The Measured Cost of Conservative Garbage Collection

BENJAMIN ZORN

Department of Computer Science, Campus Box 430, University of Colorado, Boulder CO 80309-0430, U.S.A.

SUMMARY

Because dynamic memory management is an important part of a large class of computer programs, high-performance algorithms for dynamic memory management have been, and will continue to be, of considerable interest. Experience indicates that for many programs, dynamic storage allocation is so important that programmers feel compelled to write and use their own domain-specific allocators to avoid the overhead of system libraries. As an alternative to explicit storage management techniques, conservative garbage collection has been suggested as an important algorithm for dynamic storage management in C programs. In this paper, I evaluate the costs of different dynamic storage management algorithms, including domain-specific allocators, widely-used general-purpose allocators, and a publicly available conservative garbage collection algorithm. Surprisingly, I find that programmer enhancements often have little effect on program performance. I also find that the true cost of conservative garbage collection is not the CPU overhead, but the memory system overhead of the algorithm. I conclude that conservative garbage collection is a promising alternative to explicit storage management and that the performance of conservative collection is likely to improve in the future. C programmers should now seriously consider using conservative garbage collection instead of explicitly calling free in programs they write.

KEY WORDS Garbage collection Dynamic storage allocation Performance evaluation Conservative collection Dynamic memory management

INTRODUCTION

Dynamic storage allocation is an important part of many kinds of programs, including simulators, interpreters, program and logic optimizers, and window systems. The increasing use of the object-oriented programming paradigm will substantially increase the number of programs that heavily use dynamic storage allocation. The importance of dynamic storage allocation has prompted extensive research on the topic, including design, implementation and evaluation of many different approaches. Techniques for dynamic storage management can be divided into two broad categories: explicit storage management, where the user explicitly allocates and deallocates objects; and automatic storage management, where objects no longer needed are automatically deallocated. Although automatic storage management has been used extensively with certain programming languages, such as Lisp, only recently has the technique been suggested as an extension of C language environments. Automatic storage management has been greatly beneficial in the languages that support it

because it eliminates the need for the programmer to correctly maintain information about all the objects allocated. The only factor preventing this technique from being adopted in all programming languages is its associated performance cost.

The C language commonly provides library routines malloc and free for explicit dynamic storage management. Because standard C allows untagged union types and interconversions of pointer and integer types, the classic garbage collection algorithms, which assume that all program pointers can be correctly identified, cannot be used in all existing C implementations. Conservative garbage collection, an approach to garbage collection that conservatively traces pointers, allows C programs to use automatic storage management. Another common approach to automatic storage management, reference counting, has not been successfully applied in a conservative setting and will not be considered further in this paper.

Explicit and automatic storage management algorithms are very similar and many of the costs present in automatic storage management algorithms are also present in explicit management algorithms. Likewise, many techniques that improve explicit management algorithms can be applied to automatic algorithms, and vice versa.

Given that these two approaches to storage management exist, it is important to compare them and determine their particular strengths and weaknesses. While implementations of conservative garbage collection for C are publicly available, many C programmers are currently not aware that they exist, or perhaps, feel that their performance is inadequate. The purpose of this paper is to increase the awareness of conservative garbage collection and to present a thorough evaluation of its performance relative to explicit storage management techniques.

I compare the performance of four explicit storage management algorithms with the performance of a conservative garbage collection algorithm. A major motivation for this measurement is the belief that programmers have misperceptions about the relative performance of different dynamic storage management algorithms. In particular, there is a widespread belief that garbage collection algorithms are both slow and memory-intensive and that there is little or no cost to explicit storage management.

In an effort to clarify the relative costs of different storage management methods, I have performed a careful and extensive performance evaluation of each algorithm. I not only measure the CPU overhead of the programs, but also the memory they used and the locality of reference exhibited by the programs. Only by a thorough evaluation of all relevant performance metrics can a fair comparison be made.

I evaluate the algorithms' performances in six large C programs that were chosen because they are both widely-used and they allocate a large amount of dynamic data. In four of these programs the programmer intentionally avoided using the system-provided allocators and implemented domain-specific allocation routines for the most common object types. This extra effort suggests that the programmers believed that the library storage management routines were slow.

I decided to test the intuition of the programmer by comparing the performance of the domain-specific enhancements with the general-purpose storage management algorithms. My results show that programmer intuition about storage management performance is often poor. In the programs measured, programmer enhancements had little and sometimes negative effect on the performance of dynamic storage management.

My results also show that the CPU overhead of conservative garbage collection

is comparable to that of explicit storage management techniques. I find that the CPU overhead is more dependent on the application program than the particular storage management algorithm used. Conservative garbage collection performs faster than some explicit algorithms and slower than others, the relative performance being largely dependent on the program. Measurements of memory usage and locality of reference indicate that conservative garbage collection requires substantially more memory than explicit storage management, and also interferes significantly with the program's locality of reference, increasing the number of cache misses and page faults dramatically. Nevertheless, the overall results indicate that conservative garbage collection performs sufficiently well that programmers should strongly consider using the technique instead of explicit storage allocation.

ALGORITHM DESCRIPTIONS

Explicit storage management implementations

There are many different strategies for explicit dynamic storage management, including first-fit, best-fit, buddy, and segregated-storage algorithms. It is not the intent of this paper to exhaustively compare all possible algorithms for explicit storage management. Instead, this paper considers four carefully-crafted and widely-used implementations in an effort to characterize the variation in performance exhibited by these implementations. Comparing the performance of conservative garbage collection with these implementations provides insight into the relative cost of explicit and automatic storage management. The four explicit implementations considered are: an implementation of Knuth's boundary-tag first-fit algorithm; a Cartesian-tree, better-fit algorithm (as implemented in the Sun Operating System); a fast segregated-storage algorithm (as implemented in BSD Unix); and a hybrid first-fit, segregated-storage algorithm (the GNU publicly available malloc/free implementation).

Knuth's boundary-tag first-fit algorithm

This algorithm is a straightforward implementation of a first-fit strategy with several optimizations.[1] I measured a publicly available implementation of the classic Knuth algorithm written by Mark Moraes. In this algorithm, free blocks are connected together in a doubly-linked list. During allocation the list is scanned for the first free block that is large enough. This block is split into two blocks, one of the appropriate size, and returned. As an optimization, if the extra piece is too small (in this case less than 24 bytes), the block is not split. The free list pointer is implemented as a 'roving' pointer, which eliminates the aggregation of small blocks at the front of the free list.

In this algorithm, allocated blocks contain two extra words of space overhead, one word at each end of the block. Each word contains the size of the block and its current status (allocated or free). These 'boundary-tags' are used to allow efficient freeing and coalescing of storage. When a block of storage is freed, the algorithm examines the blocks before and after it to determine if they are also free. It coalesces the allocated block with these blocks if possible and the freed block is linked into the free-list data structure. Maintaining the boundary tags reduces deallocation to a constant time operation.

