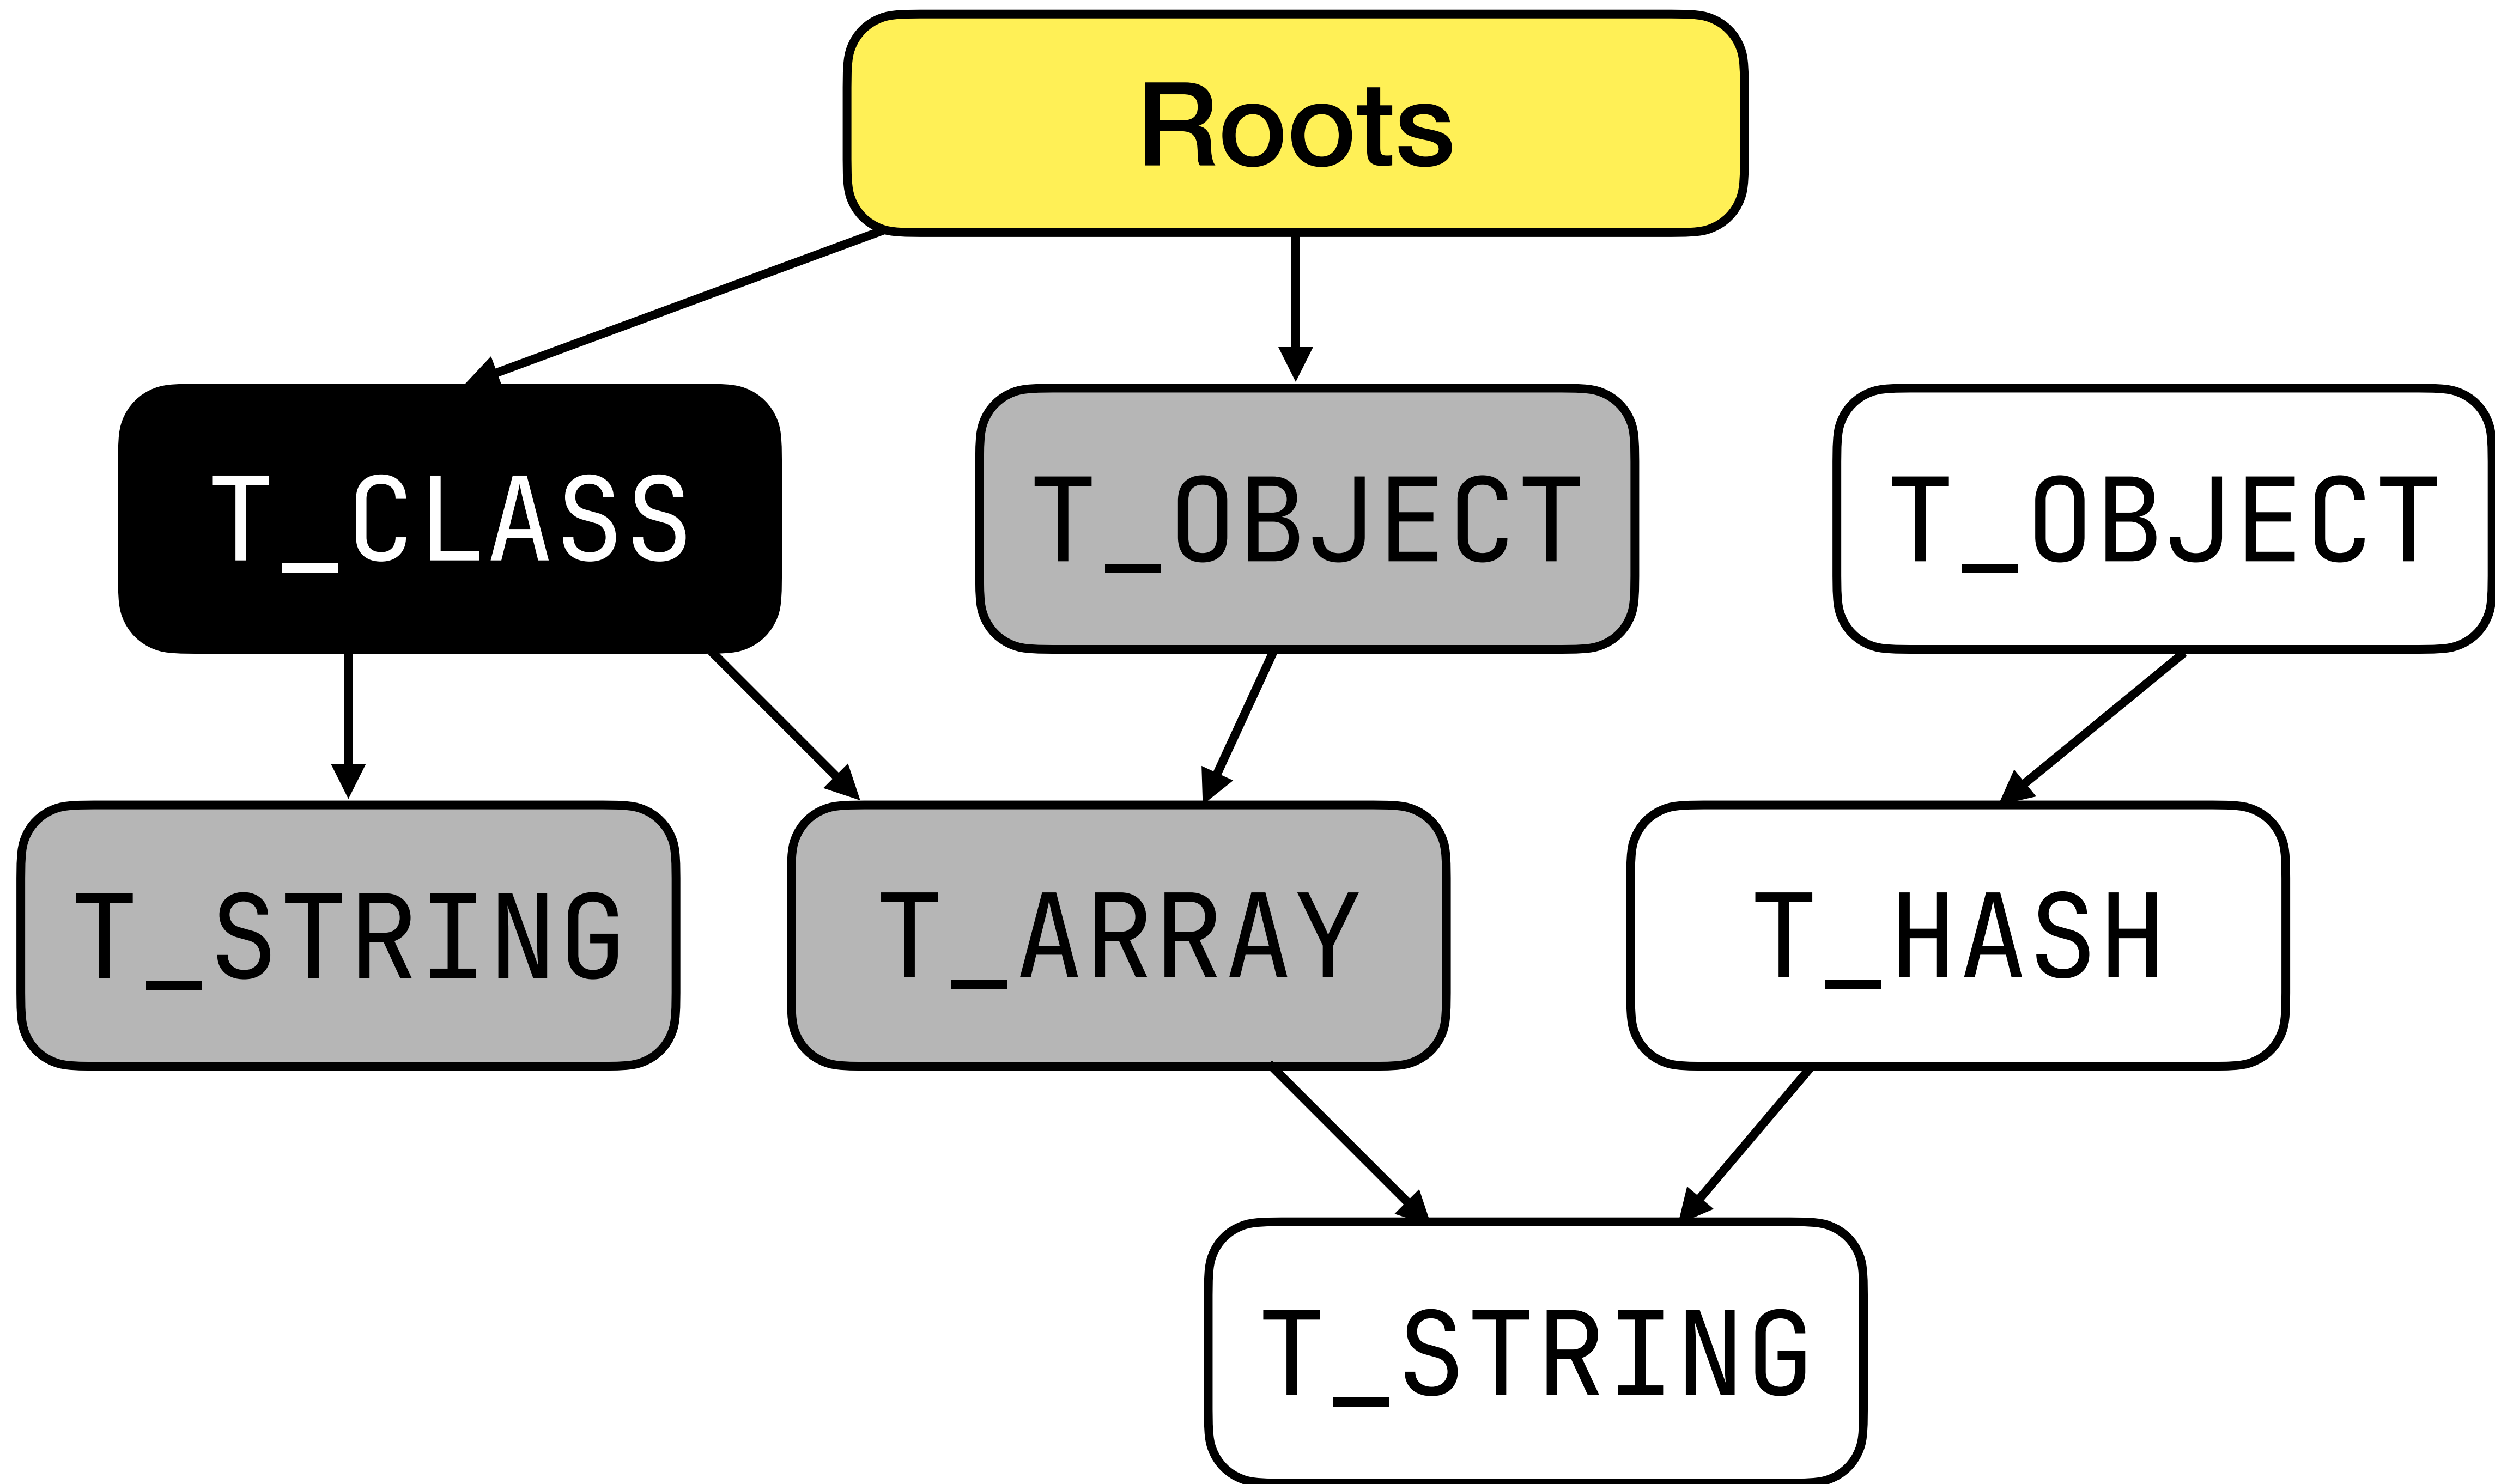
Liveliness

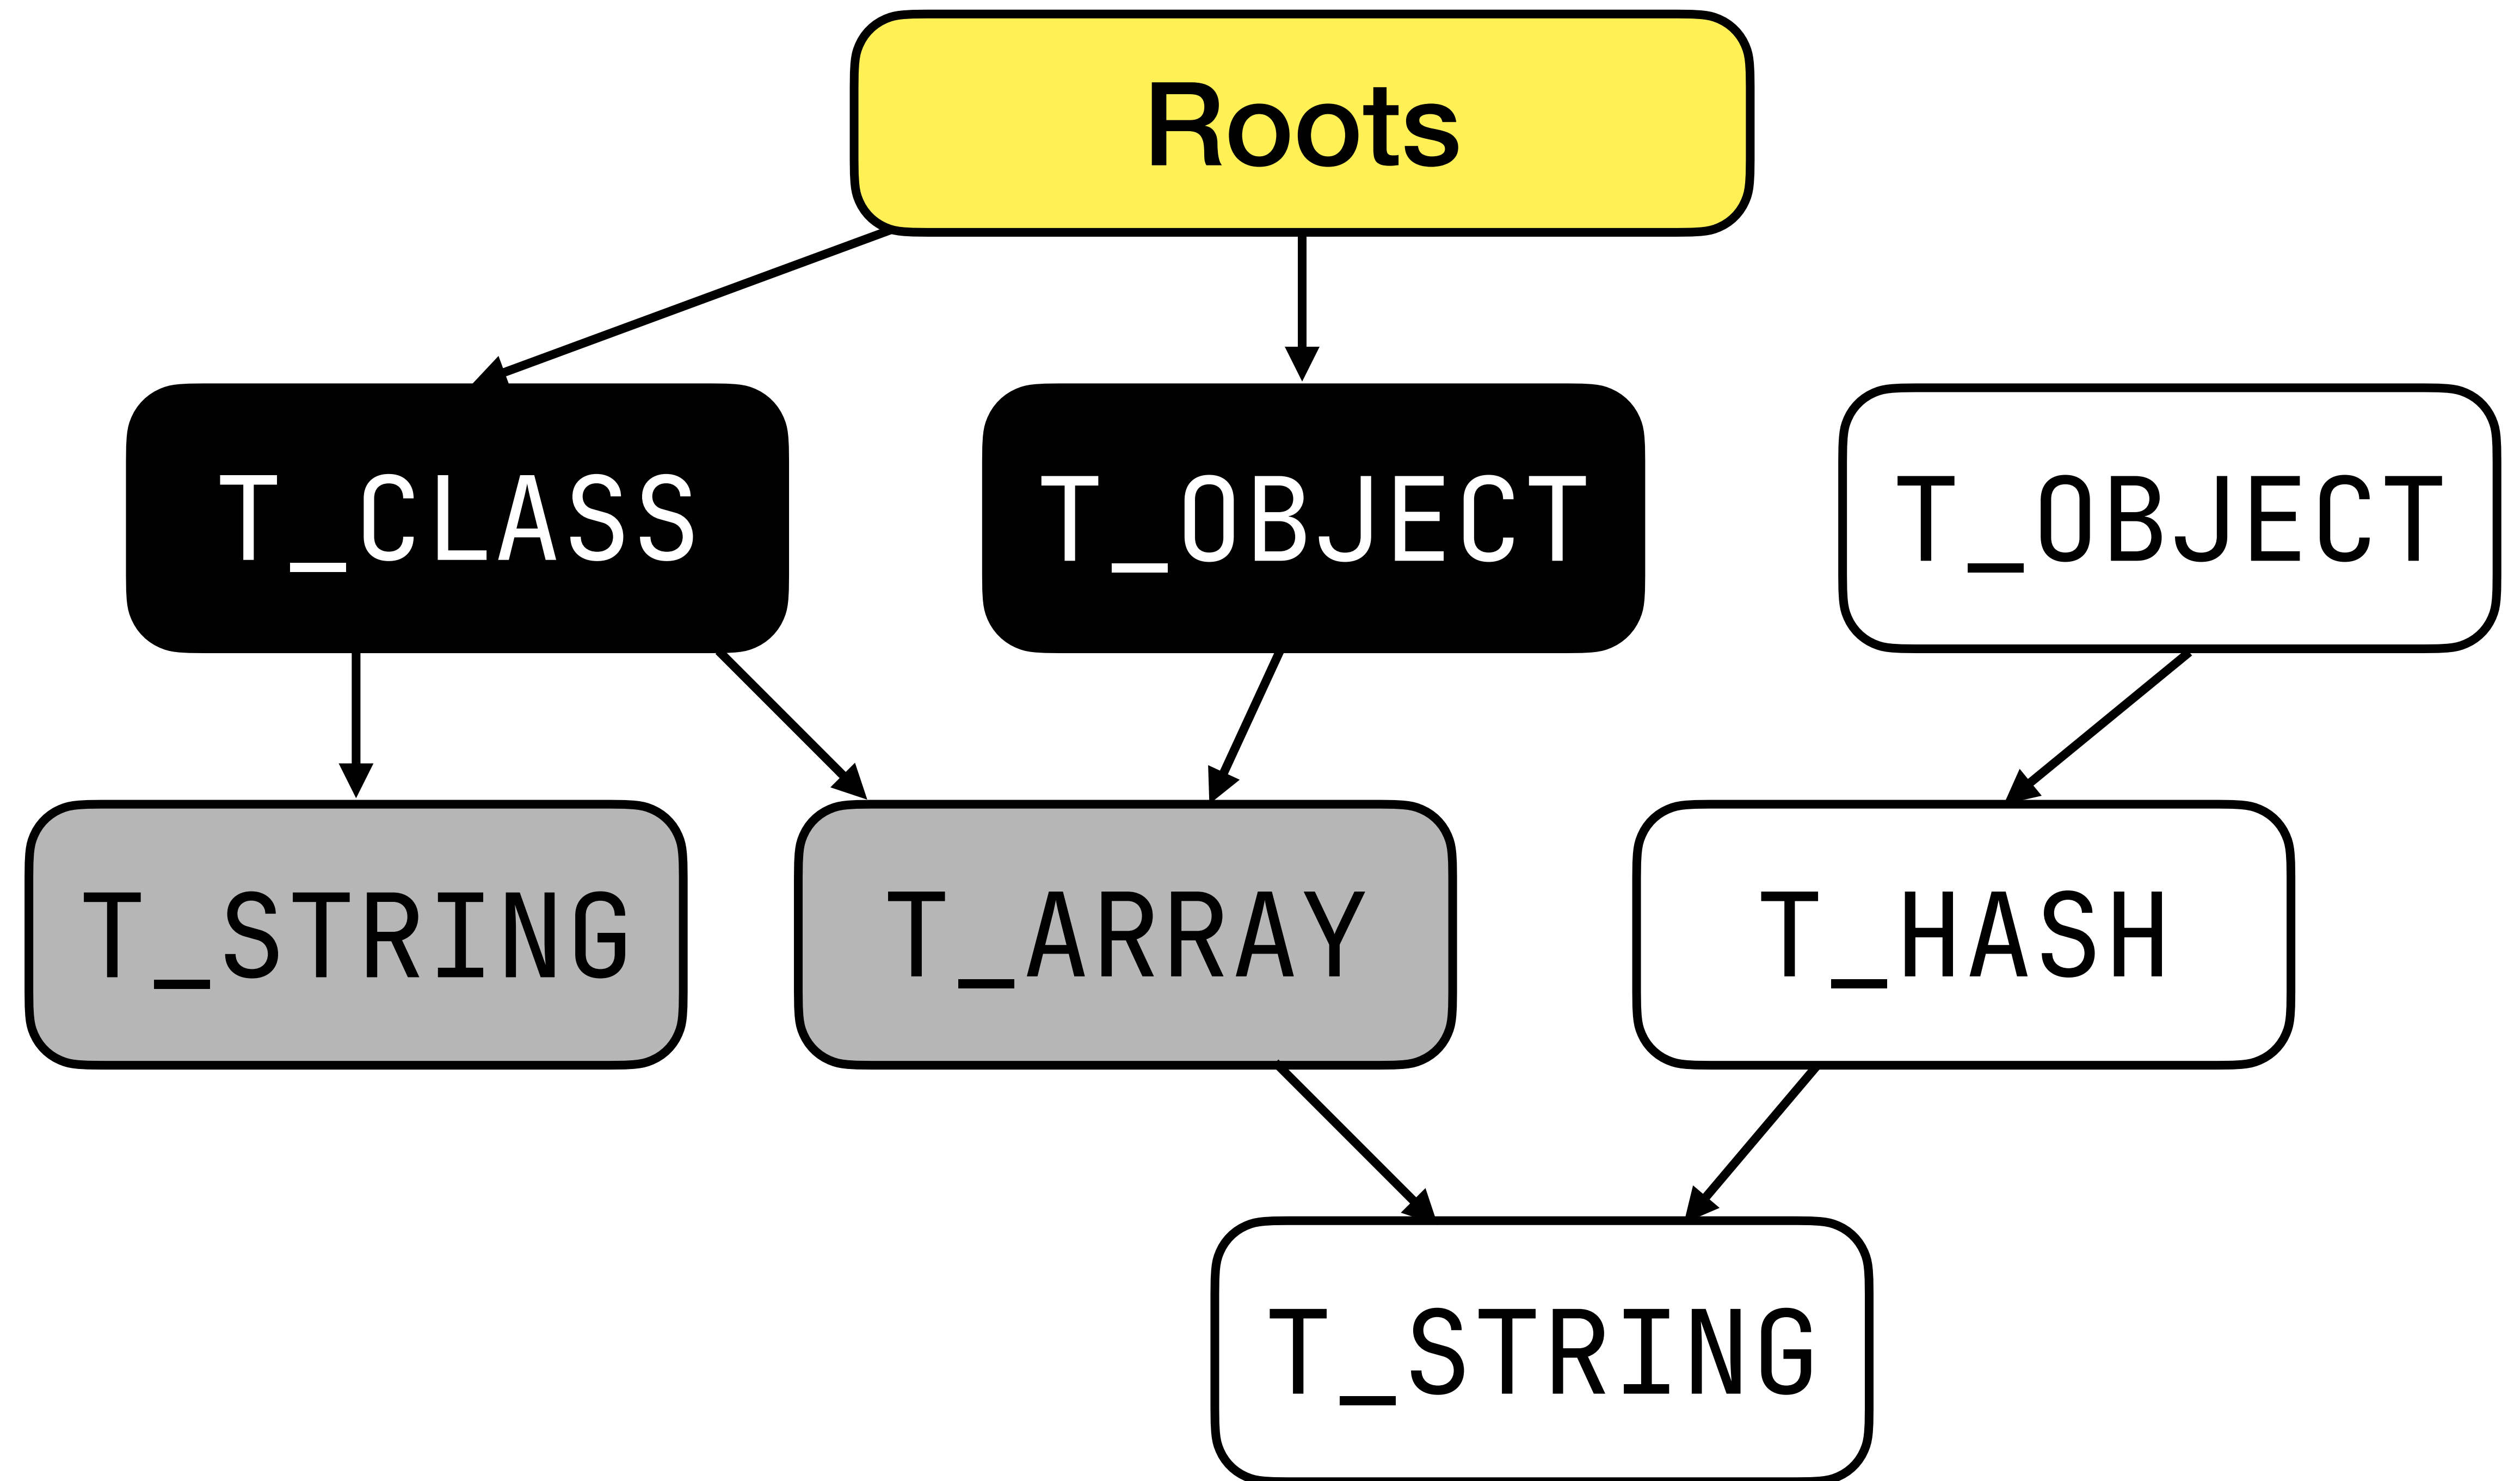
Mark-Sweep-Compact

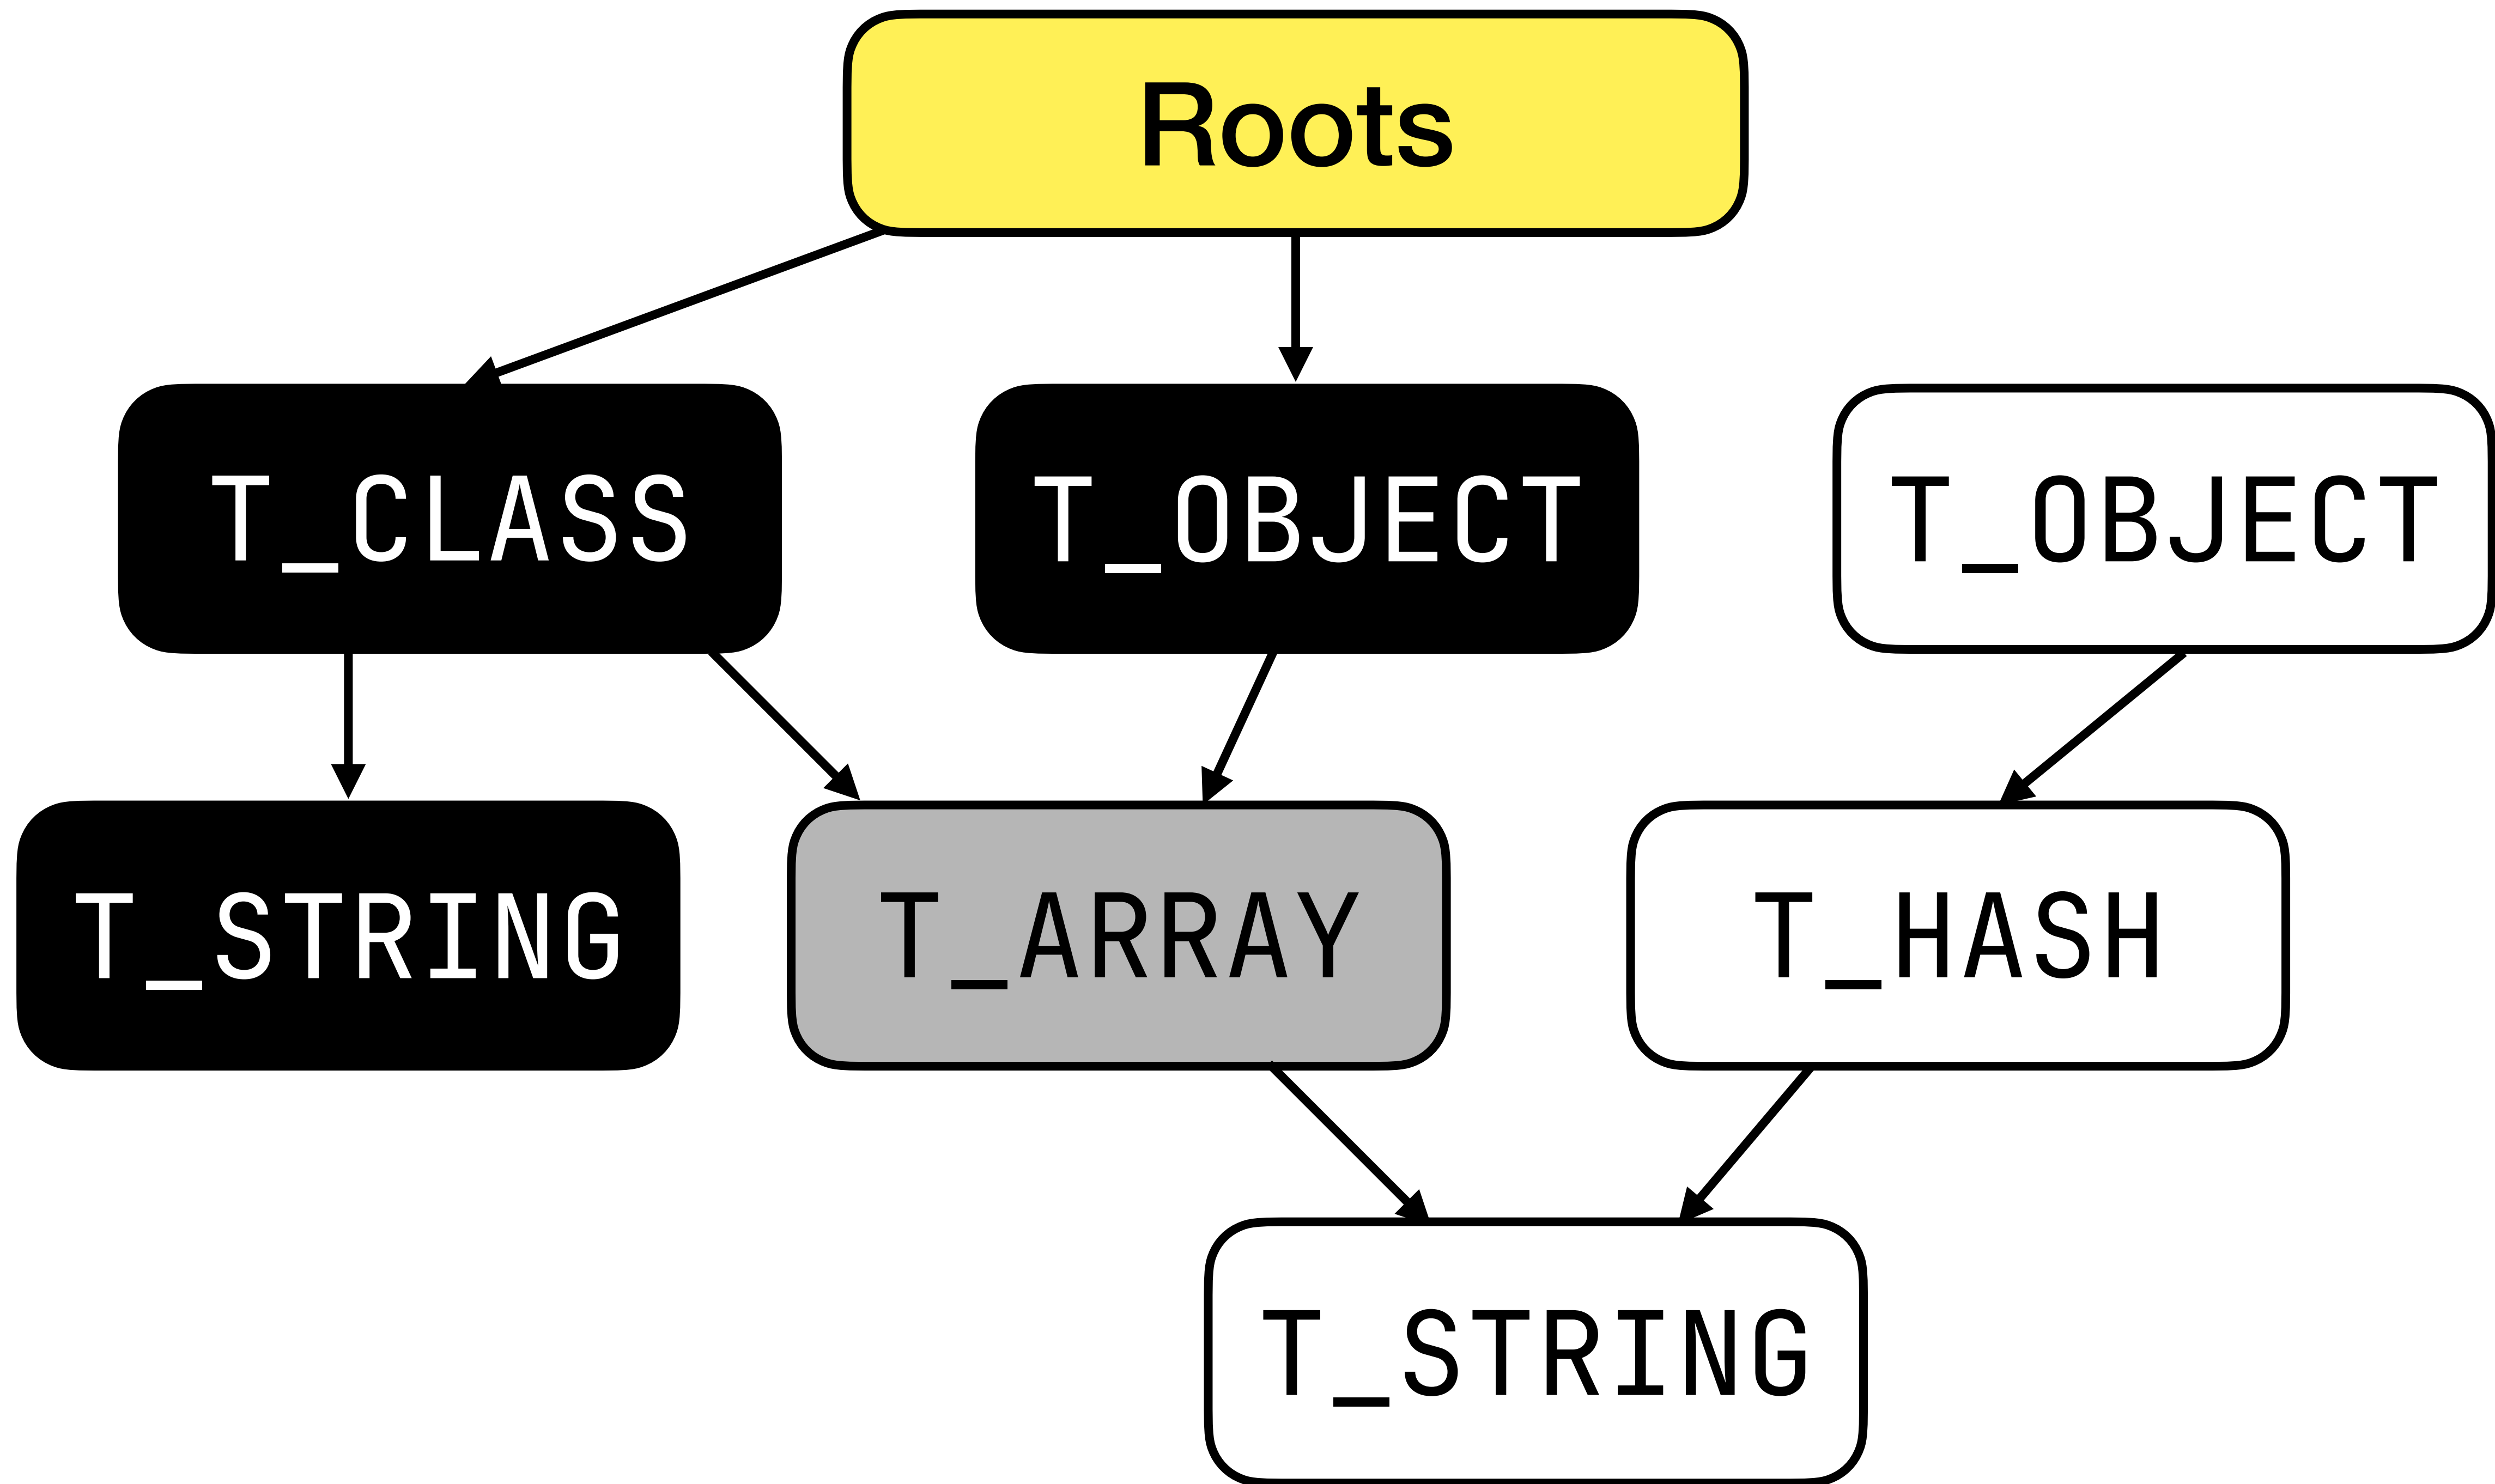
Mark Phase

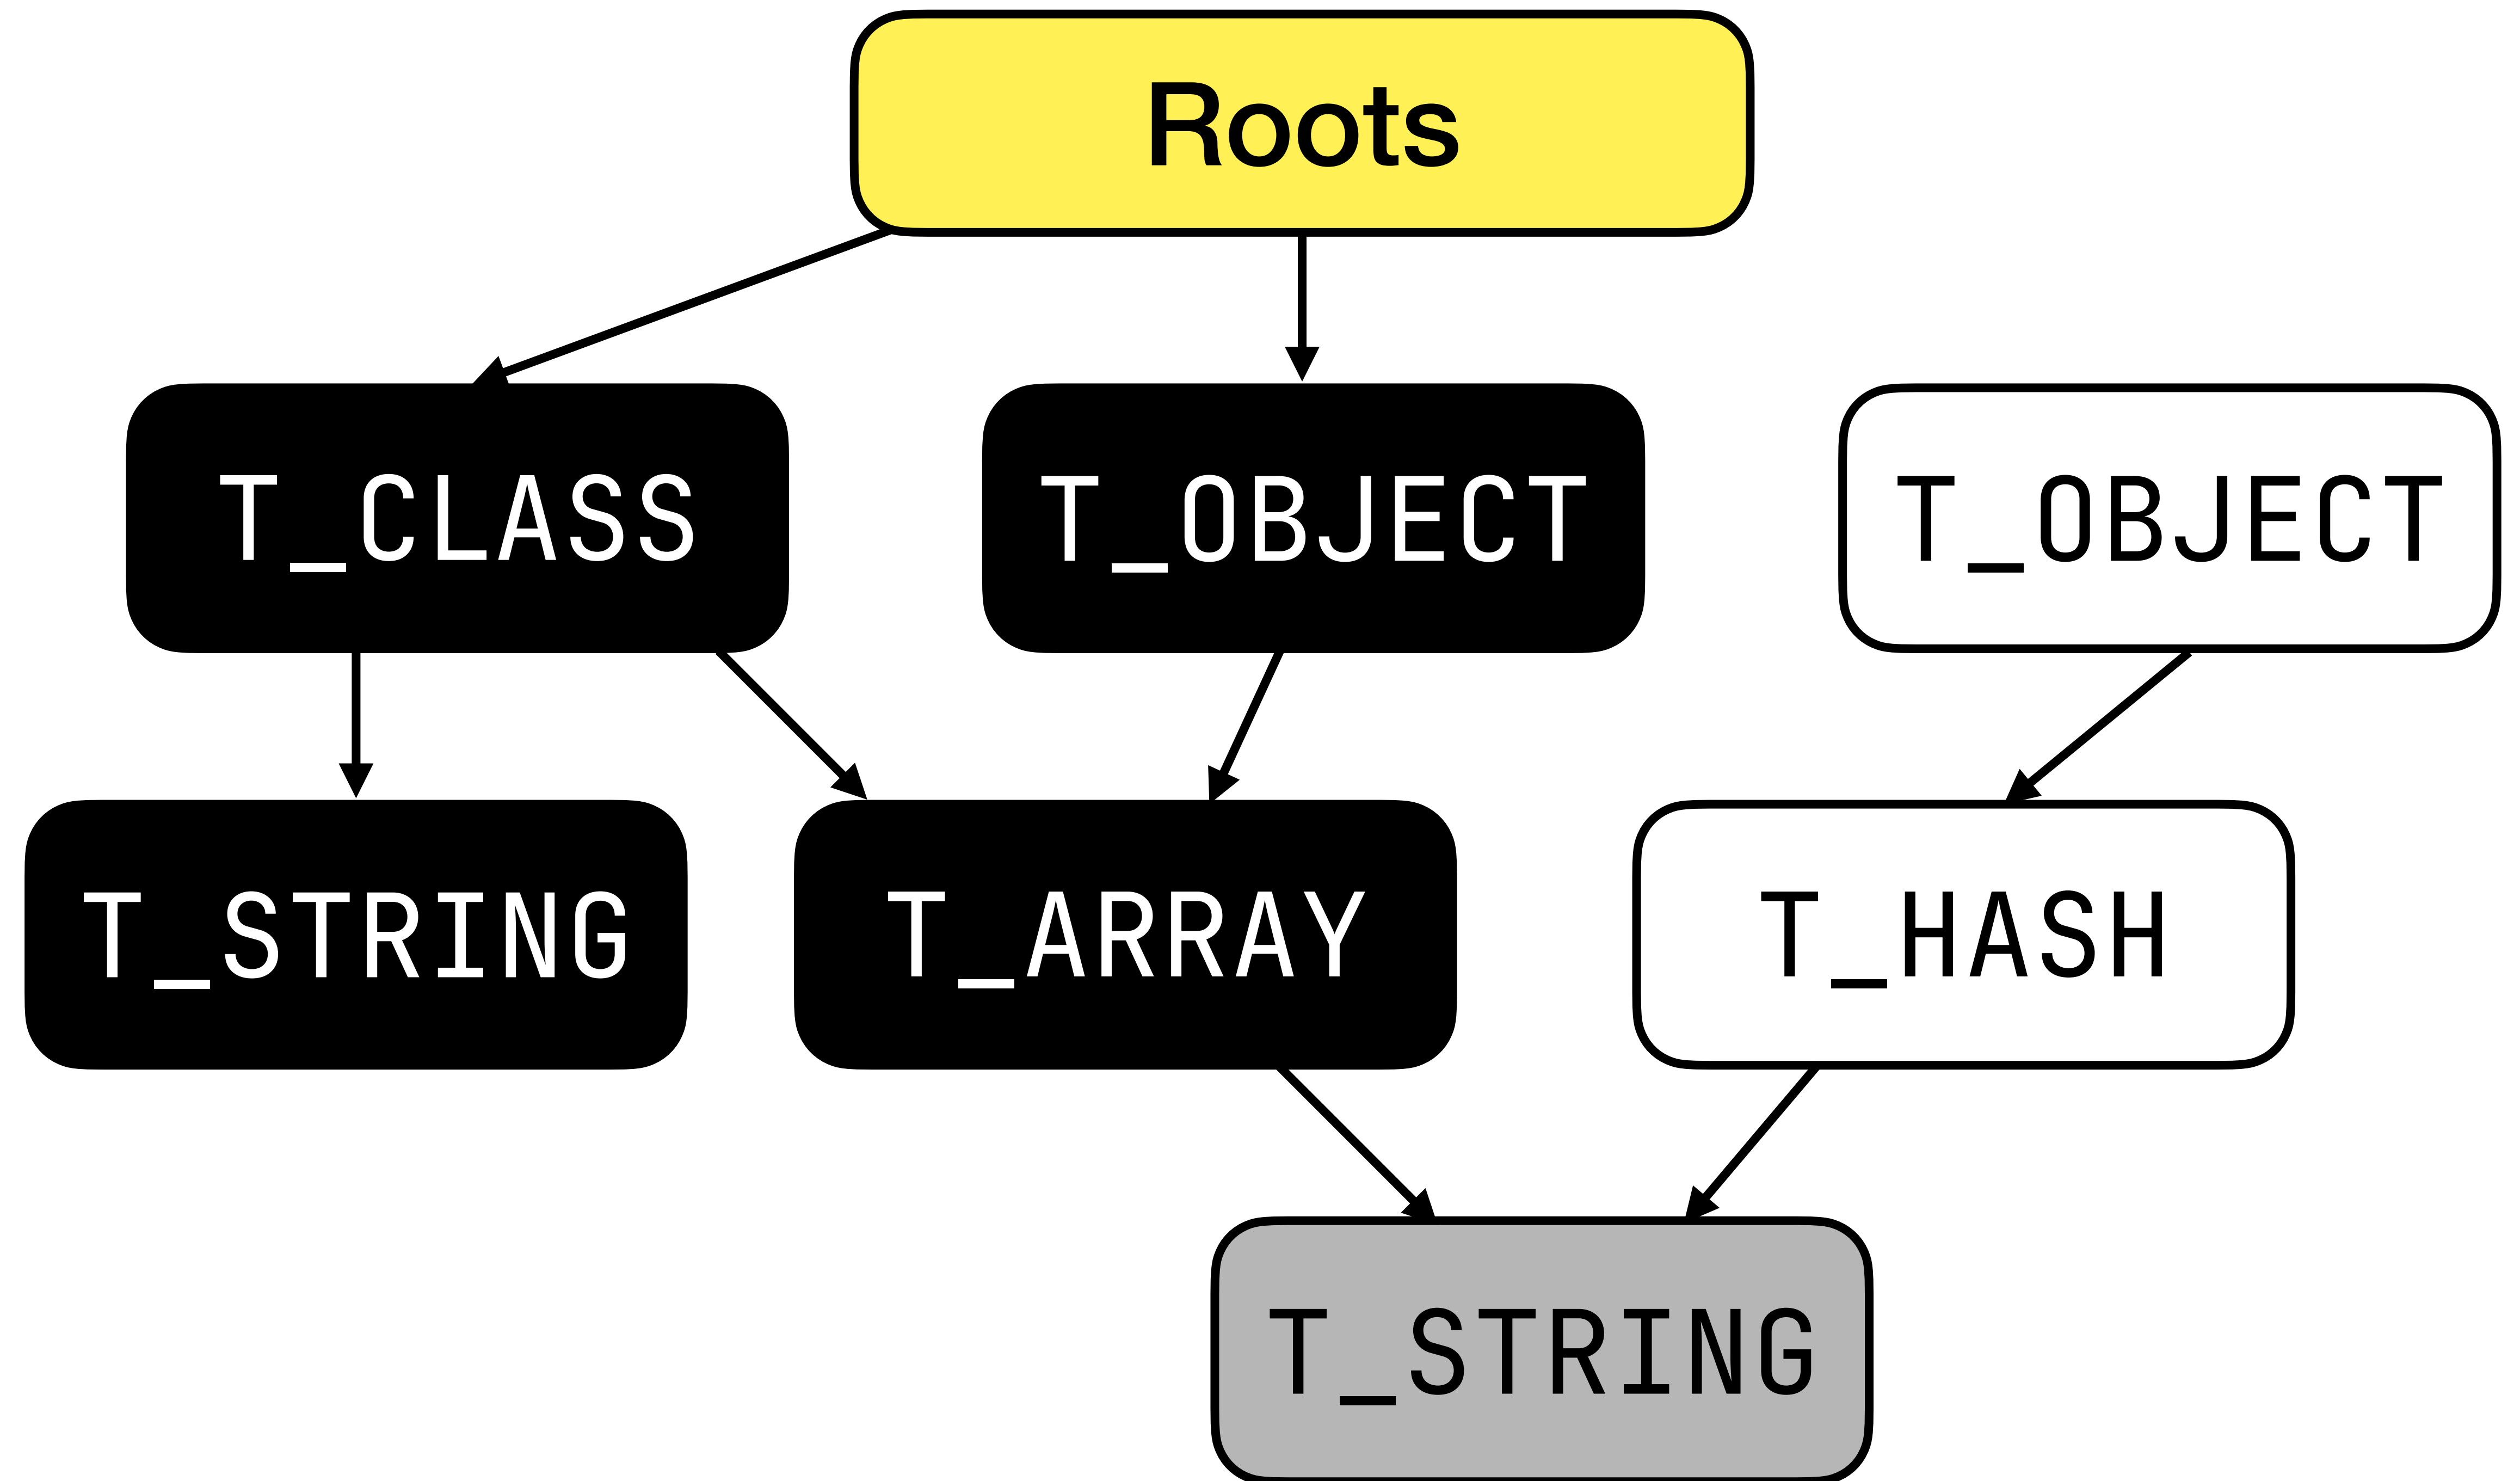


