# **Combining Machine Learning and Lifetime-based Resource Management for Memory Allocation and Beyond**

Martin Maas **Google Research** mmaas@google.com

Mohammad Mahdi Javanmard Meta mjavanmard@meta.com

David G. Andersen Carnegie Mellon University dga@cs.cmu.edu

Kathryn S. McKinley Google ksmckinley@google.com

Michael Isard Google Research misard@google.com

Colin Raffel University of North Carolina craffel@cs.unc.edu

# ABSTRACT

Memory management is fundamental to the performance of all applications. On modern server architectures, an application's memory allocator needs to balance memory utilization against the ability to use 2 MB huge pages, which are crucial for achieving high performance. This paper shows that prior C++ memory allocators are fundamentally limited because optimizing this trade-off depends on the knowledge of object lifetimes, which is information allocators lack.

We introduce a two-step approach to attain high memory utilization in huge pages. We first introduce a novel machine learning approach that predicts the lifetime of freshly allocated objects using the stack trace at the time of allocation and treats stack traces as natural language. We then present a fundamentally new type of memory allocator that exploits (potentially incorrect) object lifetime predictions to achieve high memory utilization at full huge page usage. In contrast to prior memory allocators that organize their heap around size classes and free lists, our allocator organizes the heap based on predicted *lifetime classes* and adjusts to mispredictions on the fly. We demonstrate experimentally that this learned lifetime-aware memory allocator (LLAMA) reduces fragmentation with huge pages by up to 78%.

Our approach gives rise to a new methodology for applying ML in computer systems. In addition, similar space-time bin packing problems abound in computer science and we discuss how this approach has applications beyond memory allocation to a wide range of problems.

#### **INTRODUCTION** 1.

Memory management is a decades-old research area [24]

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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 2008 ACM 0001-0782/08/0X00 ...\$5.00.



Figure 1: Overview of C++ memory allocation and how long-lived objects lead to wasted memory. Had the red and yellow objects been swapped, page 2 could be freed and the memory footprint would have halved.

that is fundamental to the performance of all applications. On modern architectures, memory managers determine a workload's ability to use 2 MB (and 1 GB) huge pages instead of traditional 4 KB pages. The use of huge pages is crucial for performance on modern servers, since they substantially reduce the cost of address translation by producing a wider reach in Translation Lookaside Buffers (TLB), reducing misses on the CPU's critical path [5].

Current huge page-aware memory managers [13] trade off huge page usage with memory utilization, breaking up huge pages when they become inefficient. Figure 1 visualizes the source of this trade-off: When a C++ program allocates memory, it calls into a memory allocator library (e.g., TC-Malloc [13]), which places the object at a particular address in memory until the program deletes it. The object may not move. The memory allocator has two goals: use as little memory as possible while placing as many objects as possible into consecutive 2 MB ranges of memory, which enables the utilization of huge pages. The latter is crucial for performance, since these 2 MB ranges can be represented by one (instead of 512) TLB entries in the CPU, reducing the number of TLB misses that slow down the application.

The original version of this paper is entitled "Learning-based Memory Allocation for C++ Server Workloads" and was published in ASPLOS 2020: Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems. \*Work done while at Google.

Challenges arise because memory allocators can only request memory from the operating system at the granularity of a page. As long as a page contains at least one object, it cannot be returned to the operating system. Poor object placement can therefore lead to mostly empty huge pages assigned to an application (Figure 1). These huge pages represent wasted memory or need to be broken up into 4 KB pages, reducing performance.

We demonstrate that wasted pages present challenges for long-running server workloads whose memory footprints shrink and grow over time depending on user demand. Many web services exhibit such highly variable memory consumption [16]. Most objects allocated by a server are short-lived but a small fraction live for a very long time, such as session state, logs, or in-memory data. This lifetime distribution is not a major problem with a 4 KB page size: If, conservatively, 99.99% of objects are short-lived and their average size is 64 B, then using 4 KB pages, the probability that any given page contains a long-lived object is less than 1%  $(1-(0.9999)^{4096/64})$ . With 2 MB huge pages, the corresponding probability is 96%.

When the footprint of a server workload shrinks, most of its huge pages cannot be released back to the operating system because they contain at least one long-lived object. Figure 2 shows this effect for a production image processing service on a synthetic workload.

Since C++ cannot move objects, solving this problem depends fundamentally on reasoning about object lifetimes and grouping objects with similar lifetimes together. In practice, the allocator thus needs to use the information it has at the time of allocation to make a placement decision that avoids long-lived objects being equally spread across pages. The most important such information is the current stack trace at the time of allocation, which is called the *allocation context*. Figure 3 shows an example of such a context and the information it contains.

Prior work on compiler and language-runtime optimization also leverages allocation contexts. Commonly, these approaches collect *profiles* of allocation contexts with associated measurements and then leverage these profiles at run time. Profiles are collected during execution (online) or in a separate profiling run (offline). Allocation context-driven optimizations are often opportunistic [9] – they do not need to provide full *coverage* and make a prediction for every possible context, but yield improvements in cases where they can. In some settings, mispredictions may be corrected (e.g., garbage collectors can move objects). In contrast, we need to be able to make predictions for every possible context we may encounter, since a single long-lived allocation may "cost" up to 2 MB and errors accumulate on long-running servers. This requirement creates a set of new challenges:

- Online profiling is challenging because of overheads. Full-context profiling adds 6% overhead [7, 20], which can be more than memory allocation itself [14]. Sampling [9] is sufficient for opportunistic optimizations but does not provide the full coverage we need.
- Offline profiling is challenging because exercising all possible application behavior ahead-of-time is infeasible and servers are configured in myriad ways with different libraries. Lifetimes are particularly difficult to profile since they require observing both an allocation and a deallocation event. We show that in practice, offline profiling does not provide full coverage either.



