


Garbage collection



Jeremy Yallop

jeremy.yallop@cl.cam.ac.uk

Algorithms

A **heap**: of one or more blocks of contiguous words

A **object**: a heap-allocated contiguous region addressed by 0+ pointers

A **mutator**: application thread, opaque to the collector except for heap operations (allocate, read, write)

A **root**: a heap pointer accessible to the mutator
(e.g. in static global storage, stack space, or registers)


An object is **live** if a mutator will access it in the future

An object is **reachable** if there is a chain of pointers to it from a root

Mark-and-sweep collection

Algorithms

```
Mark  
mark(node) =  
  if not node.marked:  
    node.marked = True  
    for c in node.children:  
      mark(c)
```



Performance


Reading



Mark-and-sweep collection


Algorithms

```
Mark  
mark(node) =  
  if not node.marked:  
    node.marked = True  
    for c in node.children:  
      mark(c)
```



Performance


Reading



Mark-and-sweep collection

Algorithms

```
Mark  
mark(node) =  
  if not node.marked:  
    node.marked = True  
    for c in node.children:  
      mark(c)
```




Performance

Reading

Mark-and-sweep collection

Algorithms

```
Mark  
mark(node) =  
  if not node.marked:  
    node.marked = True  
    for c in node.children:  
      mark(c)
```



Performance


Reading

Progress indicator for the Algorithms section, showing a sequence of seven circles. The first circle is filled, indicating the current position in the sequence.

Mark-and-sweep collection

Algorithms

```
Mark  
mark(node) =  
  if not node.marked:  
    node.marked = True  
    for c in node.children:  
      mark(c)
```




Performance

Reading

Mark-and-sweep collection

Algorithms

```
Mark  
mark(node) =  
  if not node.marked:  
    node.marked = True  
    for c in node.children:  
      mark(c)
```




Performance

Reading

Mark-and-sweep collection

Algorithms

```
Mark  
mark(node) =  
  if not node.marked:  
    node.marked = True  
    for c in node.children:  
      mark(c)
```



Performance


Reading



Mark-and-sweep collection

Algorithms

```
Mark  
mark(node) =  
  if not node.marked:  
    node.marked = True  
    for c in node.children:  
      mark(c)
```




Performance

Reading

Mark-and-sweep collection

Algorithms

```
Mark  
mark(node) =  
  if not node.marked:  
    node.marked = True  
    for c in node.children:  
      mark(c)
```



Performance

Reading



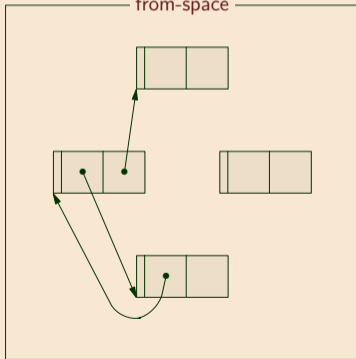
Copying collection

Algorithms


Collect

- copy live blocks to to-space (starting at the root)
- leave **forwarding addresses** in from-space
- switch **roles** of spaces

from-space



to-space



Performance


Reading

Copying collection

Algorithms

Collect

- copy live blocks to to-space (starting at the root)
- leave **forwarding addresses** in from-space
- switch roles** of spaces



Performance


Reading

Copying collection

Algorithms

Collect

- copy live blocks to to-space (starting at the root)
- leave forwarding addresses in from-space
- switch roles of spaces



Performance


Reading

Copying collection

Algorithms

Collect

- copy live blocks to to-space (starting at the root)
- leave **forwarding addresses** in from-space
- switch roles** of spaces



Performance

Reading

Copying collection

Algorithms




Performance

Reading

Collect

- copy live blocks to to-space (starting at the root)
- leave forwarding addresses in from-space
- switch roles of spaces




Copying collection

Algorithms

Collect

- copy live blocks to to-space (starting at the root)
- leave forwarding addresses in from-space
- switch roles of spaces



Performance

Reading


Copying collection

Algorithms

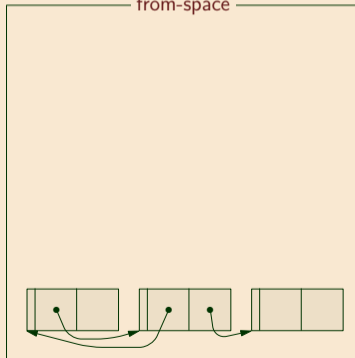
Collect

copy live blocks to to-space (starting at the root)
leave **forwarding addresses** in from-space
switch roles of spaces

to-space



from-space




Performance

Reading

The **reference count** tracks the number of pointers to each object.

