

Garbage Collection Hints

Dries Buytaert, Kris Venstermans, Lieven Eeckhout, and Koen De Bosschere

ELIS Department, Ghent University – HiPEAC Member,
St.-Pietersnieuwstraat 41, B-9000 Gent, Belgium
{dbuytaer, kvenster, leeckhou, kdb}@elis.UGent.be

Abstract. This paper shows that Appel-style garbage collectors often make suboptimal decisions both in terms of *when* and *how* to collect. We argue that garbage collection should be done when the amount of live bytes is low (in order to minimize the collection cost) and when the amount of dead objects is high (in order to maximize the available heap size after collection). In addition, we observe that Appel-style collectors sometimes trigger a nursery collection in cases where a full-heap collection would have been better.

Based on these observations, we propose *garbage collection hints (GCH)* which is a profile-directed method for guiding garbage collection. Offline profiling is used to identify favorable collection points in the program code. In those favorable collection points, the VM dynamically chooses between nursery and full-heap collections based on an analytical garbage collector cost-benefit model. By doing so, GCH guides the collector in terms of *when* and *how* to collect. Experimental results using the SPECjvm98 benchmarks and two generational garbage collectors show that substantial reductions can be obtained in garbage collection time (up to 30X) and that the overall execution time can be reduced by more than 10%.

1 Introduction

Garbage collection (GC) is an important subject of research as many of today's programming language systems employ automated memory management. Popular examples are Java, C# and .NET. Before discussing the contributions of this paper, we revisit some garbage collection background and terminology.

1.1 Garbage Collection

An Appel-style generational copying collector divides the heap into two generations [2], a variable-size *nursery* and a *mature generation*. Objects are allocated from the nursery. When the nursery fills up, a *nursery collection* is triggered and the surviving objects are copied into the mature generation. When the objects are copied, the size of the mature generation is grown and the size of the nursery is reduced accordingly. Because the nursery size decreases, the time between consecutive collections also decreases and objects have less time to die. When

the nursery size drops below a given threshold, a *full-heap collection* is triggered. After a full-heap collection all free space is returned to the nursery.

In this paper we consider two flavors of generational copying collectors, namely *GenMS* and *GenCopy* from JMTk [3]. GenMS collects the mature generation using the mark-sweep garbage collection strategy. The GenCopy collector on the other hand, employs a semi-space strategy to manage its mature generation. The semi-space collector copies scanned objects, whereas the mark-sweep collector does not.

To partition the heap into generations, the collector has to keep track of references between different generations. Whenever an object in the nursery is assigned to an object in the mature generation—i.e. there is a reference from an object in the mature generation to an object in the nursery—this information is tracked by using a so-called *remembered set*. When a nursery collection is triggered the remembered set must be processed to avoid erroneously collecting nursery objects that are referenced only from the mature generation.

1.2 Paper Contributions

While garbage collection offers many benefits, the time spent reclaiming memory can account for a significant portion of the total execution time [1]. Although garbage collection research has been a hot research topic for many years, little research has been done to decide *when* and *how* garbage collectors should collect.

Garbage is typically collected when either the heap, or a generation is full. Ideally, the heap should be collected when the live ratio is low: the fewer live objects, the fewer objects need to be scanned and/or copied, the more memory there is to be reclaimed, and the longer we can postpone the next garbage collection run. In this paper, we show that collecting prior to a full heap, at points where the live ratio is low, can yield substantial reductions in GC time.

In addition, when using an Appel-style collector with two generations, a decision needs to be made whether to trigger a full-heap or nursery collection. We found that triggering nursery collections until the nursery size drops below a certain threshold is sometimes suboptimal. In this paper, we show how to trade off full-heap collections and nursery collections so that performance improves substantially.

The approach presented in this paper to decide *when* and *how* to collect, is called *garbage collection hints (GCH)* and works as follows. GCH first determines *favorable collection points (FCPs)* for a given application through offline profiling. A favorable collection point is a location in the application code where the cost of a collection is relatively cheap. During program execution a cost function is then computed in each FCP to determine the best GC strategy: postpone GC, perform a nursery GC, or perform a full-heap GC. Our experimental results using the SPECjvm98 benchmarks and two generational collectors show that GCH can reduce the garbage collector time by up to 30X and can improve the overall execution time by more than 10%.