In the worst case, this algorithm, while simple, can also be quite slow ($O(n)$), where n is the length of the free list). Likewise, heap fragmentation associated with a first-fit approach can be considerable. In practice, however, this algorithm is often efficient both in time and space. The other algorithms discussed attempt to improve on this simple, yet effective, approach.

Stephenson's Cartesian tree algorithm

This approach is implemented in the Sun Operating System library routines malloc and free.[2] This algorithm, instead of placing the free blocks in a linear list, places them in a Cartesian tree[3] (illustrated in Figure 1). The numbers in the blocks of free storage indicate their size, and the numbers in the tree indicate the size of the largest free block available below that node in the tree. Descendants in this tree are ordered by address, so that left descendants have lower addresses than right descendants, which allows adjacent free blocks to be merged. Such an organization is attractive because the worst-case cost of all operations on the tree (allocation, deallocation, and moving blocks around) is $O(d)$, where d is the depth of the tree.

This algorithm (sometimes called 'better-fit') allows for more precise allocation than first-fit because the tree is examined starting at the root until a block that is close to the necessary size is found. If one of the children of a node is a better fit than the other, that subtree is used to allocate the storage. The leaf node used for allocation is split and whatever remains is placed appropriately in the tree. When an object is deallocated, it is located in the tree based on its address, and coalesced with adjacent free blocks if possible.

The Cartesian tree algorithm appears attractive, but in practice it can have drawbacks. The first and most important is that the tree may become unbalanced, in the worst case degenerating to a linear list of free blocks. Because the objects in the tree completely determine its shape, no balancing algorithms can eliminate this worst-case behavior. Also, while the better-fit nature of allocation would appear to reduce fragmentation, both Stephenson and Knuth warn that under some circumstances, a better/best-fit approach may in fact increase fragmentation due to the algorithm's tendency to increase the occurrence of very small blocks. Empirical comparisons, however, suggest that the memory overhead of the Stephenson algorithm is close to that of best-fit algorithms.[4]

Figure 1. Illustration of how the Cartesian tree organizes free blocks in memory

Kingsley's powers-of-two segregated-storage algorithm

Chris Kingsley implemented a fast segregated-storage algorithm that was distributed with the 4.2 BSD Unix release. Kingsley's algorithm allocates objects in a limited number of different size classes, namely powers of two minus a constant. Allocation requests are rounded up to the nearest size class and a free list of objects of each size class is maintained. If no objects of a particular size class are available, more storage is allocated. No attempt is made to coalesce objects.

Because this algorithm is so simple, it is also easy to provide a fast implementation, and it is consistently the fastest of the algorithms compared. On the other hand, it also wastes considerable space, especially if the size requests are often slightly larger than the size classes provided. This algorithm illustrates one extreme of the time/space trade-offs possible in dynamic storage management. Interestingly, its widespread use would suggest that users often consider CPU performance more important than memory usage in these systems (or, perhaps, are unaware of the penalty).

Haertel's hybrid algorithm

Mike Haertel has implemented a hybrid of the first-fit and segregated-storage algorithms. This implementation is available to the public as the Free Software Foundation implementation of malloc/free. This implementation is available on the Internet via anonymous FTP from prep.ai.mit.edu in the file pub/gnu/malloc-0.1.tar.Z. In Haertel's algorithm, requests larger than a specific size (e.g. 4096 bytes) are managed using a linked-list, first-fit strategy, whereas objects smaller than this are allocated in specific size classes (powers of two), as in the Kingsley algorithm.

An important goal of Haertel's approach is to improve the locality of reference during allocation. He does this by dividing allocated storage up into page-sized chunks and storing information about these chunks in small, highly localized chunk headers. Instead of traversing the entire heap attempting to find a fit, only the information in the chunk headers must be traversed. This algorithm also reduces the per-object space overhead required by other algorithms (such as the boundary-tags in Knuth's algorithm) in the following way. Chunks are allocated so that all objects in a chunk are the same size. The address of any object can be used to locate the chunk header, which contains information about the object size associated with the chunk. The chunk header also contains a count of the number of free blocks within a chunk and the algorithm deallocates entire chunks when all the objects in the chunk have been freed.

The Boehm–Weiser conservative collection algorithm

Any garbage collection algorithm (as distinguished from reference counting algorithms—for a complete discussion of the two approaches to automatic storage management, see the survey article by Cohen[5]) identifies garbage objects by first visiting all objects still in use by the program (or some subset of all objects if generation collection is used). These objects are defined as still in use if they are reachable by a pointer traversal starting from the root set of pointers (typically those on the run-time stack and in registers).

Adding garbage collection to C introduces the difficult problem of identifying

heap pointers without run-time type information. While the ANSI C standard is sufficiently lenient that precise tracking of pointers may be possible within the standard, current C compiler practice prevents precise collection for the following reasons. Because C is weakly typed, a pointer to an object can be stored in an integer variable. Thus, even if the locations of all pointer variables were known to a garbage collector, the collector would be unable to identify all heap pointers precisely. Furthermore, there is no way for a collector to distinguish an integer variable that 'appears' to contain a heap pointer (i.e. has a value that is an address in the heap) from one that really does. Such an integer variable is called an ambiguous pointer. The existence of ambiguous pointers leads to conservative collection algorithms, which treat the root set as a series of potential pointers (ignoring type information present in the source program), and assume that any memory location whose value is an address in the heap is a pointer. Furthermore, because these ambiguous pointers could potentially be integer variables, objects they point to cannot be relocated by the collector since doing so would require modifying the variable's value.

Several different conservative collection algorithms have been implemented and described. These include early work by Caplinger,[6] an unpublished implementation similar to the Boehm–Weiser collector by Doug McIlroy, another unpublished implementation by Paul Hilfinger, and a mostly-copying garbage collector by Bartlett.[7,8] The conservative collection algorithm that is considered in this paper is the one described by Boehm and Weiser[9] (version 1.6 of the software). Alan Demers also played a major role in implementing this collector. Of the implementations mentioned, this implementation is easy to obtain and requires the least effort to use. This implementation is available on the Internet via anonymous FTP from arisia.xerox.com in the file pub/russell/gc.tar.Z.

The Boehm–Weiser conservative collector is similar in many ways to explicit management algorithms, and in particular to Haertel's malloc/free implementation. Both algorithms use different approaches for small and large objects; both algorithms allocate smaller objects in distinct size classes; both algorithms allocate similarly sized objects in aligned, page-sized chunks; and both algorithms reclaim entire chunks if all objects in them are free.

There are also notable differences between Haertel's and Boehm and Weiser's algorithms. Unlike the BSD algorithm described above, the size classes take on values that are not powers of two, but range from 8 to 2048 bytes, concentrating on a variety of smaller sizes (including 16, 24, 32 and 48 byte objects). Chunk headers in the Boehm–Weiser algorithm include bitmaps used to mark which objects are allocated and free in each chunk.

Because of the similarities, allocation proceeds in the Boehm–Weiser algorithm much as it does in the Haertel algorithm: a free list of objects in each size class is maintained, and an object is removed from the appropriate free list when it is allocated. If the free list is empty, the Boehm–Weiser collector can either request more memory from the operating system, or more commonly, can invoke a garbage collection to reclaim unused space.

