



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
steve.blackburn@anu.edu.au

Kathryn S. McKinley

Google
United States
ksmckinley@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 simpler 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 Immix heap structure, reclaims most memory without any copying. Occasional concurrent tracing identifies cyclic garbage. LXR introduces: i) RC remembered sets for judicious copying of mature objects; ii) a novel low-overhead write barrier that combines coalescing reference counting, concurrent tracing, and remembered set maintenance; iii) object reclamation while performing a concurrent trace; iv) lazy processing of decrements; and v) novel survival rate triggers that modulate pause durations.

LXR combines excellent responsiveness and throughput, improving over production collectors. On the widely-used Lucene search engine in a tight heap, LXR delivers 7.8× better throughput and 10× better 99.99% tail latency than Shenandoah. On 17 diverse modern workloads in a moderate heap, LXR outperforms OpenJDK's default G1 on throughput by 4% and Shenandoah by 43%.

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 13–17, 2022, San Diego, CA, USA

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9265-5/22/06.

<https://doi.org/10.1145/3519939.3523440>

ACM Reference Format:

Wenyu Zhao, Stephen M. Blackburn, and Kathryn S. McKinley. 2022. Low-Latency, High-Throughput Garbage Collection. In *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation (PLDI '22)*, June 13–17, 2022, San Diego, CA, USA. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3519939.3523440>

1 Introduction

Modern concurrent garbage collectors are surprisingly expensive. Growth in memory footprints and latency-sensitive applications led vendors to focus on low pause time collectors, such as C4 [39], Shenandoah [20], and ZGC [27]. While they achieve low pause times, they do so at memory and CPU costs. Furthermore, we show that their low pause times do not always translate into low latency for latency-sensitive applications. This paper identifies why they are expensive and proposes a very different approach. We introduce LXR (Latency-critical ImmiX with Reference counting), implement it in OpenJDK, and compare it against these widely-used collectors on diverse contemporary workloads.

The early Garbage First (G1) collector is a copying collector with concurrent tracing [17]. Each collection chooses a set of profitable fixed size regions to evacuate. It is optimized for throughput and low pause times. C4, Shenandoah, and ZGC built on the G1 foundation, seeking to further reduce pause times, believing lower pause times would translate to improved application latency. These collectors all rely on i) concurrent tracing to identify live objects, ii) evacuation alone to reclaim space, and iii) expensive read and/or write barriers. These choices have fundamental implications on i) application (mutator) performance, ii) timeliness of reclamation, and iii) collector performance. They reclaim memory only with concurrent copying, an intrinsically expensive approach that relocates every object using expensive barriers to maintain coherence of heap references [34].

Table 1 shows the tradeoff Shenandoah makes to achieve low pause times and that those low pause times *do not* translate to low latency on the widely used Lucene enterprise search engine (lusearch). It compares G1—the OpenJDK default, optimized for throughput; Shenandoah—optimized for latency; and LXR—optimized for both, using a tight heap 1.3× the minimum required by G1 on our Zen 3. The workload is challenging because it is highly parallel and has a very high allocation rate. We report *throughput* using queries per second (QPS) and total time, and *latency* using query

Table 1. Throughput, latency, and GC pauses on lusearch at a 1.3× heap. Short GC pauses *do not* assure low latency.

Algorithm	Throughput		Query Latency (ms)				GC Pauses (ms)			
	QPS	Time (s)	50%	99%	99.9%	99.99%	50%	99%	99.9%	99.99%
G1	112 K	4.7	0.1	12.0	14.6	15.9	0.4	0.9	1.1	1.2
Shenandoah	15 K	34.5	0.3	78.0	116.1	127.8	0.1	0.3	2.2	3.0
LXR	119 K	4.4	0.1	3.0	8.0	13.1	0.9	1.4	1.8	2.9
Shenandoah^{10×}	170 K	3.1	0.1	10.8	14.4	16.6	0.2	0.6	0.8	0.9

time latency percentiles. Shenandoah has low 50% and 99% pause times, but G1 is lower for the 99.9% and 99.99% pauses. However, low pauses *do not* translate into low application latency, with G1 improving over Shenandoah. Furthermore, Shenandoah takes a 7× hit on throughput compared to G1. Shenandoah cannot keep up with lusearch’s 9.5 GB/s allocation rate on this 16-core, 2-way SMT machine. We also configured Shenandoah to run in a huge heap, 10× the minimum. Given this substantial memory headroom, Shenandoah can deliver good throughput and latency. LXR delivers better throughput and latency than G1 and Shenandoah without requiring extra headroom. Although LXR has slightly longer pauses than G1 and Shenandoah, it improves query latency over both, while also outperforming them on throughput.

G1, C4, Shenandoah, and ZGC all **use tracing** and must perform a trace of all live objects to reclaim any memory. G1 and C4 mitigate this delay with generational collection, reducing the scope of the trace. (Shenandoah and ZGC have generational variants under development.) Tracing is unscalable in the limit. Its worst case is a live singly-linked list which defeats parallel tracing and copying [5]. LXR’s tracing for cyclic garbage detection also experiences this worst case. The scalability of reference counting is limited only when a large list dies. LXR achieves immediacy and scalability of collection with high performance parallel reference counting as its primary mechanism for garbage identification.

C4, Shenandoah and ZGC **use concurrent evacuation** to optimize pause times, adding high overheads (Table 1). They reclaim space only when a region is empty: when all objects eventually die or they evacuate live objects. Evacuation perturbs the memory hierarchy, using caches and DRAM bandwidth. Worse, concurrent evacuation requires expensive barriers to prevent mutator and collector races, which is intrinsically more expensive than stop-the-world evacuation [34]. LXR avoids these overheads with the Immix heap structure [11] and reference counting with remembered sets, reclaiming most memory without copying, only judicious stop-the-world copying to combat fragmentation.

These collectors **depend on expensive barriers**, limiting best-case application performance. C4, Shenandoah, and ZGC use a load value barrier (LVB) [39], which filters every reference load. LXR does not require a read barrier. It uses a novel low-overhead write barrier for i) concurrent tracing to reclaim cyclic garbage [44], for ii) building remembered sets,

and for iii) high performance coalescing reference counting to reclaim both young and old objects promptly [26, 37].

Design. LXR’s design premise is that regular, brief stop-the-world collections will yield sufficient responsiveness and far greater efficiency than concurrent evacuation. LXR builds on Reference Counting Immix [37] (RCImmix), taking a novel approach to low-latency performance by limiting concurrency and copying. It copies only during stop-the-world pauses. LXR employs parallelism for scalability in every collection phase. It exploits existing optimizations, such as implicitly dead young objects [37]. The Immix hierarchical heap divides *blocks* into *lines* and tracks their liveness. The application allocates using a fast bump pointer into blocks. It recycles free lines in a block, when all the objects are on the line are dead.

LXR uses a field-logging write barrier [7] with just 1.6% mutator overhead to perform coalescing reference counting, concurrent tracing, and maintain remembered sets. At each pause, LXR performs reference counting (RC) increments and reclaims free blocks and lines with dead young objects that never receive an increment [37]. It judiciously copies young surviving objects to defragment blocks [37]. LXR triggers occasional concurrent tracing using Yuasa’s Snapshot At The Beginning (SATB) tracing algorithm [44] to collect old dead cycles and objects with stuck reference counts. SATB tracing is elegantly implemented with the same write barrier used for coalescing reference counting [2]. We show how to soundly delete objects while an SATB trace is in progress, such that the trace may span multiple RC epochs. At the beginning of an SATB trace, LXR identifies candidate *evacuation sets* of blocks with high fragmentation. The trace initializes each remembered set and the write barrier keeps them up to date. After the trace completes, LXR reclaims free blocks and lines at the next RC pause and evacuates any still fragmented blocks in the evacuation sets.

LXR modulates pause times using three novel approaches. i) Collection triggers use survival rate prediction to control the expected work LXR performs in each RC epoch to reclaim dead young objects and to determine when SATB traces are likely to be profitable. ii) Lazy concurrent decrements reclaim old dead objects. iii) Incremental copying of one or more evacuation sets defragments blocks with old objects.

Techniques novel to LXR are: i) SATB tracing spanning multiple RC pauses, ii) combining RC and remembered sets,

iii) a single, simple write barrier that combines coalescing reference counting, concurrent tracing, and remembered sets for copying with RC, iv) survival rate collection triggers, and v) lazy concurrent decrements.

Implementation. We implement LXR in MMTk [8, 9, 29] on OpenJDK 11 and compare against G1, Shenandoah, and ZGC on a range of modern workloads. We demonstrate that this fresh approach to latency-sensitive collection achieves latency results that improve over these production collectors without additional hardware requirements. LXR delivers lower latency than Shenandoah and ZGC on each of the four latency-critical benchmarks we evaluate, and does so while outperforming G1 on throughput, while Shenandoah suffers 77% and 37% slowdowns on 1.3× and 2× heaps respectively. We are optimistic that this new design approach will lead to more efficient latency-sensitive collectors and a new class of low-overhead throughput-oriented collectors.

2 Background and Related Work

This paper exposes problems with state-of-the-art concurrent collector design and develops LXR, a novel collector design to solve those problems. We begin with the background that underpins LXR and prior collectors.

2.1 Reference Counting

Reference counting (RC) tracks the number of references to each object, reclaiming objects with no references. Traditional RC performs increments for every new reference to an object and decrements for every destroyed reference [15]. Because pointers change frequently, this approach is too inefficient for performance-critical applications. Temporal coarsening almost entirely removes this overhead. The underlying insight is that for any period $t_n \rightarrow t_{n+1}$, it is sufficient to identify all modified pointer fields, p , and apply a decrement to the referent of p at t_n , and an increment to the referent of p at t_{n+1} . This insight underpins both *deferral* for highly mutated stack variables [18], and *coalescing* for heap variables [26]. LXR applies

deferral by capturing all root pointers during brief mutator stop-the-world pauses (epochs) at t_n and t_{n+1} [3] and *coalescing* by capturing the to-be-overwritten referent the first time a heap pointer is written in the period $t_n \rightarrow t_{n+1}$, and the location of the pointer to establish its referent at t_{n+1} , ignoring intermediate referents [26].