Figure 1 perfectly illustrates why garbage collection hints actually work for the `_213_javac` benchmark. This graph shows the number of live bytes as a

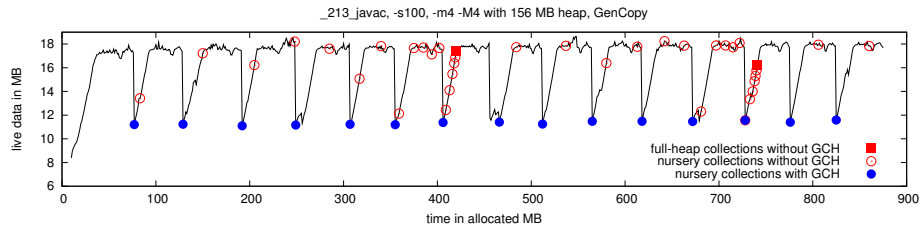


Fig. 1. Garbage collection points with and without GCH

function of the number of allocated bytes. The empty circles denote nursery collections and the squares denote full-heap collections when GCH is not enabled. Without GCH, GC is triggered at points where the number of live bytes is not necessarily low. In fact, the maximum GC time that we observed on our platform for these GC points is 225ms; and 12MB needs to be copied from the nursery to the mature generation. The GC time for a full-heap collection takes 330ms. When GCH is enabled (see the filled circles in Figure 1), garbage gets collected when the amount of live bytes reaches a minimum, i.e. at an FCP. The GC time at an FCP takes at most 4.5ms since only 126KB needs to be copied. From this example, we observe two key features why GCH actually works: (i) GCH preferably collects when the amount of live data on the heap is low, and (ii) GCH eliminates a number of full-heap collections by trading them for (cheaper) nursery collections.

The main contributions of this paper are as follows.

- We show that GC is usually not triggered when the amount of live data is low, i.e. when the amount of garbage collection work is minimal.
- We show that the collector does not always make the best decision when choosing between a nursery and a full-heap collection.
- We propose GCH which is a feedback-directed technique based on profile information that provides hints to the collector about *when* and *how* to collect. GCH tries to collect at FCPs when the amount of live data is minimal and dynamically chooses between nursery and full-heap collections. The end result is significant reductions in GC time and improved overall performance. GCH is especially beneficial for applications that exhibit a recurring phase behavior in the amount of live data allocated during program execution.

2 Experimental Setup

2.1 Java Virtual Machine

We use the Jikes Research Virtual Machine 2.3.2 (RVM) [4] on an AMD Athlon XP 1500+ at 1.3 GHz with a 256KB L2-cache, running Linux 2.6. Jikes RVM is a Java virtual machine written almost entirely in Java. Jikes RVM uses a compilation-only scheme for translating Java bytecodes to native machine instructions. For our experiments we use the *FastAdaptive* profile: all methods are

initially compiled using a baseline compiler, but sampling is used to determine which methods to recompile using an optimizing compiler.

Because Jikes RVM is written almost entirely in Java, internal objects such as those created during class loading or those created by the runtime compilers are allocated from the Java heap. Thus, unlike with conventional Java virtual machines the heap contains both application data as well as VM data. We found that there is at least 8MB of VM data that is quasi-immortal. The presence of VM data has to be taken into account when interpreting the results presented in the remainder of this work.

Jikes RVM's memory management toolkit (JMTk) [3] offers several GC schemes. While the techniques presented in this paper are generally applicable to various garbage collectors, we focus on the *GenMS* and *GenCopy* collectors. Both been used in Jikes RVM's production builds that are optimized for performance.

To get around a bug in Jikes RVM 2.3.2 we increased the maximum size of the remembered set to 256MB. In order to be able to model the shrinking/growing behavior of the heap accurately, we made one modification to the original RVM. We placed the remembered set outside the heap.

Performance is measured using the Hardware Performance Monitor (HPM) subsystem of Jikes RVM. HPM uses (i) the `perfctr`¹ Linux kernel patch, which provides a kernel module to access the processor hardware, and (ii) PAPI [5], a library to capture the processor's performance counters. The hardware performance counters keep track of the number of retired instructions, elapsed clock cycles, etc.