An object's reference count is 1 when the object is created:



The count is incremented when a pointer newly references the object:



The count is decremented when a pointer no longer references the object:



The object is unreachable garbage when the reference count goes to 0:



Motivation: collector has imperfect information about object layout (e.g. because language is compiled to C)

Idea: use an approximation to guess whether a value represents a pointer, e.g.:

1. does the value point into the heap?
2. does it point to valid metadata?


Drawbacks

1. (**chance**) can incorrectly classify addresses as pointers
2. (**subterfuge**) can fail to identify disguised pointers

Generational collection

Algorithms

Copying collector for minor heap / mark-and-sweep for major heap



Performance

Reading


Generational collection

Algorithms

Copying collector for minor heap / mark-and-sweep for major heap

Performance


Reading



Generational collection

Algorithms

Copying collector for minor heap / mark-and-sweep for major heap




Performance

Reading

Generational collection

Algorithms

Copying collector for minor heap / mark-and-sweep for major heap




Performance

Reading

Generational collection

Algorithms

Copying collector for minor heap / mark-and-sweep for major heap




Performance

Reading

Generational collection

Algorithms

Copying collector for minor heap / mark-and-sweep for major heap



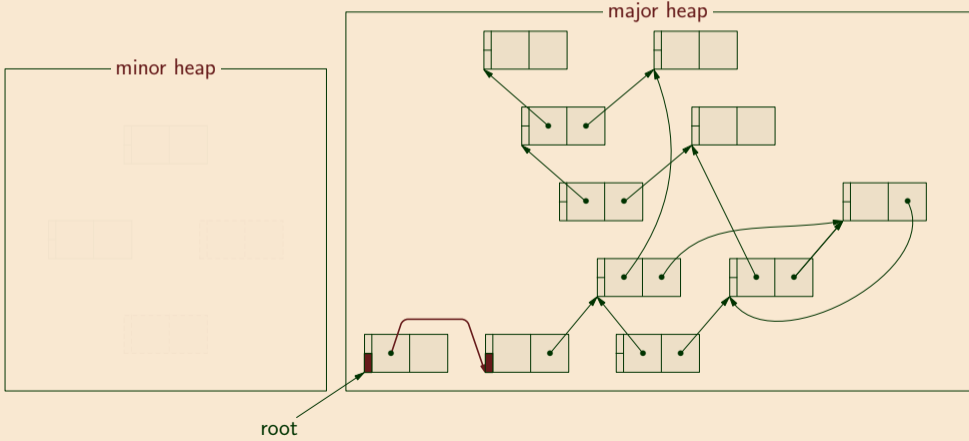
Performance

Reading

Generational collection

Algorithms

Copying collector for minor heap / mark-and-sweep for major heap



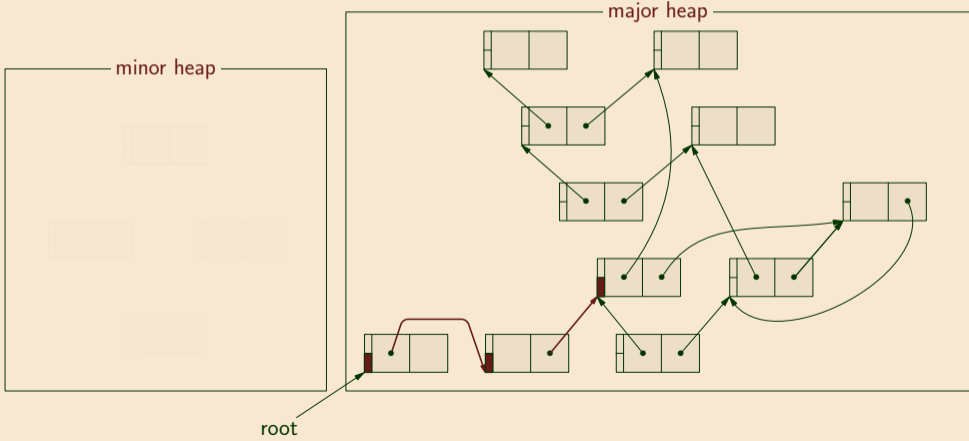
Performance

Reading

Generational collection

Algorithms

Copying collector for minor heap / mark-and-sweep for major heap




Performance

Reading

Generational collection

Algorithms

Copying collector for minor heap / mark-and-sweep for major heap




Performance

Reading

Generational collection

Algorithms

Copying collector for minor heap / mark-and-sweep for major heap




Performance

Reading

Generational collection

Algorithms

Copying collector for minor heap / mark-and-sweep for major heap




Performance

Reading

Generational collection

Algorithms

Copying collector for minor heap / mark-and-sweep for major heap




Performance

Reading

Generational collection

Algorithms

Copying collector for minor heap / mark-and-sweep for major heap




Performance

Reading

Generational collection

Algorithms

Copying collector for minor heap / mark-and-sweep for major heap



Performance

Reading

Performance

Throughput: mutator performance

Latency: pauses in mutator execution

Space overhead: e.g. due to mark bits, layout information

More (combination of program behaviour and collector design):

maximum heap size

allocation rate

collection frequency

mean object size

proportion of heap occupied by large objects

Example

Pause times alone provide little information.

A good **distribution** of pause times is needed for mutators to make progress.

Example

Compaction can slow collection but improve locality (& hence throughput)



Many mature systems combine several standard algorithms.

For example, Cedar (1985):

*“[...] provides both **a concurrent reference-counting collector** that runs in the background when needed, and **a pre-emptive conventional “trace-and-sweep” collector** that can be invoked explicitly by the user to reclaim circular data structures [...]*

*“Both collectors treat procedure-call activation records (called frames) **“conservatively”**; that is they assume that every ref-sized bit pattern found in a frame might be a ref”*

Reading



Uniprocessor Garbage Collection Techniques
Published in *ACM Computing Surveys*

Paul H. Wilson

Abstract


We cover basic automatic deletion methods, and various tasks associated and generalized collecting garbage from discrete blocks of heap-allocated memory. We discuss the relationship between garbage management, operating, debugging, and compiling. Throughout, we attempt to present a unified view based on what an internal design, addressing issues of low overhead, approximation, and consistency of collection, we can find, and a variety of implementation considerations. Such a view is expected to impact on practitioners.