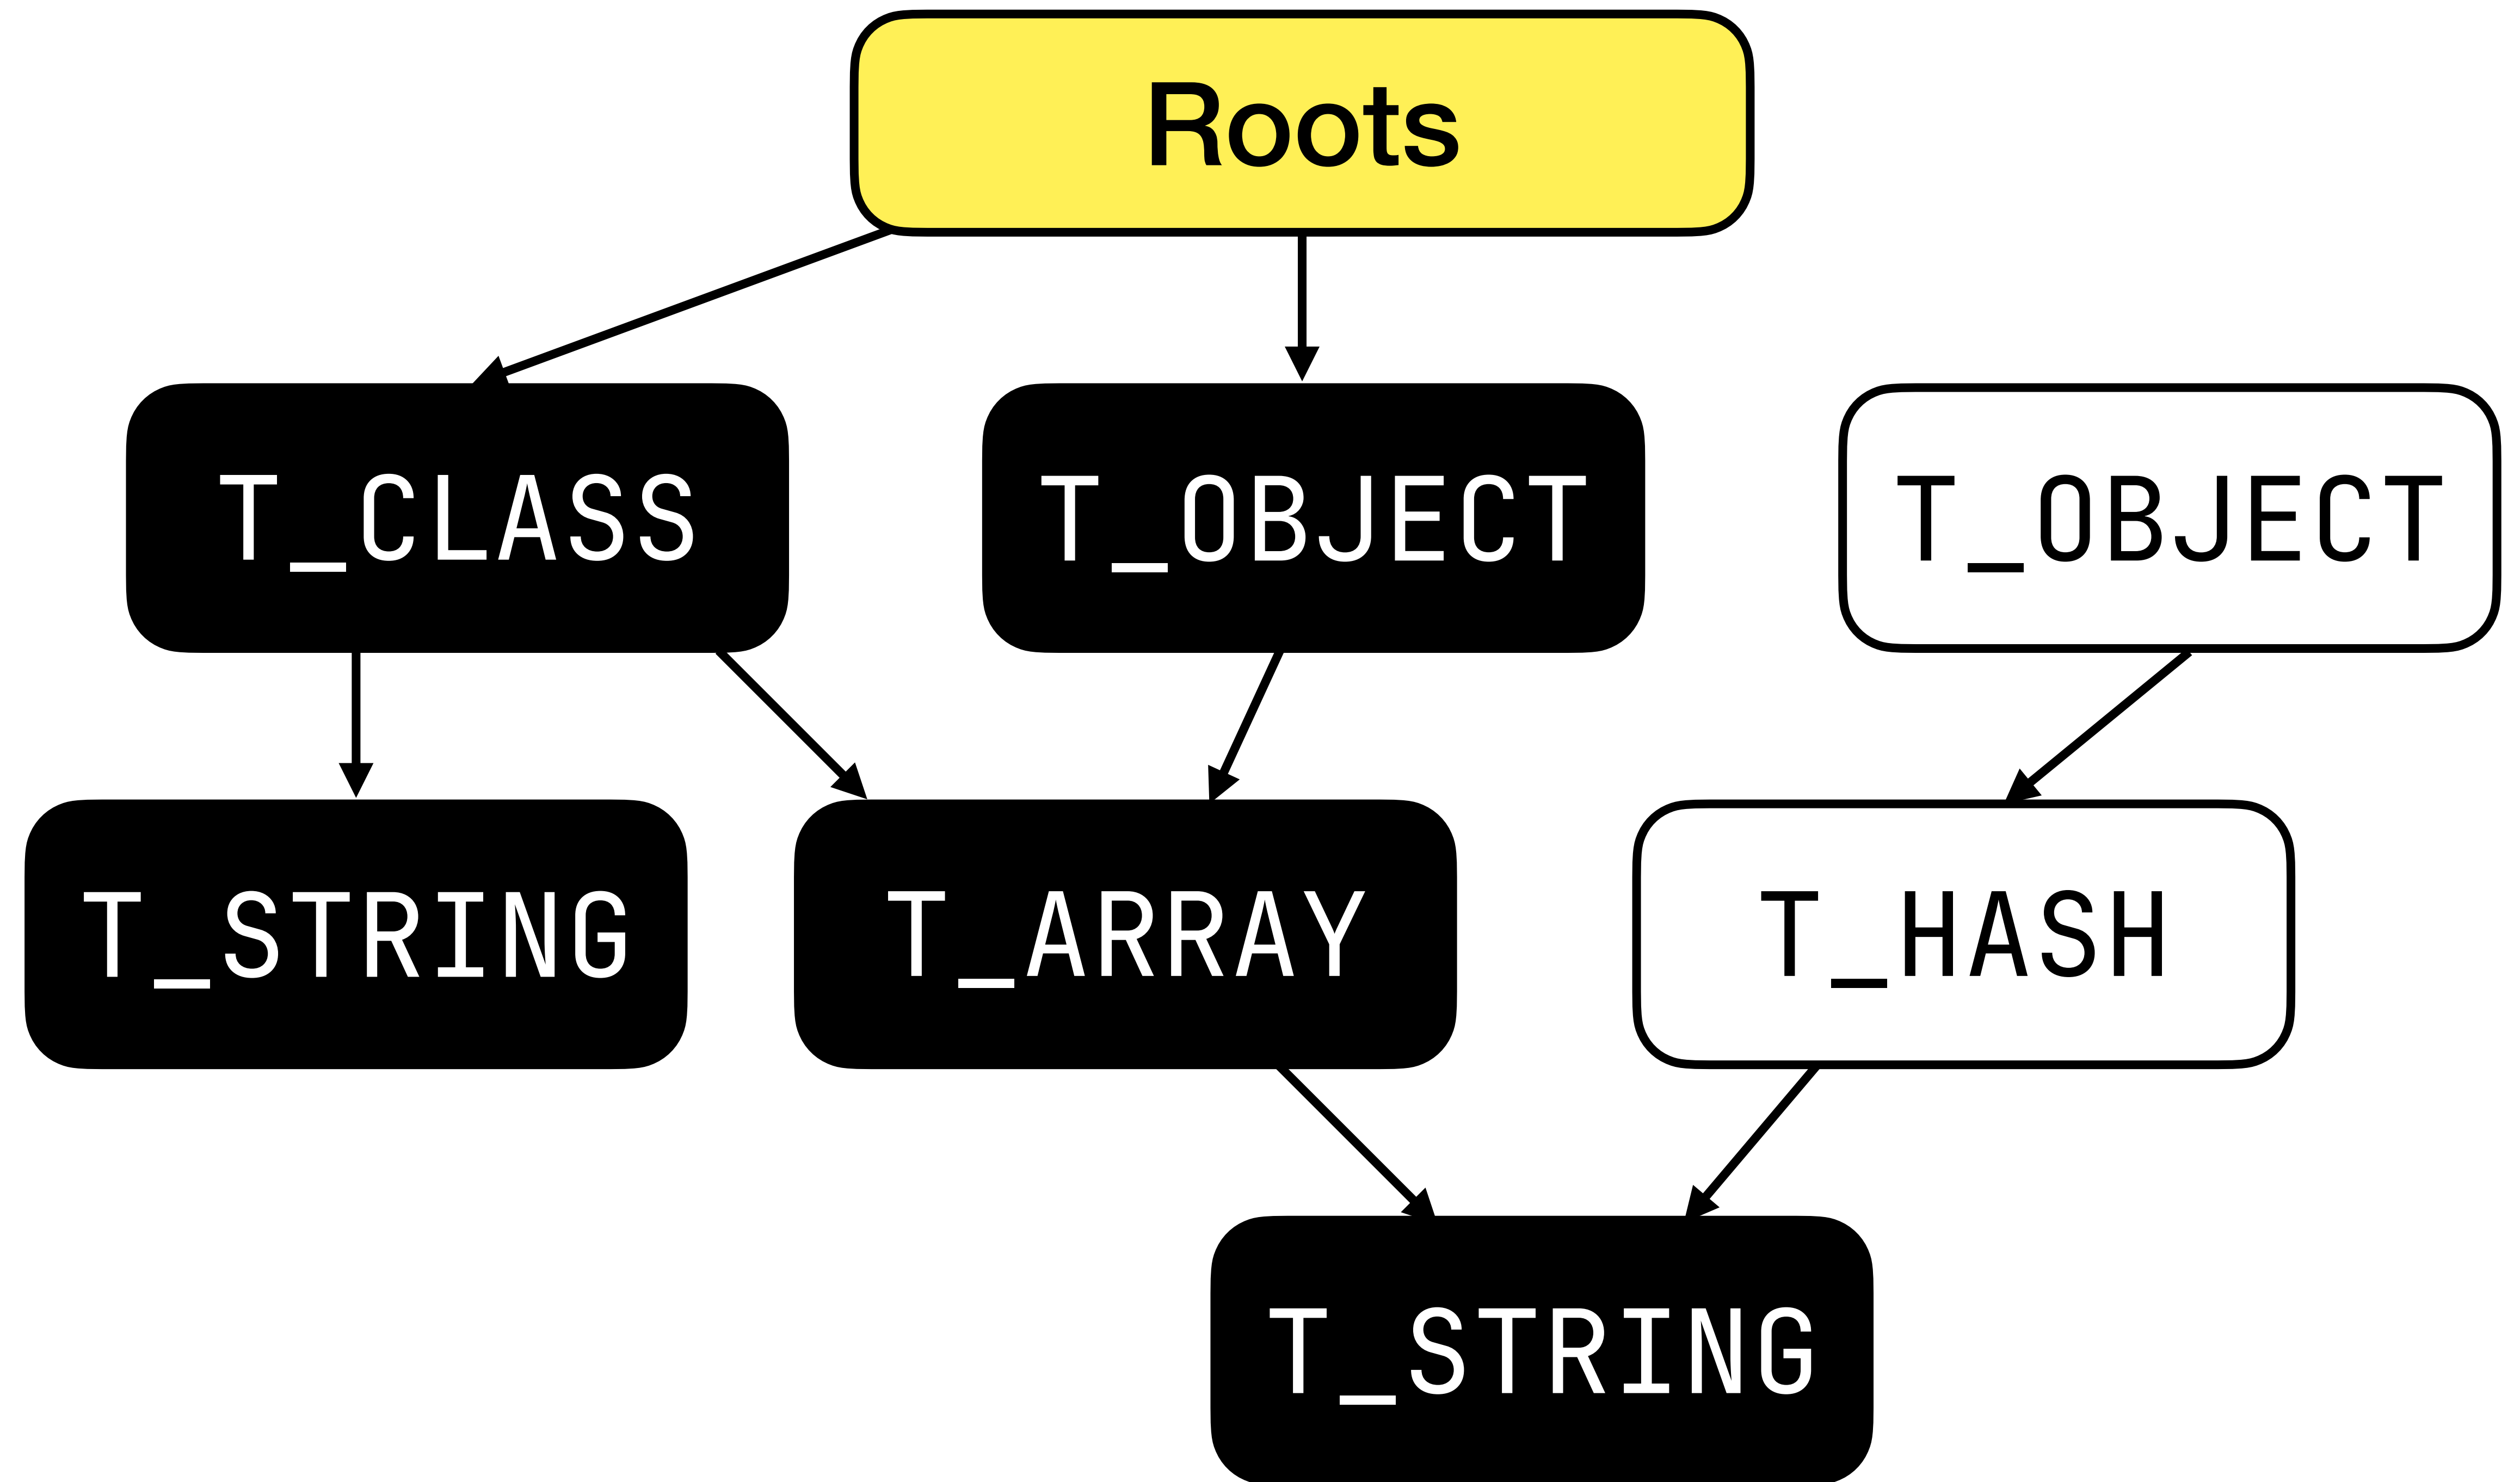




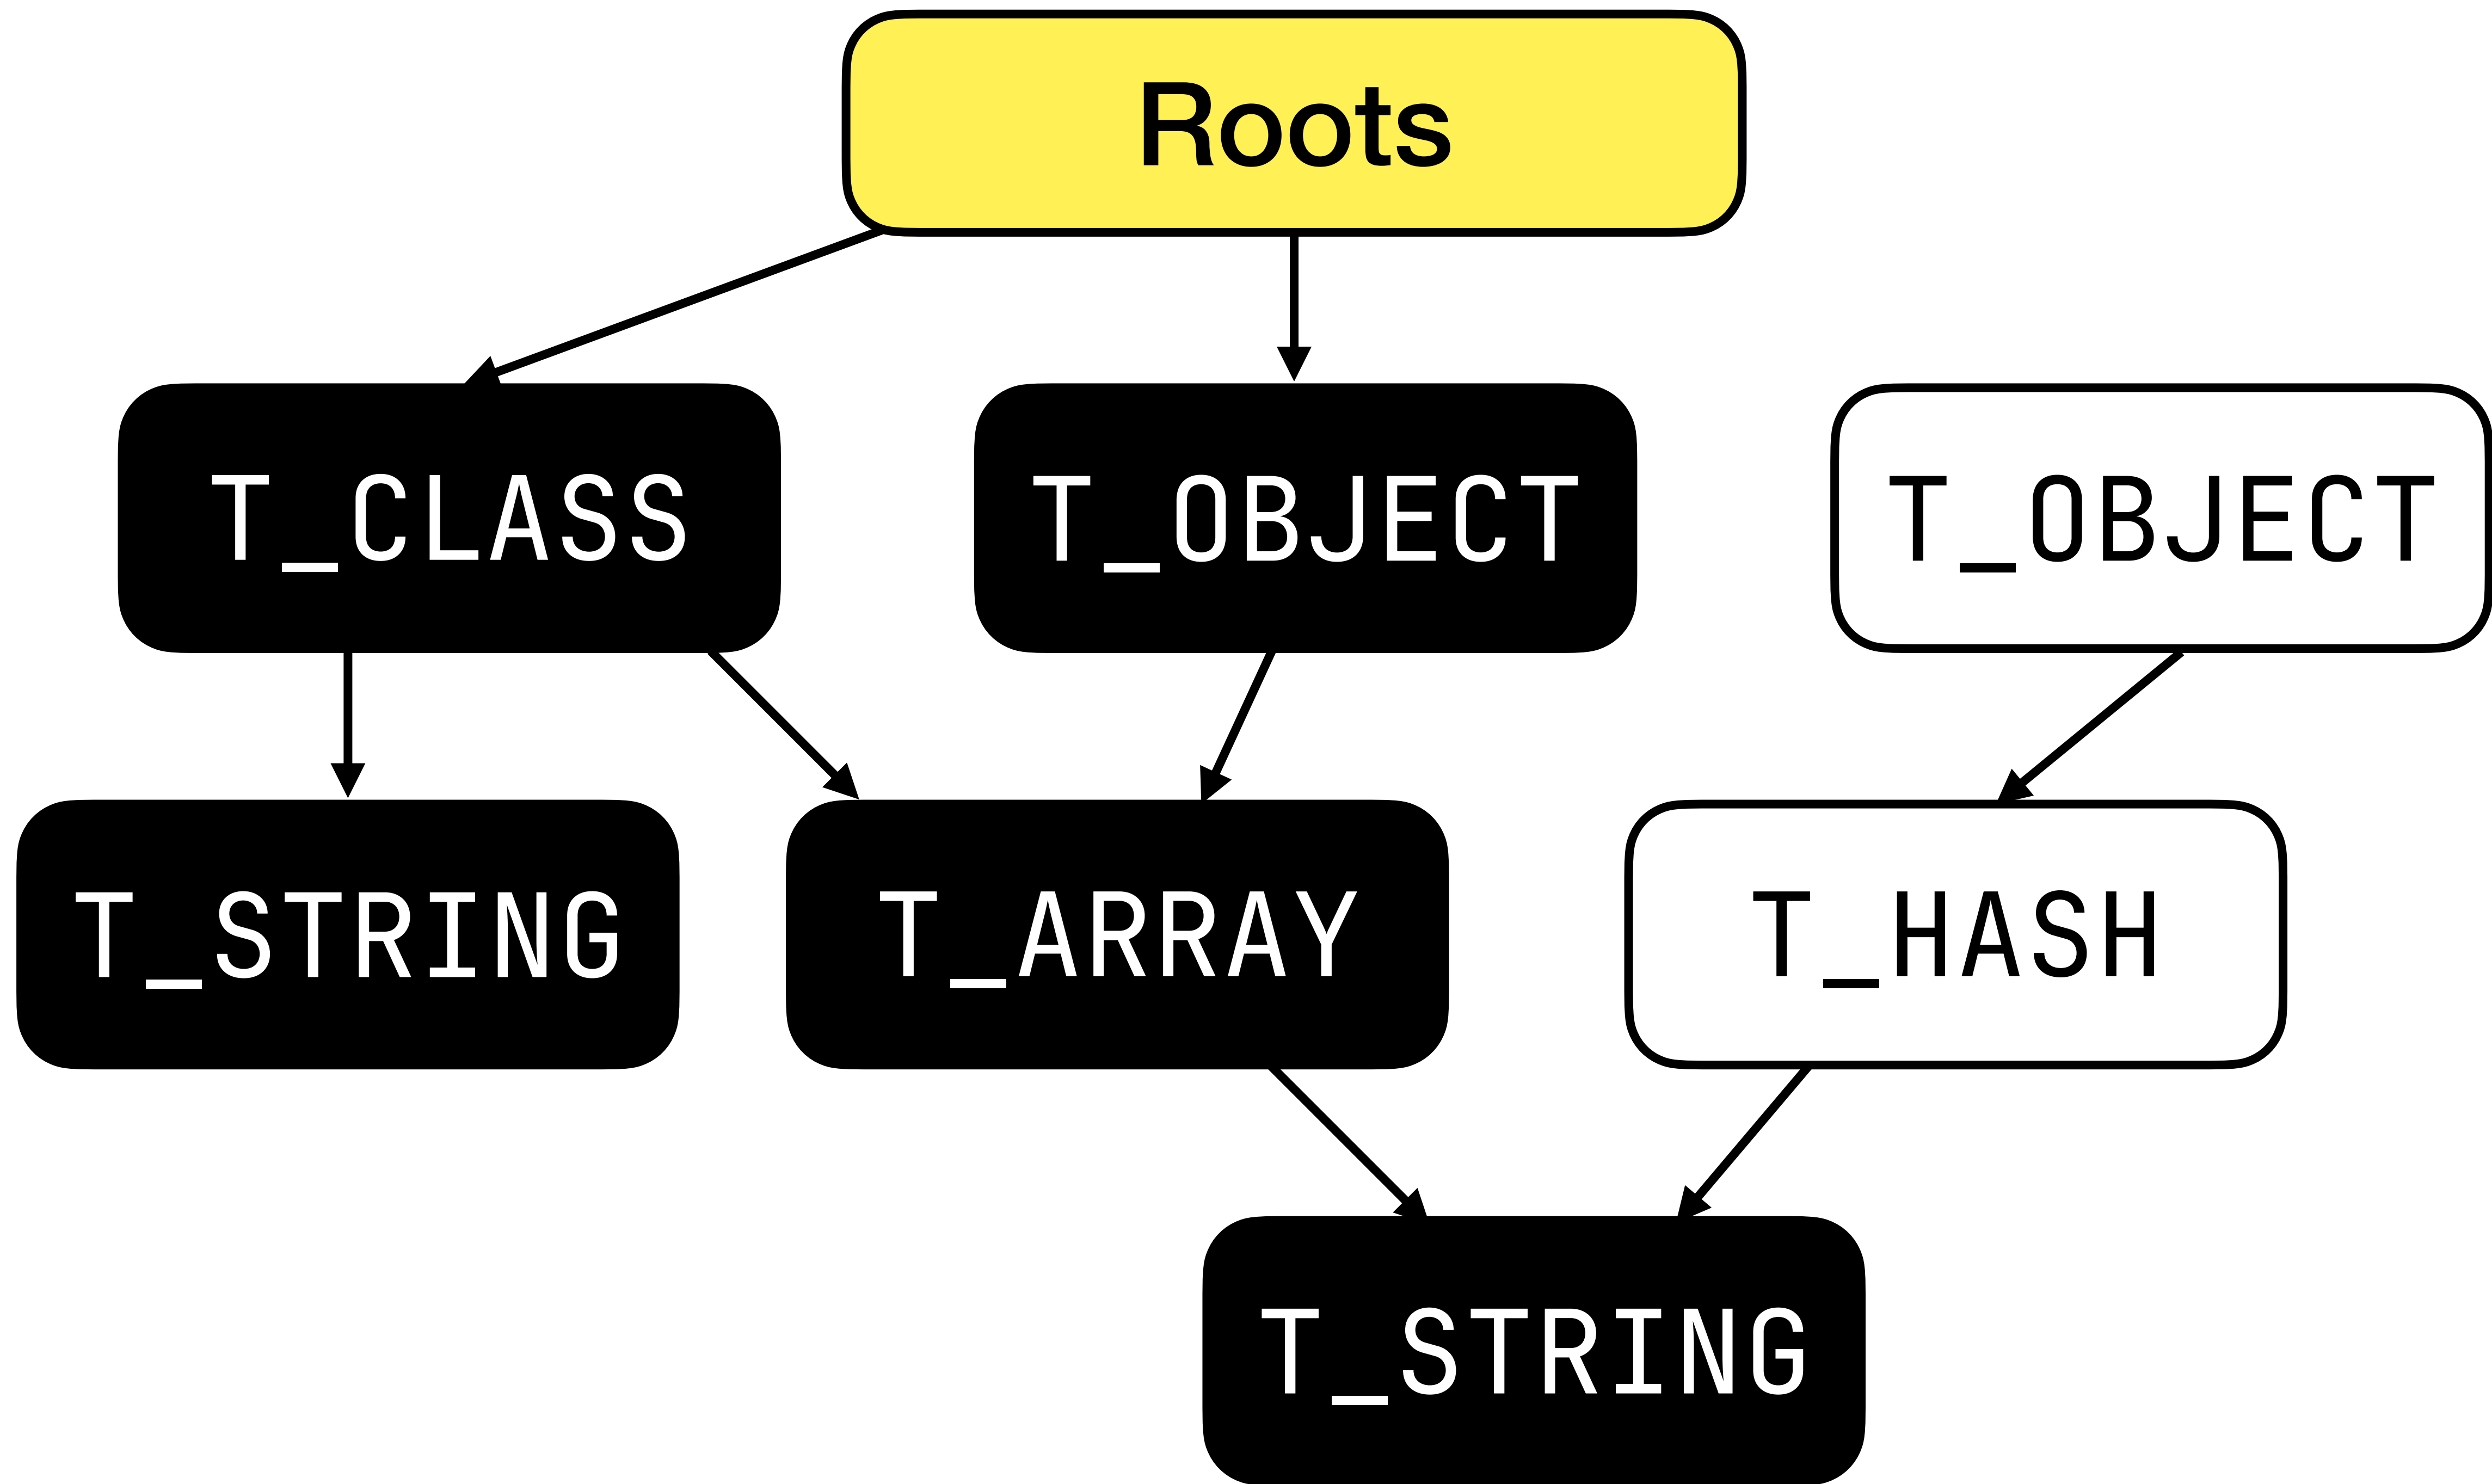


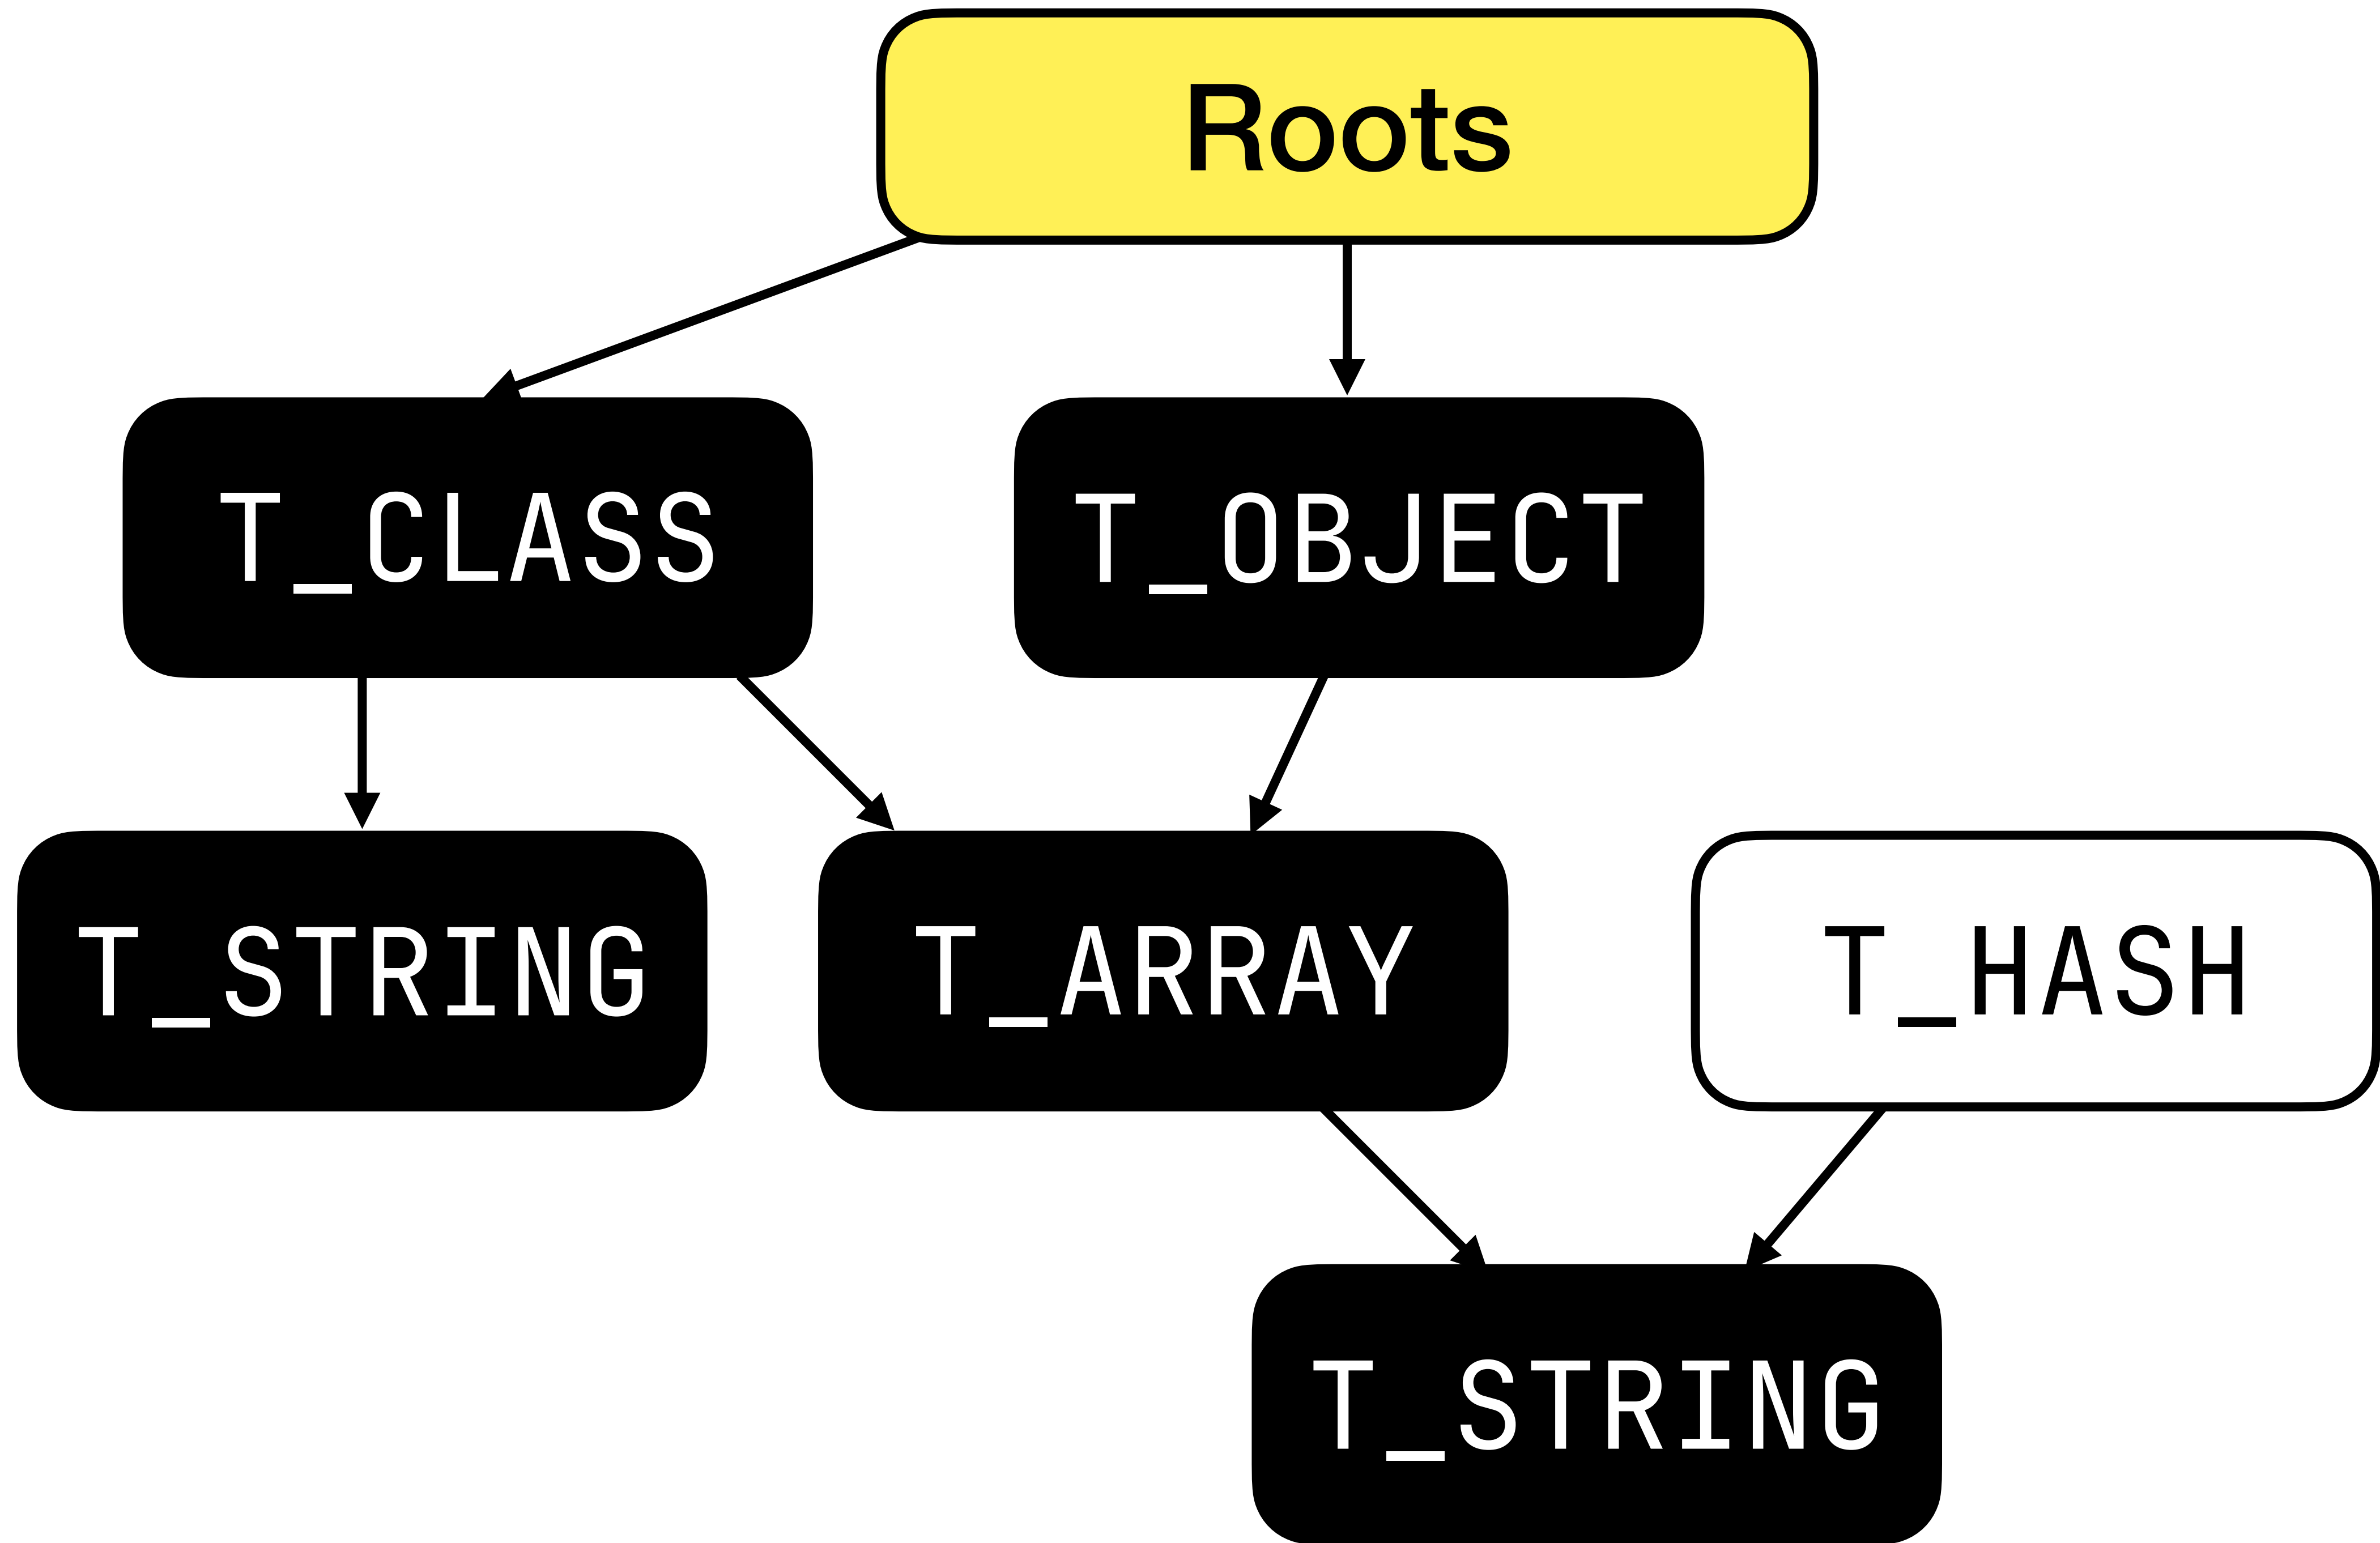


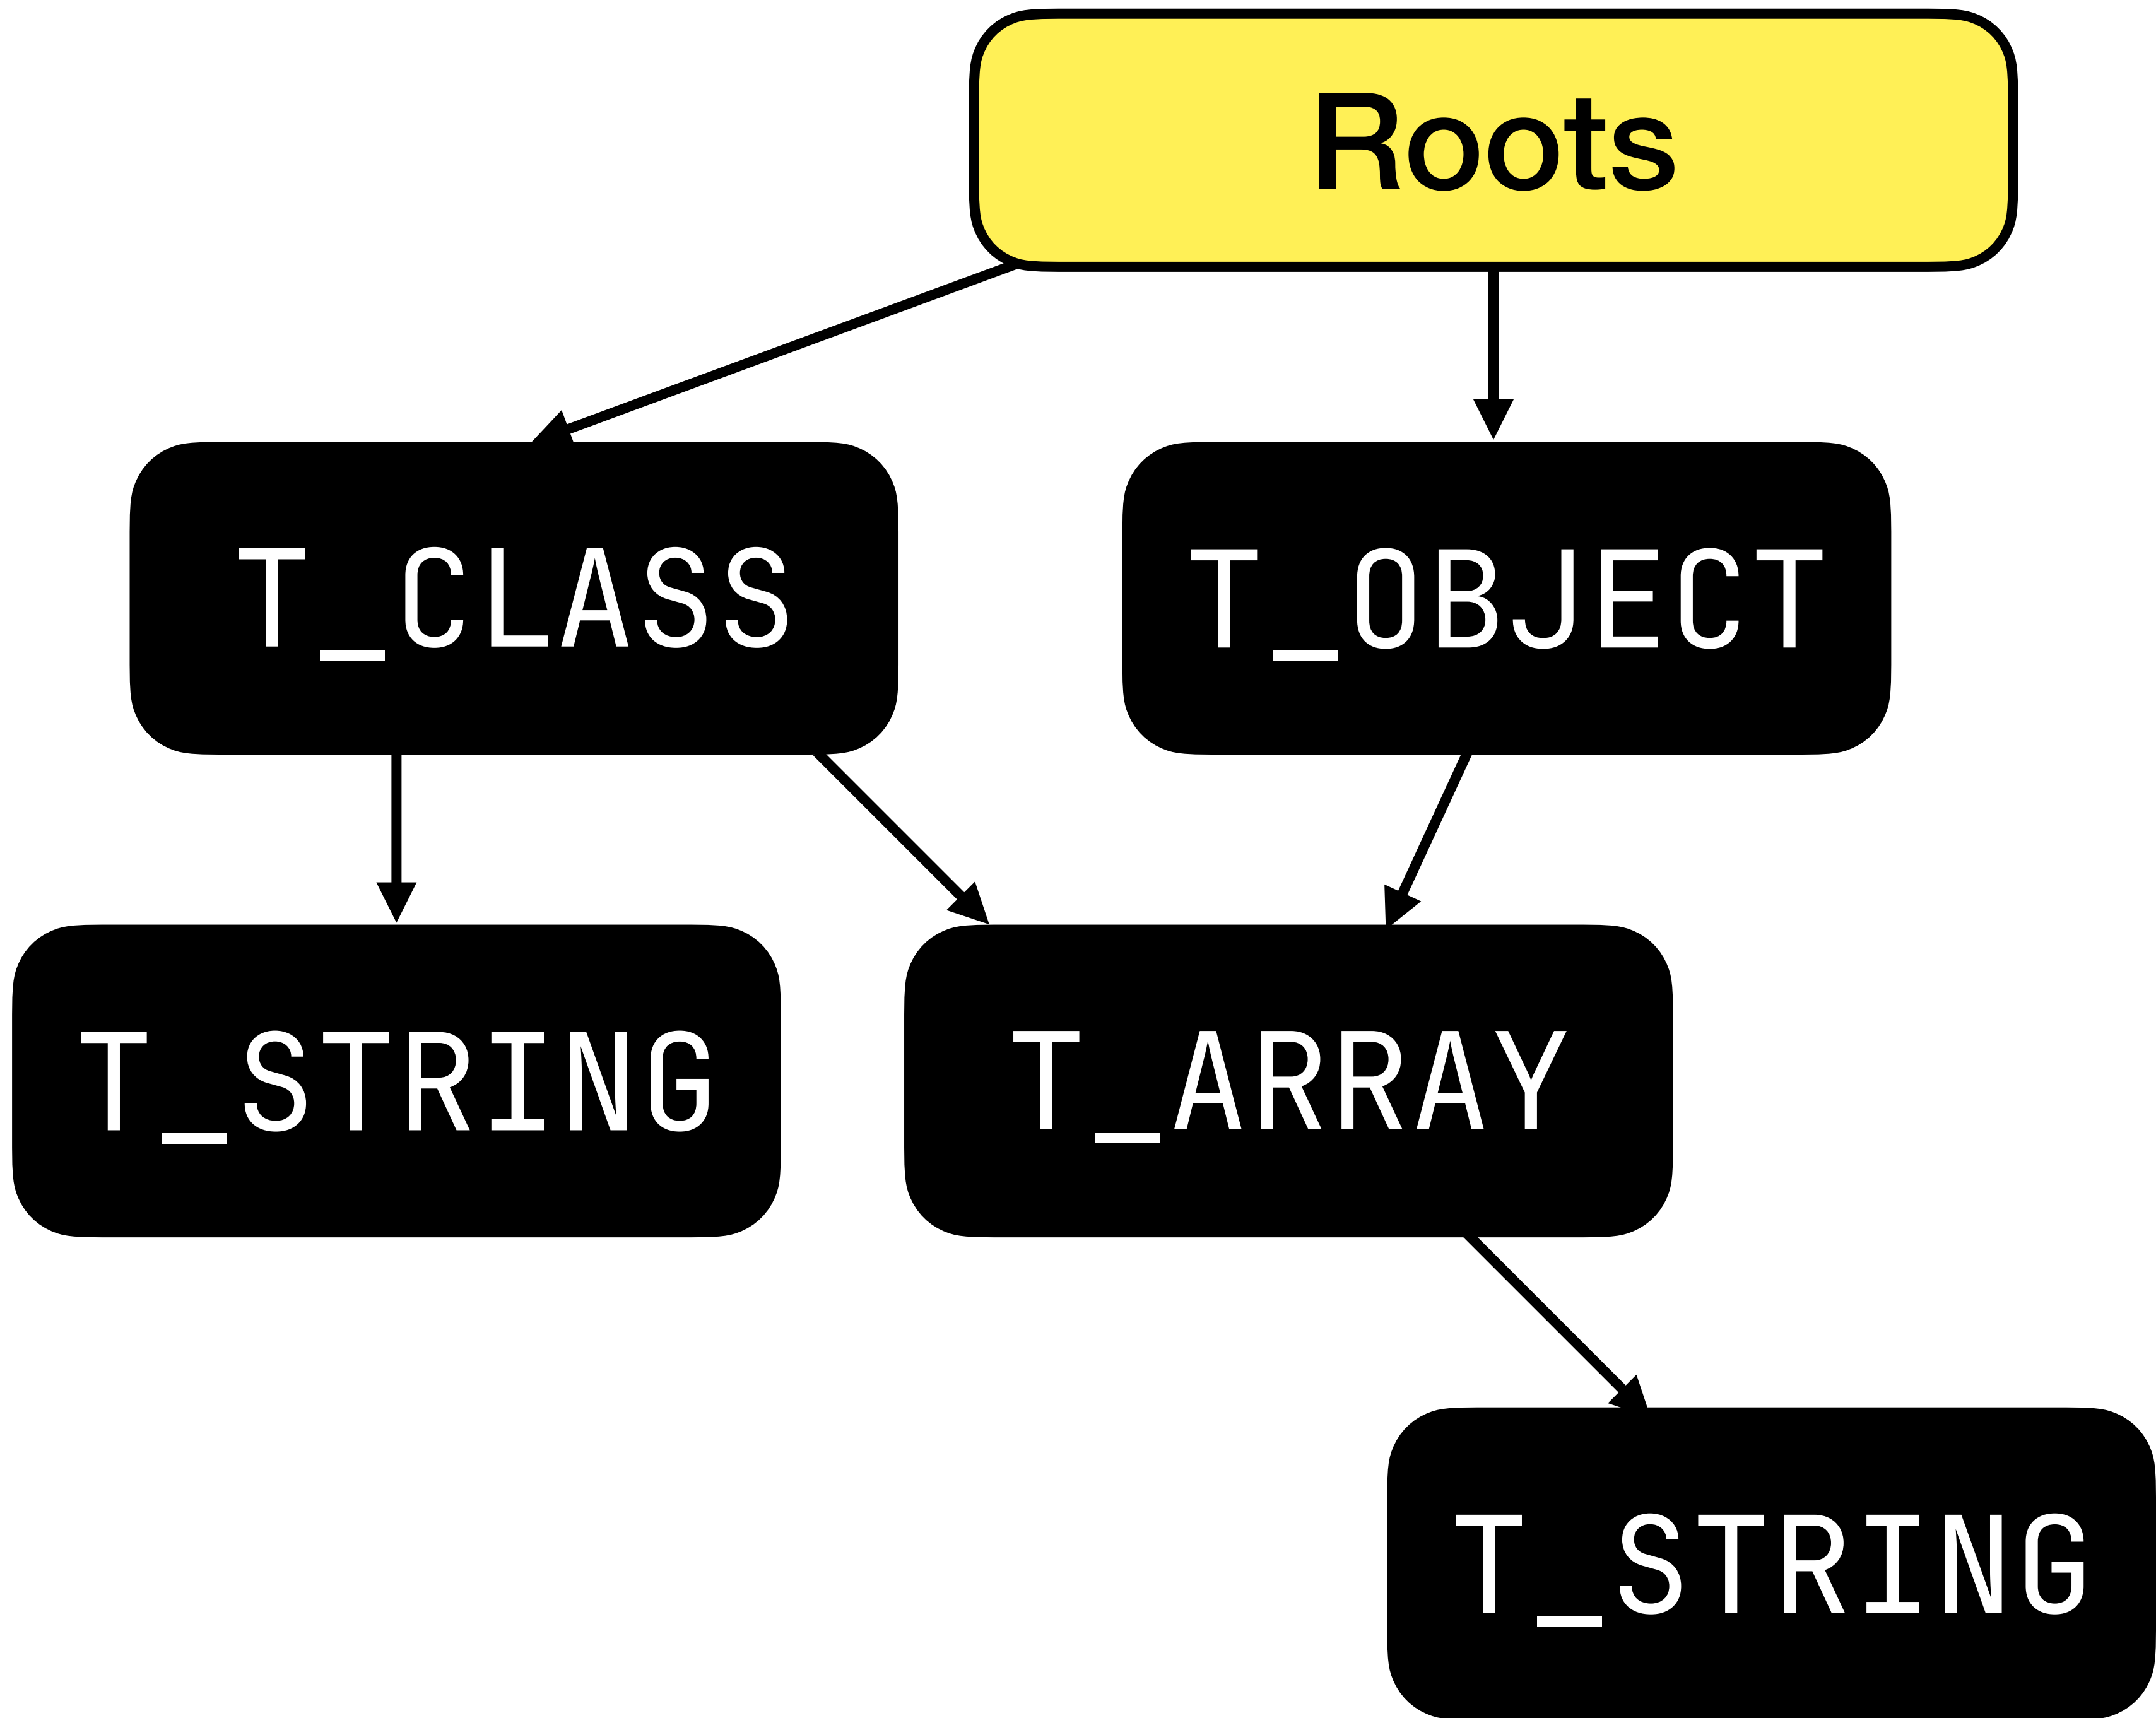




Sweep Phase







Compaction Phase

T_OBJECT		T_STRING		T_HASH
----------	--	----------	--	--------

T_OBJECT	T_STRING			T_HASH
----------	----------	--	--	--------

T_OBJECT	T_STRING	T_HASH		
----------	----------	--------	--	--



Garbage Collection in Ruby

Jul 03, 2023 (updated at Oct 02, 2023)

1. Overview

Ruby's garbage collector code lives in a single file called `gc.c`. "Garbage collector" is probably not the best term to call it because in addition to garbage collection, the code inside `gc.c` is responsible for memory allocation and management. In other words, the whole lifecycle of an object is managed by the garbage collector.

2. Primitive object types

While every type of Ruby object may appear to work the same in the Ruby source code, there