2.2 Benchmarks

To evaluate our mechanism, we use the SPECjvm98² benchmark suite. The SPECjvm98 benchmark suite is a client-side Java benchmark suite consisting of seven benchmarks, each with three input sets: `-s1`, `-s10` and `-s100`. With the `-m` and `-M` parameters the benchmark can be configured to run multiple times without stopping the VM. Garbage collection hints work well for long running applications that show some recurring phase behavior in the amount of live data. To mimic such workloads with SPECjvm98, we use the `-s100` input set in conjunction with running the benchmarks four times (`-m4 -M4`).

We used all SPECjvm98 benchmarks except one, namely `_222_mpegaudio`, because it merely allocates 15MB each run and triggers few GCs. The other benchmarks on the other hand, allocate a lot more.

All SPECjvm98 benchmarks are single-threaded except for `_227_mtrt` which is a multi-threaded raytracer. Note that because both Jikes RVM's sampling mechanism and the optimizing compiler run in separate threads the other benchmarks are not deterministic.

We ran all experiments with a range of different heap sizes. We vary the heap size between the minimum feasible heap size and the heap size at which our mechanism stops triggering or shows constant behavior. The larger the heap size, the less frequent garbage needs to be collected and the more time objects have to die.

¹ <http://user.it.uu.se/~mikpe/linux/perfctr/>

² <http://www.spec.org/jvm98/>

Some benchmarks like `_213_javac` use *forced garbage collections* triggered through calls to `java.lang.System.gc()`. For our experiments we disabled forced garbage collections unless stated otherwise.

3 Garbage Collection Hints

Our garbage collection hints approach consists of an offline and an online step, see Figure 2. The offline step breaks down into two parts: offline profiling of the application and garbage collector analysis. The offline profiling computes the live/time function of the application, i.e. the amount of live bytes as a function of the amount of bytes allocated. Based on this live/time function, favorable collection points (FCPs) can be determined. Determining the FCPs is a one-time cost per application. The garbage collector analysis characterizes the collection cost for a particular GC and application, i.e. the amount of time needed to process a given amount of live bytes. This is dependent on the collector and the platform on which the measurements are done. In the online part of GCH, the methods that have been identified as FCPs are instrumented to invoke a cost-benefit model that helps the garbage collector make decisions about *when* and *how* to collect. This decision making is based on the amount of heap space available, the live/time function of the application and the characteristics of the GC. The following subsections discuss GCH in more detail.

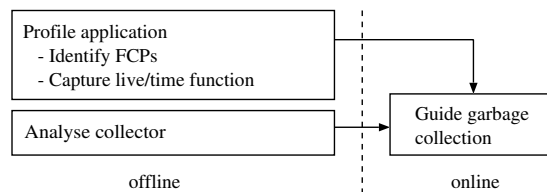


Fig. 2. An overview of the GCH methodology

3.1 Program Analysis

Live/Dead Ratio Behavior. The first step of the offline profiling is to collect the live/time function which quantifies the number of live bytes as a function of the bytes allocated so far. Moreover, we are interested in linking the live/time function to methods calls. We modified Jikes RVM to timestamp and report all method entries and exits. For each method invocation, we want to know how many objects/bytes died and how many objects are live. Therefore, a lifetime analysis is required at every point an object could have died. There are two reasons for an object to die: (i) an object's last reference is overwritten as a result of an assignment operation, or (ii) an object's last reference is on a stack frame and the stack frame gets popped because the frame's method returns or because an exception is thrown. To avoid having to do a lifetime analysis for

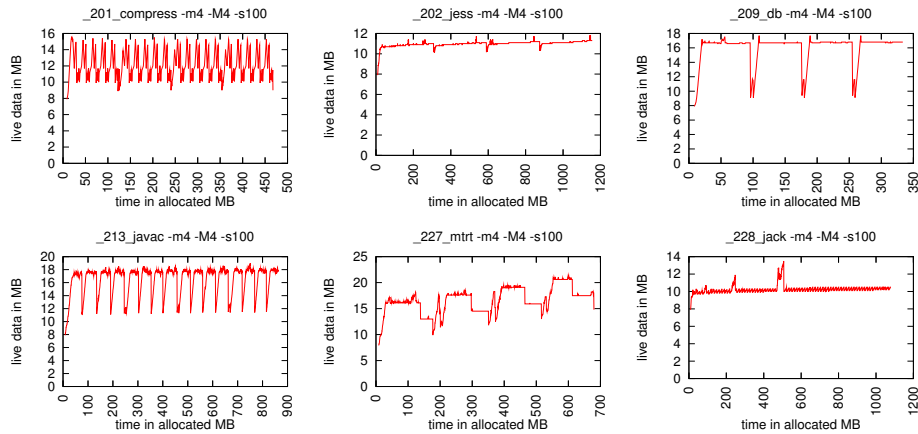


Fig. 3. The live/time function for the various benchmarks: number of live bytes as a function of the number of bytes allocated

every assignment operation, method return and exception, we used a modified version of the Merlin trace generator [6] that is part of Jikes RVM. Merlin is a tool that precisely computes every object’s last reachable time. It has been modified to use our alternative timestamping method to correlate object death with method invocations.