Every overwritten referent in the before image (t_n) receives a decrement. Every fresh referent in the after image (t_{n+1}) receives an increment. For root pointers, LXR implements deferral by applying an increment to all root-reachable objects at t_n and remembering them in a buffer so that it can apply a matching decrement at t_{n+1} [3]. Prior to coalescing [26], applying RC to heap objects required expensive synchronization and was considered an impediment to adoption in

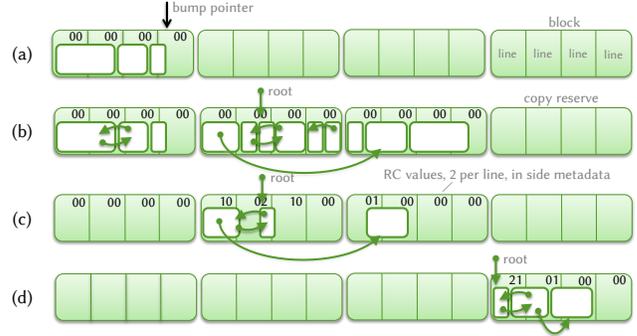


Figure 1. RC in an Immix heap with 2 bit RC values stored in line metadata: 0 is free, 1 is 1 RC to the object starting here, 2 is 2 RCs, and 3 is 3 or more RCs, simplified to a maximum of two objects per line, ignoring alignment. (a) Immix heap allocation. (b) Object graph before collection. (c) After collection. (d) The implicitly dead optimization evacuates young objects, creating 3 free blocks instead of 2.

parallel systems. Coalescing significantly reduces the synchronization work of RC systems. Coalescing with buffering also opened up optimizations such as parallel processing of the increment and decrements. Further optimizations and using a hierarchical mark-region (Immix) heap have delivered reference counting collectors that can outperform the fastest tracing collectors [2, 6, 36, 37].

LXR uses the extremely effective *implicitly dead* optimization which applies temporal RC coarsening to young objects [36]. First, young objects allocated since t_n never generate decrements in the period t_n to t_{n+1} since the initial state of their pointer fields is null. Second, they cannot receive any decrements from the heap or stack since, by definition, there were no pointers to them at t_n . Only their pointer state at t_{n+1} is thus relevant and it will only create increments. Young objects that receive no increments at t_{n+1} are dead and most do die, following the weak generational hypothesis [28, 40]. It is therefore correct to i) ignore impossible decrements from young objects, and ii) generate increments from young objects only if and when the object receives its first increment at t_{n+1} . A coalescing write barrier thus correctly ignores mutations to young objects.

Since the first GC after allocation will establish *all* references to surviving young objects, they can be correctly copied [37]. Figure 1 illustrates these optimizations (see Section 2.6). Figure 1(a) and (b) show that the initial RC values are zero in the young object metadata. Figure 1(c) shows the RC values after increments for each root reachable young object. Note that the first block of only dead young objects is free (some in cycles), no RC was ever changed, and LXR can reclaim the block before processing decrements. Figure 1(d) shows combining copying live objects while establishing their RC values, efficiently produces more free blocks.

Reference counting is fundamentally a local property. The liveness of an object can usually be determined without a global operation such as a heap trace and increases the immediacy of reclamation. RC increment and decrement processing grows only as a function of the rate of graph mutations and are both easily parallelized in general. However, if the head of a singly-linked list dies, recursively decrementing the list defeats parallelism. LXR limits the impact of this case by performing decrements concurrently. In contrast, tracing parallelism is defeated by a singly-linked list every time it is traced. We observe this worst case degrading G1 and Shenandoah throughput compared to LXR in Section 5.2.

Reference counting is incomplete because it cannot identify cycles of garbage and previously was considered incompatible with copying, lacking global information for identifying and redirecting references to relocated objects. LXR reclaims dead cycles via a backup trace. LXR uses the implicitly dead optimization to copy young objects. LXR adds novel techniques to combine the immediacy of RC with the defragmentation and locality benefits of evacuation for mature objects. We introduce remembered sets for RC and use them to combine block evacuation with RC for the first time.

2.2 Read and Write Barriers

Read and write barrier code intercepts mutator reads and writes to enforce invariants and monitor object graph mutations [42]. Coalescing reference counting uses a write barrier that, upon the first update to a pointer field, enqueues the overwritten referent for a future decrement and remembers the address of the updated field so that at the next collection, the eventual referent held at that address can receive an increment [26]. Azatchi and Petrank [2] observed that this barrier may also implement SATB tracing [44], which we exploit in LXR. The barrier is easily implemented at the object or the field granularity [7]. Section 5.3 shows the field-logging barrier used by LXR adds 1.6% overhead. G1 [17] uses several write barriers that it enables and disables during different phases of the collection. C4 [39], Shenandoah [20], and ZGC [27] all use the loaded value barrier (LVB), which filters every reference load, clearing bits if necessary, and conditionally takes further action. Shenandoah has an additional SATB write barrier for concurrent marking. Because applications typically execute field loads an order of magnitude more frequently than field stores (e.g. $64.3/\mu\text{s}$ v $4.3/\mu\text{s}$), read barriers, such as the LVB, are on the order of five times more expensive than an object remembering barrier [42].

2.3 Concurrent Tracing

Concurrent tracing algorithms identify all live objects concurrently with mutators, which may also be modifying the object graph [19, 38, 44]. G1, Shenandoah, and LXR use Yuasa’s Snapshot At The Beginning (SATB) algorithm [17, 20, 44]. SATB exploits the observation that once an object is unreachable, it remains unreachable, and it is therefore

correct to perform a trace over a stale heap snapshot. SATB captures program roots at the start of the trace. It intercepts the first time any pointer is overwritten during the trace and traces the about to be destroyed reference, along with all objects created since the start of the trace. LXR combines the SATB trace with the implicitly dead optimization, such that the trace only considers objects that survive an RC collection, significantly reducing the scope of its trace.

2.4 Concurrent Copying

As well as ensuring liveness is correctly established despite a concurrently mutated object graph, concurrent copying collectors (C4, Shenandoah, and ZGC) incur additional complexity in order to maintain referential integrity dynamically, when mutators dereference pointers which refer to objects concurrently being moved. The literature contains many such algorithms, including Brooks’s algorithm [12], initially used by Shenandoah [20], and Baker’s algorithm [4], which is the basis for the loaded value barrier (LVB) used by C4 [39], ZGC [27], and recent versions of Shenandoah [20]. Another approach is page protection and double mapping to perform concurrent evacuation while capturing any stale references [1, 14, 23, 39]. Pizlo et al. evaluated a range of concurrent copying collectors and found that the best of them incurred a throughput and latency overhead of 20% [34].

LXR’s design is based on the observation that *concurrent copying is intrinsically expensive*. LXR instead minimizes copying and eliminates concurrent copying completely by copying only during mutator pauses.

2.5 Region-Based Evacuating Collectors

G1, C4, Shenandoah, and ZGC are region-based and strictly-copying [47]. They partition the heap into fixed size regions for independent collection (e.g., G1 uses 1–32 MB depending on heap size). Prior to each GC cycle, they use heuristics such as age and fragmentation to select regions to collect, evacuating any live objects as they trace to another region. Concurrent tracing identifies garbage and fixes pointers to moved objects. Because tracing is their primary mechanism for reclamation, they are subject to the inherent scalability limitations suffered by all strictly-tracing collectors [5]. As the heap grows, the collector is not guaranteed to trace the entire heap in a timely manner, e.g., because of lack of parallelism in the heap graph or when parallel tracing threads compete with the mutator.

G1 uses a write barrier and remembered sets to identify pointers to evacuated objects, redirecting them as it moves objects during a mutator pause, which can cause long pauses [17]. C4, ZGC, and Shenandoah all target low (≤ 10 msec) pauses, necessitating concurrent evacuation. They use a loaded value barrier (LVB) on every pointer load to maintain referential integrity and evacuate concurrently with the mutator [39]. As Table 1 highlights, their designs

impose substantial costs on the mutator. C4 uses page protection to enforce referential integrity while lazily forwarding pointers to moved objects [14, 39]. These algorithms suffer because they i) *exclusively trace* to identify garbage, limiting scalability and timeliness, ii) *exclusively copy* to reclaim space, iii) *use expensive read and write barriers*, iv) *evacuate concurrently* to meet low pause time targets, and v) *have long collection cycle times* from concurrent whole-heap tracing, leading to relatively large memory requirements.

2.6 Immix Hierarchical Heap

LXR builds on RC-Immix [37] and its heap structure. Immix and RC-Immix combine the locality and simplicity of evacuating collectors without a strict copying requirement [11, 37]. Immix structures the heap in a hierarchy of 32 KB blocks, which are composed of 256 B lines. As shown in Figure 1(a), thread-local allocators use a bump pointer to allocate into a block. Objects may span lines but not blocks. Allocators reuse partially free blocks by skipping over live lines and allocating into free lines. For example the last two lines in the third block in Figure 1(c) are candidates for allocation after a collection. Allocators acquire a new free or partially-free block from a global pool when they exhaust their current block. Collectors may trace [11] or reference count [37] the Immix heap, identifying completely free lines and blocks.

The RC-Immix collector tracks the number of live objects on a line and a block. LXR uses a different approach, storing RC bits for each object (see Section 3.2.1). Figure 1(c) shows these RCs, abstracted for illustrative purposes to two per line, at the end of a reference counting epoch, when no object is evacuated. LXR reclaims old objects using reference counting without moving the live objects. Immix collectors may opportunistically copy objects to reclaim highly fragmented blocks. Figure 1(d) shows how LXR uses the implicitly dead optimization to evacuate young objects, copying them into a free block, partially free blocks are preferred. Evacuation eliminates the fragmentation shown in Figure 1(c). LXR also targets blocks with old objects and very few live lines to evacuate during cycle tracing, producing completely free blocks at low cost, similar to Immix (see Section 3.3).

2.7 Orthogonal Optimizations

Many optimizations target reducing collection and pause times, such as escape analysis [13, 33], concurrent stack scanning [22, 24, 32, 35, 45], barrier elision [41], compressed pointers [25], and large object spaces. They are orthogonal to our approach and we do not evaluate them here.

3 The LXR Algorithm

A key design insight for LXR is that *regular brief STW pauses provide a highly-efficient, low-latency approach to garbage collection*. We next present LXR’s algorithms and optimizations for allocation, identification, and reclamation. Figure 2

overviews LXR’s timeline of stop-the-world (STW) and concurrent collection activities.