Figure 2: Image server memory usage resizing groups of large and small images either backed by huge (red) or small (yellow) pages in the OS, derived from analyzing an allocation trace in a simulator. Huge pages waste systemically more memory (red) and increasingly more over time.

This paper addresses these problems by sampling a *subset* of allocation contexts and using machine learning (ML) to generalize from these contexts to previously unobserved contexts. In particular, our novel treatment of symbolized allocation contexts (stack traces) as natural language produces a model that extracts the meaning of function names and how they appear in stack traces. The model accurately predicts object lifetimes, even for previously unobserved contexts.

This new ability to predict object lifetime classes for every allocation inspires LLAMA, a fundamentally new lifetimepredicting allocator design for huge pages that substantially reduces fragmentation on C++ servers. On allocation, LLAMA predicts N *lifetime classes*, where each class differs by an order of magnitude ( $\leq 10 \text{ ms}$ , 100 ms, 1 s, 10 s, etc.). LLAMA organizes the heap by assigning each huge page to a lifetime class. It subdivides huge pages into blocks and lines, where each block lifetime is predicted less than or equal to its huge page. LLAMA handles mispredictions by observing *actual* object lifetimes and using them to reclassify huge pages.

Figure 4 shows an overview of the approach. A subset of allocation lifetimes are sampled from prior runs and versions of a workload. A model is trained against these samples to provide predictions for *all* allocations. This model is compiled into the application for use by the novel LLAMA memory manager algorithm. We also introduce a new caching approach to make these predictions fast.

This paper makes contributions to memory management and the emerging area of ML for Systems [17]. While many prior uses of ML in computer systems focused on tuning existing heuristics with ML, we instead use ML to reconsider the algorithmic context of memory management altogether. Our approach delivers the first allocator that substantially reduces fragmentation for modern C++ server workloads compared to a free-list allocator and *only* uses huge pages.

Since publication, this work has inspired a general methodology for leveraging ML in systems, with applications in operating systems and computer architecture [17]. Instead of learning a systems problem end-to-end, cheap ML techniques are used to predict a previously unknown property (in this case, object lifetimes) and these predictions are then used to fundamentally redesign the algorithm around leveraging this property while tolerating mispredictions. Some insights from this work have also inspired separate (non-ML) TCMalloc optimizations that are deployed in production, leading to an estimated 1% throughput improvement across Google's fleet [18].

While this paper focuses on C++ memory allocation, the algorithm we present applies to any *space-time bin pack*-

- 1 -\_gnu\_cxx :: \_\_g :: \_\_string\_base char , std :: \_\_g :: char\_traits char , std :: \_\_g :: allocator char :: \_M\_reserve ( unsigned long )
- 2 proto2 :: internal :: InlineGreedyStringParser ( std :: --g :: basic\_string char , std :: --g :: char\_traits char , std :: --g :: allocator char\* , char const\* , proto2 :: internal :: ParseContext\* )
- 3 proto2 :: FileDescriptorProto :: \_InternalParse ( char const\* , proto2 :: internal :: ParseContext\* )
- 4 proto2 :: MessageLite :: ParseFromArray ( void const\* ' int )
- 5 proto2 :: DescriptorPool :: TryFindFileInFallbackDatabase ( std :: \_\_g :: basic\_string char , std :: \_\_g :: char\_traits char , std :: \_\_g :: allocator char const ) const
- 6 proto2 :: DescriptorPool :: FindFileByName ( std :: \_\_G :: basic\_string char , std :: \_\_G :: char\_traits char , std :: \_\_G :: allocator char const ) const proto2 :: internal :: AssignDescriptors ( proto2 :: internal :: AssignDescriptorsTable\* )
- 7 system2 :: Algorithm\_descriptor ( )
- 8 system2 :: init\_module\_algorithm\_parse ( )
- 9 Initializer :: TypeData :: RunIfNecessary ( Initializer\* )
- 10 Initializer :: RunInitializers ( char const\* ) 11 RealInit ( char const\* , int\* , char\*\*\* , bool , bool )
- 11 RealInit ( char const\* , int\* , char\*\*\* , bool
  12 main

ization (9-12) of a larger system (7-8).

Figure 3: An example of an altered but representative allocation context, with colored tokens. A string (1) is allocated within a protocol buffer function (2-6) as part of the initial-



Figure 4: Overview of our ML-based Allocator

*ing problem*: Items (in our case, objects) are assigned to resources (in our case, pages) and a resource can only be released once all objects within it have disappeared. Often, the items' lifetimes are predictable, which enables the use of the LLAMA algorithm. For example, we showed in a later paper that file lifetimes in storage systems are predictable [25]. The general LLAMA approach applies to areas such as storage systems, OS process management, and potentially operations research.

#### 2. OBJECT LIFETIME PREDICTION

This section describes how we predict lifetimes of C++ objects at the time of allocation, which has several challenges: (1) Lifetime depends on the entire calling context at allocation, not only the *allocation site* where the allocator was called. (2) The overhead of online profiling is impractical, because it costs 6% CPU performance [7, 20], which would be more than allocation alone [14]. (3) Full coverage of calling contexts and perfect accuracy are not achievable with offline profiling. Because servers evolve and are configured in myriad ways with different libraries, an offline profiler will thus only ever see a subset of contexts.

We address overhead and coverage challenges by sampling a subset of allocations across multiple executions (Section 2.1). We connect to a given application for a sample period and collect lifetimes for a small fraction of all allocations that occur during this period. Sampling is suitable for both server applications in datacenters and multiple runs of a popular application (e.g., a web browser) on a client.