When a collection is required, objects are visited transitively, starting from the root set, and marked. For this algorithm, the root set includes all data in the stack, static data, and registers. Any potential pointer is followed and the object pointed to is marked (details about how ambiguous pointers are resolved are included in

Reference 9). After visiting reachable objects, the sweep phase reclaims objects that are no longer in use and places them on the appropriate free list.

One of the biggest differences between an explicit storage management algorithm and a garbage collection algorithm is the size and structure of the heap. Garbage collection, because it visits all the live data, is a time-consuming operation. The efficiency of collection is increased if more garbage can be collected during each collection phase. The obvious way to allow more garbage to be collected is to wait longer between collection phases, which results in more data having a chance to become garbage. If collection efficiency were the only consideration, this line of reasoning would suggest that collections should take place infrequently. Unfortunately, the longer one waits between collections, the more address space is required to hold all the allocated and not yet collected garbage. Thus, the prime trade-off in collection algorithm design is between efficiency of collection (how much garbage is reclaimed per collection) and program address space size.

The Boehm–Weiser collector seeks to balance these concerns: one heuristic it uses is that if a collection phase reclaims less than 25 per cent of the heap, then the heap is expanded by 50 per cent. Thus, in these collection algorithms, the heap will grow until each collection phase consistently reclaims more than 25 per cent of the heap. Of course, these percentages are quite heuristic and in many collectors are tunable by the programmer. The heap configuration used in the experiments presented started the heap at 256 kilobytes and allowed the heap to expand based on this heuristic.

RELATED WORK

There have been several substantial efforts to compare a broad range of dynamic storage management algorithms, but none of the existing performance surveys has also evaluated a conservative collection algorithm.

In Knuth's *Fundamental Algorithms text*,[1] he compares the time and space overheads of first-fit, best-fit and buddy system methods. The performance evaluation is based on a simulation where the object sizes and object lifetimes were calculated using probability distributions. Knuth concludes that the first-fit method, with optimizations, works surprisingly well and in all cases better than the best-fit method. The buddy system algorithm he defines also does well in comparison with the other methods. Knuth also compares the expected overhead of a mark-and-sweep algorithm, and concludes that under the best circumstances, garbage collection can be competitive with the other algorithms (an observation confirmed here).

A more recent comparison of explicit management algorithms was conducted by Korn and Vo.[4] They implemented and tested 11 different storage management algorithms, including a doubly-linked first-fit algorithm, the BSD malloc and the Stephenson Cartesian tree algorithm. Their results indicate that the BSD algorithm consistently used the most space and the least time of all the algorithms compared. The Stephenson algorithm, although approaching the best-fit algorithms in space wasted, also had an execution time that was higher than many of the other algorithms.

Bozman *et al.* measured the object allocation behavior of the operating system on two large IBM mainframes with hundreds of users.[10] In their paper, they compare a variety of algorithms, including several variations of buddy systems, first-fit, Cartesian tree algorithms, and subpool caching algorithms. Subpool algorithms are

much like Haertel's hybrid algorithm—pools of specific object sizes are maintained, and requests for larger objects are handled as free-list searches. Bozman *et al.* investigate a variety of subpool algorithms and identify a two-level subpool algorithm that is most appropriate for their multi-user operating systems environment.

The Bozman *et al.* paper contains complete information about the observed interarrival and holding times of objects in the systems measured. Empirical information from actual systems has since been used by other researchers to investigate and compare the performance of their algorithms. In particular, Brent uses their data to compare the performance of a first-fit, balanced binary tree algorithm with Knuth's first-fit boundary tag algorithm,[11] and Oldehoeft and Allan also use empirical data to study the performance of an adaptive enhancement to standard storage management algorithms.[12]

The research presented in this paper differs substantially from the related work in several ways. First, none of the papers discusses or compares the performance of a conservative garbage collection algorithm with explicit storage management algorithms. Secondly, like the Bozman *et al.* study, this paper measures the performance of the algorithms in actual programs; however, this work differs from Bozman *et al's* in that we investigate the performance of individual programs and not multi-user operating systems. Thirdly, this paper presents not only the gross performance measures of execution time and memory used, but also presents detailed information about the reference locality of different algorithms. Finally, this paper considers the performance effect of programmer-written, domain-specific enhancements to general-purpose allocation routines.

METHODS

I evaluate the effect of the storage management algorithms on the performance of six large C programs. Although this approach has been taken in countless performance studies, including the empirical studies mentioned above, the use of actual programs to compare performance has the drawback that only a small sample of the actual space of programs is investigated. Ideally, hundreds of programs would be measured in this study, but such a complete sampling of the space of programs is impractical. Another approach taken in related work is to compare performance based on synthetic benchmarks that use standard distributions of object lifespan, size, and birth rate. Synthetic benchmarks have the drawback that they may be completely unrepresentative of actual program execution. My view is that algorithms are mainly of interest if they improve the performance of programs that people actually use. As such, I have collected a group of programs that are in widespread use and that allocate large amounts of dynamic memory.

The programs considered, presented in Table I, are drawn from different application areas, including language interpreters, number factoring, logic optimization, and VLSI layout tools. Table II shows the fraction of time each program spent doing dynamic storage management, with numbers ranging from 23–38 per cent. In general, the programs spent more time freeing objects than allocating them, except in the case of yacr, where the program seldom called free. Yacr is a special case among these programs because, in this version, the programmer made almost no attempt to deallocate dead objects.

In four other programs (cfrac, gawk, ghostscript, and perl), the programmer wrote

Table I. General information about the test programs

cfrac	Cfrac, version 2.00, is a program that factors large integers using the continued fraction method. The input is a 33-digit number that is the product of two primes.
espresso	Espresso, version 2.3, is a logic optimization program. The input file was one of the larger example inputs provided with the release code.
gawk	Gnu Awk, version 2.11, is a publicly available interpreter for the AWK report and extraction language. The input script formats the words in a dictionary into filled paragraphs.
ghostscript	GhostScript, version 2.1, is a publicly available interpreter for the PostScript page-description language. The input used is a large system documentation manual. For this study, GhostScript did not run as an interactive application as it is often used, but instead was executed with the NODISPLAY option that simply forces the interpretation of the Postscript without displaying the results.
perl	Perl 4.10, is a publicly available report extraction and printing language, commonly used on UNIX systems. The input script formats the words in a dictionary into filled paragraphs.
yacr	YACR (yet another channel router), version 2.1, is a channel router for printed circuit boards. The input file was one of the larger example inputs provided with the release code.