Figure 3 shows the live/time function for the various benchmarks. As can be seen from these graphs, the live/time function shows recurring phase behavior. This recurring phase behavior will be exploited through GCH. Applications that do not exhibit a phased live/time function are not likely to benefit from GCH. Next, the live/time function is used to select FCPs and to compute the FCP live/time patterns.

Favorable Collection Points. For a method to represent a favorable collection point (FCP), it needs to satisfy three criteria:

1. An FCP’s invocation should correspond to a local minimum in terms of the number of live bytes. In other words, we need to select methods that are executed in the minima of the live/time function. This will allow GCH to collect garbage with minimal effort.
2. An FCP should not be executed frequently. To minimize the overhead of the instrumentation code, FCPs that represent cold methods are preferred. A method that gets executed *only* in local minima is an ideal FCP.
3. The live/time pattern following the execution of the FCP should be fairly predictable, i.e. each time the FCP gets executed, the live/time function should have a more or less similar shape for some range after the FCP.

Given the live/time function, selecting FCPs is fairly straightforward. Table 1 shows the selected FCPs that we selected for the SPECjvm98 benchmarks. Some

Table 1. The selected FCPs for each of the benchmark applications

Benchmark	Favorable collection points
_201_compress	spec.io.FileInputStream.getLength()I
_202_jess	spec.benchmarks._202_jess.jess.undefrule.<init>()V
_209_db	spec.harness.BenchmarkTime.toString()Ljava/lang/String; spec.harness.Context.setBenchmarkRelPath(Ljava/lang/String;)V
_213_javac	spec.io.FileInputStream.getCachingtime()J
_227_mtrt	spec.benchmarks._213_javac.ClassPath.<init>(Ljava/lang/String;)V spec.io.TableOfExistingFiles.<init>()V
_228_jack	spec.harness.Context.clearIOtime()V spec.io.FileInputStream.getCachingtime()J spec.benchmarks._228_jack.Jack.the.Parser.GeneratorInternals.-compare(Ljava/lang/String;Ljava/lang/String;)V

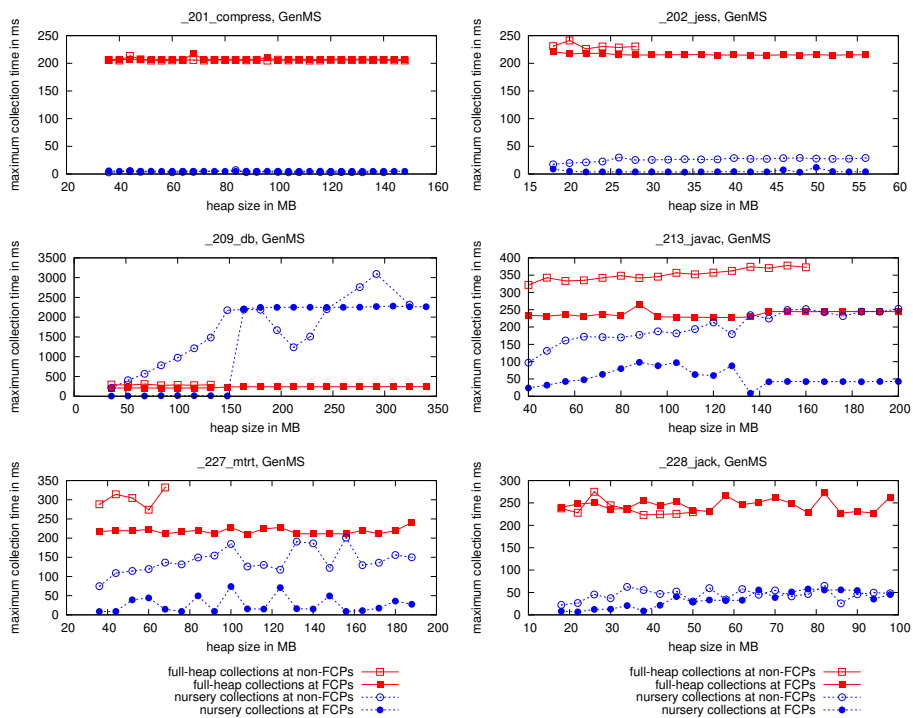


Fig. 4. The maximum garbage collection times across different heap sizes for each of the different scenarios

benchmarks have only one FCP (see for example Figure 1 for `_213_javac`); others such as `_227_mtrt` have three FCPs.