Sampling may not observe all allocation contexts and we must combine samples from a heterogeneous set of different software versions, while the code bases are constantly updated. Our solution uses ML on observed samples of *tokenized* (subdivided) contexts to predict object lifetimes. We train a supervised model (Section 2.2) that maps from calling context to lifetime and generalizes to unseen contexts.

Another challenge is to perform prediction without significant overhead. For example, TCMalloc's allocation fast path is 8.3 ns (Table 1), which is too short to obtain a prediction from an ML model. In fact, it is not even sufficient to gather all the required features, since collecting a deep stack trace takes 400 ns. We address this problem with a hashingbased cache (Section 2.3) that identifies previously seen contexts by using values that are already in registers (the return address and stack pointer) to index a hash table and execute the model only if the lookup fails. We thus amortize model executions over the lifetime of a long-running server. We next explain each component in more detail.

#### 2.1 Sampling-Based Data Collection

Always-on collection of allocation lifetimes incurs a substantial overhead, e.g., stack tracing adds 14% end-to-end overhead and writing to disk further increases the cost, making continuous profiling infeasible in production. We thus introduce a cheap sample-based continuous profiling mechanism and implement it in TCMalloc [13], similar to other production profiling tools [14]. Our sampling approach periodically connects to servers (for a short duration such as  $\approx 5$  minutes) and samples a subset of all memory allocations within the process. Each sample includes stack trace, object size, and address at allocation and deallocation time.

We implement this profiler with TCMalloc *hooks* that are called periodically, based on the number of allocated bytes. These hooks incur virtually no overhead when they are disabled. We also assign each sampled object an identifier at allocation time and match it at deallocation time to compute lifetimes. For each sampled allocation, we keep a running tally of the distribution of lifetimes, by storing the maximum, minimum, count, sum, and sum of squares. We calculate mean and variance of the lifetimes during post processing. At the end of a sampling period, we store the result into a protocol buffer for later analysis using pprof [11].

#### 2.2 Lifetime Prediction Model

Our sampling approach collects allocation contexts and their associated lifetimes. The simplest way to use these samples would be to store the values in a lookup table that maps allocation contexts to a lifetime class. Within a binary, an allocation context can be stored as a sequence of 64bit pointers representing the locations of the instructions on the call stack. Building this table online is prohibitive due to large overheads of always-on profiling. We therefore use data from prior executions of the application.

However, stack traces are brittle when used across executions. Even stack traces on the exact same binary may differ due to address layout randomization. Using symbol information, it is possible to reconstruct the original method name for each stack frame, but different builds of the same

TCMalloc Fast Path (new/delete)		8.3 ns				
TCMalloc Slow Path (central list)		81.7 ns				
Capture full stack trace		396 ns $\pm$ 364 ns				
Look up stack hash (Section 2.3)		22.5 ns				
Table 1: Timescale comparisons						
Version Difference	Matching/Total # Traces					

Revisions 1 week apart	20,606 / 35,336 (58.31%)				
Revisions 5 months apart	127 / 33,613 (0.38%)				
Opt. vs. non-opt. build	43 / 41,060 (0.10%)				
Table 2. Exaction of individual stack traces that match he					

Table 2: Fraction of individual stack traces that match between different binary versions using exact match of symbolized function names.

binary may still differ. For example, changing libraries can affect inlining decisions, different compiler settings lead to slightly different symbol names, and function names and interfaces change over time. This problem also occurs when collecting traces across a large number of instances of the same server with different build configurations and software versions. Table 2 shows that the fraction of matching stack traces between builds with even minor changes is low and decreases over time. This shows that a lookup table would not be suitable, particularly since our predictor needs to provide a high-quality prediction for every allocation context, which includes unseen contexts.

To address this problem, we design an ML-based predictor that learns on calling contexts of tokenized class and method names to produce accurate predictions for unobserved contexts. We train this model using supervised learning. Training data is generated by grouping samples by allocation context and calculating the distribution of observed lifetimes for each context. We use the 95th percentile  $T_{95}^{i}$  of observed lifetimes of context *i* to assign a label  $L_i \in \{1, ..., 7, \infty\}$  such that  $T_{95}^i < T(L_i) = (10)^{L_i}$  ms. Objects the program never frees get a special long-lived label  $\infty$ . This produces lifetime classes of 10 ms, 100 ms, 1 s, 10 s, 100 s, 21000 s, and  $\infty$ . Our model classifies stack traces according to these labels. To ensure that our model assigns greater importance to stack traces that occur more often, we weight each stack trace according to the number of times it was observed and sample multiple copies for frequently occurring traces. The resulting datasets for our applications contain on the order of tens of thousands of elements.

The use of wallclock time for lifetime prediction is a departure from prior work that expresses lifetime with respect to allocated bytes [4], which can be more stable across environments at short timescales. We experimented with logical time measured in bytes, but believe wallclock time works better because 1) our lifetime classes are very coarse-grained  $(10\times)$  and absorb variations, 2) if the speed difference between environments is uniform, nothing changes (lifetime classes are still a factor of  $10 \times$  apart). Meanwhile, variations in application behavior make the bytes-based metric very brittle over long time ranges. For example, in our image server, the sizes of submitted images, number of asyn-



Figure 5: LSTM-based model architecture

chronous external events, etc. dilate logical time.

We use a model similar to text models. In recent years, there has been an explosion of work that learns on text representations of code [2], and our paper represents an example of this approach. In contrast to much of this work, we use code to reason about dynamic program properties rather than static code, an area that has seen less attention [8].