3.1 Allocation and Heap Structure

LXR uses the Immix hierarchical heap structure and a bump point allocator. It uses two global lock-free data structures to maintain lists of free and partially free blocks respectively. Thread local allocators obtain blocks from these lists. Following Immix, LXR uses the partially free blocks first to maximize the availability of free blocks for large allocations which are managed separately.

When allocating into a partially free block, LXR uses a *reference count table* (Section 3.2), to identify unavailable lines, skipping over them to find free lines. Because objects may straddle lines but not blocks, following Immix [11], LXR conservatively assumes that the first line after a used line is unavailable, side-stepping computations to track objects that straddle two lines. When LXR applies the first RC increment to an object, if the object is larger than a line, it writes a non-zero value into the reference count table entry for each trailing line except the last, to ensure trailing lines are not reused. The last line does not need an entry because the conservative treatment of straddling objects already accounts for it correctly, as shown in block 3 of Figure 1(c).

LXR employs the dynamic overflow optimization from Immix to fill lines in partially free blocks. If there are one or more free lines between the cursor and the limit, and the current object does not fit, i.e., it is bigger than a line, then we allocate that object in an initially completely free overflow block for medium objects. Dynamic overflow avoids wasting lines. In the common case, allocation is fast and the allocator provides good spatial locality for contemporaneously allocated objects. Objects larger than half a block in size (16 KB) are delegated to a large object allocator.

If the runtime requires that LXR zero memory before use, it zeros each free region immediately before allocating into it. Thus, LXR zeroes free blocks in bulk [43]. For a partially free block, it zeros contiguous free lines immediately prior to allocating into them. Most managed languages require memory to be zeroed. OpenJDK optionally provides compiler-injected per-allocation zeroing. In this case, LXR does not need to perform zeroing at allocation time.

3.2 Identification

LXR promptly identifies dead objects using reference counting with deferral [18], coalescing [26], and implicitly dead [36] optimizations, with a back-up concurrent trace [44] to collect dead cycles and dead objects with stuck reference counts.

LXR uses regular, brief stop-the-world (STW) pauses to coherently scan runtime roots including mutator threads and stacks. Compared to taking a more aggressive on-the-fly approach to almost entirely remove pauses [26], LXR instead uses STW pauses as opportunity for highly efficient, targeted evacuation (Section 3.3) and increment processing.

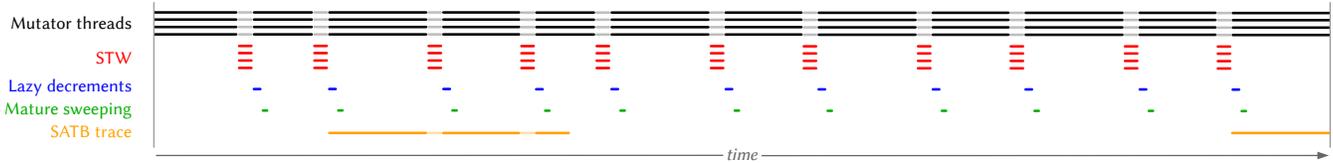


Figure 2. A simplified LXR timeline with four mutator threads. Four parallel STW GC threads increment object RCs, perform evacuation, and sweep young blocks. A concurrent LXR thread performs lazy decrements, old block sweeping, and the SATB trace. Unshown mutator activities: allocation and write barriers. We show 11 RC epochs, a complete and partial SATB trace.

3.2.1 RC Identification. Each RC epoch applies all increments first and then decrements. For each store to a reference field during mutator execution, LXR’s write barrier (Section 3.4) puts the overwritten reference in a *decrement buffer* and puts the address of the updated field in a *modified fields buffer*. Each RC epoch, LXR performs root scanning and performs an increment for each root-reachable object and buffers a corresponding decrement for the next epoch. LXR processes the modified fields buffer by applying an increment to the current referent of each modified field. After applying all increments, it processes the decrement buffer, which includes the root targets from the prior epoch and all overwritten references from the last mutator epoch, applying a decrement to each target object.

Whenever a decrement transitions a count from $1 \rightarrow 0$, LXR enqueues the dead object for recursive decrements, since all pointers in the dead object are about to be destroyed. LXR processes the recursive decrement queue by scanning each object and applying a decrement to each of its referents. Once the decrement queue and the recursive decrement queue are exhausted, the RC epoch is complete. In principle, recursive decrement processing could apply decrements to most objects, leading to arbitrarily long pauses, e.g., when the last reference to a large object graph is destroyed. We address this limitation, which is common to all RC algorithms, by restarting the mutators and performing concurrent lazy decrement processing (Section 3.2.1).

LXR is implemented in OpenJDK, which lacks sufficient object header bits for a reference count. LXR therefore stores reference counts in a *reference count table*, where each count is reachable from an object address via simple address arithmetic. Given a minimum object size of 16 bytes, we allocate an N -bit reference count for each 16 bytes of allocated memory. Shahriyar et al. [36] found $N_{rc} = 3$ was sufficient. We select two bits for simple address arithmetic and to limit memory overhead. With $N_{rc} = 2$, each 256 B line of allocated data consumes four bytes of metadata. The LXR allocator (Section 3.1) scans these densely packed reference counts. Free lines and blocks have all zero counts.

If a reference count for an object reaches 3, then we *stick* its count, and do not apply further decrements or increments to it. We rely on the concurrent trace to identify dead objects with stuck reference counts. In our analysis, 16 of 17

benchmarks have very few stuck counts (0 to 1% in Table 7). Batik is an outlier with 6.6%.

Optimization: Lazy Decrements. LXR processes RC decrements lazily, concurrent with the mutator. Because LXR reclaims a lot of memory, i.e., all the memory consumed by dead young objects [36], after applying RC increments and *without* processing any decrements (Section 2.1), lazy processing of decrements affects the promptness only of mature object reclamation. Lazy decrements are processed immediately after each brief STW pause with higher priority than SATB trace operations, thus the impact on prompt mature object reclamation is small. Lazy decrement processing has the further advantage of addressing the problem of recursive decrements taking arbitrarily long. If the next RC epoch starts and LXR still has decrements to process, it finishes them first. In our analysis, among 17 benchmarks only xalan (21%), lusearch (5%), and avrora (1%) have pauses that must process unfinished lazy decrements (Table 7).

Heuristic: RC Triggers. LXR triggers an RC pause if the heap is full or if increment or survival thresholds are reached. The *increment threshold* protects against long pauses due to increment processing. The *survival threshold* uses a prediction of young object survival rates to protect against long pauses due to recursive increments and copying of young objects. By using a predictor, the trigger targets controlling the *expected* case RC pause time, favoring throughput, in contrast to controlling worst case RC pause times with an allocation trigger. We use an exponential decay predictor, conservatively biasing it to higher survival rates. After each RC, if the survival rate is higher than predicted, then the new prediction uses 3/4 of the latest observation and 1/4 of the old prediction. Otherwise we reverse the weights, 1/4 (latest) : 3/4 (previous). Although survival rates vary greatly between workloads (see Table 3), this predictor is very effective.

3.2.2 SATB Identification. LXR periodically uses Yuasa’s SATB concurrent tracing algorithm [44] to collect dead objects RC cannot identify: objects in cycles and objects with stuck reference counts. SATB is an elegant match to coalescing RC since both depend on capturing the referents of overwritten pointer fields starting at some time t . RC enqueues the overwritten referents so that it can later decrement their reference counts, while SATB must include the

overwritten referents in its trace, since they comprise part of the initial heap snapshot at time t .

An SATB tracing epoch starts during a regular RC collection pause, seeded with the root set from the RC collection. The trace continues concurrently with the mutator, adding additional edges as they are provided by the write barrier. Section 3.4 describes how SATB piggybacks on the RC write barrier with no additional mutator overhead.

Optimization: SATB with Interruptions. Note that it is correct to add new referents to the SATB trace. We exploit this property to allow, as needed, a single SATB trace to span multiple RC epochs, t_n to t_{n+k} , giving LXR the freedom to perform RC collections as frequently as necessary, independent of the time taken to perform a full heap trace. SATB tracing time is governed by the size and shape of the reachable object graph and may not be amenable to parallelism [5]. Since the SATB trace must visit all objects that were reachable at the *start* of the trace, RC deleting objects during the SATB trace presents a challenge to correctness. However, deferring RC until the trace completes would be a barrier to timely reclamation.

We generalize SATB tracing to handle *interruptions* (RC collections) during the trace. The correctness of SATB depends on visiting all objects reachable at the time the SATB began. We note that we can achieve this goal by enforcing the invariant that RC may never delete an unmarked object while an SATB trace is underway. We implement this invariant with mark bits stored in side metadata and by immediately marking and scanning any mature object that RC determines is dead if the SATB has not already marked it. Note that RC will reclaim space promptly regardless of whether the SATB mark bit is set. However, once the object's mark bit is set, the SATB will not attempt to visit the object, avoiding the possibility of the SATB following a reference to a deleted object. These potential references come both from objects no longer reachable and references held in the SATB's mark stack. The collector clears these mark bits only after the SATB epoch is finished. Reclamation of unmarked objects happens during the first RC epoch that starts after the SATB finishes.

Optimization: Mature-Only SATB. Yuasa's SATB algorithm conservatively keeps new objects live. Using the implicitly dead optimization, we correctly apply SATB tracing only to mature objects that survive an RC collection. LXR implements this optimization by ignoring objects with a zero reference count when it performs the SATB trace. This optimization substantially reduces the SATB's working set and eliminates its prior conservative treatment of objects allocated during the SATB trace.

Heuristic: SATB triggers. SATB traces place a burden on the mutator, but are essential to defragmenting mature objects and freeing cyclic garbage and objects with stuck counts.

To maximize their effectiveness and limit their costs, we trigger SATB traces judiciously. LXR has an available free block trigger and a predicted heap *wastage* trigger for SATB. We define wastage as uncollected dead mature objects and fragmentation. If an RC yields fewer clean blocks than a prescribed threshold, the next pause will initiate an SATB trace. We also trigger an SATB if predicted wastage reaches a threshold percentage of the heap. Each RC epoch, we use a live block predictor to predict the number of live blocks that will be in the heap if we perform an SATB trace. This predictor is driven by live block observations after each SATB trace. Similar to the RC trigger, the predictor uses a 1:3 ratio conservatively biased exponential decay. We use the predictor to estimate wastage after each RC by taking the difference between the current live blocks and the predicted live blocks.

3.3 Reclamation

At its simplest, LXR is a non-moving reference counting collector. We augment this algorithm with reclamation performed as a result of the SATB trace and two optimizations that opportunistically copy objects to defragment the heap and improve allocation locality.