Table II. Test program performance information. Execution times were measured on a Solbourne S4000 Desktop Workstation computer (SPARC architecture) with 13 SPECMarks performance and 16 megabytes of memory. Execution times provided are for the SunOS implementation of malloc/free including programmer enhancements, which is used as a base case throughout this paper. The fraction of time spent in malloc/free was measured with the SunOS malloc/free and programmer allocator enhancements removed. The fraction of time in free includes time spent in realloc

Program	Lines of code	Objects allocated ($\times 10^6$)	Bytes allocated ($\times 10^6$)	Execution time (seconds)	Percentage of time in malloc	Percentage of time in free
cfrac	6000	3·81	65·0	216	10·9	14·4
espresso	15,500	1·66	105	237	8·0	16·4
gawk	8500	4·47	170	144	12·0	26·7
ghostscript	29,500	0·92	89	147	9·7	20·6
perl	34,500	0·36	8·3	174	8·7	14·7
yacr	9000	0·88	28·8	92	28·0	0·1

domain-specific allocation routines for the most common object types allocated by the program. I call programs that use special routines to allocate certain object types the 'optimized' versions of the programs. In performing this research, each program was modified so that it no longer made use of these special-purpose routines but instead always called the general-purpose allocator, and I call this form of the program the 'unoptimized' version. The unoptimized programs were also modified so that calls to malloc invoked the Boehm–Weiser conservative garbage collection implementation and calls to free became null calls, resulting in the garbage-collected version of the programs.

It is important to note that the programs needed to be changed very little to use the Boehm–Weiser collector instead of malloc/free, for which they were written. The program requiring the most complex conversion to garbage collection was cfrac, in which the programmer had implemented a version of reference counting. The reference counting code becomes unnecessary when garbage collection is used, and the removal of that code from cfrac required less than 30 minutes. The ease with which programs can be converted to use garbage collection suggests that programmers should at least experiment with this approach, since the experiment is easy to perform.

To collect the data for this paper, each program was executed in optimized and unoptimized form with each of the four different explicit storage management implementations and also with garbage collection. To ensure repeatability, five executions of each program were performed, the results averaged, and a high consistency was observed (the range of values observed was less than 5 per cent of the mean in all cases).

In addition to executing the programs and observing the CPU overhead and memory usage, the programs were instrumented with an address-trace extraction tool. The tool, AE,[13] extracts all memory references the programs perform, allowing the memory system performance of the programs (and storage management implementations) to be studied using trace-driven simulation. Among other metrics, this approach allows me to collect information about the cache miss rate and virtual memory page fault rate for all cache and memory sizes.

RESULTS

Performance comparison

Figure 2 (numeric data in Table III) shows the CPU performance of the four explicit management algorithms and the conservative collection algorithm in each of the six test programs. From the Figure we see that the CPU performance of the algorithms varies widely. The Figure illustrates that the CPU performance of the conservative garbage collection algorithm is often comparable and sometimes better than the performance of the explicit storage management algorithms. In cfrac, the conservative collector did better than all explicit management algorithms by about 20 per cent because the reference counting code, necessary to perform explicit deallocation in the program, was disabled. The Figure also shows that the CPU performance among explicit management algorithms varies widely. Furthermore, the relative performance of algorithms depends on the application. For example, in four cases the BSD algorithm performs substantially better than the SunOS algorithm, whereas in the other two cases it performs worse (up to 28 per cent worse in the ghostscript program).

The most important conclusion to draw from this comparison is that the cost of conservative garbage collection is comparable to that of explicit memory management. Also, the CPU performance of different dynamic storage management algorithms is highly program dependent—unless there is prior knowledge that garbage collection will show much worse performance than explicit management, potential CPU performance should not be a factor in choosing between explicit and automatic algorithms.

CPU time (seconds)

Figure 2. CPU overhead of explicit and automatic forms of storage management

Table III. Absolute and relative CPU performance of five different dynamic storage management algorithms for the six programs measured. For each measurement, the absolute time (user time plus system time) and the time relative to the SunOS implementation is given (smaller is faster). Execution times were measured on a Solbourne S4000 Desktop Workstation computer (SPARC architecture) with 13 SPECMarks performance and 16 megabytes of memory

Program	SunOS (seconds)/ SunOS = 1	B–W GC (seconds)/ SunOS = 1	GNU (seconds)/ SunOS = 1	Knuth (seconds)/ SunOS = 1	BSD (seconds)/ SunOS = 1
cfrac	216·2/1·00	175·9/0·81	217·2/1·00	225·3/1·04	221·6/1·03
espresso	236·9/1·00	165·9/0·70	206·9/0·87	170·4/0·72	158·8/0·67
gawk	144·0/1·00	156·6/1·09	147·8/1·03	117·8/0·82	105·2/0·73
ghostscript	147·1/1·00	187·3/1·27	176·1/1·20	178·3/1·21	187·8/1·28
perl	173·7/1·00	165·6/0·95	149·2/0·86	145·1/0·84	137·7/0·79
yacr	91·6/1·00	91·2/1·00	99·2/1·08	78·5/0·86	85·0/0·93

Figure 3 (numeric data in Table IV) shows the memory usage of the six test programs using the different storage management techniques. The Figure shows that conservative garbage collection does not come without cost. In programs where memory is correctly deallocated by the programmer (i.e. all but yacr), the conservative garbage collection algorithm requires 30 to 150 per cent more memory than the SunOS implementation. Of course, in the case where the programmer did not correctly deallocate free storage, the garbage collected implementation performs better than the explicit algorithms.

There are three reasons why the conservative collector requires more space: there

Memory Usage (kilobytes)

Figure 3. Memory usage of explicit and automatic forms of storage management. For clarity of exposition, the memory usage numbers for the yacr program have been scaled down by a factor of ten

Table IV. Absolute and relative memory overhead of five different dynamic storage management algorithms for the six programs measured. For each measurement, the absolute memory overhead and the overhead relative to the SunOS implementation are given

Program	SunOS (kilobytes)/ SunOS = 1	B–W GC (kilobytes)/ SunOS = 1	GNU (kilobytes)/ SunOS = 1	Knuth (kilobytes)/ SunOS = 1	BSD (kilobytes)/ SunOS = 1
cfrac	736/1·00	1807/2·45	748/1·02	760/1·03	761/1·03
espresso	1387/1·00	3166/2·28	1315/0·95	1448/1·04	1974/1·42
gawk	1006/1·00	1581/1·57	1086/1·08	1018/1·01	1134/1·13
ghostscript	4053/1·00	9167/2·26	4206/1·04	4908/1·21	4756/1·17
perl	1422/1·00	1901/1·34	1469/1·03	1470/1·03	1597/1·12
yacr	13,841/1·00	10,944/0·79	14,148/1·02	13,697/0·99	14,245/1·03

is additional internal fragmentation caused by rounding object sizes, the heap size is increased to decrease the frequency of collection (and increase CPU performance), and the conservative aspect of the collector is retaining more data than necessary. Measurements of the rounding algorithm used by the conservative collector indicate that internal fragmentation typically accounts for less than 5 per cent of the space allocated by the algorithm. This fragmentation is comparable and often better than that of the other allocators measured. Additional measurements show that increasing the frequency of collection can reduce the total memory usage of the conservative collector, but even with more frequent collections the conservative collector often requires twice as much total memory. Unfortunately, measuring the amount of

additional data retained due to conservatism is difficult because the behavior is not easily reproducible. Later studies will attempt to quantify this effect.