Contents

1	Introduction	1	3	Incremental Tracing Collection	37
1.1	Motivation	1	3.1	Garbaron and Garbaron	37
1.2	The Tracing Abstraction	2	3.2	Tracing Methods	38
1.3	Basic Terminology	3	3.3	Incremental Approaches	38
1.4	Design of the Tracer	3	3.4	Write Barrier Algorithms	38
2	Basic Garbage Collection Techniques	4	3.5	Superspace/Overlapped Algorithms	39
2.1	Reference Counting	4	3.6	Incremental Hybrid Algorithms	39
2.2	The Problem with Spills	5	3.7	Hybrid of Tracing Algorithms	39
2.3	The Stop-and-Copy Problem	7	3.8	Incremental Copying	39
2.4	Block and Reference Counting	8	3.9	Basic Incremental Memory Management Algorithms to Transform	39
2.5	Scavenging in Real-time Systems	8	3.10	Characteristics of Basic Heap Tracing	39
2.6	Mark-and-Sweep Algorithms	10	3.11	Variation in the Real Barrier	39
2.7	Garbage Collection Optimization	10	3.12	Real-time Garbage Collection	39
3.1	A Space-Copying Garbage Collector: "Stop-and-Copy" Under Stress	10	3.13	Garbaron and Garbaron Hybrid	39
3.2	The Stop-and-Copy Garbage Collector	11	3.14	Garbaron and Garbaron in Memory Management	39
3.3	Scavenging Garbage Collection	12	3.15	Garbaron and Garbaron in Garbage Collection	39
3.4	Chaitin Automata and Tracing Techniques	12	3.16	Hybrid of Tracing and Reference Counting	39
3.5	Tracing with Write Barriers	12	3.17	Garbaron and Garbaron in Real-time Systems	39
3.6	Tracing with Write Barriers	12	3.18	Garbaron and Garbaron in Garbage Collection	39
3.7	Tracing with Write Barriers	12	3.19	Garbaron and Garbaron in Garbage Collection	39
3.8	Tracing with Write Barriers	12	3.20	Garbaron and Garbaron in Garbage Collection	39
3.9	Tracing with Write Barriers	12	3.21	Garbaron and Garbaron in Garbage Collection	39
3.10	Tracing with Write Barriers	12	3.22	Garbaron and Garbaron in Garbage Collection	39
3.11	Tracing with Write Barriers	12	3.23	Garbaron and Garbaron in Garbage Collection	39
3.12	Tracing with Write Barriers	12	3.24	Garbaron and Garbaron in Garbage Collection	39
3.13	Tracing with Write Barriers	12	3.25	Garbaron and Garbaron in Garbage Collection	39
3.14	Tracing with Write Barriers	12	3.26	Garbaron and Garbaron in Garbage Collection	39
3.15	Tracing with Write Barriers	12	3.27	Garbaron and Garbaron in Garbage Collection	39
3.16	Tracing with Write Barriers	12	3.28	Garbaron and Garbaron in Garbage Collection	39
3.17	Tracing with Write Barriers	12	3.29	Garbaron and Garbaron in Garbage Collection	39
3.18	Tracing with Write Barriers	12	3.30	Garbaron and Garbaron in Garbage Collection	39
3.19	Tracing with Write Barriers	12	3.31	Garbaron and Garbaron in Garbage Collection	39
3.20	Tracing with Write Barriers	12	3.32	Garbaron and Garbaron in Garbage Collection	39
3.21	Tracing with Write Barriers	12	3.33	Garbaron and Garbaron in Garbage Collection	39
3.22	Tracing with Write Barriers	12	3.34	Garbaron and Garbaron in Garbage Collection	39
3.23	Tracing with Write Barriers	12	3.35	Garbaron and Garbaron in Garbage Collection	39
3.24	Tracing with Write Barriers	12	3.36	Garbaron and Garbaron in Garbage Collection	39
3.25	Tracing with Write Barriers	12	3.37	Garbaron and Garbaron in Garbage Collection	39
3.26	Tracing with Write Barriers	12	3.38	Garbaron and Garbaron in Garbage Collection	39