First, we treat each frame in the stack trace as a string and tokenize it by splitting based on special characters such as, and ::. We separate stack frames with a special token: **©**. We take the most common tokens and create a table that maps them to IDs. One special ID is reserved for unknown or rare tokens, denoted as UNK. The table size is a configuration parameter (e.g., 5,000 covers most common tokens).

We use a long short-term memory (LSTM) recurrent neural network model [12]. LSTMs are typically used for sequence prediction, e.g., for next-word prediction in natural language processing. They capture long-term sequential dependencies by applying a recursive computation to every element in a sequence and outputting a prediction based on the final step. In contrast, feed-forward neural networks like multi-layer perceptrons [10] or convolutional neural networks [15] can recognize local patterns, but require some form of temporal integration in order to apply them to variable-length sequences.

Our choice of an LSTM is informed by stack trace structure. Figure 3 shows an example. Sequentially processing a trace from top to bottom conceptually captures the nesting of the program. In this case, the program is creating a string, which is part of a protocol buffer (proto) operation, which is part of another subsystem. Each part on its own is not meaningful: A string may be long-lived or short-lived, depending on whether it is part of a temporary data structure or part of a long-lived table. Similarly, some operations in the proto might indicate that a string constructed within it is temporary, but others make the newly constructed string part of the proto itself, which means they have the same lifetime. In this case, the enclosing context that generates the proto indicates whether the string is long or short-lived.

To learn these patterns, our model must step through the stack frames, carrying through information, and, depending on the context, decide whether or not a particular token is important. This capability is a particular strength of LSTMs (Figure 5). We feed the stack trace into the LSTM as a sequence of tokens (ordered starting from the top of the trace) by first looking up an "embedding vector" for each token in a table represented as a matrix **A**. The embedding matrix A is trained as part of the model. Ideally, A will map tokens with a similar meaning close together in embedding space, similar to word2vec embeddings [19] in natural language processing. Here lies an opportunity for the model to generalize. If the model can learn that tokens such as ParseFromArray and InternalParse appear in similar contexts, it can generalize when it encounters stack traces that it has not seen before.

Note that our approach is not specific to LSTMs. We chose the LSTM architecture since it is one of the simplest sequence models, but future work could explore more sophisticated model architectures that incorporate more details of the underlying program, e.g., Graph Neural Networks trained on program code [3].

We implement and train our model using TensorFlow [1]. Calling into the full TensorFlow stack to obtain a lifetime



Figure 6: High-level overview of low-latency prediction. We use the model only when the hash of the current stack trace is not in the cache. Discarding cache entries periodically helps dynamically adapting to workload changes.

prediction would be prohibitively expensive for a memory allocator, so after training, we use TensorFlow's XLA compiler to transform the trained model into C++ code that we compile and link into our allocator directly.

# **2.3 Speeding Up Predictions**

The allocator must predict object lifetimes quickly to meet latency requirements. TCM alloc allocation times are <100cycles - recording the complete calling context and invoking even a simple neural network takes microseconds, and both are thus too costly. Table 1 shows recording the calling stack for an allocation alone can take an order of magnitude longer than the allocation. We solve these problems by cheaply caching predictions, using a hash of values that are already in registers. We cache predictions as shown in Figure 6 by computing a hash of the return address, stack height and object size, and index a thread-local hashmap. Prior work shows that stack height identifies C/C++ stack traces with 68% accuracy [20]. We add object size to increase the accuracy further. If the hash hits, we use the cached prediction. Otherwise, we run the compiled model, which takes hundreds of  $\mu s$ , and store the result in the cache.

When stack hashes with very different lifetimes alias or workloads change, prediction accuracy suffers. We found that 14% of predictions disagreed with the currently cached value. To address this problem, we periodically discard cached entries. Every, e.g., 1,000 cache hits, we run prediction again. If the result agrees with the current entry, we do nothing. Otherwise, we set the cache entry to the maximum lifetime of the old and new prediction. We use maximum because the allocator is more resilient to overpredicted lifetimes than under-predicted lifetimes.

# 3. LIFETIME-AWARE ALLOCATOR

This section introduces LLAMA, a fundamentally new design for C/C++ memory managers based on predicted object lifetimes. Instead of building an allocator around segmenting allocations into size classes [13, 24], we directly manage huge pages and segment object allocation into predicted lifetime classes. We further divide, manage, and track huge pages and their liveness at a block and line granularity to limit fragmentation. We implement our allocator from scratch. LLAMA is a fundamentally different approach for heap management, although its hierarchical heap organization has similarities to Immix [6]. LLAMA is an untuned research prototype, but demonstrates the potential of a lifetime-based memory allocation approach.

#### 3.1 High-Level Structure

LLAMA organizes the heap into huge pages. To limit

fragmentation, we divide huge pages into 8 KB blocks and track their liveness. LLAMA assigns each active huge page one of N lifetime classes (LC), separated by an order of magnitude (e.g., 10 ms, 100 ms, 1000 ms, ...,  $\infty$ ).

LLAMA's global allocator manages huge pages and their blocks. It acquires and releases huge pages from the OS as needed. It directly manages large objects (>= 8 KB), placing them into contiguous free blocks in partially free huge pages or in new huge pages. Small objects are handled by thread-local allocators that request ranges of blocks from the global allocator. A huge page may contain large and small objects. We will first describe the global allocator's algorithm for large objects and then describe the handling of small objects in Section 3.5.

#### 3.2 Lifetime-Based Huge Page Management

LLAMA stores a small amount of metadata for each huge page. Huge pages have three states: *open, active,* and *free.* Open and active huge pages are *live* and consume 2 MB of memory each. Each huge page has an associated LC and only one huge page per LC is open at a time. While a huge page is open, LLAMA only assigns its blocks to objects with the same predicted LC as the huge page. LLAMA transitions a huge page from open to active after filling all its constituent blocks for the first time. The huge page remains active for the rest of its lifetime. A huge page is free when all its blocks are free and is immediately returned to the OS.