The Figure shows that the fast BSD algorithm also requires more space (up to 42 per cent) than the SunOS algorithm, which showed the best memory performance. In any event, the memory overhead of even the worst explicit storage management algorithm is substantially lower than the conservative garbage collection algorithm.

Given that the memory overhead of conservative garbage collection is so large, it is important to understand the implications of this measurement. Assuming that a program uses virtual memory, the actual program performance cost of such a large address space depends on the locality of reference in accessing it. Before investigating this aspect of performance, however, I present results showing the effect of programmer-written, domain-specific enhancements to the general-purpose allocation algorithms.

Domain-specific allocator enhancements

In four of the programs investigated, the programmer felt compelled to avoid using the general-purpose storage allocator by writing type-specific allocation routines for the most common object types in the program. In this section, I compare the CPU performance and the memory size of the optimized (enhanced) and unoptimized versions of these programs.

Figure 4 (numeric data in Table V) shows the relative performance of the optimized and unoptimized versions of the four programs using each of the explicit storage management algorithms. For example, the BSD optimized bar shows the performance of the program using domain-specific allocation routines for the most common object types and the BSD allocation routines for the rest. The BSD unoptimized bar shows what the performance is when all allocations are performed using the general-purpose BSD allocator. The Figure shows that in some programs, such as cfrac, programmer tuning did improve performance over the standard storage management algorithms, sometimes by as much as 60 per cent. In the ghostscript application, however, the optimized BSD algorithm is much slower than the optim-

Table V. Relative CPU performance of programs using domain-specific allocators for the most common object types (optimized) versus using a general-purpose allocator for all objects (unoptimized). All measurements are relative to the performance of the optimized SunOS implementation of the program (smaller is faster)

| | Relative CPU overhead (SunOS, optimized = 1·0) | | | |
	cfrac	gawk	ghostscript	perl
Sun OS, optimized	1·00	1·00	1·00	1·00
Sun OS, unoptimized	1·62	1·22	1·15	1·24
GNU, optimized	1·00	1·03	1·20	0·86
GNU, unoptimized	1·26	1·10	0·93	0·95
Knuth, optimized	1·04	0·82	1·21	0·84
Knuth, unoptimized	1·32	0·91	1·04	0·91
BSD, optimized	1·03	0·73	1·28	0·79
BSD, unoptimized	1·10	0·75	0·85	0·82

ized version. The next section shows that this behavior is caused by cache miss anomalies in the optimized version that are part of the measured CPU cost.

The general conclusion that must be reached from the Figure is that programmer optimizations in these programs were mostly unnecessary. Although the programmer optimizations can improve the performance of the program for a particular storage management algorithm, simply using a different algorithm (and especially the BSD algorithm) appears to improve performance even more. In three of the applications, programmer enhancements affected the performance of the BSD algorithm only minimally (2–7 per cent). In the fourth application, ghostscript, programmer enhancements actually decreased the program performance in most cases.

Given that the CPU performance of the programs was not significantly improved by programmer optimizations, the possibility exists that the programmer's intent in optimizing was to reduce the storage required by the program. Table VI shows the relative memory sizes required in the optimized and unoptimized versions of the programs for each of the algorithms considered. The table shows that optimizations had little (and sometimes negative) effect on the memory requirements of the programs. Because little space was actually saved in the optimized storage allocators, the results in the table suggest that space savings were not the intent of program optimizations.

The strongest argument one can make for the optimizations is that they reduce the program overhead most when the SunOS algorithm is used. Because the SunOS algorithm often required the least memory overhead of the algorithms considered, optimizing the CPU performance of that algorithm provides both CPU and memory efficiency. Thus the program optimizations sometimes provide performance comparable to the fast BSD algorithm with only the space requirements of the SunOS algorithm. Such a benefit did not always occur, and the BSD algorithm, even without optimization, sometimes reduced program execution time by 20 per cent, with an average 10 per cent additional memory overhead.

Table VI. Relative memory overhead of programs using domain-specific allocators for the most common object types (optimized) versus using a general-purpose allocator for all objects (unoptimized). All measurements are relative to the performance of the optimized SunOS implementation of the program

| | Relative memory overhead (SunOS, optimized = 1·0) | | | |
	cfrac	gawk	ghostscript	perl
Sun OS, optimized	1·00	1·00	1·00	1·00
Sun OS, unoptimized	1·01	1·00	0·92	1·00
GNU, optimized	1·02	1·08	1·04	1·03
GNU, unoptimized	1·01	1·05	1·05	1·03
Knuth, optimized	1·03	1·10	1·21	1·03
Knuth, unoptimized	1·09	1·03	1·56	1·03
BSD, optimized	1·03	1·13	1·17	1·12
BSD, unoptimized	1·03	1·09	1·15	1·12

Reference locality

In this section, I investigate the program locality of reference in three cases: the BSD algorithm with programmer enhancements, the BSD algorithm without enhancements, and the Boehm–Weiser conservative garbage collection algorithm. Using trace-driven simulation in combination with cache and virtual memory simulation tools, I compare the page fault rate and the cache miss rate of these three cases for each of the test programs. In these measurements, the cache miss rates are calculated using the all-associativity cache simulator tycho.[14] Page fault rates were computed using a modified stack simulation algorithm,[15] in which an LRU page replacement policy is assumed. For these measurements, only data references were measured because the dynamic memory management algorithm used has little effect on instruction reference locality. In all cases over 100 million data references were used to measure the miss rate and page fault rate of the programs.

I have already shown that conservative garbage collection increases the memory requirements of the programs measured from 30–150 per cent. In this section, I show that the increased memory requirements translate to a decrease in the locality of reference.

Figures 5 and 6 compare the virtual memory localities of the BSD and garbage collection algorithms. Figure 5 compares the total page faults of the different algorithms in each of the test programs, assuming a fixed memory size of two megabytes. Two megabytes was chosen because it represents a memory size for which there is a significant difference between the allocators in the applications chosen. If smaller memory sizes are measured, all allocators show higher fault rates; if larger memory sizes are chosen, all allocators show low fault rates (illustrated in Figure 6).

Figure 5 shows conclusively that conservative garbage collection significantly decreases the locality of reference in all the programs, typically causing one to two orders of magnitude more page faults for the given memory size. Interestingly, even in the yacr program, where conservative garbage collection reduced the total memory requirements by 20 per cent, the garbage collection algorithm still showed far less locality of reference. This can be attributed to the non-locality of this gargage collection algorithm: the mark phase visits all reachable objects and the sweep phase sequentially references substantial parts of the heap.

Figure 6 shows the page fault rate (faults per reference) as a function of the memory size for the gawk program. This Figure better illustrates the relative reference localities of the optimized and unoptimized versions of the BSD algorithm. The page fault rates of the unoptimized algorithm are slightly higher than those of the optimized BSD algorithm, for all memory sizes. The fault rate of the garbage collection algorithm is substantially higher than either BSD algorithm for memories up to two megabytes.