Modular Garbage Collectors

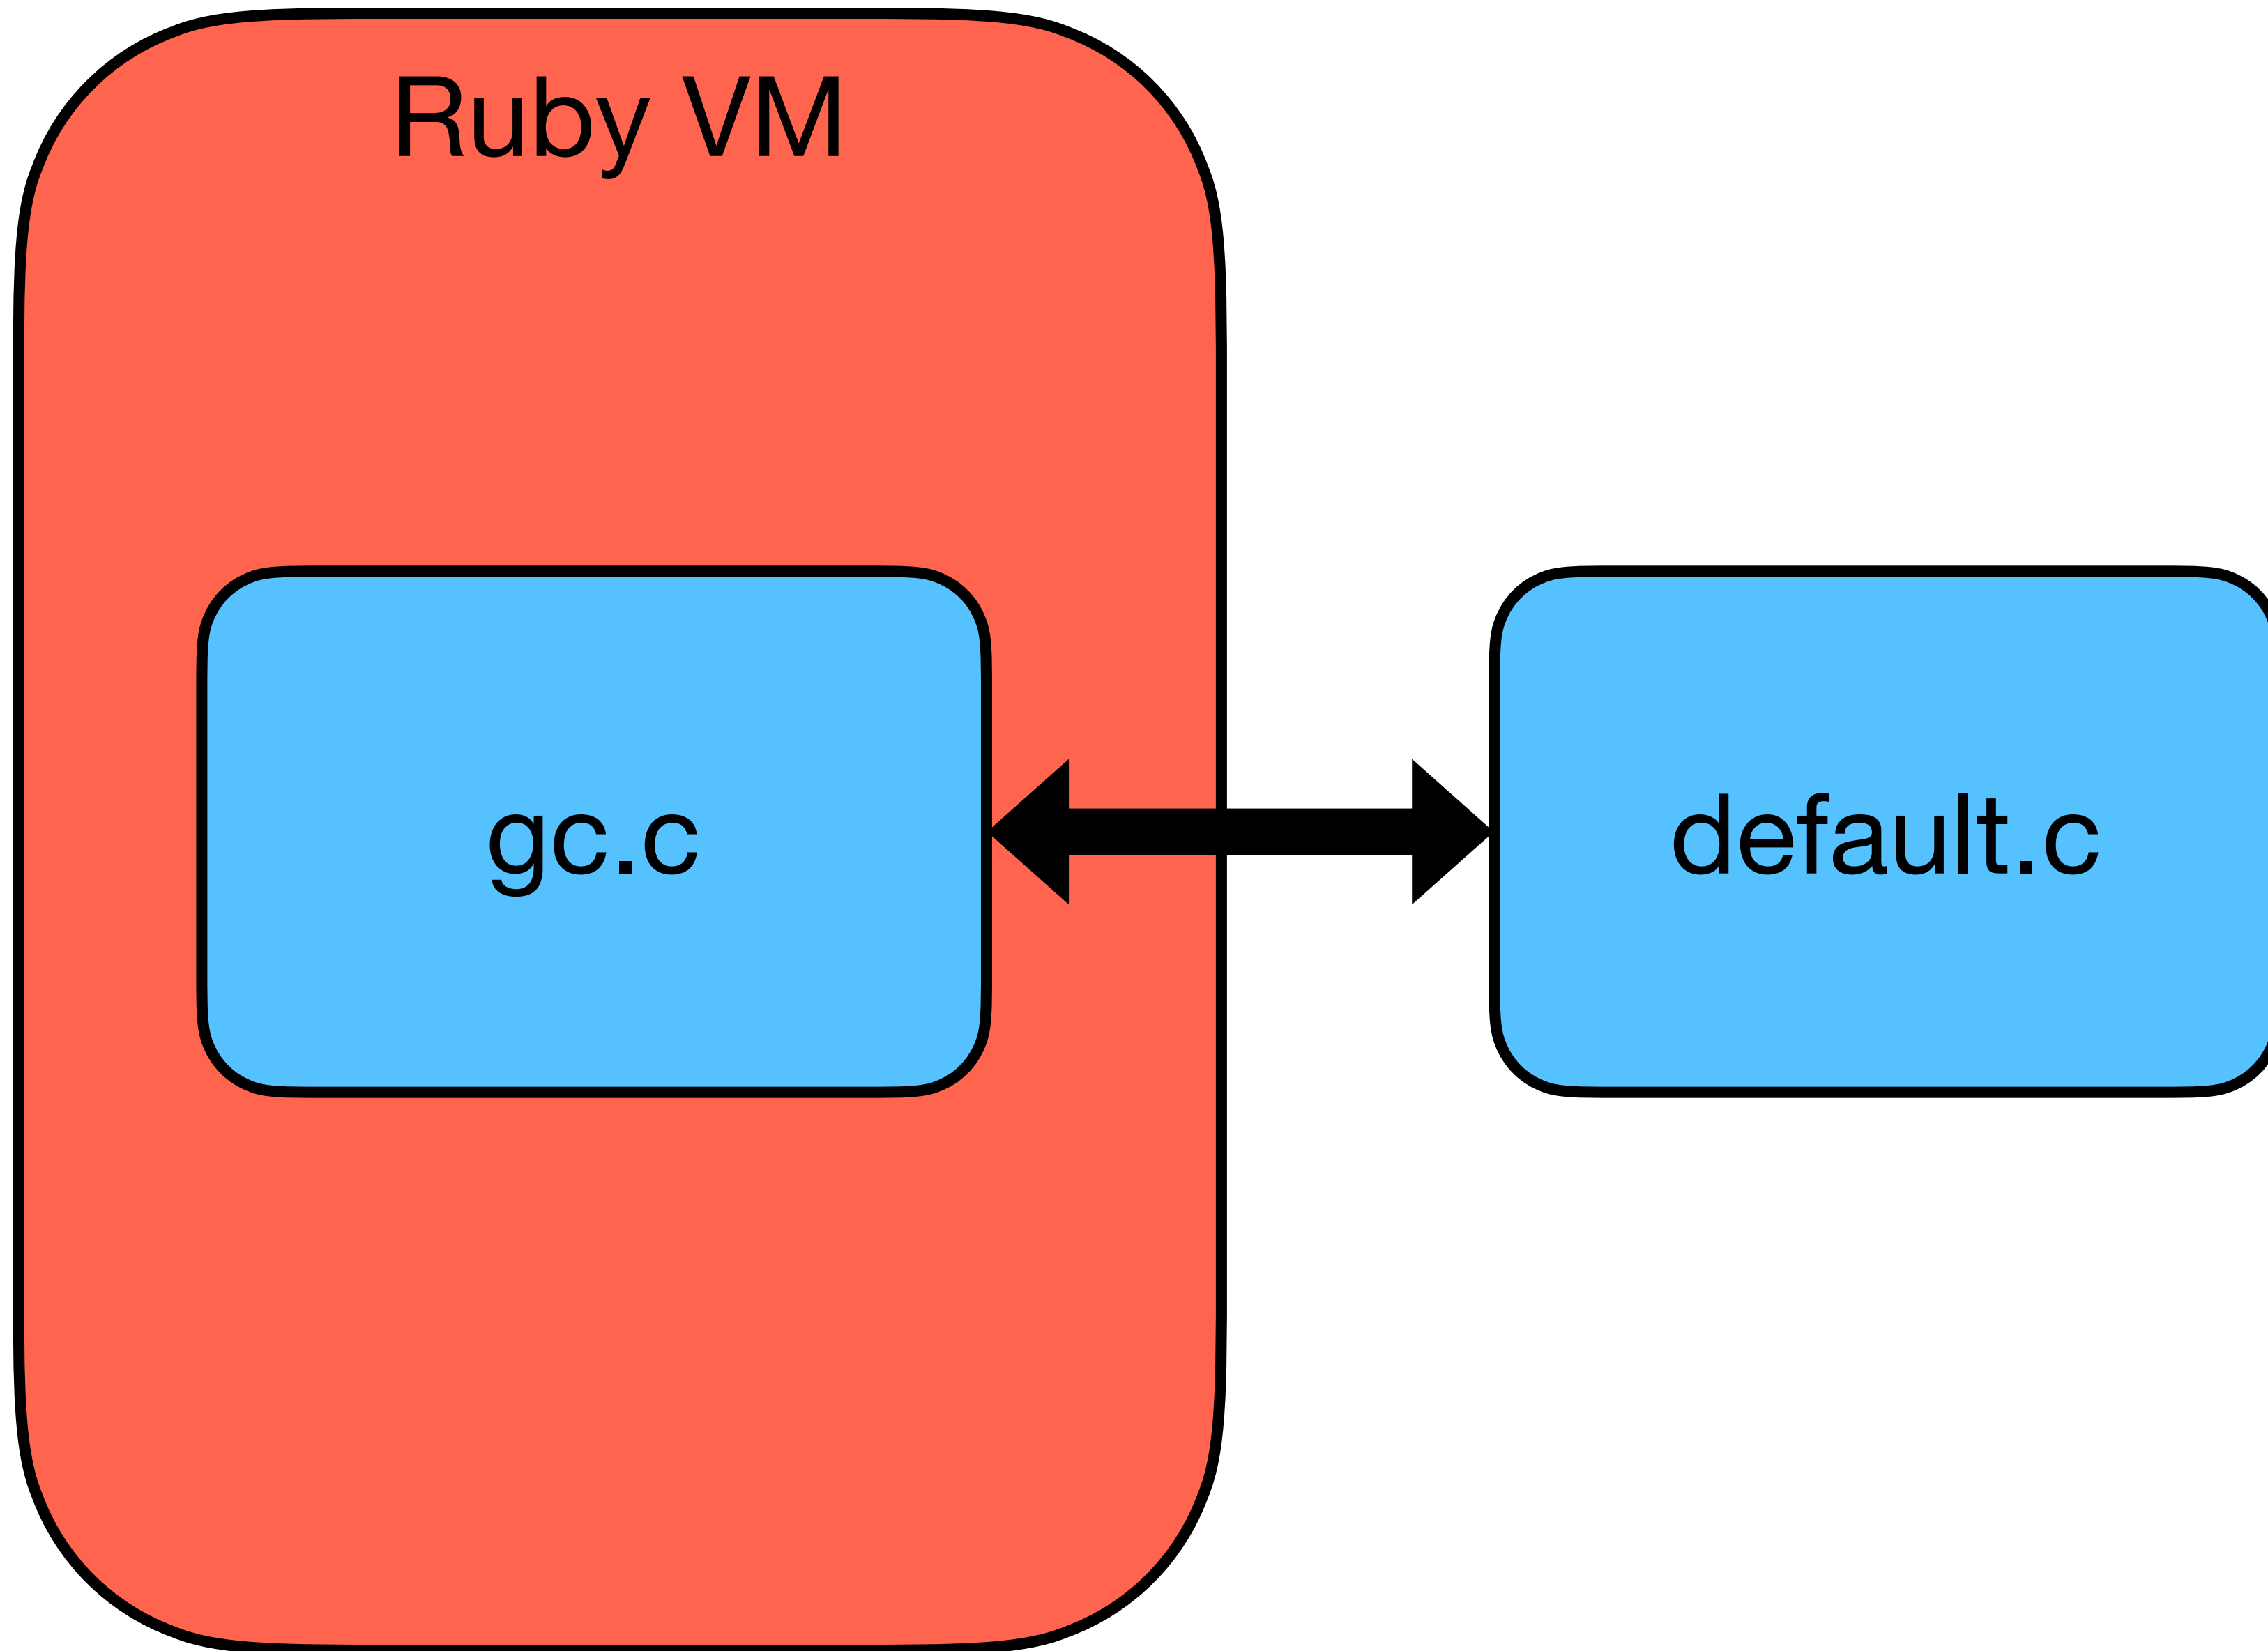
Ruby VM

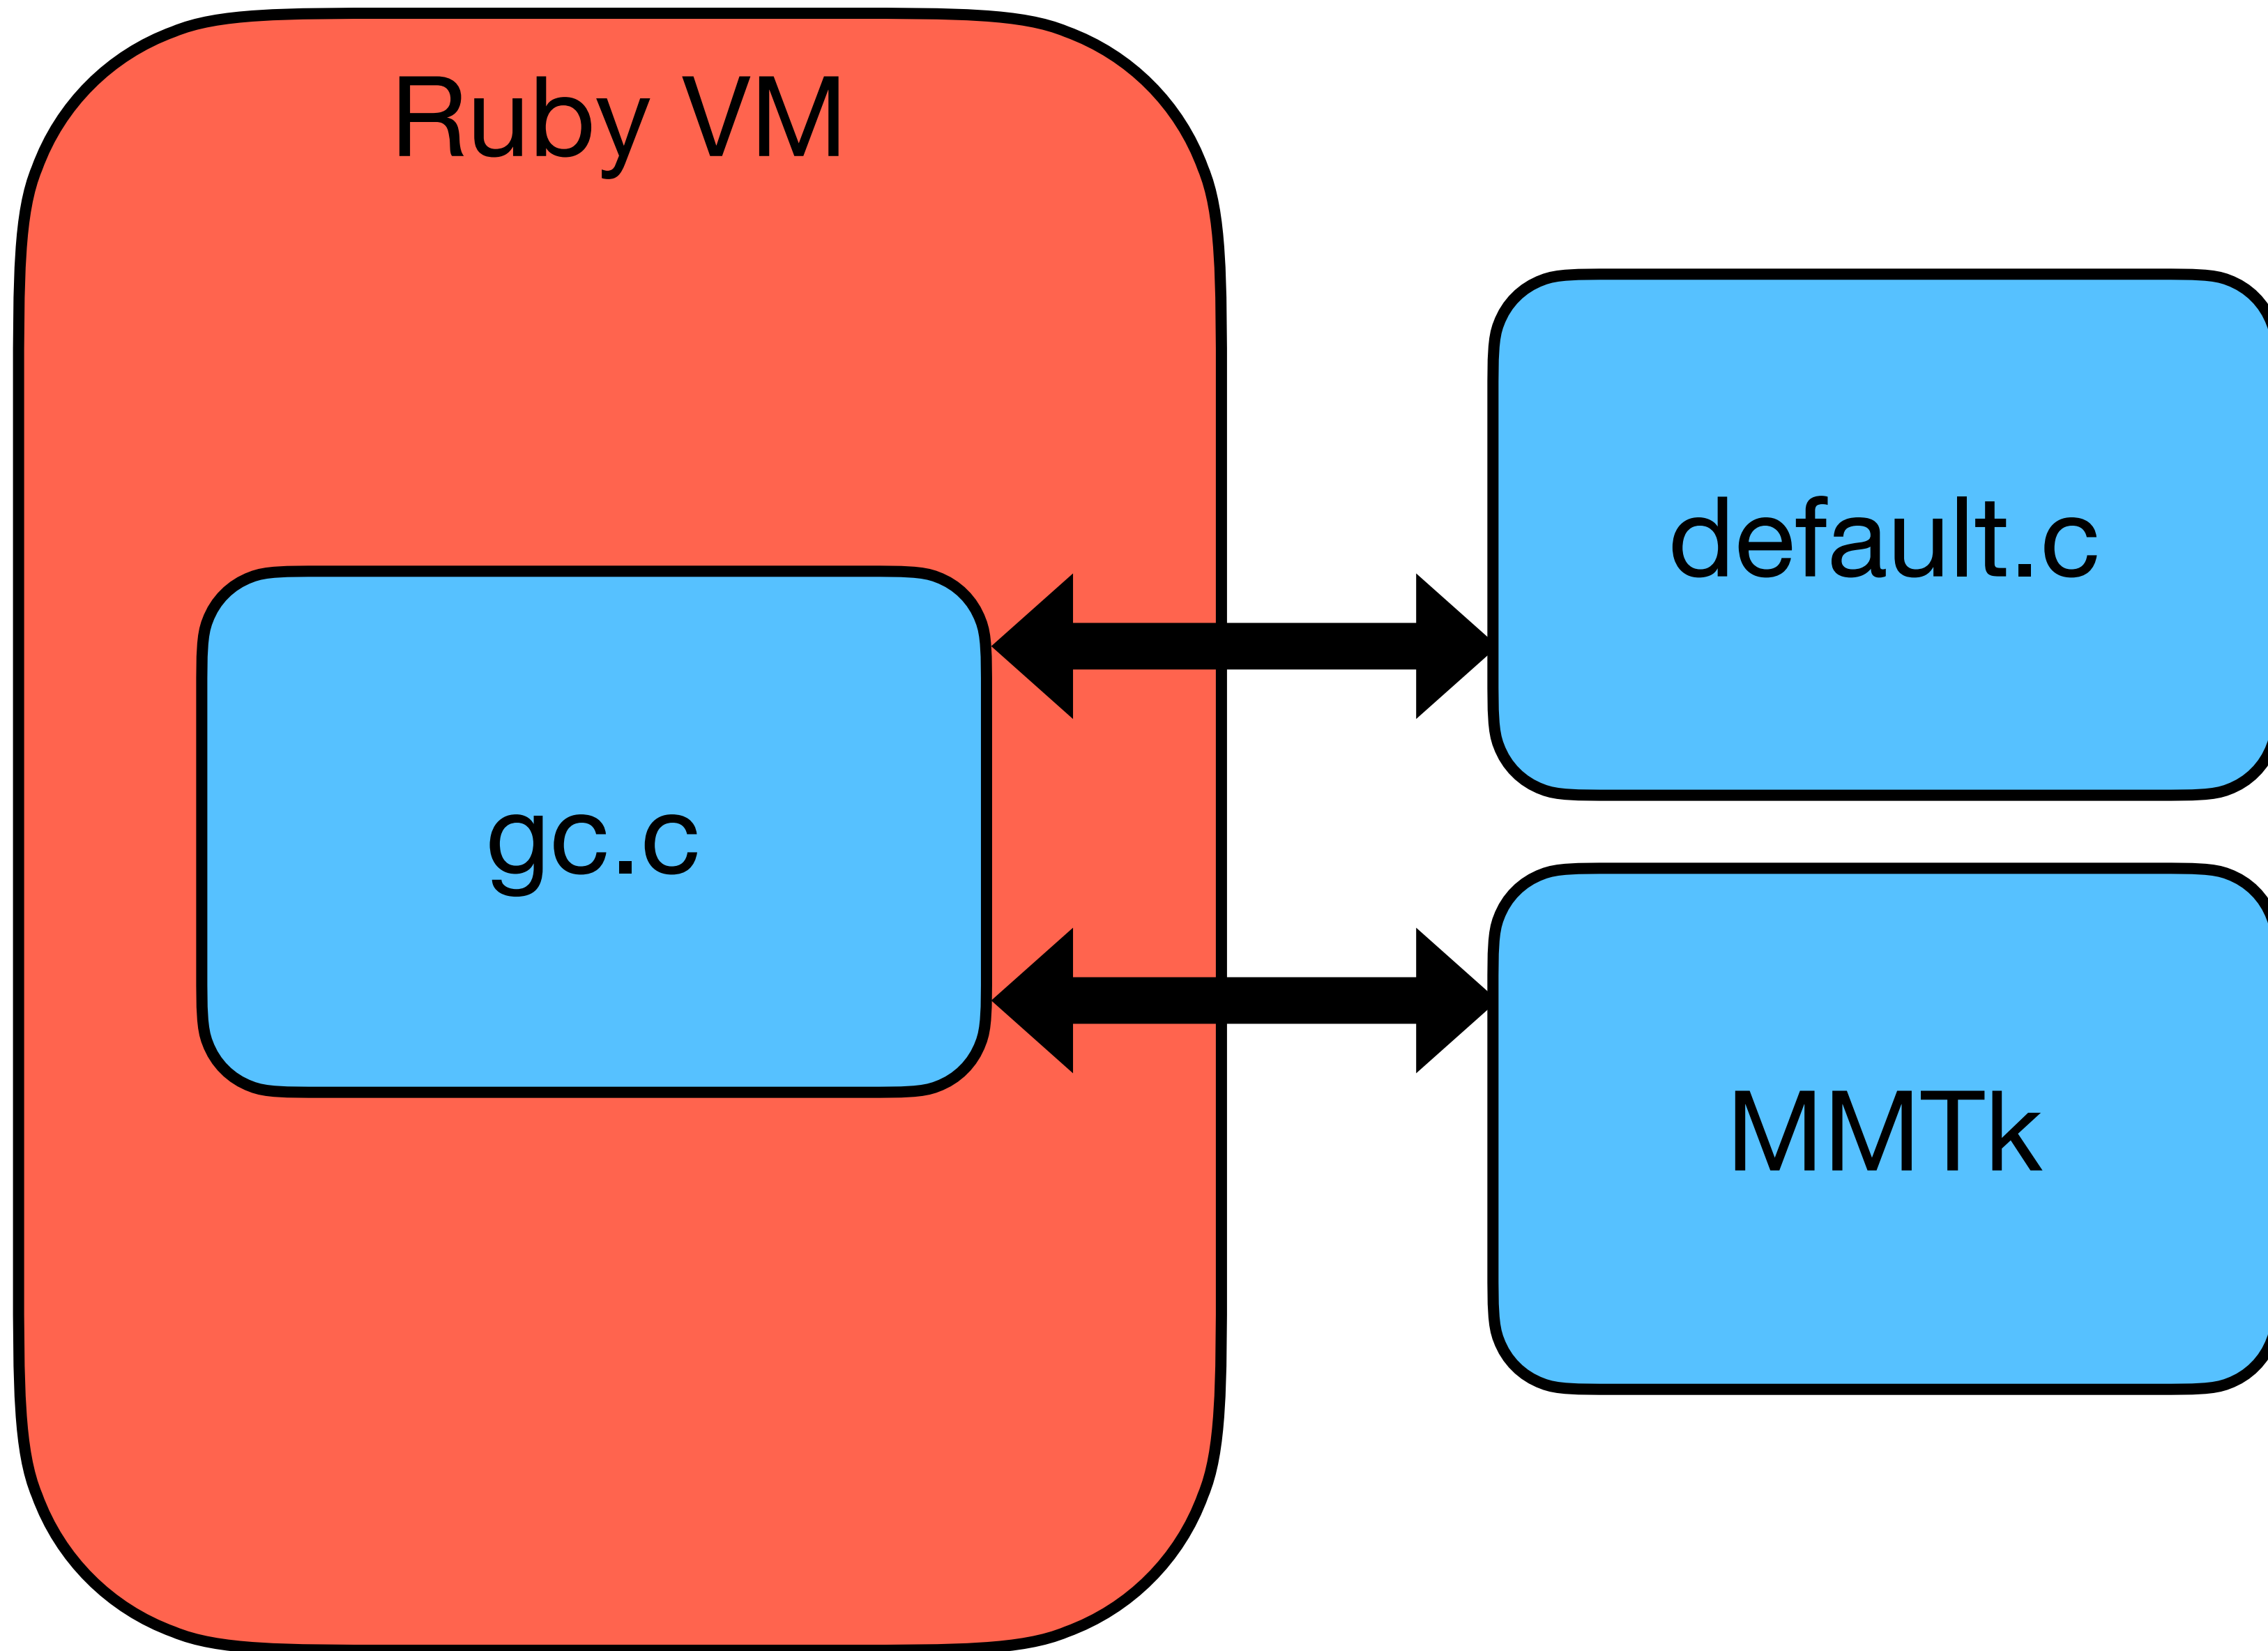
gc.c

Ruby VM

gc.c

default.c






youtu.be/04axm4JcaT4

Modular Garbage Collectors in Ruby

Peter Zhu

Ruby Core Committer
Senior Developer, Shopify



RubyKaigi 2025

Sponsored by



Memory Management Toolkit

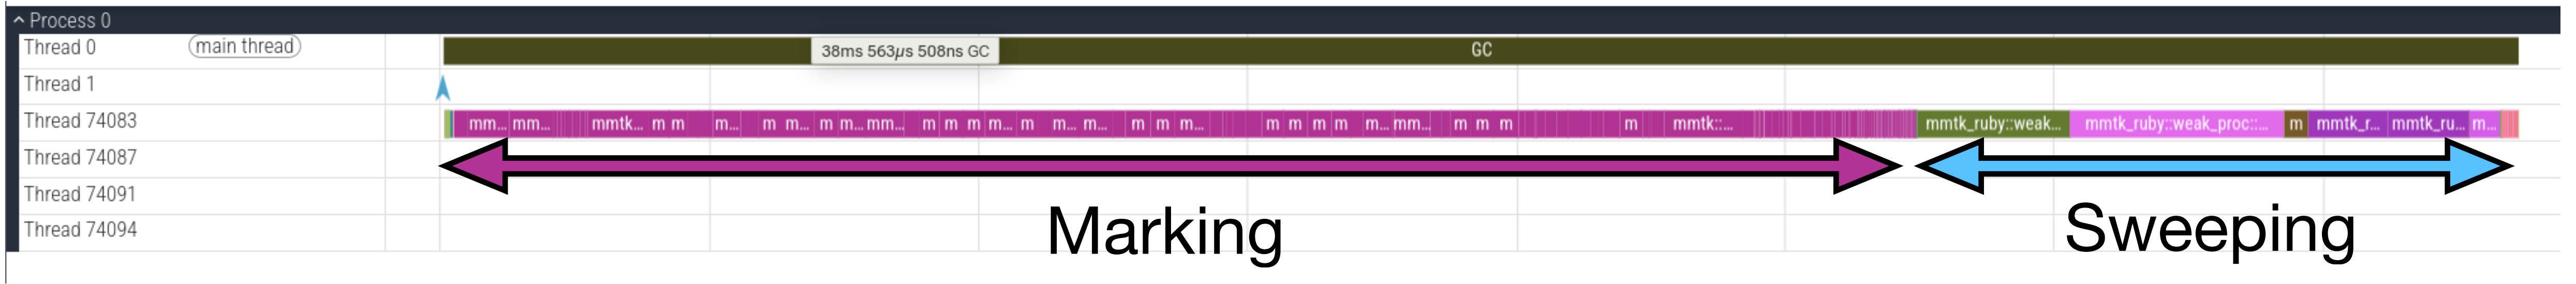
Memory Management Toolkit

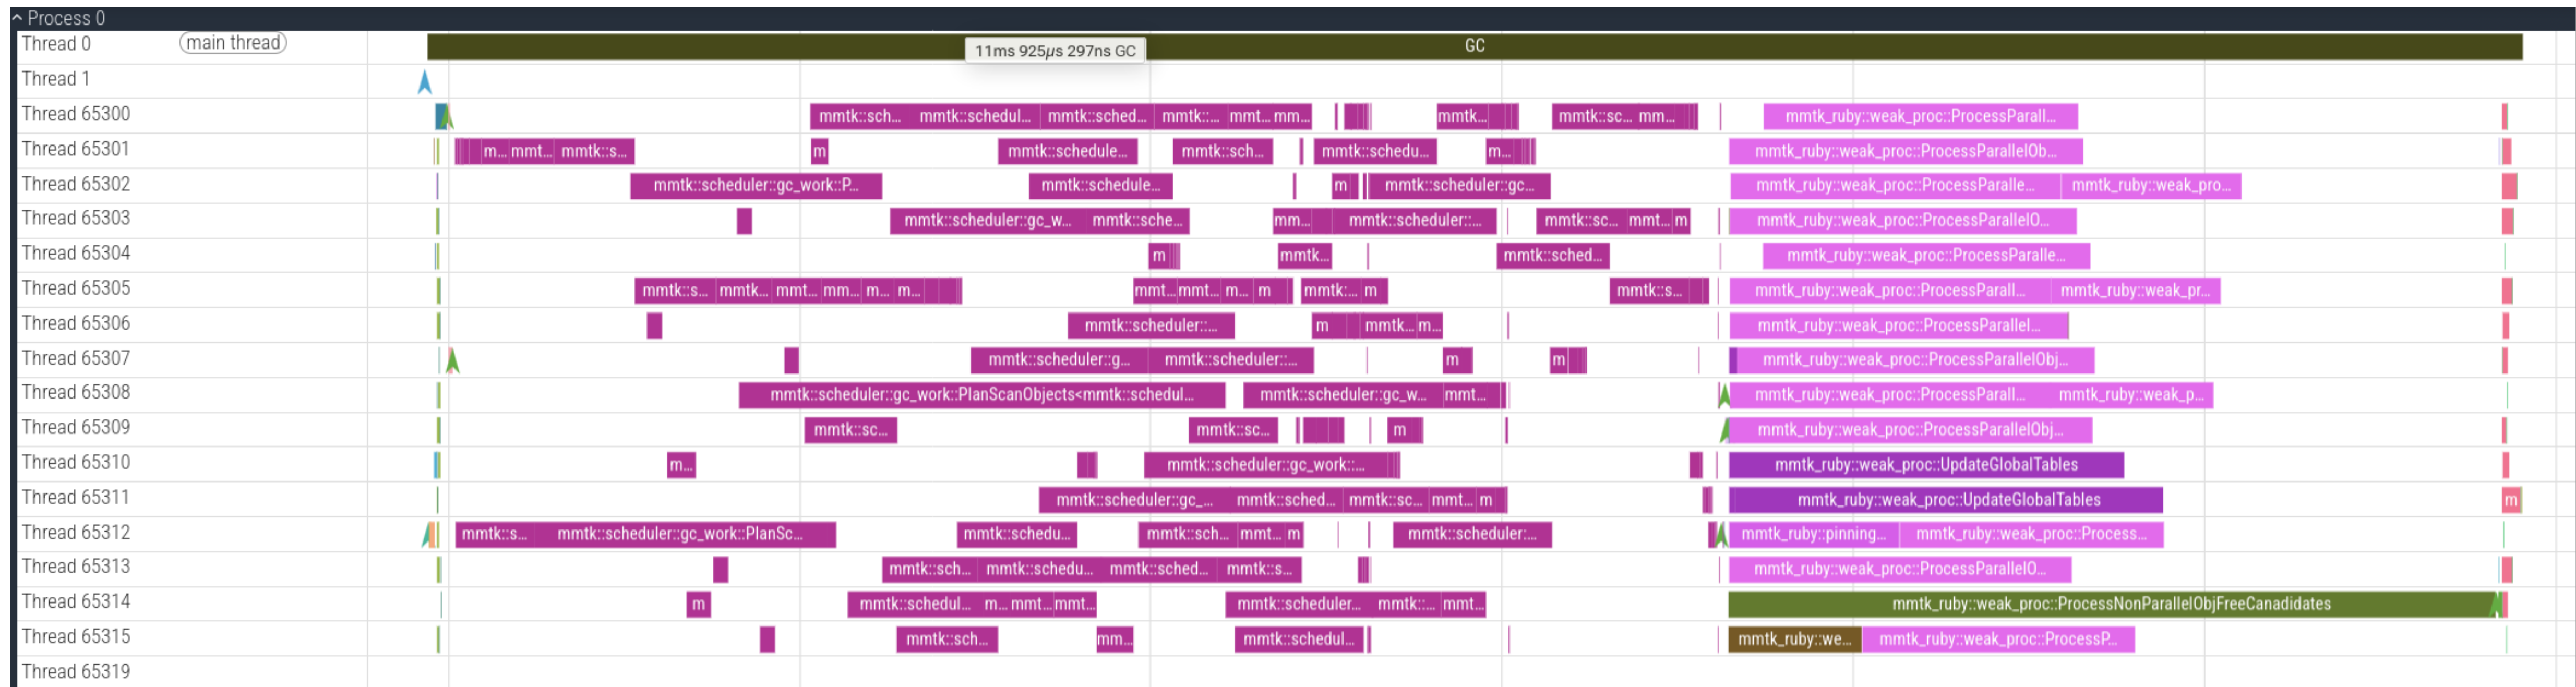
Alternative GC Algorithms

Memory Management Toolkit

Alternative GC Algorithms

Parallelism





38ms 563 μ s 508ns GC

11ms 925 μ s 297ns GC

h

mmtk''schedul

mmtk''sched

mmtk''

Memory Management Toolkit

Alternative GC Algorithms

Parallelism

Memory Management Toolkit

Alternative GC Algorithms

Parallelism

Dynamic Object Sizes

shopify.engineering/ruby-variable-width-allocation



Solutions ▾ Pricing Resources ▾ What's new ▾


Engineering Blog AI & Machine Learning Mobile Infrastructure Culture Latest More topics ▾

Search 🔍

BLOG | DEVELOPMENT

Optimizing Ruby's Memory Layout: Variable Width Allocation

Shopify is improving CRuby's performance in Ruby 3.2 by optimizing the memory layout in the garbage collector through the Variable Width Allocation project.



Memory Management Toolkit

Alternative GC Algorithms

Parallelism

Dynamic Object Sizes

Improvements to MMTk Since RubyKaigi 2025

Moving Immix

Immix

Immix Combines Mark and Compact