All blocks in a huge page are *free* or *live*; and *residual* or *non-residual*. These properties are tracked via bitmaps. When blocks on an *open huge page* are assigned, these blocks are marked as *residual*, which means that they are predicted to match the LC of their huge page. An *active huge page* may also contain other live (non-residual) blocks, but these blocks will contain objects of a shorter lifetime class, as explained below.

LLAMA initially places objects in the open pages corresponding to their predicted LC and transitions these pages from open to active once they are full. At this point, the huge page contains residual blocks and maybe free blocks. Figure 7 shows a simple example with three lifetime classes, separated by orders of magnitude. A large amount of initial allocations are placed in open pages (Figure 7a), including a large object in huge pages 11 and 12. Figure 7b then shows many frees, which cause LLAMA to return free huge pages 2 and 6 to the OS.

#### 3.3 Recycling Blocks to Limit Fragmentation

As shown in Figure 7b, active huge pages contain free blocks and live residual blocks of the same LC. If the free blocks were never reused, the allocator would waste significant amounts of memory to fragmentation. LLAMA limits fragmentation by aggressively *recycling* such free blocks for objects in shorter LCs. Given a request for LC lr, the global allocator prefers to use free blocks from a longer-lived active huge page (LC > lr). These recycled blocks are marked as non-residual, as illustrated in Figure 7c. If no such recyclable blocks exist, the global allocator uses block(s) from the open huge page of the same LC = lr.

Intuitively, if the predictor is accurate, all objects in a huge page with lifetime class LC will be freed within  $1.1 \times$  LC. All residual objects have a lifetime of at most LC and all non-residual objects are of the next-lower lifetime class, i.e.,  $0.1 \times$  LC. Because lifetime classes are separated by an

order of magnitude, the allocator may reuse the non-residual blocks many times while the longer-lived objects on the huge page are in use, reducing the maximum heap footprint.

For example, given the heap state in Figure 7b and a request for a two-block large object with lr < 10 ms, the global allocator allocates it into huge page 7 with LC < 100 ms and marks the blocks non-residual, as illustrated in Figure 7c. Once the program has freed all the objects on residual blocks within a huge page, all remaining (non-residual) blocks have a lifetime class at least one less than the huge page's current lifetime class. At this point, the huge page is reclassified as the next-lower lifetime class and all the current live blocks are set to residual (Figure 7d). Allocation then proceeds as before, filling the gaps between these blocks with even shorter-lived blocks and repeatedly reducing the lifetime class of the page until it is free.

## 3.4 Tolerating Prediction Errors

Since not all predictions made by the model are correct, LLAMA needs the ability to handle both over-predicted and under-predicted lifetimes.

Over-predicted lifetimes are handled using the mechanism that reclassifies huge pages once all residual objects are gone. For example, if residual blocks are freed before their lifetime has expired, the page will be reclassified earlier, which may result in it freeing up earlier.

Under-predicted lifetimes are more difficult. For example, if a huge page in the lowest lifetime class contains a longlived object, it can result in a large amount of fragmentation since the page does not become free and also cannot be used for allocating new objects. We detect under-prediction of lifetimes using deadlines. When a huge page becomes full for the first time, the global allocator transitions it from *open* to *active* and assigns it a deadline as follows:

$$deadline = current\_timestamp + K \times LC_{Huge Page}$$

When LLAMA changes the LC of a huge page, it assigns the huge page a new deadline using the same calculation. The intuition is that when all predictions are correct, the page would be free after  $1.1 \times LC$ . If we observe that a huge page has been active without reclassification for much longer than this (e.g., setting K = 2), we know that one or more of its constituent blocks were under-predicted.

When a huge page's deadline expires, then the predictor made a mistake. To recover, LLAMA increases the huge page's lifetime class and gives it a new deadline. The huge page remains in the *active* state. Figure 7e depicts this case. The residual blocks in huge page 1 outlive their deadline and LLAMA increases its LC to 100 ms. A huge page may also contain non-residual blocks, which are left unchanged. If blocks live for even longer than this LC, this process will repeat until the blocks are freed or reach the longest-lived LC. This policy ensures that huge pages with under-predicted objects eventually end up in the correct lifetime class, tolerating mispredictions.

#### 3.5 Handling Small Objects

LLAMA achieves scalability on multicore hardware by using mostly unsynchronized *thread-local* allocation for small objects ( $\leq =8$  KB). The global allocator gives 16 KB *block* spans to local allocators upon request. Local allocators hold one or two block spans for each LC. LLAMA further subdivides block spans into 128 B *lines* and *recycles* lines in



(a) Initial allocations. Huge pages are bump-pointer allocated into LC regions. Each huge page is first filled with same LC blocks, marked residual with a dot.



(b) After object frees, some blocks and huge pages are free (white). LLAMA immediately returns free huge pages to the OS to control maximum heap size.



(c) Subsequent allocations of shorter LC small objects first fill free blocks in the highest LC in A(ctive) huge pages 9 and 10, and then blocks in huge page 7. These blocks are not residual (no dot) and expected to be freed before the residual blocks. O(pen) pages 5, 8, and 12 are ineligible for such allocation.



(d) Huge page 9 only contains non-residual blocks and consequently, LLAMA decreases its LC. It marks all live blocks residual since they match or are less than the huge page's LC.