Figures 7 and 8 compare the total cache miss rate of the three algorithms assuming a direct-mapped cache with 32-byte blocks. Just as the previous figures showed that garbage collection has much worse locality of reference at the page level, these figures show that the cache locality of the Boehm–Weiser garbage collection algorithm is also much worse than the BSD algorithm, except in the ghostscript and yacr applications. Ghostscript shows anomalous behavior when the programmer enhancements are combined with the BSD algorithm, resulting in extremely high cache miss rates. This overhead is also reflected in the measured CPU time of this implementation,

Figure 4. Comparison of CPU performance of four programs using domain-specific allocators for the most common object types (optimized) versus using a general-purpose allocator for all objects (unoptimized). For each program, four different general-purpose allocators are measured
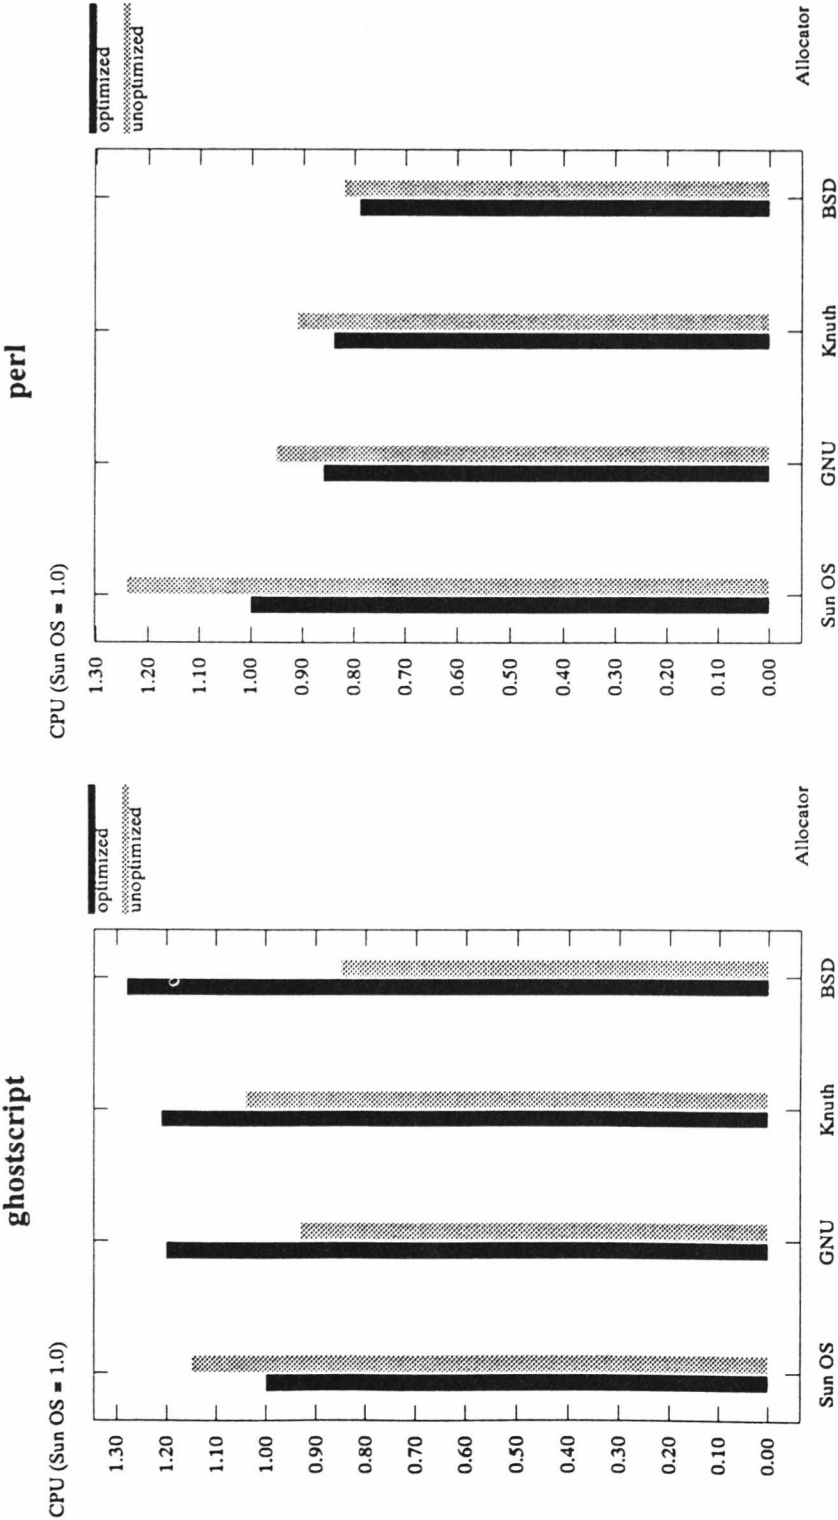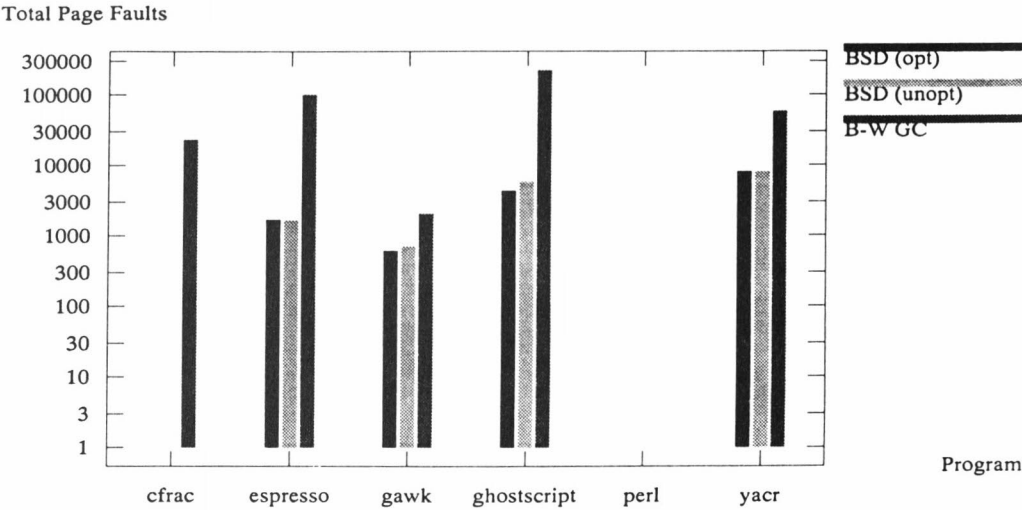
Figure 4. Continued

Total Page Faults

Figure 5. Page fault rates for a memory size of two megabytes in the six test programs. For several of the programs (e.g. perl) and allocators, the address space required was less than two megabytes, and no page faults were required

Page Fault Rate

Figure 6. Page fault rates for different memory sizes in the gawk application

as shown in Table III, where the BSD algorithm uncharacteristically takes longer to execute than the SunOS algorithm. Although I have not investigated the source of this cache contention, I surmise that it results from access conflicts between the domain-specific allocator and the underlying general-purpose allocator.