3.3.1 RC Reclamation. LXR reclaims objects with a reference count of zero after applying all increments. Because young objects start with zero reference counts and die quickly, fully and partially free blocks with young objects are common. LXR reclaims free lines and blocks by sweeping the blocks that contain young objects and examining their reference count table. If all counts are zero, then the entire block is free and LXR returns the block to the global free block list, where it is available for use by any part of the collector, including large object allocation. Otherwise, if lines are free in a block, it places the partially free block on the global partially free block list.

Optimization: Lazy Reclamation. After sweeping blocks containing young objects, LXR resumes the mutator threads and processes decrements concurrently. As soon as it finishes all the decrements, LXR selectively sweeps those blocks containing objects which received a decrement, returning any partially or completely free blocks to the global allocator.

3.3.2 SATB Reclamation. At the start of an RC epoch, LXR checks whether an SATB trace has completed. If a trace is complete, SATB-identified dead objects will be unmarked, but will have a non zero reference count. Note these objects cannot have outstanding RC increments since the SATB algorithm guarantees that they were dead at the start of the SATB, which was at least one epoch earlier. When LXR finishes its increments and decrements, the decrement sweep will include clearing the reference counts of all dead unmarked objects from the SATB trace and will reclaim objects dead due to decrements and dead due to SATB tracing.

Optimization: Young Evacuation. Shahriyar et al. [37] observe that since all surviving young objects will receive an increment from every referent during their first collection, the first collection may move them, safely redirecting (forwarding) each of the referents in the process. Furthermore, they show that the Immix heap structure allows this movement to be performed *opportunistically* on a per-object basis, according to space and time remaining. We apply this optimization to LXR and note that it can be slightly improved by observing that young objects can be trivially identified as any object that receives a $0 \rightarrow 1$ increment.

Heuristic: All Young Evacuation. In our implementation, we mark the blocks which only contain young objects, i.e., these blocks are empty when assigned to a thread-local allocator prior to this RC epoch. We then target all these blocks for young object evacuation. This optimization exploits the generational hypothesis. Since few objects survive, LXR will copy only a few objects, creating completely free blocks without much copying, and creating locality by compacting the survivors into the partially free blocks with other survivors from this collection and prior collections. It will copy these survivors into completely free blocks when the partially free blocks are exhausted or when it applies dynamic overflow for medium size objects [11]. Figure 1(d) shows this optimization when all the blocks contain only young objects. If there are no free or partially free blocks, it can stop copying young objects and increment their reference counts in place.

Optimization: Mature Evacuation. Shahriyar et al. note that copying and RC can be combined during a stop-the-world full heap trace for collecting cycles [37]. This opportunity is not available to LXR since it does not perform a stop-the-world full heap trace. We observe that the write barrier LXR uses for RC and SATB (Section 3.4) is sufficient to maintain remembered sets, allowing portions of the heap to be independently evacuated during RC collection pauses. Each LXR remembered set tracks all pointers into an *evacuation set*. An evacuation set is a set of blocks within a contiguous region of the heap that will be evacuated together. LXR uses either 4 MB regions, leading to many remembered sets, or the whole heap, yielding a single remembered set. Prior to each SATB, LXR identifies and marks *target blocks*, which will form the evacuation set(s). We use the reference count table as an upper bound on the number of live bytes in each block. For all blocks that have less than 50% occupancy, we sort them from the lowest to highest occupancy, and mark the N lowest occupancy blocks as *target blocks*.

The remembered sets are not maintained continuously, but only from the start of the SATB to the RC pause in which the set is evacuated. Each set is bootstrapped by piggybacking on the first SATB trace, which must traverse every pointer into the evacuation set. Each such pointer is appended to the region’s remembered set. The remembered set is kept up to date until the set is evacuated, via the write barrier,

```

1 void putfield(ref* field, ref value) {
2   if (isUnlogged(field)) /* non-atomic */
3     logField(field);
4   *field = value; /* store value */
5 }
6 void logField(ref* field) {
7   if (attemptToLog(field)) { /* synchronized */
8     ref old = *field; /* remember old */
9     markAsLogged(field);
10    decbuf.push(old); /* log old value */
11    modbuf.push(field); /* log address */
12  }
13 }

```

Figure 3. LXR’s field write barrier has 1.6% overhead.

which remembers each subsequently-created reference into the evacuation set.

If remembered sets need to be maintained for more than one RC epoch, entries can become stale, leading to a correctness concern. Each remembered set entry is a pointer to the location of an incoming reference. It is benign for the reference to be overwritten, however if the source object dies and the space is reused, the remset entry may point to a non-pointer. In that case, it could be possible for a non-pointer value to be treated as a pointer, and in the unlikely event that the value appeared to point to an object being evacuated, the collector would incorrectly update it. We avoid this problem by maintaining a reuse counter for each line which is reset at each SATB. Each remset entry is tagged with the reuse count for the source line. If at evacuation time the source line is newer than the remset entry, we discard the entry.

LXR evacuates each set by tracing through the blocks using the current collection roots and the remembered set as roots. When the evacuation encounters an object in the evacuation set, it copies the object into a destination in a free or partially free block, updates the incoming reference, and leaves a forwarding pointer. If the object has already been copied, the incoming pointer is updated using the forwarding pointer stored in the stale object. When the trace encounters a pointer outside the evacuation set, it ignores it, limiting the scope of the evacuation to within the evacuation set. At each STW pause, LXR evacuates one or more evacuation sets and may use a time budget.

3.4 Write Barrier

For each field the mutator overwrites during $t_n \rightarrow t_{n+1}$, LXR needs to know the field’s value at t_n and at t_{n+1} . It uses the former for SATB tracing and to generate decrements for destroyed references. It uses the latter to create remembered set entries and increments for new references. The barrier may operate at one of two granularities. It can remember *objects* containing fields that are overwritten or with slightly higher mutator overhead, but greater precision, it can remember just overwritten *fields* [7, 26, 36]. We implemented both, but our results use the field barrier shown in Figure 3.

LXR’s field-remembering barrier uses one unlogged bit per field. LXR uses side metadata to hold the unlogged bit used by both `isUnlogged()` and `logField()`. At each pointer store, the barrier checks the field’s unlogged bit. If the field is unlogged, the slow path adds the referent of the field to the decrement buffer and adds the address of the field to the *modified field buffer*. Note that only this slow path, that records the field the first time it is written, is synchronized. The `attemptToLog()` function will block if another thread is logging, returning only after the other thread has captured the to-be-overwritten value. Since objects are zeroed prior to allocation, they are naturally initialized with the unlogged bit/s cleared. The barrier therefore does not log changes to new objects, cleanly implementing the implicitly dead optimization (Section 2.1, [36]).

When LXR processes the modified fields buffer, it generates an increment for referents of each field and resets its unlogged bit. LXR uses entries in the `decbuf` for RC decrements and to create the SATB snapshot. During RC pauses, LXR dereferences entries in the `modbuf` to find final referents, which receive RC increments. The geometric mean overhead of this write barrier compared to no write barrier on Immix is 1.6%, with a worst case for h2o at 4.6% (Table 7).

3.5 Parallelism

LXR is designed for scalability. For mutator scalability, LXR uses lock free block allocators to issue clean and recycled blocks to thread-local allocators with minimal contention, even at very high allocation rates. For collector scalability, coalescing reference counting is naturally scalable. Reference counts are inherently local, so applying increments and decrements is embarrassingly parallel. We found partitioning very large reference arrays necessary for scalability of parallelized increment processing. We use work stealing within the collector to maximize load balancing. Only the concurrent SATB trace has significant potential for scalability issues [5]. Because it is a tertiary collection mechanism (following young object treatment and reference counting), this concern is not a major one and rarely encountered in practice. Most decrements are processed concurrently with limited impact on pause times, immediacy, or mutator performance.

4 Methodology

We implement LXR in the new MMTk [8, 9, 29] on OpenJDK, fork `jdk-11.0.11+6`.¹ We use G1, Shenandoah, and ZGC on the same fork. C4 is not available in OpenJDK. Because our implementation is missing class unloading, compressed pointers, and weak references (see Section 6), we disable them in all systems. The implementation we evaluate here is publicly available [46], and we expect it to become integrated and maintained as part of MMTk’s public release [29]. We

Table 2. We use three contemporary stock processors. We report results from the Zen 3 by default.

Name	Model	Cores	Clock	LLC	Memory
Zen 3	AMD 5950X	16/32	3.4 GHz	64 MB	64 GB DDR4 3200MHz
Zen 2	AMD 3900X	12/24	3.8 GHz	64 MB	64 GB DDR4 2133MHz
C’Lake	Core i9-9900K	8/16	3.6 GHz	16 MB	128 GB DDR4 3200MHz

use three hardware platforms described in Table 2. The Zen 3 uses a solid-state drive while the Zen 2 and C’Lake use hard disk drives. Unless otherwise stated, we report AMD Zen 3 results. We use a Intel Coffee Lake (C’Lake) to assess the robustness of our findings to microarchitecture. All machines run Ubuntu 18.04. The remainder of this section describes LXR configurations, benchmarks, and measurement methodologies.

LXR Configuration. Unless otherwise stated, we use the following configuration of LXR for performance measurements: a 4 MB lock-free global block allocation buffer (Section 3.5), a two bit reference count, a 128 MB survival threshold, no increment threshold (Section 3.2.1), a 5% mature wastage threshold (Section 3.2.2), and a single evacuation set (Section 3.3.2).

Benchmarks. We use 17 diverse up-to-date benchmarks from the Chopin development branch of the DaCapo benchmark suite [10, 21].² It contains four latency-sensitive workloads that report request latency percentiles and throughput (total time). We omit tradebeans and tradesoap because they frequently fail on this version of JDK 11 unless the C2 compiler is disabled. ZGC fails on many benchmarks at 2× their nominal heap size because it requires a substantial minimum heap. JDK 17 fixes this limitation.

Table 3 lists key benchmark characteristics. It presents the minimum G1 heap size, allocation characteristics, and object demographics. Allocation rates are relative to the mutator time of G1, excluding pause times. High allocation rates and high ratios of allocation to minimum heap sizes make most of these applications challenging workloads for collectors. Ten of 17 benchmarks allocate over a GB per second. The second-from-last column indicates the percentage of allocated bytes for objects larger than 16 KB. The last column shows the percentage of bytes that survive a fixed 32 MB nursery. Nine benchmarks are highly generational with 0-5% survival rates. High survival rate is typically coupled with low ratios of allocation to heap size, with the exception of `xalan`. In contrast, `Batik` has a 51% survival rate.