(e) When huge page 1's deadline expires, residual blocks are still live (misprediction). LLAMA increases the huge page's LC by one, from 10 to 100 ms. Residual blocks remain residual; their expected lifetime is now at least 100 ms. Figure 7: LLAMA's logical heap organization with three lifetime classes (< 10 ms, < 100 ms, < 1 s). Each live huge page is A(ctive) or O(pen) and divided into blocks. Block color depicts predicted LC or free (white). Residual blocks

partially free block spans for small objects. It tracks line and block liveness using *counters* that describe how many objects are live within this span or line.

are marked with a dot. Deadlines and lines are omitted.

Small objects occupy one or more contiguous lines within the same span. Once a span is closed (filled at least once), subsequent frees may create a fully or partially free span. Fully free spans are returned to the global allocator. Partially free spans are recycled, but only after the deadline of their huge page expires. On huge page expiration, the global allocator scans the huge page and adds any closed partially free spans to a list, to be reassigned to thread-local allocators in the future. When a span is assigned to a thread-local allocator, it is marked as open.

A local allocator may have one or two open spans per

LC: one initially partially free and one initially fully free. LLAMA sequentially allocates small objects into partially free spans until it encounters an occupied line or the end of the span. When it encounters an occupied line, it skips to the next free line(s). If an object still does not fit, the allocator uses the fully free span, similar to Immix [6].

## 4. EVALUATION

We evaluate LLAMA on four workloads. Except for Redis, they are large production code bases. Experiments are run in a research setup on a workstation with a 6-core Intel Xeon E5-1650 CPU running at 3.60GHz with 64 GB of DRAM and Linux kernel version 4.19.37.

- **Image Processing Server.** A Google-internal production image processing server that filters and transforms images, using synthetic inputs for measurement that produce fragmentation consistent with production.
- **TensorFlow.** The open-source TensorFlow Serving framework [21] running the InceptionV3 [23] image recognition model. This workload exercises libraries with complex memory allocation behavior, such as the *Eigen* linear algebra library. It runs 400 batches of requests in a harness. While running an old model, the benchmark is representative of modern workloads as well.
- **Data Processing Pipeline.** A Google-internal data processing workload running word count on a 1 GB file with 100 M words. We run the entire computation in a single process, which creates very high allocator pressure, resulting in 476 parallel threads and 5M allocations per second.
- **Redis.** The open-source Redis key-value store (4.0.1) running its standard **redis-benchmark**, configured with 5K concurrent connections and 100K operations of 1000 B.

These workloads stress every part of our allocator. They use 10s to 100s of threads, a mix of C++ and C memory allocation, object alignment, have a large ratio of allocation to live objects, and a large amount of thread sharing. They frequently communicate objects between threads, causing the free lists to be "shuffled" and leading to fragmentation. We believe these workloads are representative of modern C/C++ server application. They stress the memory allocator significantly more than workloads used in some prior C/C++ memory manager evaluations, such as SPEC CPU. These patterns are similar to Java applications, illustrating the evolution of C/C++ applications and how they now heavily rely on their memory managers.

The goal of the evaluation is to 1) demonstrate this approach is promising and works on large production code bases; 2) understand trade-offs, such as the model's generalization abilities; and 3) characterize LLAMA. The original paper contains more details and a comparison to Mesh [22].

#### 4.1 End-to-end Evaluation

Table 3 shows end-to-end fragmentation improvements over TCMalloc for the four workloads, ranging from 19% to 78%. Figure 8 shows the image processing server's fragmentation as a function of time. Since vanilla TCMalloc did not support huge pages at the time the paper was written (it does now [13]), we reconstructed the number of occupied



Figure 8: LLAMA reduces huge page (HP) fragmentation compared to TCMalloc on the Image Processing Server. TC-Malloc numbers optimistically assume all free spans are immediately returned to the OS, which is not the case.



(a) Sampling Rate (b) Workload Variations Figure 9: The lifetime model generalizes to unobserved allocation sites from different versions and compiler settings. Blue shows accuracy per stack trace, green weighted by allocations. Light/dotted data shows off-by-one accuracy.

and free huge pages from its bookkeeping information. This method is a lower bound because it does not take into account that TCMalloc does not immediately (or sometimes ever) release pages to the OS. TCMalloc's actual occupancy will be between this amount and the largest peak in the trace, depending on page release rate. Even when compared with the most optimistic variant, we eliminate 43% of the fragmentation introduced by TCMalloc for the image server (in steady state and at termination). Note these results include the memory overheads of our model.

#### 4.2 Model Evaluation

This section evaluates the accuracy of our model. Figure 9 shows that classification accuracy remains high when training our model on one version of the image server and applying it to another, and at relatively low sampling rates. The same configuration in Table 2 shows almost no matching stack traces with a lookup table. In contrast, the model achieves upwards of 80% accuracy when applied to the other revision, and increases to 95% when ignoring errors where the prediction is off by at most one lifetime class.

We see an interesting effect for the non-optimized build. This example achieves few exact matches but higher accuracy for off-by-one errors. We hypothesize that because the non-optimized version of the code runs slower, lifetimes are consistently in a higher class than optimized code.

#### 4.3 Performance Overheads

This section explores overheads compared to TCM alloc. The original paper contains more details. The model takes 100-500  $\mu \rm s$  per prediction for common stack sizes, which would be too expensive to run on every allocation but is acceptable since the model only runs when missing in the cache. The allocator consumes 56 MB for our first workload, less than 2% of the maximum heap size. As we show in Section 4.1, LLAMA recoups this memory easily.