In four of the applications, however, the Boehm–Weiser collector results in much higher miss rates than the other algorithms. Boehm and Weiser note that their
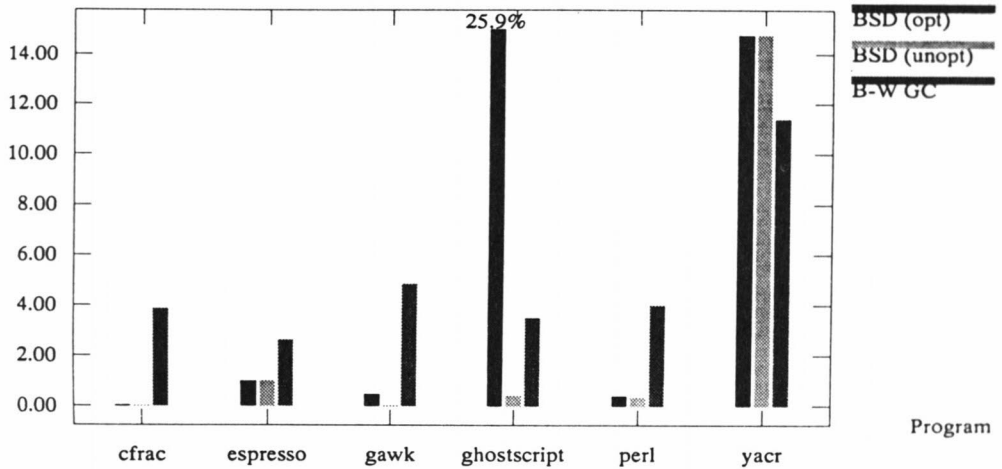
Cache Miss Rate

Figure 7. *Cache miss rates for a 128-kilobyte direct-mapped cache in the six test programs. The BSD (opt) value for the ghostscript program is shown in the graph as 15·0 per cent, but is actually 25·9 per cent*

Cache Miss Rate (%)

Figure 8. *Cache miss rates for different direct-mapped cache sizes in the gawk application*

collector may result in extra cache misses because chunk headers, which are frequently accessed, are also page aligned, and thus contain the same least significant bits in their addresses.[9] My measurements confirm their observation.

Figure 7 shows the miss rates of the algorithms and programs in a direct-mapped, 128-kilobyte cache with a block-size of 32 bytes. In practice, the miss rate of the conservative algorithm appears to be 3–10 times higher than that of the BSD

algorithm. Figure 8 shows the cache miss rate as a function of cache size in the gawk application, assuming a direct-mapped cache. The conservative garbage collection algorithm shows a relatively high miss rate for cache sizes as large as two megabytes.

Based on these measures of reference locality, I reach two important conclusions. First, the Boehm–Weiser conservative garbage collection algorithm not only requires a larger address space, but also disrupts the memory reference locality of the test programs far more than explicit storage management methods. Secondly, I conclude that programmer optimizations do not significantly increase the reference locality of the programs studied in the Solbourne system I have used to compare performance. In computer systems with smaller caches and main memories the impact of programmer optimizations may be more significant.

SUMMARY

In this paper, I have compared a conservative garbage collection algorithm with four explicit storage management algorithms. I have compared CPU performance, memory overhead, and reference locality of the different algorithms to fully evaluate the effect of algorithm choice on program performance. To perform the comparison, I used six widely-used, memory-intensive test programs. In four of these programs, the programmers implemented domain-specific allocators for the most common object types in their programs. I measured the impact of programmer enhancements to determine what programmers consider to be a substantial performance improvement (at least substantial enough to warrant the extra effort of coding the optimizations).

From these measurements, I reach the following conclusions. First, the CPU overhead of conservative garbage collection is often not significantly higher than that of different explicit storage management algorithms. The CPU overhead of storage management is application dependent, and conservative garbage collection compares well with other explicit storage management algorithms. These test programs, which spend up to 30 per cent of their total execution time in storage management, represent the worst case overhead that might be expected to be associated with garbage collection. Even so, the Boehm–Weiser conservative garbage collection algorithm, when compared with four well-implemented, widely-used allocator implementations, was the slowest allocator in only one of the six programs, and the faster allocator in one other.

Conservative garbage collection does not come without a cost. In the programs measured, the garbage collection algorithm used 30–150 per cent more address space than the most space efficient explicit management algorithm. In addition, the conservative garbage collection algorithm significantly reduced the reference locality of the programs, greatly increasing the page fault rate and cache miss rate of the applications for a large range of cache and memory sizes. This result suggests that not only does the conservative garbage collection algorithm increase the size of the address space, but also frequently references the entire space it requires. For systems with enough cache and memory, this disadvantage of the conservative garbage collection algorithm does not significantly degrade its performance, but in systems with smaller caches and memories, some performance degradation is inevitable.

Finally, I have evaluated the effect of domain-specific allocator enhancements in four of the test programs. With some algorithms, programmer enhancements increase

performance significantly, whereas with other algorithms (and most notably the BSD algorithm), enhancements have little, and sometimes negative effect. I conclude that programmers, instead of spending time writing domain-specific storage allocators, should consider using other publicly available implementations of storage management algorithms if the one they are using performs poorly. It appears that choice of algorithm has more impact on performance than programmer enhancement of a particular algorithm. Perhaps operating systems should provide users with a variety of dynamic storage management algorithms instead of providing a single choice. A colleague and I have investigated the possibility of automatically generating an appropriate allocator.[16]

Given that conservative garbage collection has a cost, when should it be considered as an alternative to explicit storage management? This paper has compared conservative garbage collection head-to-head with explicit storage management, until now treating them as if they are equally desirable alternatives. Garbage collection, however, has a tremendous advantage over explicit management methods because it is far easier for a programmer to use. With programs that do substantial explicit storage management, programmers spend a significant amount of time finding and eliminating storage management bugs (memory leaks and duplicate frees). Rovner estimates that developers using the Mesa language spent 40 per cent of the development time implementing memory management procedures and finding bugs related to explicit storage reclamation.[17] In addition, storage management bugs that are not found can greatly contribute to the unreliability of software. Bartlett has noted that a large fraction of software-caused total system failures are caused by memory management errors.

Although the advantages of freeing the programmer from explicit storage management are difficult to quantify, they are obvious and substantial. Keeping this in mind, along with the result that the CPU overhead of conservative garbage collection is not significant, I would recommend that conservative garbage collection be used in almost every case. As a default, the algorithm has low overhead and is far easier to use. In cases where the memory available is known to be quite limited, explicit storage management may be necessary, but should only be considered if garbage collection has been shown to require excessive memory or CPU overhead.

Algorithm improvements

Explicit storage management algorithms have had decades to be tuned and improved, whereas conservative garbage collection algorithms have only been implemented in the past few years. Trends in garbage collection technology suggest that improvements to the Boehm–Weiser conservative collection algorithm are possible and likely in the near future. In particular generation techniques, already successfully applied in Lisp environments, can be applied to conservative collection algorithms.