3.27	Tracing with Write Barriers	12	3.39	Garbaron and Garbaron in Garbage Collection	39
3.28	Tracing with Write Barriers	12	3.40	Garbaron and Garbaron in Garbage Collection	39
3.29	Tracing with Write Barriers	12	3.41	Garbaron and Garbaron in Garbage Collection	39
3.30	Tracing with Write Barriers	12	3.42	Garbaron and Garbaron in Garbage Collection	39
3.31	Tracing with Write Barriers	12	3.43	Garbaron and Garbaron in Garbage Collection	39
3.32	Tracing with Write Barriers	12	3.44	Garbaron and Garbaron in Garbage Collection	39
3.33	Tracing with Write Barriers	12	3.45	Garbaron and Garbaron in Garbage Collection	39
3.34	Tracing with Write Barriers	12	3.46	Garbaron and Garbaron in Garbage Collection	39
3.35	Tracing with Write Barriers	12	3.47	Garbaron and Garbaron in Garbage Collection	39
3.36	Tracing with Write Barriers	12	3.48	Garbaron and Garbaron in Garbage Collection	39
3.37	Tracing with Write Barriers	12	3.49	Garbaron and Garbaron in Garbage Collection	39
3.38	Tracing with Write Barriers	12	3.50	Garbaron and Garbaron in Garbage Collection	39
3.39	Tracing with Write Barriers	12	3.51	Garbaron and Garbaron in Garbage Collection	39
3.40	Tracing with Write Barriers	12	3.52	Garbaron and Garbaron in Garbage Collection	39
3.41	Tracing with Write Barriers	12	3.53	Garbaron and Garbaron in Garbage Collection	39
3.42	Tracing with Write Barriers	12	3.54	Garbaron and Garbaron in Garbage Collection	39
3.43	Tracing with Write Barriers	12	3.55	Garbaron and Garbaron in Garbage Collection	39
3.44	Tracing with Write Barriers	12	3.56	Garbaron and Garbaron in Garbage Collection	39
3.45	Tracing with Write Barriers	12	3.57	Garbaron and Garbaron in Garbage Collection	39
3.46	Tracing with Write Barriers	12	3.58	Garbaron and Garbaron in Garbage Collection	39
3.47	Tracing with Write Barriers	12	3.59	Garbaron and Garbaron in Garbage Collection	39
3.48	Tracing with Write Barriers	12	3.60	Garbaron and Garbaron in Garbage Collection	39
3.49	Tracing with Write Barriers	12	3.61	Garbaron and Garbaron in Garbage Collection	39
3.50	Tracing with Write Barriers	12	3.62	Garbaron and Garbaron in Garbage Collection	39
3.51	Tracing with Write Barriers	12	3.63	Garbaron and Garbaron in Garbage Collection	39
3.52	Tracing with Write Barriers	12	3.64	Garbaron and Garbaron in Garbage Collection	39
3.53	Tracing with Write Barriers	12	3.65	Garbaron and Garbaron in Garbage Collection	39
3.54	Tracing with Write Barriers	12	3.66	Garbaron and Garbaron in Garbage Collection	39
3.55	Tracing with Write Barriers	12	3.67	Garbaron and Garbaron in Garbage Collection	39
3.56	Tracing with Write Barriers	12	3.68	Garbaron and Garbaron in Garbage Collection	39
3.57	Tracing with Write Barriers	12	3.69	Garbaron and Garbaron in Garbage Collection	39
3.58	Tracing with Write Barriers	12	3.70	Garbaron and Garbaron in Garbage Collection	39
3.59	Tracing with Write Barriers	12	3.71	Garbaron and Garbaron in Garbage Collection	39
3.60	Tracing with Write Barriers	12	3.72	Garbaron and Garbaron in Garbage Collection	39
3.61	Tracing with Write Barriers	12	3.73	Garbaron and Garbaron in Garbage Collection	39
3.62	Tracing with Write Barriers	12	3.74	Garbaron and Garbaron in Garbage Collection	39
3.63	Tracing with Write Barriers	12	3.75	Garbaron and Garbaron in Garbage Collection	39