Immix Performs Selective Compaction

Moving Immix

github.com/ruby/ruby/pull/15744

ruby / ruby

Search: Type / to search

Code Pull requests 568 Agents Actions Wiki Security Insights

Implement moving Immix in MMTk #15744

Code

Merged peterzhu2118 merged 4 commits into ruby:master from peterzhu2118:register-pinning-obj on Dec 29, 2025

Conversation 2 Commits 4 Checks 91 Files changed 25 +470 -71



peterzhu2118 commented on Dec 26, 2025

Member

This commit implements moving Immix in MMTk, which allows objects to move in the GC.

The performance of this implementation is not yet amazing. It is very similar to non-moving Immix in many of them and slightly slower in others.

The benchmark results is shown below.

bench	Moving Immix (ms)	stddev (%)	RSS (MiB)
activerecord	241.9	0.5	86.6
chunky-png	447.8	0.8	74.9
erubi-rails	1183.9	0.8	136.1
hexapdf	1607.9	2.6	402.3
liquid-c	45.4	6.7	44.9
liquid-compile	44.1	9.3	53.0
liquid-render	105.4	4.5	55.9
lobsters	650.1	9.7	418.4
mail	115.4	2.1	64.4
psych-load	1656.8	0.8	43.6
railsbench	1653.5	1.3	149.8
rubocop	127.0	15.6	142.1
ruby-lexer	120.7	10.5	80.4

Reviewers

No reviews

Assignees

No one assigned

Labels

None yet

Projects

None yet


Milestone

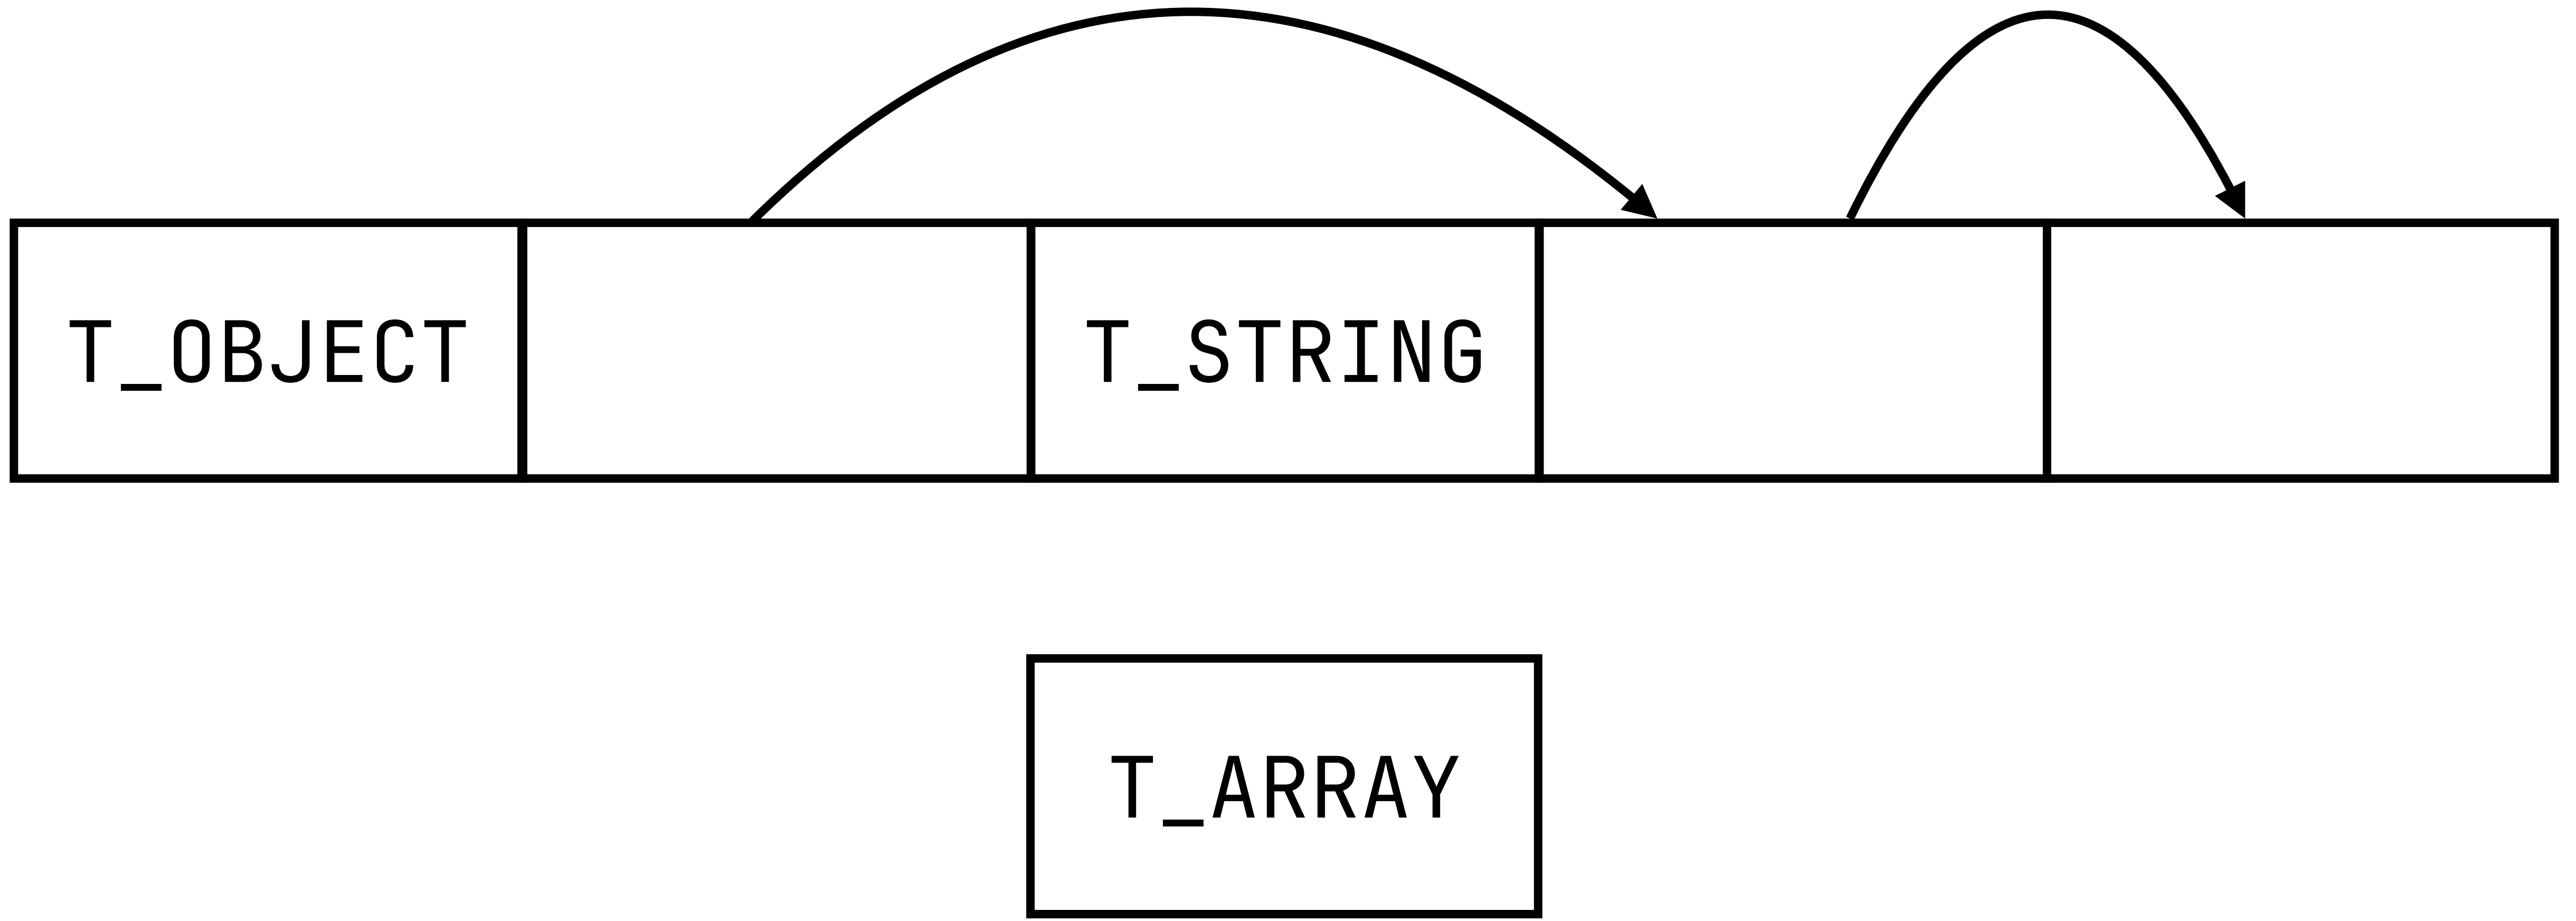
No milestone

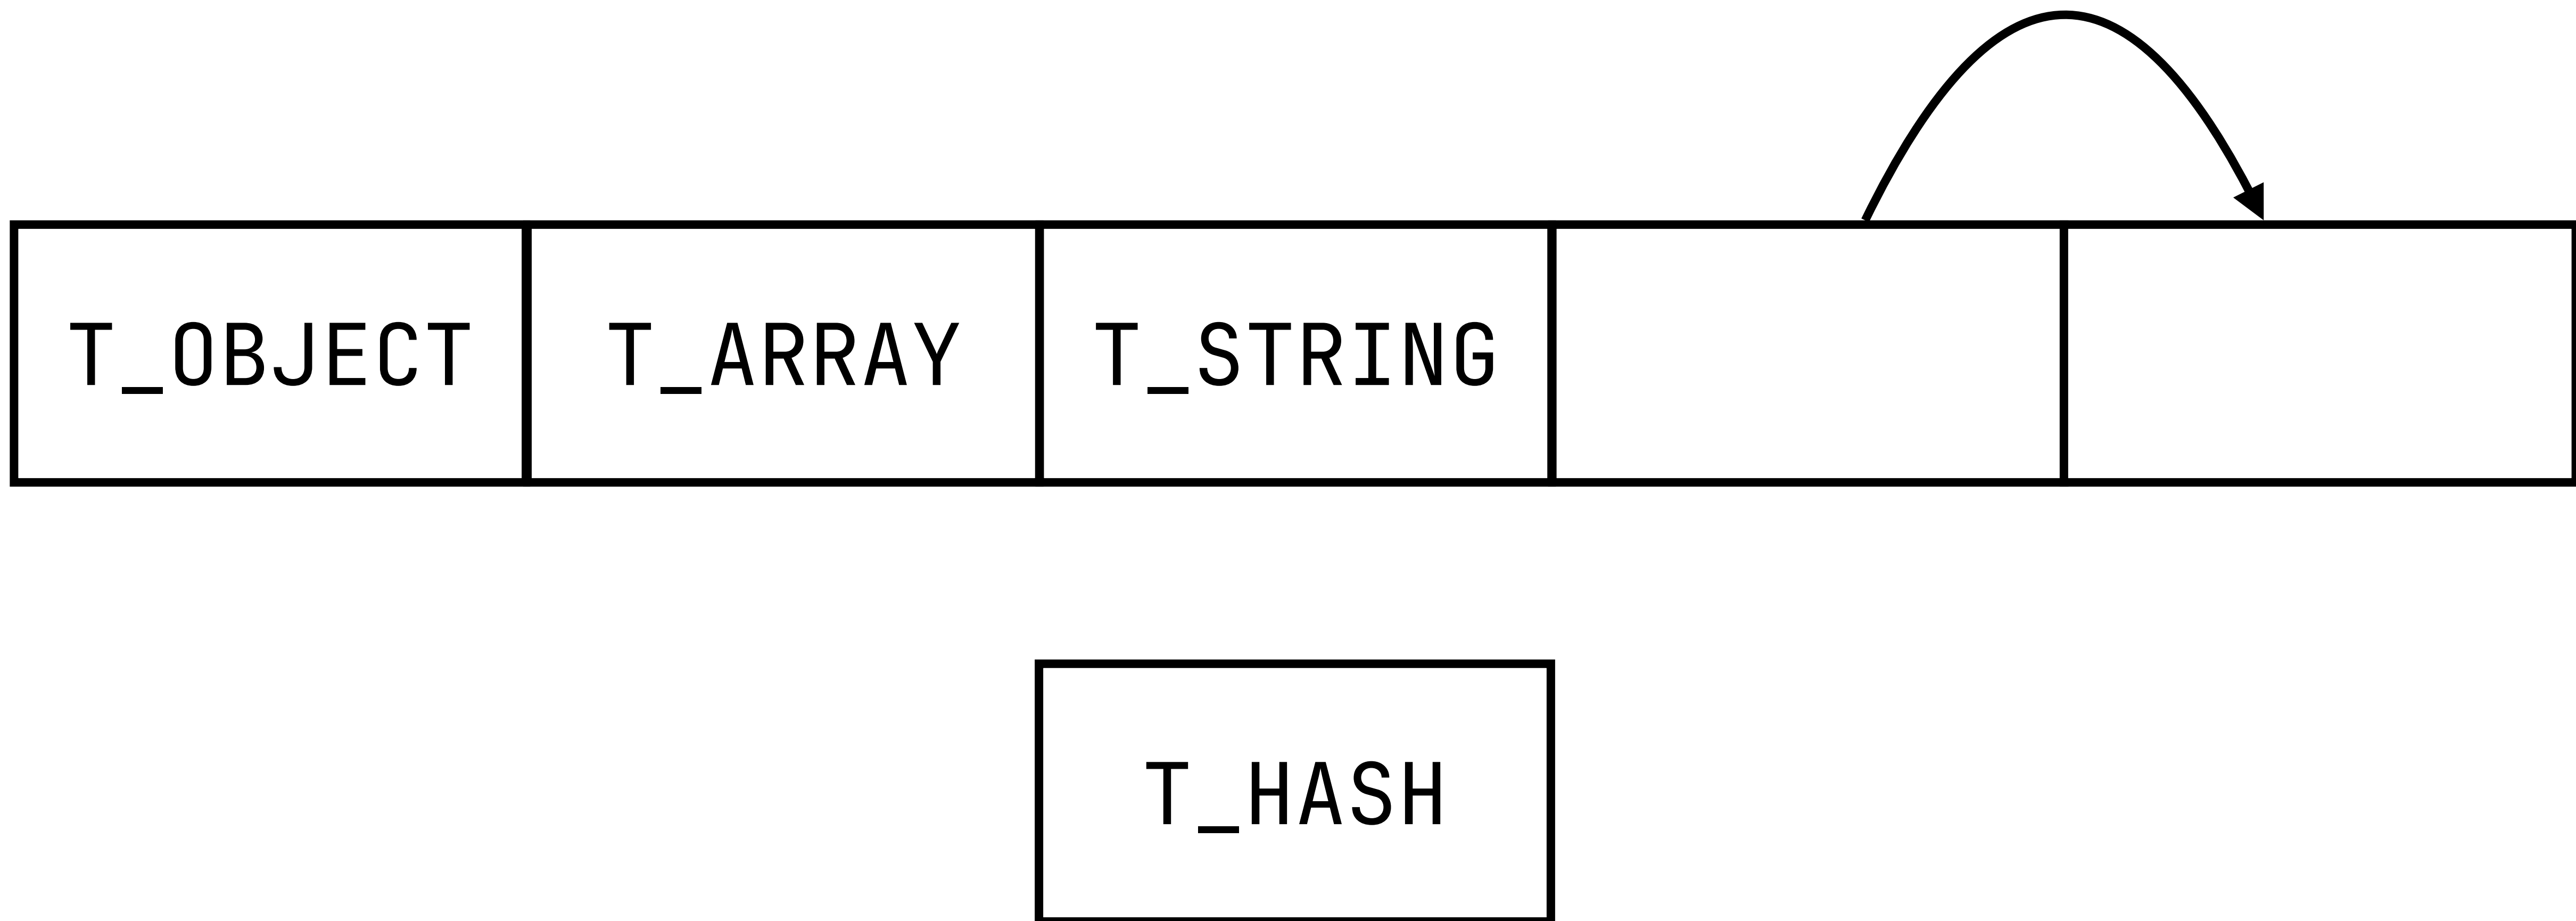
Development



Successfully merging this pull request may close these issues.