Workload	Prediction Accuracy		Final Steady-state Memory			Fragmentation
	Weighted	Unweighted	TCMalloc	LLAMA	Live	reduction
Image Processing Server	96%	73%	664 MB	446  MB	153  MB	43%
TensorFlow InceptionV3 Benchmark	98%	94%	282 MB	269  MB	214  MB	19%
Data Processing Pipeline	99%	78%	1964 MB	481  MB	50  MB	78%
Redis Key-Value Store	100%	94%	832 MB	$312 \mathrm{~MB}$	$115 \mathrm{~MB}$	73%
	1		' <u> </u>			

Table 3: Summary of Model Accuracy and End-to-end Fragmentation Results

We next evaluate our stack hashing approach. For the image server, 95% of predictions hit in the cache, which shows that stack hashing reduces model evaluations. To evaluate accuracy, we sample predictions and measure how often they disagreed with the cached value. They disagree 14% of the time, but only require updates to longer lifetime classes for 1.6% of allocation sites.

Finally, we characterize LLAMA's overall performance using a microbenchmark that stress tests the allocator. The average latency for global allocations hitting in the cache is 88.0 ns (vs. 81.7 ns for TCMalloc) while fast path allocations take 48.8 ns (vs. 8.3 ns for TCMalloc), much of it because of predictions. In practice, the memory allocator receives much less pressure than in this stress test and the difference is thus less pronounced. For example, the image server slows down  $\approx 12.5\%$  per query compared to TCMalloc.

Our allocator is largely unoptimized. The global allocator is protected by a central lock that is currently the main performance bottleneck. We believe the prototype's bottlenecks can be addressed in a production implementation. Production allocator optimizations include rigorous tuning of every instruction on the fast path, software prefetch instructions, use of restartable sequences to reduce synchronization overheads, size class tuning, and fine-grained locking.

#### 5. **DISCUSSION**

This section discusses how we believe this work has broader implications for other space-time bin packing problems and the emerging field of ML for Systems.

#### 5.1 Generalization to Other Problems

While this paper focuses on 2 MB huge pages, 1 GB huge pages are already available on current hardware. However, to our knowledge, they are not widely relied upon for C++ workloads. LLAMA could feasibly be extended to handle this use case as well, but may need to adjust some of its policies.

While this paper focuses on memory management, the LLAMA algorithm applies to any space-time bin packing problem where items are assigned to resources and a resource can only be released once all items within it are gone. These kinds of resource scheduling problems abound in computer science. For example, a component of a system can only be switched off if nothing is running on it, a storage system may require that a block can only be freed once it is fully empty, and load balancing in distributed systems often cannot move workloads and can only free resources once all running jobs on them have finished. Another example are NAND blocks in flash-based SSDs: When an SSD runs out of empty blocks, it needs to reclaim existing blocks and relocate all remaining data within them, incurring an overhead for any block that is not completely empty.

In many of these cases, lifetimes of items are predictable, e.g., how long it will take to process a particular request or how long a particular file will persist. In this case, the LLAMA algorithm may provide an effective solution to these problems. The prerequisite is that there is some additional information that enables predictions, such as stack traces, memory addresses, or request metadata [25].

## 5.2 A General Pattern of ML for Systems

Generalizing some of the insights from this paper led to a new methodology for applying Machine Learning in computer systems [17]. Prior work on ML for Systems focused on learning a problem end-to-end, which can be inefficient and requires complex models. This paper instead shows an approach of using ML to learn only a particular piece of information that was previously unavailable – in this case, the lifetime of an allocated object. By applying ML to a more limited problem, simpler models can be used and learning can be integrated into a traditional system in a more practical way. At the same time, exploiting the new information may require redesigning the rest of the system, both a cost and an opportunity, and it certainly requires adding the ability to tolerate mispredictions.

Isolating the portion of the problem that needs to be learned also means that the community now can establish best practices, benchmarks and tools for these specific subproblems. We found that most of these problems fall into a small number of categories [17].

# 5.3 Follow-Up Work Since the Original Paper

While LLAMA focuses on memory allocation, we also worked on storage systems and showed that similar patterns repeat in this area [25]. Since the publication of the LLAMA paper, the TCMalloc team has published a paper about Temeraire, a new hugepage-aware allocator [13]. In follow-up work, we applied insights from this work to make better decisions about when to break up huge pages in this allocator, which led to an estimated 1% throughput improvement across Google's fleet [18].

#### 6. CONCLUSION

We show that modern ML techniques can be effectively used to address fragmentation in C++ server workloads that is induced by long-lived objects allocated at peak heap size. We use language models to predict lifetimes *for unobserved allocations contexts*, a problem unexplored in prior lifetime prediction work. We introduce LLAMA, a novel memory manager that organizes the heap using huge pages and lifetime classes, instead of size classes.

LLAMA packs objects with similar lifetimes into the same huge pages, tracks actual lifetimes, and uses them to correct for mispredictions. It limits fragmentation by filling gaps created by frees with shorter-lived objects. In this context, this work solves challenges related to applying ML to systems problems with strict resource and latency constraints. We believe that the LLAMA approach applies to a wide range of other space-time bin packing problems.

# 7. ACKNOWLEDGEMENTS

We would like to thank Harry Xu who was the shepherd of the original paper. We would also like to thank Ana Klimovic, Chris Kennelly, Christos Kozyrakis, Darryl Gove, Deniz Altinbuken, Jae W. Lee, Jeff Dean, Khanh Nguyen, Mark Hill, Martin Abadi, Mike Burrows, Milad Hashemi, Paul Barham, Paul Turner, Sanjay Ghemawat, Steve Blackburn, Steve Hand and Vijay Reddi, as well as the anonymous reviewers, for their feedback. Finally, we would like to give credit to Rebecca Isaacs and Amer Diwan for the initial implementation of the stack hashing mechanism, and to Snehasish Kumar for open-sourcing the lifetime profiler.