To illustrate the potential benefit of FCPs, Figure 4 plots the maximum time spent in GC when triggered at an FCP and when triggered in a non-FCP. We make a distinction between full-heap and nursery collections, and plot data for a range of heap sizes. For most benchmarks we observe that the maximum GC time spent at an FCP is substantially lower than the GC time in a non-FCP. This

reinforces our assumption that collecting at an FCP is cheaper than collecting in a non-FCP. However, there are two exceptions, `_201_compress` and `_228_jack`, for which GC time is insensitive to FCPs. For `_201_compress`, this is explained by the fact that the live/time function shown in Figure 3 is due to a few objects that are allocated in the Large Object Space (LOS). Because objects in LOS never get copied, GCH cannot reduce the copy cost. Furthermore, because there are only a few such objects it will not affect the scan cost either. For `_228_jack`, the height of the live/time function’s peaks is very low, see Figure 3. Because `_201_compress` and `_228_jack` are insensitive to FCPs we exclude them from the other results that will be presented in this paper. (In fact, we applied GCH to these benchmarks but the result was a neutral impact of overall performance. Due to space constraints, we do not to include these benchmarks in the rest of this paper.)

It is also interesting to note that for `_209_db`, a nursery collection can be more costly than a full-heap collection. This is due to the fact that the remembered set needs to be scanned on a nursery collection. As such, for `_209_db` a full-heap collection can be more efficient than a nursery collection. This is exploited through GCH.

FCP’s Live/Time Pattern. For each unique FCP we have to capture the *live/time pattern* following the FCP. This is a slice of the live/time function following the FCP that recurs throughout the complete program execution. We sample the FCP’s live/time pattern at a frequency of one sample per 0.5MB of allocated memory and use it as input for the cost-benefit model. An FCP’s live/time pattern is independent of the heap size (the same information is used for all heap sizes) and is independent of the collection scheme (the same information is used for both GenMS and GenCopy). It only needs to be computed once for each benchmark. For completeness, we add that in order to be independent of the collection scheme, we differentiate between data that is eligible to be copied and data that cannot be copied.

3.2 Collector Analysis

So far, we discussed the offline profiling that is required for GCH. We now discuss the characterization of the garbage collector. This characterization will be used in the cost model that will drive the decision making in GCH during program execution. The characterization of the garbage collector quantifies the cost of a collection. We identify three cost sources: the cost of a full-heap collection, the cost of a nursery collection and the cost of processing the remembered set. The cost functions take as input the amount of live data and output the estimated collection time. These cost functions are dependent on the application, the collector and the given platform (VM, microprocessor, etc.).

Figure 5 shows how the cost functions are to be determined for the GenMS and GenCopy collectors. The graphs are obtained by running the benchmarks multiple times with different heap sizes using instrumented collectors. In these graphs we make a distinction between nursery collections, full-heap collections