3.64	Tracing with Write Barriers	12	3.76	Garbaron and Garbaron in Garbage Collection	39
3.65	Tracing with Write Barriers	12	3.77	Garbaron and Garbaron in Garbage Collection	39
3.66	Tracing with Write Barriers	12	3.78	Garbaron and Garbaron in Garbage Collection	39
3.67	Tracing with Write Barriers	12	3.79	Garbaron and Garbaron in Garbage Collection	39
3.68	Tracing with Write Barriers	12	3.80	Garbaron and Garbaron in Garbage Collection	39
3.69	Tracing with Write Barriers	12	3.81	Garbaron and Garbaron in Garbage Collection	39
3.70	Tracing with Write Barriers	12	3.82	Garbaron and Garbaron in Garbage Collection	39
3.71	Tracing with Write Barriers	12	3.83	Garbaron and Garbaron in Garbage Collection	39
3.72	Tracing with Write Barriers	12	3.84	Garbaron and Garbaron in Garbage Collection	39
3.73	Tracing with Write Barriers	12	3.85	Garbaron and Garbaron in Garbage Collection	39
3.74	Tracing with Write Barriers	12	3.86	Garbaron and Garbaron in Garbage Collection	39
3.75	Tracing with Write Barriers	12	3.87	Garbaron and Garbaron in Garbage Collection	39
3.76	Tracing with Write Barriers	12	3.88	Garbaron and Garbaron in Garbage Collection	39
3.77	Tracing with Write Barriers	12	3.89	Garbaron and Garbaron in Garbage Collection	39
3.78	Tracing with Write Barriers	12	3.90	Garbaron and Garbaron in Garbage Collection	39
3.79	Tracing with Write Barriers	12	3.91	Garbaron and Garbaron in Garbage Collection	39
3.80	Tracing with Write Barriers	12	3.92	Garbaron and Garbaron in Garbage Collection	39
3.81	Tracing with Write Barriers	12	3.93	Garbaron and Garbaron in Garbage Collection	39
3.82	Tracing with Write Barriers	12	3.94	Garbaron and Garbaron in Garbage Collection	39
3.83	Tracing with Write Barriers	12	3.95	Garbaron and Garbaron in Garbage Collection	39
3.84	Tracing with Write Barriers	12	3.96	Garbaron and Garbaron in Garbage Collection	39
3.85	Tracing with Write Barriers	12	3.97	Garbaron and Garbaron in Garbage Collection	39
3.86	Tracing with Write Barriers	12	3.98	Garbaron and Garbaron in Garbage Collection	39
3.87	Tracing with Write Barriers	12	3.99	Garbaron and Garbaron in Garbage Collection	39
3.88	Tracing with Write Barriers	12	4.00	Garbaron and Garbaron in Garbage Collection	39
3.89	Tracing with Write Barriers	12	4.01	Garbaron and Garbaron in Garbage Collection	39
3.90	Tracing with Write Barriers	12	4.02	Garbaron and Garbaron in Garbage Collection	39
3.91	Tracing with Write Barriers	12	4.03	Garbaron and Garbaron in Garbage Collection	39
3.92	Tracing with Write Barriers	12	4.04	Garbaron and Garbaron in Garbage Collection	39
3.93	Tracing with Write Barriers	12	4.05	Garbaron and Garbaron in Garbage Collection	39
3.94	Tracing with Write Barriers	12	4.06	Garbaron and Garbaron in Garbage Collection	39
3.95	Tracing with Write Barriers	12	4.07	Garbaron and Garbaron in Garbage Collection	39
3.96	Tracing with Write Barriers	12	4.08	Garbaron and Garbaron in Garbage Collection	39
3.97	Tracing with Write Barriers	12	4.09	Garbaron and Garbaron in Garbage Collection	39
3.98	Tracing with Write Barriers	12	4.10	Garbaron and Garbaron in Garbage Collection	39
3.99	Tracing with Write Barriers	12	4.11	Garbaron and Garbaron in Garbage Collection	39
4.00	Tracing with Write Barriers	12	4.12	Garbaron and Garbaron in Garbage Collection	39