`Cassandra`, `h2`, `lusearch`, and `tomcat` are request-based, latency-sensitive, and widely-used workloads, appropriate for evaluating latency optimized collectors. `Lusearch`, `h2`, and `tomcat` have high allocation rates, ranging from 1–9 GB/s on the Zen 3. `Cassandra` runs the YCSB workloads [16] on the Apache Cassandra NoSQL database management system (v.

¹Commit 1007c9c1f46644d325bb4d4bd3a3d6dc718c713e, 2021/3/10.

²Commit b00bfa96b6db296bdc6f57a57e56a9a34b2d2d89, 2022/04/09

Table 3. Benchmark characteristics: minimum G1 heap size; total bytes allocated; ratio of allocation to minimum heap; allocation rate; mean object size; percentage of large object bytes to total bytes; and percentage of survivor bytes to total bytes for a 32 MB nursery.

Benchmark	Heap		Allocation		Obj Size	% Lrg. Srv.	
	MB	GB	/heap	MB/s		Lrg	Srv.
cassandra	263	5.6	22	596	50	0	4
h2	1 191	13.0	11	1 534	64	0	17
lusearch	53	31.2	603	9 520	97	1	1
tomcat	71	6.9	100	1 440	95	21	1
avro	7	0.2	28	46	45	0	5
batik	1 076	0.5	0	257	71	10	51
biojava	191	11.8	63	800	37	3	2
eclipse	534	8.3	16	595	100	29	17
fop	73	0.5	7	557	58	3	10
graphchi	255	11.9	48	1 117	134	3	4
h2o	3 689	11.8	3	3 065	168	23	14
python	325	5.2	16	1 038	60	4	0
luindex	41	2.2	54	335	288	75	3
pmd	637	7.0	11	3 952	46	2	14
sunflow	87	20.5	241	6 267	45	0	3
xalan	43	3.9	92	4 265	122	41	17
zxing	153	1.5	10	1 750	183	50	23

3.11.10). H2 performs a TCP-C-like workload on the Apache Derby database (v. 10.14.2.0). Lusearch performs half a million search queries on the Apache Lucene search framework (v. 8.8.2). It is highly parallel and has a small heap footprint. Tomcat runs the sample web application workload provided with the Apache Tomcat web server (v. 9.0.45).

Throughput Measures. We report throughput as total time. We invoke each collector on each benchmark 20 times, perform five iterations, and report the time for the fifth iteration on an otherwise empty machine. This methodology gives the compiler and runtime time to warm up. We calculate the geometric mean and 95% confidence intervals on 20 measurements. Our measurements interleave collectors to minimize bias due to systemic interference.

Latency Measures. The Chopin development version of DaCapo reports *simple* and *metered* latency measurements. Three significant factors impact request latency: i) the uninterrupted time to *compute* the request, ii) the time consumed by *interruptions* such as scheduling and garbage collection, and iii) the time consumed by *request queuing*. The simple measurement captures the first two, but not the third. DaCapo’s metered measurement captures all three by modeling request queuing, which is standard in cloud-based services. Requests arrive at servers at some *metered* rate. The server enqueues these requests if there are insufficient resources to process them immediately. When an interruption arises

(such as a garbage collection event), it delays both active requests and any enqueued requests. DaCapo models an unbounded queue and arbitrarily long delays, whereas some systems drop requests. We present DaCapo’s metered 50, 90, 99, 99.9 and 99.99 percentile tail latencies in tabular form and the full data in latency response curves (e.g., Figure 5). Because they are so sensitive to noise, we use 40 invocations for all of our latency measurements.

5 Evaluation

We first compare LXR to G1, Shenandoah, and ZGC with respect to request latency on the four latency-sensitive applications. We then present throughput and the sensitivity results. Because garbage collection is a time-space tradeoff, we explore sensitivity to heap size. Our results show that LXR combines low latency with high throughput. Compared to the low latency collectors (Shenandoah, ZGC) and G1, LXR matches or improves over their latency on all four benchmarks. On throughput, LXR improves over G1 by 4% and Shenandoah by 43% on our benchmarks.

5.1 Request Latency

We evaluate G1, LXR, Shenandoah, and ZGC on the four latency-sensitive workloads described in Section 4. Table 4 and Figure 5 report results for a tight 1.3× heap. This version of ZGC cannot run in small heaps so produced results only for cassandra and h2. Table 4 shows the median, 90, 99, 99.9 and 99.99 percentile request latencies and the 95% confidence intervals, which are expressed as a fraction: ± 0.500 means the 95% confidence interval extends 50% over the reported result in each direction. Figure 5(a) graphs the mean at all request percentiles with their 95% confidence intervals shaded.

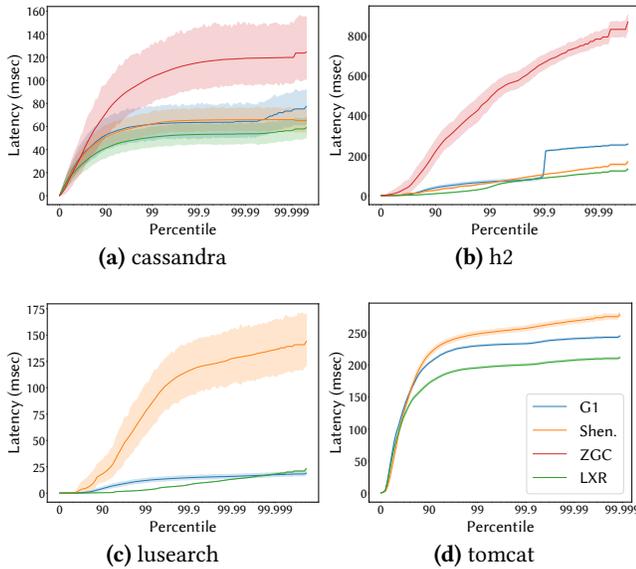
LXR performs better or as well as the other collectors at essentially all percentiles. Recall the three components of a latency measure (Section 4) and that variability in the computation of the requests may dominate. For cassandra, noise, regardless of the collector, reduces the differences between LXR, Shenandoah, and G1, as shown in Figure 5(a).

For h2, LXR outperforms both latency-targeting collectors at all percentiles. G1 performs similar to Shenandoah until the 99.9th percentile, when its latency roughly doubles, but it substantially outperforms ZGC at every percentile. The performance on h2 with a minimum heap size of 1.2 GB demonstrates that LXR performs well in large heaps running traditional workloads such as TPC-C, as well as performing well on the other workloads with smaller minimum heap sizes (Table 3).

For lusearch, LXR and G1 are similar, but Shenandoah performs poorly, as it currently lacks an effective mechanism for young objects and very high allocation rates. Lusearch has a small heap size, a very high allocation rate, and a very low survival rate in contrast to h2. LXR’s excellent performance on both h2 and lusearch shows its design at its

Table 4. Median and tail request latencies in a 1.3× heap on the Zen 3 with 95% confidence intervals. Best results in green.

Benchmark	G1				LXR				Shenandoah				ZGC			
	50	99	99.9	99.99	50	99	99.9	99.99	50	99	99.9	99.99	50	99	99.9	99.99
cassandra	24.2 ±0.332	62.0 ±0.244	63.6 ±0.237	64.5 ±0.749	19.9 ±0.342	50.8 ±0.180	53.1 ±0.172	53.7 ±0.755	23.2 ±0.286	62.4 ±0.161	65.5 ±0.156	65.9 ±0.749	21.7 ±0.571	103.1 ±0.253	115.5 ±0.235	119.133 ±0.922
h2	1.1 ±0.112	69.7 ±0.134	104.6 ±0.089	246.1 ±0.996	0.6 ±0.009	42.9 ±0.024	87.7 ±0.013	113.6 ±0.995	1.3 ±0.014	63.3 ±0.009	105.9 ±0.012	140.3 ±0.991	16.3 ±1.481	500.0 ±0.127	677.1 ±0.078	813.2 ±1.010
lusearch	0.1 ±0.458	12.0 ±0.170	14.6 ±0.142	15.9 ±0.997	0.1 ±0.003	3.0 ±0.020	8.0 ±0.010	13.1 ±0.996	0.3 ±1.673	78.0 ±0.313	116.1 ±0.226	127.8 ±1.001				
tomcat	89.3 ±0.040	230.1 ±0.007	233.4 ±0.007	241.1 ±0.644	77.8 ±0.045	195.7 ±0.009	200.4 ±0.009	207.6 ±0.642	68.2 ±0.088	248.7 ±0.017	257.5 ±0.016	268.7 ±0.768				

**Figure 5.** Request latency response curves on the Zen 3 with a 1.3× heap. Each y-axis shows latency in milliseconds and the x-axis shows the percentile of requests experiencing that latency. ‘90’ on the x-axis represents the latency observed by the 90th percentile of requests. Each solid line represents the average from 40 invocations and the shaded area shows its 95-percentile variance. LXR improves over the other collectors at nearly every point of every curve.

best, delivering excellent results on such different workloads. On the other hand, while G1 does well on lusearch, it does not do well in the tail on h2, and Shenandoah does well on h2, but struggles with lusearch’s high allocation rate.

For tomcat, LXR delivers large improvements at the 90th percentile and higher. The narrow confidence intervals and consistency in the result among the three collectors suggests that request latency is, like cassandra, dominated by the intrinsic computational cost of each of benchmark’s request more so than interruptions or request queuing. The better responsiveness of LXR is nonetheless visible in Figure 5(d).

These four applications exhibit diverse behaviors, such as high and low survival rates. These differences require the range of optimizations in LXR. Table 3 shows the high

Table 5. Geometric mean of 99.99% latency for cassandra, h2, lusearch, and tomcat and throughput for all benchmarks. Results relative to G1 for each collector on three heap sizes.

Heap	99.99% Latency/G1				Time/G1			
	G1	LXR	Shen	ZGC	G1	LXR	Shen	ZGC
1.3×	1.00	0.72	1.51	–	1.00	0.97	1.77	–
2×	1.00	0.92	2.54	–	1.00	0.96	1.37	–
6×	1.00	0.85	1.41	1.44	1.00	1.01	1.09	1.26

survival rate of h2 at 17%, whereas the others have low survival rates: cassandra at 4%, lusearch at 1%, and tomcat at 1%. LXR’s RC and implicitly dead optimization are extremely effective on the three generational workloads. Mature space optimizations are however critical for all benchmarks, even lusearch and tomcat, because their ratio of allocation to minimum heap size is high (Table 3).