Bump Pointer Fast Path Allocator












```
def allocate_from_freelist
  obj = @freelist
  if obj
    @freelist = obj.next
    obj
  else
    gc_slowpath
  end
end
```



limit



cursor




T_ARRAY



CURSOR




limit




T_ARRAY

T_HASH




cursor



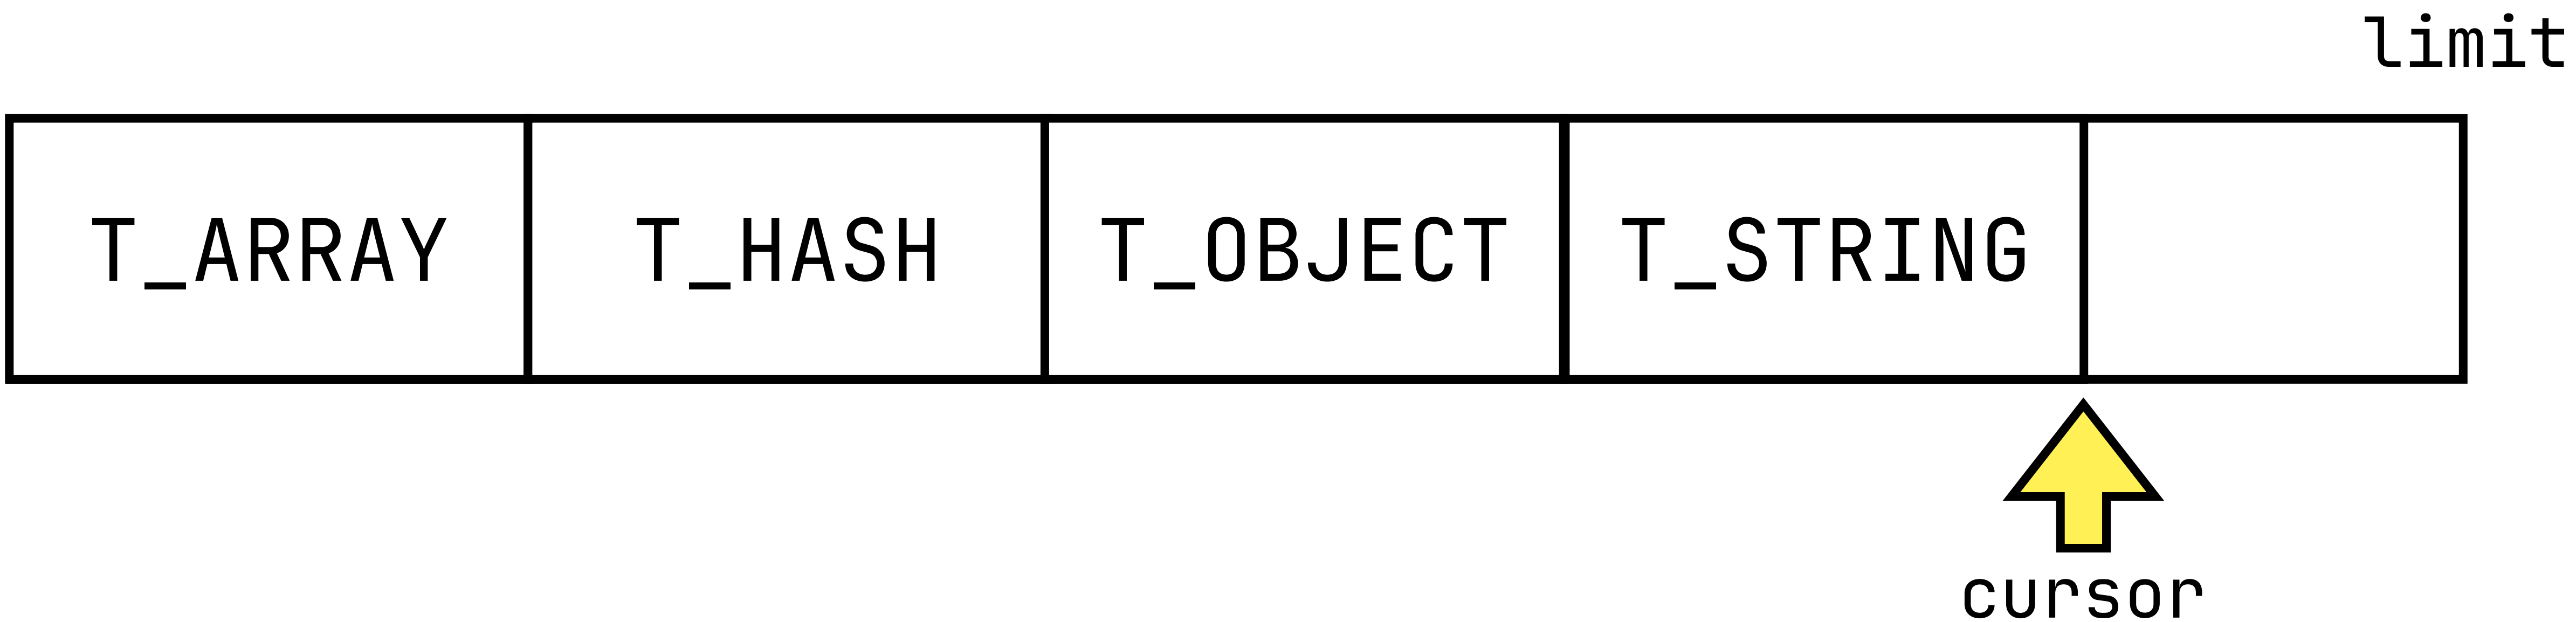
T_OBJECT


limit



↑
cursor







```
def allocate_from_bump_pointer(size)
  obj = @cursor
  @cursor += size
  if @cursor <= @limit
    obj
  else
    gc_slowpath
  end
end
```

```
def allocate_from_freelist
  obj = @freelist
  if obj
    @freelist = obj.next
    obj
  else
    gc_slowpath
  end
end
```



```
def allocate_from_bump_pointer(size)
  obj = @cursor
  @cursor += size
  if @cursor <= @limit
    obj
  else
    gc_slowpath
  end
end
```

Bump Pointer Fast Path Allocator

gc/mmtk/mmtk.c

```
static VALUE
rb_mmtk_alloc_fast_path(struct objspace *objspace, struct MMTk_ractor_cache *ractor_cache, size_t size)
{
    MMTk_BumpPointer *bump_pointer = ractor_cache->bump_pointer;
    if (bump_pointer == NULL) return 0;

    uintptr_t new_cursor = bump_pointer->cursor + size;

    if (new_cursor > bump_pointer->limit) {
        return 0;
    }
    else {
        VALUE obj = (VALUE)bump_pointer->cursor;
        bump_pointer->cursor = new_cursor;
        return obj;
    }
}
```

github.com/ruby/mmtk/pull/53

ruby / mmtk

Type / to search

Code Issues 2 Pull requests Agents Actions Projects Security Insights

Implement fast path for bump pointer allocator #53

Code ▾

Merged peterzhu2118 merged 1 commit into main from alloc-fast-path on Dec 20, 2025

Conversation 0 Commits 1 Checks 22 Files changed 4 +63 -3

peterzhu2118 commented on Dec 20, 2025 Member

Adding a fast path for bump pointer allocator can improve allocation performance.

For the following microbenchmark with MMTK_HEAP_MIN=100MiB:

```
10_000_000.times { String.new }
```

Before:

```
810.7 ms ± 8.3 ms [User: 790.9 ms, System: 40.3 ms]
```

After:

```
777.9 ms ± 10.4 ms [User: 759.0 ms, System: 37.9 ms]
```

Reviewers
No reviews

Assignees
No one assigned

Labels
None yet

Projects
None yet

Milestone
No milestone

Development
Successfully merging this pull request may close

Ruby Heap

Heaps in MMTk

Heaps in MMTk

Fixed Heap

Heaps in MMTk

Fixed Heap

Dynamic Heap

dl.acm.org/doi/10.1145/3563323

Optimal Heap Limits for Reducing Browser Memory Use

MARISA KIRISAME, University of Utah, United States

PRANAV SHENOY, University of Utah, United States

PAVEL PANCHEKHA, University of Utah, United States

Garbage-collected language runtimes carefully tune heap limits to reduce garbage collection time and memory usage. However, there's a trade-off: a lower heap limit reduces memory use but increases garbage collection time. Classic methods for setting heap limits include manually tuned heap limits and multiple-of-live-size rules of thumb, but it is not clear when one rule is better than another or how to compare them.

We address this problem with a new framework where heap limits are set for multiple heaps at once. Our key insight is that every heap limit rule induces a particular allocation of memory across multiple processes, and this allocation can be sub-optimal. We use our framework to derive an optimal “square-root” heap limit rule, which minimizes total memory usage for any amount of total garbage collection time. Paradoxically, the square-root heap limit rule achieves this coordination without communication: it allocates memory optimally across multiple heaps without requiring any communication between those heaps.

To demonstrate that this heap limit rule is effective, we prototype it for V8, the JavaScript runtime used in Google Chrome, Microsoft Edge, and other browsers, as well as in server-side frameworks like node.js and Deno. On real-world web pages, our prototype achieves reductions of approximately 16.0% of memory usage while keeping garbage collection time constant. On memory-intensive benchmarks, reductions of up to 30.0% of garbage collection time are possible with no change in total memory usage.

CCS Concepts: • **Software and its engineering** → **Runtime environments**; • **Information systems** →

Ruby Heap



Used: 60%

Empty: 40%



Used: 100%

Empty: 0%



Run GC



Used: 90%

Empty: 10%



Used: 60%

Empty: 40%



Used: 100%

Empty: 0%



Run GC



Used: 20%

Empty: 80%



Used: 60%

Empty: 40%

Ruby Heap

github.com/ruby/mmtk/pull/54

ruby / mmtk

Q Type / to search

<> Code Issues 2 Pull requests Agents Actions Projects Security Insights

Implement Ruby heap #54

Code ▾

Merged peterzhu2118 merged 1 commit into main from ruby-heap on Dec 22, 2025

Conversation 0 Commits 1 Checks 22 Files changed 6 +176 -5



peterzhu2118 commented on Dec 22, 2025

Member ...

This heap emulates the growth characteristics of the Ruby default GC's heap. By default, the heap grows by 40%, requires at least 20% empty after a GC, and allows at most 65% empty before it shrinks the heap. This is all configurable via the same environment variables the default GC uses (`RUBY_GC_HEAP_FREE_SLOTS_GOAL_RATIO`, `RUBY_GC_HEAP_FREE_SLOTS_MIN_RATIO`, `RUBY_GC_HEAP_FREE_SLOTS_MAX_RATIO`, respectively).

The Ruby heap can be enabled via the `MMTK_HEAP_MODE=ruby` environment variable.

Compared to the dynamic heap in MMTk (which uses the MemBalancer algorithm), the Ruby heap allows the heap to grow more generously, which uses a bit more memory but offers significant performance gains because it runs GC much less frequently.

We can see in the benchmarks below that this Ruby heap heap gives faster performance than the dynamic heap in every benchmark, with over 2x faster in many of them. We see that memory is often around 10-20% higher with certain outliers that use significantly more memory like hexapdf and erubi-rails. We can also see that this brings MMTk's Ruby heap much closer in performance to the default GC.

Ruby heap benchmark results:

Reviewers

No reviews

Assignees

No one assigned

Labels

None yet

Projects

None yet

Milestone

No milestone

Development

Successfully merging this pull request may close these issues.

activerecord

**Dynamic
Heap**

845ms

76.7MB

**Ruby
Heap**

233ms

85.9MB

**Default
GC**

148ms

67.9MB

3.63x



1.12x



1.57x



1.26x



railsbench

**Dynamic
Heap**

3218ms $\xrightarrow{1.89x}$

124MB $\xrightarrow{1.18x}$

**Ruby
Heap**

1707ms $\xrightarrow{1.35x}$

146MB $\xrightarrow{1.27x}$

**Default
GC**

1263ms

115MB

github.com/ruby/mmtk/pull/54

ruby / mmtk

Q Type / to search

<> Code Issues 2 Pull requests Agents Actions Projects Security Insights

Implement Ruby heap #54

Code ▾

Merged peterzhu2118 merged 1 commit into main from ruby-heap on Dec 22, 2025

Conversation 0 Commits 1 Checks 22 Files changed 6 +176 -5



peterzhu2118 commented on Dec 22, 2025

Member ...

This heap emulates the growth characteristics of the Ruby default GC's heap. By default, the heap grows by 40%, requires at least 20% empty after a GC, and allows at most 65% empty before it shrinks the heap. This is all configurable via the same environment variables the default GC uses (`RUBY_GC_HEAP_FREE_SLOTS_GOAL_RATIO`, `RUBY_GC_HEAP_FREE_SLOTS_MIN_RATIO`, `RUBY_GC_HEAP_FREE_SLOTS_MAX_RATIO`, respectively).

The Ruby heap can be enabled via the `MMTK_HEAP_MODE=ruby` environment variable.

Compared to the dynamic heap in MMTk (which uses the MemBalancer algorithm), the Ruby heap allows the heap to grow more generously, which uses a bit more memory but offers significant performance gains because it runs GC much less frequently.

We can see in the benchmarks below that this Ruby heap heap gives faster performance than the dynamic heap in every benchmark, with over 2x faster in many of them. We see that memory is often around 10-20% higher with certain outliers that use significantly more memory like hexapdf and erubi-rails. We can also see that this brings MMTk's Ruby heap much closer in performance to the default GC.

Ruby heap benchmark results:

Reviewers

No reviews

Assignees

No one assigned

Labels

None yet

Projects

None yet

Milestone

No milestone

Development

Successfully merging this pull request may close these issues.

Future Roadmap

Sticky Immix

Generational GC

Weak Generational Hypothesis:
Most Objects Die Young

Generational GC

Young & Old Objects

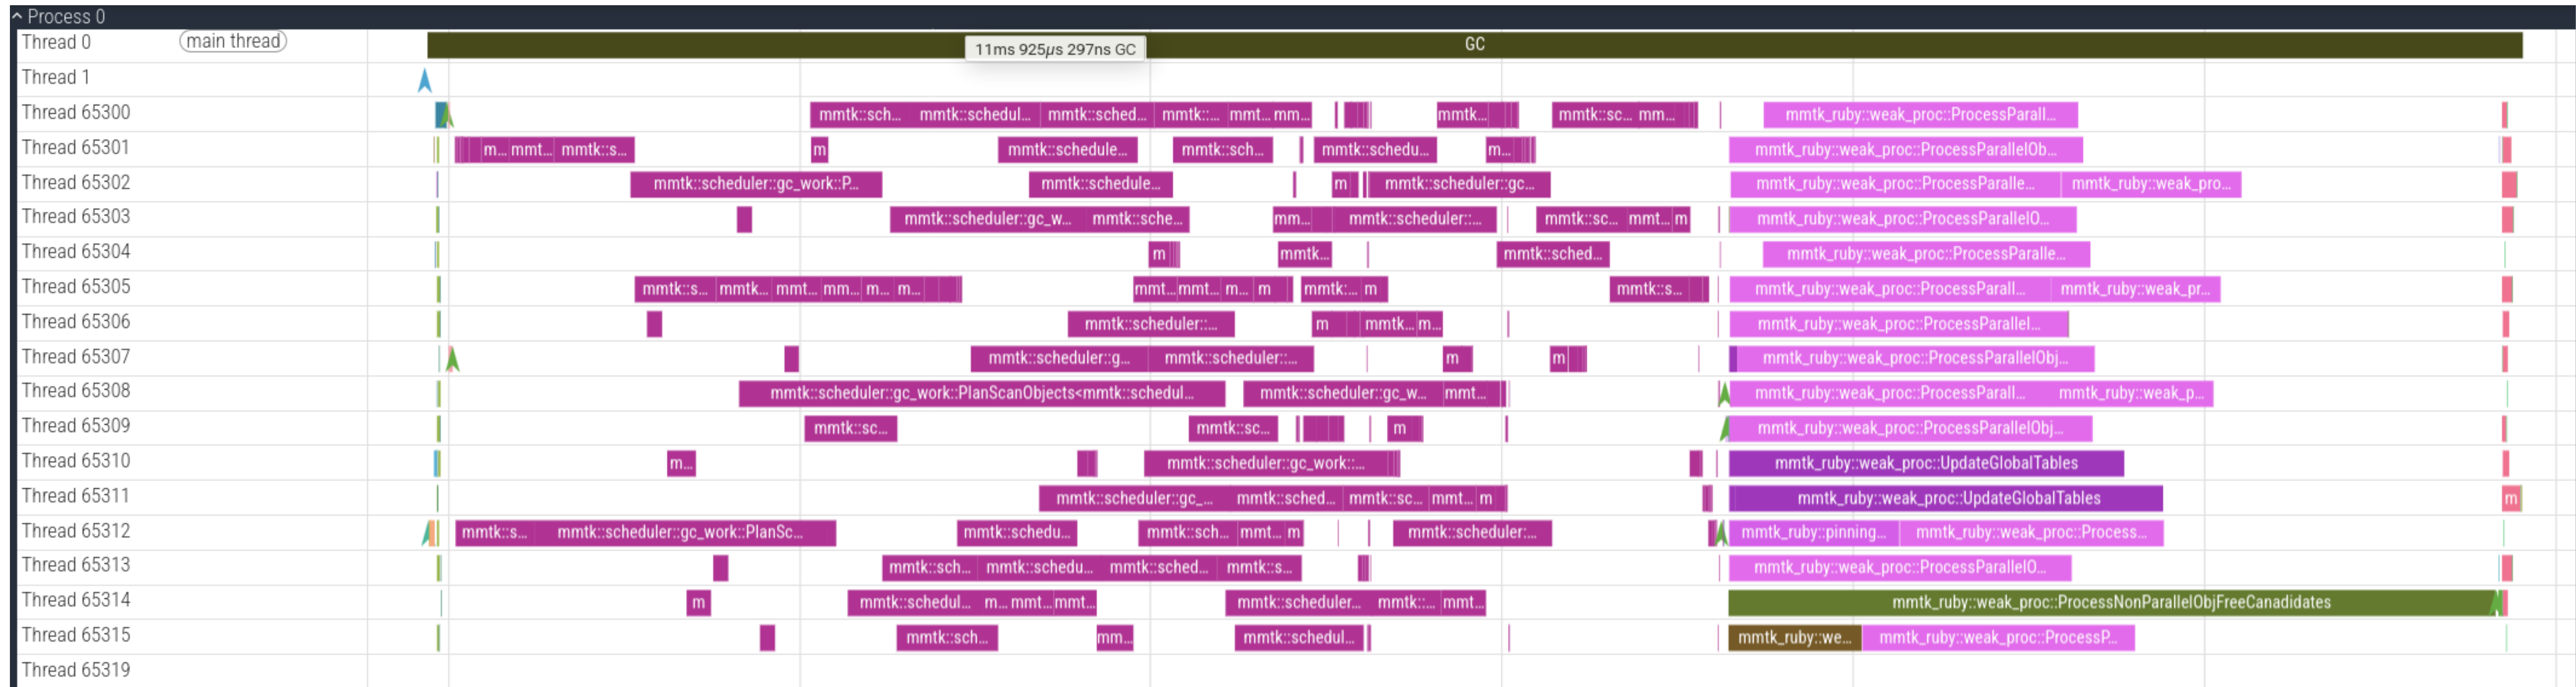
Young & Old Objects
Minor & Major GC

Sticky Immix

Sticky Mark Bit

Sticky Immix

Improve Parallelism



Implement Buffer Objects for String & Array

Inline GC Fast Path in JIT Compiler

- Introduce Ruby's Garbage Collector
- Modular Garbage Collectors
- Memory Management Toolkit
 - Improvements since RubyKaigi 2025
 - Moving Immix
 - Bump Pointer Allocator Fast Path
 - Ruby Heap
 - Future Roadmap

Thank You!



blog.peterzhu.ca



peter@peterzhu.ca




[@peterzhu2118](https://twitter.com/peterzhu2118)



[@peterzhu2118@ruby.social](https://ruby.social/@peterzhu2118)



[@peterzhu.ca](https://x.com/peterzhu.ca)



[@peterzhu.photos](https://www.instagram.com/peterzhu.photos)