Algorithms



A Real-time Garbage Collector with Low Overhead and Consistent Utilization

David F. Bacon
dfb@watson.ibm.com

Perry Cheng
perychc@us.ibm.com

V.T. Rajan
vt@us.ibm.com

IBM T.J. Watson Research Center
P.O. Box 704
Yorktown Heights, NY 10598

ABSTRACT

Now that the use of garbage collection in languages like Java is becoming widely accepted due to the safety and software engineering benefits it provides, there is significant interest in applying garbage collection to hard real-time systems. Past approaches have generally suffered from one of two major flaws: either they were not provably real-time, or they imposed large space overheads to meet the real-time bounds. We present a mostly non-moving, dynamically defragmenting collector that combines both of these features. By avoiding copying in most cases, space requirements are kept low, and by fully incrementalizing the collector we are able to meet real-time bounds. We implemented our algorithm in the RiscV JVM and show that at real-time resolution we are able to obtain mutator utilization rates of 45% with only 1.6–2.5 times the actual space required by the application, a factor of 4 improvement in utilization over the best previously published results. Defragmentation causes no more than 4% of the traced data to be copied.

General Terms
Algorithms, Languages, Measurement, Performance

Categories and Subject Descriptors
C.1 [Special-Purpose and Application-Based Systems]: Real-time and embedded systems; D.3.2 [Programming Languages]: Java; D.3.4 [Programming Languages]: Processors—Memory management (garbage collection)

Keywords
Real barrier, defragmentation, real-time scheduling, utilization

1. INTRODUCTION

Garbage collection languages like Java are making significant inroads into domains with hard real-time constraints, such as aerospace command-and-control systems. However, the engineering and product life-cycle advantages consequent from the simplicity of

programming with garbage collection remain unavailable for use in the core functionality of such systems, where hard real-time constraints must be met. As a result, real-time programming requires the use of multiple languages, or at least (in the case of the Real-Time Specification for Java [9]) two programming models within the same language. Therefore, there is a pressing practical need for a system that can provide real-time guarantees for Java without imposing major penalties in space or time.

We present a design for a real-time garbage collector for Java, an analysis of its real-time properties, and implementation results that show that we are able to run applications with high mutator utilization and low variance in pause times.

The target is supervisor embedded systems. The collector is therefore concurrent, but not parallel. This choice both simplifies and simplifies the design: the design is complicated by the fact that the collector must be interleaved with the mutator, instead of being able to run on a separate processor; the design is simplified since the programming model is sequentially consistent.

Previous incremental collectors either attempt to avoid overhead and complexity by using a non-copying approach (and are therefore subject to potentially unbounded fragmentation), or attempt to prevent fragmentation by performing concurrent copying (and therefore require a minimum of a factor of two overhead in space, as well as requiring barrier-on-read and/or writes, which are costly and tend to make response time unpredictable).