Generation techniques[18] focus the attention of garbage collection on the most recently allocated objects, which empirical evidence shows are the objects most likely to become garbage. Focusing garbage collection on these younger objects serves three purposes. First, the efficiency of collection is increased because a higher percentage of the objects visited are garbage. Secondly, because fewer objects are visited during each collection, program pauses associated with garbage collection

are shortened. The final and most important purpose of generation garbage collection (based on my measurements) is that the reference locality of garbage collection is substantially increased.[19] This increase occurs because only a small part of the program's address space is visited during a collection.

Generation garbage collection has been successfully used in languages including Lisp,[20] Smalltalk[21] and ML.[22] Because generation collection relies on the behavior that most objects live a relatively short time, generation collection will only be effective for C if C programs also display this behavior. Figure 9 shows the survival curves for objects in five of the applications used in this paper. The survival curve plots the fraction of objects surviving past a particular age as a function of age (e.g., in the Figure, approximately 50 per cent of the objects live beyond 10,000 cycles in the perl application). The Figure shows that in all cases except yacr, fewer than 10 per cent of objects allocated live beyond 100,000 cycles (or approximately 3 ms on a 33 MHz CPU). Object lifespans were measured as the number of instructions between when they were allocated and freed. Object lifespans in yacr are quite long because the programmer failed to free garbage objects. I conclude from the data in the Figure that generation techniques can be successfully applied to C programs. A generation collector that collected objects less frequently than every 100,000 cycles (e.g. every 5–10 seconds) would find mostly dead objects and thus perform very efficiently.

Generation techniques have already been applied to conservative garbage collection algorithms. Boehm et al. have a technique called 'sticky-mark-bit', currently available in the Portable Common Runtime system, that adds generations to their conservative collection algorithm.[23] Bartlett has also extended a mostly-copying conservative collection algorithm for C++ to use generation techniques.[7] From these research initiatives and others that may develop as the true potential of conservative garbage collection becomes clear, it is likely that much better conserva-


Figure 9. Fraction of objects surviving as a function of age in five test programs

tive collection algorithms, with increased locality of reference and decreased address space needs, will be discovered.

In this paper I have measured the performance of a conservative garbage collection algorithm and found it to be comparable to that of the best explicit storage management algorithms. As object-oriented programming becomes widely used, techniques that reduce the complexity of managing heap-allocated storage will become more important. Efficient conservative garbage collection algorithms are available for use right now and, in the future, are likely to increase in efficiency as better algorithms are discovered. C programmers should now seriously consider using conservative garbage collection instead of explicitly calling free in programs they write.

REFERENCES

1. D. E. Knuth, *Fundamental Algorithms, Vol. 1 of The Art of Computer Programming*, Addison Wesley, Reading, MA, 2nd edn, 1973, Chap. 2, pp. 435–451.
2. C. J. Stephenson, 'Fast fits: new methods for dynamic storage allocation', in *Proc. Ninth ACM Symposium on Operating System Principles*, Bretton Woods, NH, October 1983, pp. 30–32.
3. J. Vuillemin, 'A unifying look at data structures', *Commun. ACM*, **23**, 229–239 (1980).
4. D. G. Korn and K.-P. Vo, 'In search of a better malloc', *Proc. Summer 1985 USENIX Conference*, 1985, pp. 489–506.
5. J. Cohen, 'Garbage collection of linked data structures', *ACM Computing Surveys*, **13**, 341–367 (1981).
6. M. Caplinger, 'A memory allocator with garbage collection for C', *Proc. Winter 1988 USENIX Conference*, Dallas, TX, February 1988, pp. 325–330.
7. J. Bartlett, 'Mostly-copying garbage collection picks up generations and C++', *Tech. Rep. TN-12*, Digital Equipment Corporation Western Research Laboratory, Palo Alto, CA, October 1989.
8. J. F. Bartlett, 'Compacting garbage collection with ambiguous roots', *Tech. Rep. 88/2*, Digital Equipment Corporation Western Research Laboratory, Palo Alto, CA, February 1988.
9. H. Boehm and M. Weiser, 'Garbage collection in an uncooperative environment', *Software—Practice and Experience*, pp. 807–820 (1988).
10. G. Bozman, W. Buco, T. P. Daly and W. H. Tetzlaff, 'Analysis of free-storage algorithms—revisited', *IBM Systems Journal*, **23**, (1), 44–64 (1984).
11. R. P. Brent, 'Efficient implementation of a first-fit strategy for dynamic storage allocation', *ACM Trans. Programming Languages and Systems*, **11**, 388–403 (1989).
12. R. R. Oldehoeft and S. J. Allan, 'Adaptive exact-fit storage management', *Communications of the ACM*, **28**, 506–511 (1985).
13. J. R. Larus, 'Abstract execution: a technique for efficiently tracing programs', *Software—Practice and Experience*, **20**, 1241–1258 (1990).
14. M. D. Hill, 'Aspects of cache memory and instruction buffer performance', *Ph.D. Thesis*, University of California at Berkeley, Berkeley, CA, November 1987. Also appears as *Tech. Report UCB/CSD 87/381*.
15. B. Zorn, 'Comparative performance evaluation of garbage collection algorithms', *Ph.D. Thesis*, University of California at Berkeley, Berkeley, CA, November 1989. Also appears as *Tech. Report UCB/CSD 89/544*.

16. D. Grunwald and B. Zorn, 'CustoMalloc: efficient synthesized memory allocators', *Software—Practice and Experience*, **23**, (1993)—to appear.
17. P. Rovner, 'On adding garbage collection and runtime types to a strongly-typed, statically checked, concurrent language', *Tech. Rep. CSL-84-7*, Xerox Palo Alto Research Center, Palo Alto, California, July 1985.
18. H. Lieberman and C. Hewitt, 'A real-time garbage collector based on the lifetimes of objects', *Communications of the ACM*, **26**, 419–429 (1983).
19. B. Zorn, 'The effect of garbage collection on cache performance', *Tech. Rep. CU-CS-528-91*, Department of Computer Science, University of Colorado, Boulder, Boulder, CO, May 1991.
20. D. A. Moon, 'Garbage collection in a large Lisp system', *Conference Record of the 1984 ACM Symposium on LISP and Functional Programming*, Austin, Texas, August 1984, pp. 235–246.
21. D. Ungar, 'Generation scavenging: a non-disruptive high performance storage reclamation algorithm', *SIGSOFT/SIGPLAN Practical Programming Environments Conference*, April 1984, pp. 157–167.
22. A. W. Appel, 'Simple generational garbage collection and fast allocation', *Software—Practice and Experience*, **19**, 171–183 (1989).
23. A. Demers, M. Weiser, B. Hayes, H. Boehm, D. Bobrow and S. Shenker, 'Combining generational and conservative garbage collection: framework and implementations', *Conference Record of the Seventeenth ACM Symposium on Principles of Programming Languages*, January 1990, pp. 261–269.