Mature space RC is an effective optimization for these benchmarks. Table 7 shows that LXR’s mature object RC (Old) reclaims many mature objects. Mature object RC reclaims 60% of mature objects for cassandra, 48% for h2, 40% for lusearch, and 71% for tomcat. The SATB trace reclaims the other mature objects. We note that LXR’s SATB trigger successfully tunes itself to the effectiveness of the RC mature reclamation behavior (see Table 7 under GC Pauses with SATB%), running less frequently (14% of RC pauses) on tomcat versus 20-21% for the others, since LXR’s mature RC reclaims most mature objects for tomcat.

Heap Size Sensitivity. The above analysis uses a tight, 1.3× heap to highlight the time-space tradeoff inherent to garbage collection and how that affects application latency. Table 5 shows the heap size sensitivity of request latency and performance with a 2× heap and a generous 6× heap. It presents the geometric mean latency for the four latency benchmarks and time for all benchmarks. For latency, LXR outperforms G1 and Shenandoah at each size, and at 1.3× and 6× LXR is substantially better. On throughput, LXR is better than G1 in tighter heaps and substantially better than Shenandoah. These results illustrate that LXR successfully pairs low latency and high throughput on a range of heap sizes.

Table 6. Benchmark throughput with a 2× heap. We report G1 time in milliseconds and relative performance of LXR, Shenandoah, and ZGC, normalized to G1. Best results for each collector are shown in green, worst are shown in orange.

Benchmark	G1	LXR	Shen.	ZGC
cassandra	9 718	0.961	1.015	0.990
h2	8 902	1.101	1.246	1.764
lusearch	3 810	0.909	6.596	
tomcat	5 082	1.024	1.220	
avroa	4 336	0.881	1.014	
batik	1 852	1.015	1.031	0.989
biojava	15 215	1.032	1.107	2.464
eclipse	14 596	1.019	0.992	1.046
fop	934	0.999	1.159	
graphchi	11 032	0.991	1.079	1.139
h2o	3 972	1.099	1.002	0.999
ython	5 244	1.018	1.814	2.793
luindex	6 638	0.995	1.071	
pmd	2 088	0.992	1.084	1.124
sunflow	3 717	0.779	1.450	
xalan	1 430	0.643	4.784	
zxing	871	0.952	1.030	
geomean		0.958	1.373	

5.2 Throughput

Table 6 reports throughput using a moderate heap, 2× the minimum heap for G1. Column two reports running time for G1 in milliseconds. The next three columns show relative performance for LXR, Shenandoah, and ZGC. The confidence intervals (unshown) are less than 1% for most systems, with the highest value on sunflow at 4%. Table 5 shows that the throughput results discussed below hold up on a wide range of heap sizes. ZGC cannot run some workloads because the heap size is too small.

LXR’s average speedup is 4% over G1 and 43% over Shenandoah. Among the latency-sensitive workloads, LXR offers a 4% and 9% improvement over G1 on cassandra and lusearch. On h2, G1 has the best performance, whereas LXR is slower by 10%, but Shenandoah and ZGC are substantially slower at 25% and 76%, respectively. On lusearch, Shenandoah is 6.6× slower than G1, as it still struggles to keep up with the extremely fast rate at which Lucene generates garbage. Shenandoah is within 3% of G1 on five benchmarks, but otherwise it pays for its short pauses with more pauses (measured but unshown) and more expensive barriers, degrading mutator and total performance. Shenandoah’s best result is on eclipse which has a high 17% survival rate.

By contrast, LXR’s young object optimizations give it an advantage on lusearch and other low survival rate benchmarks (Section 3.3.2, [36]). LXR offers substantial improvements over G1 on avroa, sunflow, and xalan: 12%, 22%, and 36% respectively, and even higher improvements over

Shenandoah. These benchmarks are less generational, with survival rates of 5%, 3%, and 17%, demonstrating LXR’s ability to adapt to diverse application behaviors. On these benchmarks LXR is working as intended, modulating the SATB tracing to about 20% of collections, where the SATB trace reclaims most of the mature dead objects (Table 7).

Examining LXR’s excellent performance on avroa in detail, we find that G1 and Shenandoah are spending significantly amounts of time in concurrent collection processing a large, long-lived linked list. This list creates a bottleneck for the tracing collectors [5]. The live list has much less effect on LXR because it only occasionally performs backup SATB tracing. Avroa demonstrates that the worst case limitations when relying only on tracing and copying are not just theoretical, but are encountered in practice. On the benchmarks where LXR performs the best (avroa, sunflow, and xalan), LXR has substantially fewer RC pauses, retired instructions, and cache misses (measured but unshown) than G1 and Shenandoah. LXR’s improvements over Shenandoah highlight the CPU and memory penalties of region-based concurrent copying.

5.3 Analysis of Overheads

This section explores the features and overheads of LXR. Table 7 presents results for LXR in a 2× heap. The first column indicates the running time for each benchmark.

Concurrency. The next three columns of Table 7 show the performance impact of turning off concurrency features by i) tracing (-SATB) in the RC pause, ii) decrementing (-LD) in the RC pause, and iii) both (STW), yielding a stop-the-world collector. It shows ratios with respect to the default LXR.

Tracing concurrently or in the pause has the same total performance on these workloads. H2 degrades the most with stop-the-world tracing, adding 4% to execution time, whereas xalan speeds up by 18%. Decrementing in the pause slows the benchmarks down by 3% on average, and slows down h2 by 26%! These results show that increases in pause time work versus concurrency with the application often make little impact on total time. Xalan is an outlier. It speeds up when the concurrency features are turned off. This result shows that concurrency features sometimes interfere with application threads and doing them when the mutators are already paused can be efficient with respect to total time.

With both SATB tracing and decrements in the pause (a fully stop-the-world collector), the benchmarks run 3% slower on average. H2 has the worst slow down at 31% and xalan improves by 9%. LXR outperforms STW, an approximation of RC-Immix [37]. However, STW lacks the carefully tuned whole heap defragmentation heuristics RC-Immix used. Summarizing, the concurrency features usually do not penalize throughput and are essential to low pause times.

Table 7. Breaking down LXR. The first column indicates execution time. The next three columns show the impact of **concurrency optimizations** relative to LXR by turning off concurrency for the trace (-SATB) and for lazy decrements (-LD), and for both (STW). The next five columns show stats for **GC Pauses**: pauses/second, 50th and 95th percentile pauses, fraction due to SATB, and percentage that occur before lazy decrements complete. The next two capture key **barrier** statistics: increments per ms and the overall field barrier overhead. The last five columns show key **reclamation** statistics: percentage of reclamation due to young RC objects, old RC objects, SATB, percentage of old objects with a stuck RCs, and the ratio of evacuated young volume over the reclaimed young clean blocks (YC).

	time ms	Concurrency			GC Pauses					Barriers		Reclamation (%)				
		-SATB	-LD	STW	/s	50% ms	95% ms	SATB%	!Lazy%	Inc/ms	o/h	Young	Old	SATB	Stuck	YC
cassandra	9 335	1.01	1.05	1.05	2.0	4.6	7.4	21	0	399	0.991	98.5	0.9	0.6	0.0	1.2
h2	9 803	1.04	1.26	1.31	1.6	9.8	20.2	19	0	1 460	1.033	96.0	1.9	2.1	0.4	4.6
lusearch	3 461	1.02	1.04	1.01	154.8	1.0	1.4	21	5	5 058	1.022	99.5	0.2	0.3	0.0	0.4
tomcat	5 203	1.01	1.10	1.09	23.2	2.3	3.9	14	0	2 426	1.013	99.4	0.5	0.2	0.0	0.5
avro	3 822	1.02	1.02	1.06	15.8	1.1	1.5	23	1	479	0.997	94.9	0.1	5.0	0.2	4.6
batik	1 880	0.99	1.00	1.01	1.1	20.7	24.5	50	0	3 511	1.022	42.9	2.6	54.5	6.6	52.6
biojava	15 695	1.00	1.00	1.00	4.5	1.3	4.5	10	0	1 596	1.021	100.0	0.0	0.0	0.0	0.0
eclipse	14 878	1.00	1.04	1.05	1.5	5.4	12.4	29	0	845	1.039	93.8	2.0	4.3	1.0	7.8
fop	933	1.00	1.02	1.02	5.2	2.9	4.4	25	0	776	1.038	97.7	0.0	2.3	0.1	2.7
graphchi	10 935	1.00	1.01	1.00	4.1	1.7	4.3	24	0	512	0.977	95.4	0.0	4.6	0.0	1.7
h2o	4 365	1.00	1.00	1.00	0.5	12.6	13.3	0	0	70	1.046	100.0	0.0	0.0	0.0	0.1
jython	5 340	1.00	1.01	1.00	1.9	2.7	3.1	0	0	219	1.041	100.0	0.0	0.0	0.0	0.0
luindex	6 602	1.01	1.01	1.02	6.6	1.1	1.6	19	0	204	1.015	99.6	0.3	0.1	0.0	0.4
pmd	2 071	1.03	1.02	1.02	5.1	13.5	20.1	33	0	9 296	1.028	93.1	0.0	6.9	0.6	10.6
sunflow	2 894	1.02	1.02	1.01	58.7	1.5	2.0	20	0	6 202	1.010	97.9	0.3	1.8	0.0	2.5
xalan	920	0.82	0.94	0.91	210.6	1.3	1.7	20	22	9 197	1.017	94.8	0.8	4.4	0.2	9.8
xzing	829	0.99	1.01	1.00	10.3	1.4	2.0	27	0	429	0.968	99.4	0.3	0.2	0.0	0.4
min		0.82	0.94	0.91	0.5	1.0	1.4	0	22	70	0.968	42.9	0.0	0.0	0.0	0.0
max		1.04	1.26	1.31	210.6	20.7	24.5	50	2	9 296	1.046	100.0	2.6	54.5	6.6	52.6
mean		1.00	1.03	1.03	29.9	5.0	7.5	21	2	2 511	1.017	94.3	0.6	5.1	0.5	5.9
geomean		1.00	1.03	1.03	6.7	3.0	4.7			1 094	1.016	92.9	0.1	0.0	0.0	1.1

GC Pauses. The next five columns characterize collector pauses, showing their frequency (/s), the 50th and 95th percentile pause time in ms, the fraction of pauses that start an SATB trace, and the fraction of pauses that occur before lazy decrements complete. The frequency of pauses varies by two orders of magnitude, with lusearch and xalan having the highest rate, at about one every 10 ms. The median and 95th percentiles of pause durations are 5.0 and 7.5 ms respectively, on average. H2, batik, and pmd have the longest pauses. The number of collections that trigger an SATB is 21% on average and as high as 50% for batik. It is rare for lazy decrement processing to run into the next pause. For most benchmarks, it never happens, but in the worst case, xalan, it reaches 22%.