Our collector is unique in that it accepts an under-optimized portion of the design space for real-time incremental collectors: it is a mostly non-copying hybrid. As long as space is available, it acts like a non-copying collector, with the consequent advantages. When space becomes scarce, it performs defragmentation with fast-copying of objects. We show experimentally that such a design is able to achieve low space and time overhead, and high and consistent mutator CPU utilization.

In order to achieve high performance with a copying collector, we have developed optimization techniques for the Brooks-style real barrier [10] using an “outer invariant” that keeps real barrier overhead to 4%, an order of magnitude faster than previous software real barriers.

Our collector can use either time- or work-based scheduling. Most previous work on real-time garbage collection, starting with Euler’s algorithm [5], has used work-based scheduling. We show both analytically and experimentally that time-based scheduling is superior, particularly at the short intervals that are typically of interest in real-time systems. Work-based algorithms may achieve short individual pause times, but are unable to achieve consistent utilization.

The paper is organized as follows: Section 2 describes previ-

“[...] there is significant interest in applying garbage collection to hard real-time systems.”

“Past approaches have generally suffered from one of two major flaws: either they were not provably real-time, or they imposed large space overheads to meet the real-time bounds.”

“We [...] show that at real-time resolution we are able to obtain mutator utilization rates of 45% with only 1.6–2.5 times the actual space required by the application”

Performance

Reading

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear the notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
POPL'03, January 13–14, 2003, New Orleans, Louisiana, USA.
Copyright © 2003 ACM 1-58113-630-9/03/0001 \$5.00



Algorithms

Concurrent GCs and Modern Java Workloads: A Cache Perspective

Maria Carpen-Amarie
Hawari Zurich Research Center
maria.carpen.amarie@hazwei.com

Georgios Vavoulas
Hawari Zurich Research Center
georgios.vavoulas@hazwei.com

Konstantinos Tsvetkov
Hawari Zurich Research Center
konstantinos.tsvetkov@hazwei.com

Boris Gost
University of Edinburgh
Hawari Zurich Research Center
boris.gost@ed.ac.uk

Rene Mueller
Hawari Zurich Research Center
rene.mueller@hazwei.com

Abstract

The garbage collector (GC) is a crucial component of language runtimes, offering correctness guarantees and high productivity in exchange for a run-time overhead. Concurrent collection can alleviate application threads (mutators) and share CPU resources. A likely point of contention between mutators and GC threads and, consequently, a potential overhead source is the shared last-level cache (LLC).

This work builds on the hypothesis that the cache pollution caused by concurrent GCs hurts application performance. We validate this hypothesis with a cache-sensitive Java micro-benchmark. We find that concurrent GC activity may slow down the application by up to 3 \times and increase the LLC misses by 3 orders of magnitude. However, when we extend our analysis to a suite of benchmarks representative for today's server workloads (Rearmance), we find that only 5 out of 23 benchmarks show a statistically significant correlation between GC-induced cache pollution and performance. Even for these, the performance overhead of GCs does not exceed 10%. Based on further analysis, we conclude that the lower impact of the GC on the performance of Rearmance benchmarks is due to their lack of sensitivity to LLC capacity.

CCS Concepts: Software and its engineering \rightarrow Runtime environments; Garbage collection; Computer systems organization \rightarrow Multicore architectures.

Keywords: JVM, garbage collection, ZGC, cache pollution

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. *ISMM '23*, June 18, 2023, Orlando, FL, USA. © 2023 Copyright held by the owner/authors. Publication rights licensed to ACM.

ACM ISBN 978-1-609-59732-9/23/06...\$15.00
<https://doi.org/10.1145/3591195.3591249>

ACM Reference Format:

Maria Carpen-Amarie, Georgios Vavoulas, Konstantinos Tsvetkov, Boris Gost, and Rene Mueller. 2023. Concurrent GCs and Modern Java Workloads: A Cache Perspective. In *Proceedings of the 2023 ACM SIGPLAN International Symposium on Memory Management (ISMM '23)*, June 18, 2023, Orlando, FL, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3591195.3591249>

1 Introduction

Automatic memory management, also known as garbage collection (GC), is a technique that provides memory access safety and reliability while significantly reducing developers' load. These aspects make garbage collection an essential component of managed runtime environments (e.g., Java, C#, JavaScript), which are intensively used by web services (e.g., Twitter), web browsers, and mobile platforms (e.g., Android). For these reasons, automatic memory management continues to be a hot topic today, even after more than half a century of active research towards its optimization.