# 8. REFERENCES

- M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, M. Kudlur, J. Levenberg, R. Monga, S. Moore, D. G. Murray, B. Steiner, P. Tucker, V. Vasudevan, P. Warden, M. Wicke, Y. Yu, and X. Zheng. Tensorflow: A system for large-scale machine learning. In *Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation*, OSDI'16, pages 265–283, Berkeley, CA, USA, 2016. USENIX Association.
- [2] M. Allamanis, E. T. Barr, P. Devanbu, and C. Sutton. A survey of machine learning for big code and naturalness. ACM Computing Surveys (CSUR), 51(4):81, 2018.
- [3] M. Allamanis, M. Brockschmidt, and M. Khademi. Learning to represent programs with graphs. In 6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings, 2018.
- [4] D. A. Barrett and B. G. Zorn. Using lifetime predictors to improve memory allocation performance. In Proceedings of the ACM SIGPLAN 1993 Conference on Programming Language Design and Implementation, PLDI '93, pages 187–196, New York, NY, USA, 1993. ACM.
- [5] A. Basu, J. Gandhi, J. Chang, M. D. Hill, and M. M. Swift. Efficient virtual memory for big memory servers. In *Proceedings of the 40th Annual International Symposium on Computer Architecture*, ISCA '13, page 237–248, New York, NY, USA, 2013. Association for Computing Machinery.
- [6] S. M. Blackburn and K. S. McKinley. Immix: A mark-region garbage collector with space efficiency, fast collection, and mutator performance. In Proceedings of the 29th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '08, pages 22–32, 2008.
- [7] R. Bruno, D. Patricio, J. Simão, L. Veiga, and P. Ferreira. Runtime object lifetime profiler for latency sensitive big data applications. In *Proceedings of the Fourteenth EuroSys Conference 2019*, EuroSys '19, pages 28:1–28:16, New York, NY, USA, 2019. ACM.
- [8] B. Chen, D. Tarlow, K. Swersky, M. Maas, P. Heiber, A. Naik, M. Hashemi, and P. Ranganathan. Learning to improve code efficiency, 2022.
- [9] D. Clifford, H. Payer, M. Stanton, and B. L. Titzer. Memento mori: Dynamic allocation-site-based optimizations. In *Proceedings of the 2015 ACM*

SIGPLAN International Symposium on Memory Management, pages 105–117, 2015.

- [10] I. Goodfellow, Y. Bengio, and A. Courville. Deep Learning. MIT Press, 2016. http://www.deeplearningbook.org.
- [11] Google. pprof, 2020.
- [12] S. Hochreiter and J. Schmidhuber. Long short-term memory. Neural computation, 9(8):1735–1780, 1997.
- [13] A. Hunter, C. Kennelly, P. Turner, D. Gove, T. Moseley, and P. Ranganathan. Beyond malloc efficiency to fleet efficiency: a hugepage-aware memory allocator. In 15th USENIX Symposium on Operating Systems Design and Implementation (OSDI 21), pages 257–273. USENIX Association, July 2021.
- [14] S. Kanev, J. P. Darago, K. Hazelwood, P. Ranganathan, T. Moseley, G.-Y. Wei, and D. Brooks. Profiling a warehouse-scale computer. In Proceedings of the 42Nd Annual International Symposium on Computer Architecture, ISCA '15, pages 158–169, New York, NY, USA, 2015. ACM.
- [15] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. Backpropagation applied to handwritten zip code recognition. *Neural Computation*, 1(4):541–551, 1989.
- [16] D. Lo, L. Cheng, R. Govindaraju, L. A. Barroso, and C. Kozyrakis. Towards energy proportionality for large-scale latency-critical workloads. In ACM International Conference on Computer Architecture (ISCA), pages 301–312, 2014.
- [17] M. Maas. A taxonomy of ml for systems problems. *IEEE Micro*, 40(5):8–16, 2020.
- [18] M. Maas, C. Kennelly, K. Nguyen, D. Gove, K. S. McKinley, and P. Turner. Adaptive huge-page subrelease for non-moving memory allocators in warehouse-scale computers. In *Proceedings of the 2021* ACM SIGPLAN International Symposium on Memory Management, ISMM 2021, page 28–38, New York, NY, USA, 2021. Association for Computing Machinery.
- [19] T. Mikolov, K. Chen, G. Corrado, and J. Dean. Efficient estimation of word representations in vector space. arXiv preprint arXiv:1301.3781, 2013.
- [20] T. Mytkowicz, D. Coughlin, and A. Diwan. Inferred call path profiling. In *Proceedings of the 24th ACM SIGPLAN Conference on Object Oriented Programming Systems Languages and Applications*, OOPSLA '09, pages 175–190, New York, NY, USA, 2009. ACM.
- [21] C. Olston, N. Fiedel, K. Gorovoy, J. Harmsen, L. Lao, F. Li, V. Rajashekhar, S. Ramesh, and J. Soyke. Tensorflow-serving: Flexible, high-performance ml serving. arXiv preprint arXiv:1712.06139, 2017.
- [22] B. Powers, D. Tench, E. D. Berger, and A. McGregor. Mesh: Compacting memory management for C/C++ applications. In ACM Conference on Programming Language Design and Implementation(PLDI), pages 333–346, 2019.
- [23] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. Rethinking the inception architecture for computer vision. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [24] P. R. Wilson, M. S. Johnstone, M. Neely, and

D. Boles. Dynamic storage allocation: A survey and critical review. In H. G. Baler, editor, *Memory Management*, pages 1–116, Berlin, Heidelberg, 1995. Springer Berlin Heidelberg.

[25] G. Zhou and M. Maas. Learning on distributed traces for data center storage systems. In A. Smola,
A. Dimakis, and I. Stoica, editors, *Proceedings of Machine Learning and Systems*, volume 3, pages 350–364, 2021.