Barriers. The next two columns of Table 7 show the number of write barriers per ms and the field barrier overhead. The rate of increments per millisecond directly corresponds to the barrier slow path take rate (Figure 3, line 3). This rate varies by more than two orders of magnitude from just 70 for h2o up to 9 296 for pmd. The total overhead of the field barrier is

measured relative to no write barrier using full heap Immix. On average the field barrier adds 1.6% overhead compared to no barrier, ranging from no measurable overhead for avro and a still modest 4.6% for h2o.

Reclamation. The last four columns of Table 7 show the percentage of objects reclaimed via the implicitly dead RC optimization (Young), mature RC (Old), and by the SATB trace; followed by the percentage of objects with stuck reference counts and the ratio of young object bytes evacuated to blocks freed (YC). At this 2× heap size, the vast majority of objects are reclaimed via the implicitly dead optimization. In the case of biojava, h2o, and jython, it reclaims 100%. Batik has the highest number of objects reclaimed by the SATB trace, at 54.5%, with pmd at 6.9%. The second last column shows that very few objects have stuck reference counts, at most 6.6%, for batik. The last column shows the number of bytes evacuated from young blocks relative to blocks reclaimed. The outlier is batik, which has a very high survival rate, so has to perform a lot of copying to yield free blocks.

However, the geometric mean is just 1.1%, illustrating that LXR reclaims most space without performing copying.

5.4 Sensitivity Analysis

This section summarizes the sensitivity of LXR to i) block size, ii) number of reference counting bits, iii) size of the lock free block allocator, and iv) architecture.

LXR uses a block size of 32 KB, and we found that halving it improved performance by 0.6% on average but one benchmark (pmd) did not run to completion, while doubling introduced a 3.9% overhead and h2o ran out of memory. LXR uses 2 reference counting bits per object. We evaluated 4 and 8 bits. Since address arithmetic maps objects to their counts, we considered only power-of-two bit counts. LXR is sensitive to this choice, with a 2.9% penalty for 4 bits and a 3.4% penalty for 8 bits. This result is notable, demonstrating that just two bits is quite effective for reference counting these workloads.

LXR uses a 32 entry lock-free free block buffer by default (Section 3.5). We evaluated both 64 and 128 entry buffers. The differences were small, just 1.1% and 1.3% slower respectively. Notably, lusearch, which is the fastest allocating workload (Table 3) and motivated us to introduce this structure, saw a 12% slowdown for a 128-entry buffer.

We also explored sensitivity to architecture on three machines, each with a different core count, microarchitecture, and clock speed. At a 2× heap, LXR’s throughput relative to G1 was robust to architecture, showing 4.2%, 3.4%, and 2.2% geometric mean performance improvements over G1 on the Zen 3, Zen 2, and C’Lake respectively. Shenandoah showed a similar trend but much greater sensitivity, reporting 37.3%, 29.4%, and 19.6% slowdowns respectively. We observed similar patterns at the other heap sizes.

6 Discussion and Threats to Validity

OpenJDK 11. Building an entirely new garbage collector in OpenJDK is a major undertaking. We implement LXR in OpenJDK 11 because it is a long term service (LTS) release, including significant back-ported improvements to Shenandoah and ZGC. On 2021-09-14, JDK 17 reached *General Availability* status and is now the long-term support (LTS) release, following JDK 11 [30]. We plan to port our work to JDK 17. Because of the substantial differences between JDK 11 and JDK 17 [31], it is not possible to compare our collector and the JDK 17 collectors directly. The JDK 17 version of ZGC removes its limitation of operating only in large heaps. To validate our most striking findings on Lucene, we ran Shenandoah from OpenJDK 17 and found that its performance was essential identical to the results with OpenJDK 11—even with the many advances included in JDK 17, it still exhibits the same pathologies we report here.

C4 is implemented in Platform Prime (formerly Zing) JDK, not OpenJDK [39]. Apples-to-apples collector comparisons

require using the same runtime. However, the design decisions that we critique and respond to in this paper are common to C4, Shenandoah, and ZGC. The notable advantage of C4 is its generational collection. For highly allocating workloads such as lusearch, the non-generational Shenandoah and ZGC are at a disadvantage with respect to C4, G1, and LXR, which all optimize for generational behavior.

Compressed Pointers, Weak References, and Class Unloading. We have not yet implemented compressed pointers so we disable it in our evaluations [25]. Although profitable, we believe it is orthogonal to LXR and do not expect its absence to change our findings. We have not completed the implementation of weak references or class unloading, so we also disable them on all systems. Turning off compressed pointers increased minimum heap sizes for most benchmarks and together with the other disabled features slowed G1 down by just 2.6% at a 3× heap size.

Implementation and Workloads. Our collector is new so has neither had the benefit of years of heuristic tuning, testing, or optimization by a production team. We believe that our ability to execute on 17 substantial workloads from the recent Chopin development branch of DaCapo demonstrates robustness and completeness.

7 Conclusion

Since 2004, G1’s region-based design with concurrent tracing and strict evacuation has dominated production garbage collectors [20, 27, 39, 47]. This paper identifies fundamental limitations in timeliness due to full heap tracing in all systems and in the concurrent copying approach in C4, Shenandoah, and ZGC. Although they deliver very short pause times, short pauses do not always translate into low-latency for request-driven time critical workloads. We introduce LXR, which takes an entirely different approach. It uses brief stop-the-world pauses and reference counting to promptly reclaim most memory without copying. It defragments the heap with limited opportunistic copying. It concurrently reclaims mature objects, identifying dead objects in cycles and dead objects with stuck reference counts. The result is a collector with moderate pause times that delivers low application latency and high throughput. Our results are robust to heap size and microarchitecture. These initial results demonstrate that this new design delivers performance without requiring additional hardware or memory. We hope that this work will provoke a rethink of modern collector design and a reinvigoration of GC research.

Acknowledgments

This material is based upon work supported by the Australian Research Council under project DP190103367.

Data Availability Statement

We made this implementation of LXR, the DaCapo benchmark suite, and instructions for reproducing our results [48] and the latest source code [46] publically available.