However, the benefits offered by the GC do not come for free. Prior work [1, 4, 13, 15, 17, 30] has shown that application performance is profoundly contingent on the effectiveness of the garbage collector. Historically, GCs harmed application performance due to unacceptably long stop-the-world (STW) pauses. To reduce the performance overhead and responsiveness issues created by STW events, significant effort was put into re-designing and implementing more efficient GCs [7, 12, 16, 22, 26, 28, 29]. Therefore, most modern GCs [10, 12] have increasingly shorter pauses, while performing most of the work concurrently with the application threads (e.g., ZGC [18]). However, concurrency comes at a price as well [5], as the application needs to share resources with the GC (e.g., cache capacity, bandwidth, CPU time) and even sometimes help with the collection itself. To the best of our knowledge, there is no work on concurrent GCs that quantifies these overhead components individually. Such information would reveal new weaknesses and opportunities and facilitate well-targeted performance improvements.

The rest of the paper we use GC to refer to the general of garbage collection as well as the garbage collector itself.

“This work builds on the hypothesis that the cache pollution caused by concurrent GCs hurts application performance.”

“We find that concurrent GC activity may slow down the application by up to 3 \times and increase the LLC misses by 3 orders of magnitude.”

“However, [...] we find that only 5 out of 23 benchmarks show a statistically significant correlation between GC-induced cache pollution and performance.”

Reading



Algorithms

Low-Latency, High-Throughput Garbage Collection

Wenyu Zhao
School of Computing
Australian National University
Australia
wenyu.zhao@anu.edu.au

Stephen M. Blackburn
School of Computing
Australian National University
Australia
stev.blackburn@anu.edu.au

Kathryn S. McKinley
Google
United States
kmcinley@google.com



Abstract

To achieve short pauses, state-of-the-art concurrent copying collectors such as C4, Shenandoah, and ZGC use substantially more CPU cycles and memory than stop-the-world collectors. They suffer from design limitations: (i) concurrent copying with inherently expensive read and write barriers, (ii) scalability limitations due to tracing, and (iii) immediacy limitations for mature objects that impose memory overheads.

This paper takes a different approach to optimizing responsiveness and throughput. It uses the insight that regular, brief stop-the-world collections deliver sufficient responsiveness at greater efficiency than concurrent evacuation. It introduces LXR, where stop-the-world collections use reference counting (RC) and judicious copying. RC delivers scalability and immediacy, promptly reclaiming young and mature objects. RC, in a hierarchical frame heap structure, reclaimsmost memory without any copying. Occasional concurrent tracing identifies cyclic garbage. LXR introduces: (i) RC re-memorized sets for judicious copying of mature objects, (ii) a novel low-overhead write barrier that combines condensing reference counting, concurrent tracing, and rememorized set maintenance, (iii) object relocation while performing a concurrent trace, (iv) lazy processing of decrements, and (v) novel survival rate triggers that modulate pause durations.

LXR combines efficient responsiveness and throughput, improving over production collectors. On the widely-used Lomene search engine in a tight heap, LXR delivers 7.8x better throughput and 10x better 99.99% tail latency than Shenandoah. On 17 diverse mature workloads in a moderate heap, LXR outperforms OpenJDK default G1 on throughput by 4% and Shenandoah by 65%.

CCS Concepts: · Software and its engineering → Garbage collection.

Keywords: Garbage collection, Reference counting



This work is licensed under a Creative Commons Attribution 4.0 International License.

PLDI '22, June 19–17, 2022, San Diego, CA, USA
© 2022 Copyright held by the author(s).
ACM ISBN 978-1-60959-522-6.
<https://doi.org/10.1145/3529919.3523189>

“To achieve short pauses, state-of-the-art concurrent copying collectors such as C4, Shenandoah, and ZGC use substantially more CPU cycles and memory than simpler collectors.”

“This paper [...] uses the insight that regular, brief stop-the-world collections deliver sufficient responsiveness at greater efficiency than concurrent evacuation.”

“[...] LXR delivers 7.8× better throughput and 10× better 99.99% tail latency than Shenandoah.”

Performance

Reading