References

- [1] Andrew W. Appel and Kai Li. 1991. Virtual Memory Primitives for User Programs. In *Proceedings of the ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, Santa Clara, CA, USA. 96–107. <https://doi.org/10.1145/106972.106984>
- [2] Hezi Azatchi and Erez Petrank. 2003. Integrating Generations with Advanced Reference Counting Garbage Collectors. In *Proceedings of the International Conference on Compiler Construction (CC)*, Warsaw, Poland, Vol. 2622. Springer, 185–199. https://doi.org/10.1007/3-540-36579-6_14
- [3] David F. Bacon, C. Richard Attanasio, Han Bok Lee, V. T. Rajan, and Stephen E. Smith. 2001. Java without the Coffee Breaks: A Nonintrusive Multiprocessor Garbage Collector. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, Snowbird, UT, USA. 92–103. <https://doi.org/10.1145/378795.378819>
- [4] Henry G. Baker, Jr. 1978. List Processing in Real Time on a Serial Computer. *Commun. ACM* 21, 4 (1978), 280–294. <https://doi.org/10.1145/359460.359470>
- [5] Katherine Barabash and Erez Petrank. 2010. Tracing garbage collection on highly parallel platforms. In *Proceedings of the ACM SIGPLAN International Symposium on Memory Management (ISMM)*, Toronto, Ontario, Canada. 1–10. <https://doi.org/10.1145/1806651.1806653>
- [6] Stephen Blackburn and Kathryn S. McKinley. 2003. Ulterior reference counting: Fast garbage collection without a long wait. In *Proceedings of the ACM SIGPLAN Conference on Object-Oriented Programming Languages, Systems, and Applications (OOPSLA)*, Anaheim, CA, USA. 344–358. <https://doi.org/10.1145/949305.949336>
- [7] Stephen M. Blackburn. 2019. Design and analysis of field-logging write barriers. In *Proceedings of the ACM SIGPLAN International Symposium on Memory Management (ISMM)*, Phoenix, AZ, USA. 103–114. <https://doi.org/10.1145/3315573.3329981>
- [8] Stephen M. Blackburn, Perry Cheng, and Kathryn S. McKinley. 2004. Myths and realities: The performance impact of garbage collection. In *Proceedings of the ACM International Conference on Measurements and Modeling of Computer Systems, (SIGMETRICS)*, New York, NY, USA. 25–36. <https://doi.org/10.1145/1005686.1005693>
- [9] Stephen M. Blackburn, Perry Cheng, and Kathryn S. McKinley. 2004. Oil and Water? High Performance Garbage Collection in Java with MMTk. In *Proceedings of the IEEE International Conference on Software Engineering (ICSE)*, Edinburgh, United Kingdom. 137–146. <https://doi.org/10.1109/ICSE.2004.1317436>
- [10] Stephen M. Blackburn, Robin Garner, Chris Hoffmann, Asjad M. Khan, Kathryn S. McKinley, Rotem Bentzur, Amer Diwan, Daniel Feinberg, Daniel Frampton, Samuel Z. Guyer, Martin Hirzel, Antony L. Hosking, Maria Jump, Han Bok Lee, J. Eliot B. Moss, Aashish Phansalkar, Darko Stefanović, Thomas VanDrunen, Daniel von Dincklage, and Ben Wiedermann. 2006. The DaCapo benchmarks: Java benchmarking development and analysis. In *Proceedings of the ACM SIGPLAN Conference on Object-Oriented Programming Languages, Systems, and Applications (OOPSLA)*, Portland, OR, USA. 169–190. <https://doi.org/10.1145/1167473.1167488>
- [11] Stephen M. Blackburn and Kathryn S. McKinley. 2008. Immix: A mark-region garbage collector with space efficiency, fast collection, and mutator performance. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, Tucson, AZ, USA. 22–32. <https://doi.org/10.1145/1375581.1375586>
- [12] Rodney A. Brooks. 1984. Trading Data Space for Reduced Time and Code Space in Real-Time Garbage Collection on Stock Hardware. In *Proceedings of the ACM Conference on LISP and Functional Programming (LFP)*, Austin, Texas, USA. 256–262.
- [13] Jong-Deok Choi, Manish Gupta, Mauricio J. Serrano, Vugranam C. Sreedhar, and Samuel P. Midkiff. 2003. Stack allocation and synchronization optimizations for Java using escape analysis. *ACM Transactions on Programming Languages and Systems* 25, 6 (2003), 876–910. <https://doi.org/10.1145/945885.945892>
- [14] Cliff Click, Gil Tene, and Michael Wolf. 2005. The pauseless GC algorithm. In *Proceedings of the ACM/USENIX SIGPLAN/SIGOPS International Conference on Virtual Execution Environments, (VEE)*, Chicago, IL, USA. 46–56. <https://doi.org/10.1145/1064979.1064988>
- [15] George E. Collins. 1960. A method for overlapping and erasure of lists. *Commun. ACM* 3, 12 (1960), 655–657. <https://doi.org/10.1145/367487.367501>
- [16] Brian F. Cooper, Adam Silberstein, Erwin Tam, Raghu Ramakrishnan, and Russell Sears. 2010. Benchmarking cloud serving systems with YCSB. In *Proceedings of the ACM Symposium on Cloud Computing (SoCC)*, Indianapolis, Indiana, USA. 143–154. <https://doi.org/10.1145/1807128.1807152>
- [17] David Detlefs, Christine Flood, Steve Heller, and Tony Printezis. 2004. Garbage-first Garbage Collection. In *Proceedings of the ACM SIGPLAN International Symposium on Memory Management (ISMM)*, Vancouver, BC, Canada. 37–48. <https://doi.org/10.1145/1029873.1029879>
- [18] L. Peter Deutsch and Daniel G. Bobrow. 1976. An Efficient, Incremental, Automatic Garbage Collector. *Commun. ACM* 19, 9 (1976), 522–526. <https://doi.org/10.1145/360336.360345>
- [19] Edsger W. Dijkstra, Leslie Lamport, Alain J. Martin, Carel S. Scholten, and Elisabeth F. M. Steffens. 1978. On-the-Fly Garbage Collection: An Exercise in Cooperation. *Commun. ACM* 21, 11 (1978), 966–975. <https://doi.org/10.1145/359642.359655>
- [20] Christine H. Flood, Roman Kennke, Andrew Dinn, Andrew Haley, and Roland Westrelin. 2016. Shenandoah: An open-source concurrent compacting garbage collector for OpenJDK. In *ACM Proceedings of the International Conference on Principles and Practices of Programming on the Java Platform: Virtual Machines, Languages, and Tools (PPPJ)*, Lugano, Switzerland, Walter Binder and Petr Tuma (Eds.). 13:1–13:9. <https://doi.org/10.1145/2972206.2972210>
- [21] DaCapo Group. 2021. DaCapo Benchmarks, Chopin Development. <https://github.com/dacapobench/dacapobench/tree/dev-chopin>
- [22] Roman Kennke. 2021. Shenandoah in OpenJDK 17: Sub-millisecond GC pauses. Blog post. <https://developers.redhat.com/articles/2021/09/16/shenandoah-openjdk-17-sub-millisecond-gc-pauses#>
- [23] Haim Kermany and Erez Petrank. 2006. The Compressor: Concurrent, incremental, and parallel compaction. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, Ottawa, Ontario, Canada. ACM, 354–363. <https://doi.org/10.1145/1133981.1134023>
- [24] Gabriel Kliot, Erez Petrank, and Bjarne Steensgaard. 2009. A lock-free, concurrent, and incremental stack scanning for garbage collectors. In *Proceedings of the ACM/USENIX SIGPLAN/SIGOPS International Conference on Virtual Execution Environments, (VEE)*, Washington, DC, USA. 11–20. <https://doi.org/10.1145/1508293.1508296>
- [25] Chris Lattner and Vikram S. Adve. 2005. Transparent pointer compression for linked data structures. In *Proceedings of the ACM Workshop on Memory System Performance (MSP)*, Chicago, IL, USA. 24–35. <https://doi.org/10.1145/1111583.1111587>
- [26] Yossi Levanoni and Erez Petrank. 2001. An On-the-Fly Reference Counting Garbage Collector for Java. In *Proceedings of the ACM SIGPLAN Conference on Object-Oriented Programming Languages,*

- Systems, and Applications (OOPSLA), Tampa, FL, USA.* 367–380. <https://doi.org/10.1145/504282.504309>
- [27] Per Lidén. 2017. CFV: New Project: ZGC. <https://mail.openjdk.java.net/pipermail/announce/2017-October/000237.html>
- [28] Henry Lieberman and Carl Hewitt. 1983. A Real-Time Garbage Collector Based on the Lifetimes of Objects. *Commun. ACM* 26, 6 (1983), 419–429. <https://doi.org/10.1145/358141.358147>
- [29] MMTk Team. 2022. MMTk: A high performance language-independent memory management framework. <https://mmtk.io>
- [30] Oracle. 2021. JDK 17. <https://openjdk.java.net/projects/jdk/17/>
- [31] Oracle. 2021. JEPs in JDK 17 integrated since JDK 11. <https://openjdk.java.net/projects/jdk/17/jeps-since-jdk-11>
- [32] Erik Österlund. 2020. JEP 376: ZGC: Concurrent Thread-Stack Processing. <https://openjdk.java.net/jeps/376>
- [33] Young Gil Park and Benjamin Goldberg. 1992. Escape Analysis on Lists. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), San Francisco, CA, USA.* ACM, 116–127. <https://doi.org/10.1145/143095.143125>
- [34] Filip Pizlo, Erez Petrank, and Bjarne Steensgaard. 2008. A study of concurrent real-time garbage collectors. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), Tucson, AZ, USA.* 33–44. <https://doi.org/10.1145/1375581.1375587>
- [35] Hideaki Saiki, Yoshiharu Konaka, Tsuneyasu Komiya, Masahiro Yasugi, and Taiichi Yuasa. 2005. Real-Time GC in JeRTyTMVM Using the Return-Barrier Method. In *IEEE International Symposium on Object-Oriented Real-Time Distributed Computing (ISORC), Seattle, WA, USA.* 140–148. <https://doi.org/10.1109/ISORC.2005.45>
- [36] Rifat Shahriyar, Stephen M. Blackburn, and Daniel Frampton. 2012. Down for the count? Getting reference counting back in the ring. In *Proceedings of the ACM SIGPLAN International Symposium on Memory Management (ISMM), Beijing, China.* 73–84. <https://doi.org/10.1145/2258996.2259008>
- [37] Rifat Shahriyar, Stephen M. Blackburn, Xi Yang, and Kathryn S. McKinley. 2013. Taking off the gloves with reference counting Immix. In *Proceedings of the ACM SIGPLAN Conference on Object-Oriented Programming Languages, Systems, and Applications (OOPSLA), Indianapolis, IN, USA.* 93–110. <https://doi.org/10.1145/2509136.2509527>
- [38] Guy L. Steele, Jr. 1975. Multiprocessing Compactifying Garbage Collection. *Commun. ACM* 18, 9 (1975), 495–508. <https://doi.org/10.1145/361002.361005>
- [39] Gil Tene, Balaji Iyengar, and Michael Wolf. 2011. C4: The continuously concurrent compacting collector. In *Proceedings of the ACM SIGPLAN International Symposium on Memory Management (ISMM), San Jose, CA, USA.* 79–88. <https://doi.org/10.1145/1993478.1993491>
- [40] David M. Ungar. 1984. Generation Scavenging: A Non-Disruptive High Performance Storage Reclamation Algorithm. In *Proceedings of the ACM SIGSOFT/SIGPLAN Software Engineering Symposium on Practical Software Development Environments, Pittsburgh, PA, USA.* 157–167. <https://doi.org/10.1145/800020.808261>
- [41] Martin T. Vechev and David F. Bacon. 2004. Write barrier elision for concurrent garbage collectors. In *Proceedings of the ACM SIGPLAN International Symposium on Memory Management (ISMM), Vancouver, BC, Canada.* 13–24. <https://doi.org/10.1145/1029873.1029876>
- [42] Xi Yang, Stephen M. Blackburn, Daniel Frampton, and Antony L. Hosking. 2012. Barriers reconsidered, friendlier still!. In *Proceedings of the ACM SIGPLAN International Symposium on Memory Management (ISMM), Beijing, China.* 37–48. <https://doi.org/10.1145/2258996.2259004>
- [43] Xi Yang, Stephen M. Blackburn, Daniel Frampton, Jennifer B. Sartor, and Kathryn S. McKinley. 2011. Why nothing matters: The impact of zeroing. In *Proceedings of the ACM SIGPLAN Conference on Object-Oriented Programming Languages, Systems, and Applications (OOPSLA), Portland, OR, USA.* 307–324. <https://doi.org/10.1145/2048066.2048092>
- [44] Taiichi Yuasa. 1990. Real-time garbage collection on general-purpose machines. *Elsevier Journal of System Software* 11, 3 (1990), 181–198. [https://doi.org/10.1016/0164-1212\(90\)90084-Y](https://doi.org/10.1016/0164-1212(90)90084-Y)
- [45] Taiichi Yuasa, Yuichiro Nakagawa, Tsuneyasu Komiya, and Masahiro Yasugi. 2002. Return Barrier. In *Proceedings of the ACM International LISP Conference (ILC), San Francisco, CA, USA,* Raymond de Lacaze (Ed.). 79–88.
- [46] Wenyu Zhao. 2022. LXR fork of MMTk. <https://github.com/wenyuzhao/mmtk-core/tree/lxr>
- [47] Wenyu Zhao and Stephen M. Blackburn. 2020. Deconstructing the garbage-first collector. In *Proceedings of the ACM/USENIX SIGPLAN/SIGOPS International Conference on Virtual Execution Environments (VEE), (virtual) Lausanne, Switzerland.* 15–29. <https://doi.org/10.1145/3381052.3381320>
- [48] Wenyu Zhao, Stephen M. Blackburn, and Kathryn S. McKinley. 2022. LXR Artifacts. <https://doi.org/10.5281/zenodo.6501899>