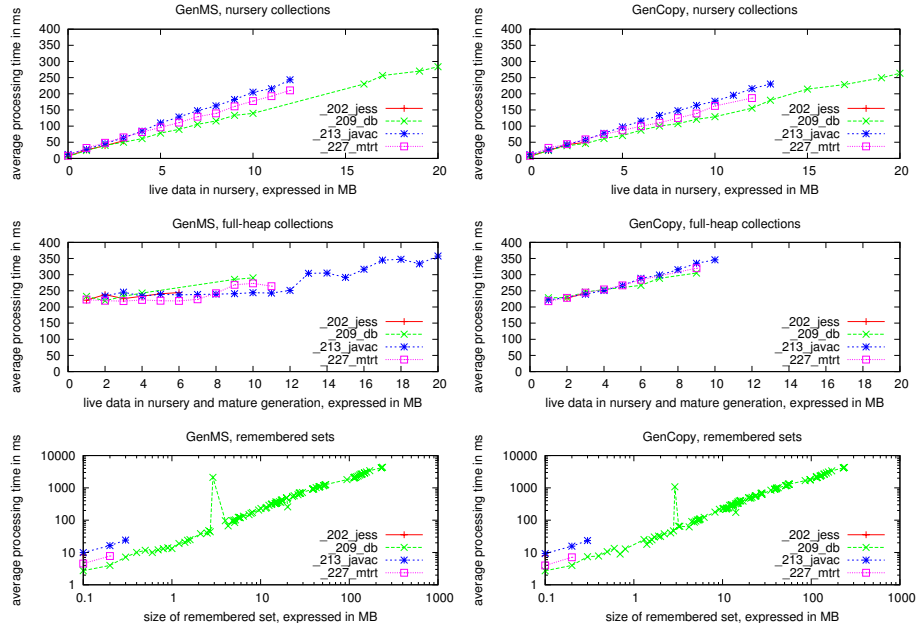


Fig. 5. The cost of a nursery and full-heap collection in terms of the amount of copied/live data

and processing of the remembered set. Hence, the processing times on the nursery collection graphs do not include the time required to process the remembered sets.

GC time appears to be a linear function of the amount of live data for both collectors. In other words, the scanning and copying cost is proportional to the amount of live bytes. Likewise, the processing cost of the remembered set appears to be a linear function of its size. In summary, we can compute linear functions that quantify the cost of a nursery collection, full-heap collection and processing of the remembered set.

In this paper we employ application-specific cost functions, i.e. we compute cost functions per application. In fact, on a specialized system with dedicated long running applications, it is appropriate to consider a cost function that is specifically tuned for the given application. Nevertheless, given the fact that the cost functions appear to be fairly similar across the various applications, see Figure 5, choosing application-independent cost functions could be a viable scenario for general-purpose environments. Such a scenario is likely to produce results that are only slightly worse compared to the ones presented in this paper.

3.3 GCH at Work

The information that is collected through our offline analysis is now communicated to the VM to be used at runtime. Jikes RVM reads all profile information

at startup. This contains (i) a list of methods that represent the FCPs, (ii) the live/time pattern per FCP, and (iii) the cost functions for the given garbage collector. Jikes RVM is also modified to dynamically instrument the FCPs. The instrumentation code added to the FCPs examines whether the current FCP should trigger a GC. The decision to collect or not is a difficult one as there exists a trade-off between reducing the amount of work per collection and having to collect more frequently. Clearly, triggering a collection will have an effect on subsequent collections. Because GC is invoked sooner due to GCH than without GCH, additional collections might get introduced. On the other hand, triggering a collection at an FCP can help reduce the GC overhead. A collection at an FCP will generally introduce modest pause times compared to collections at a non-FCPs. Moreover, triggering a full-heap collection grows the nursery size and gives objects more time to die, while triggering a nursery collection when few objects are live will result in the mature generation filling up slower, reducing the need for more expensive full-heap collections.

To make this complex trade-off, the instrumentation code in the FCPs implements an *analytical cost-benefit model*. The cost-benefit model estimates the total GC time for getting from the current FCP to the end of its FCP's live/time pattern. The cost-benefit model considers the following three scenarios: (i) do not trigger a GC in the current FCP, (ii) trigger a full-heap GC, or (iii) trigger a nursery GC. For each of these three scenarios, the cost-benefit model computes the total GC time by *analytically simulating* how the heap will evolve through the FCP's live/time pattern. This is done as follows. First, the cost-benefit model computes the GC cost *in the current FCP* under the three scenarios:

- (i) The cost for not triggering a GC is obviously zero. The available heap size remains unchanged.
- (ii) For computing the cost for triggering a full-heap collection in the current FCP, we first calculate the number of live bytes at the current FCP. We get this information from the live/time pattern. We subsequently use the full-heap GC cost function to compute the GC time given the amount of live data in the current FCP. The available heap size after the current (hypothetical) full-heap collection then equals the maximum heap size minus the amount of live data in the current FCP.
- (iii) To compute the cost for triggering a nursery GC in the current FCP, we assume that the amount of live bytes in the nursery at that FCP is close to zero. The GC cost is then computed based on the nursery-GC cost function. This GC cost is incremented by an extra cost due to processing the remembered set. This extra cost is proportional to the size of the remembered set, which is known at runtime at an FCP. The heap size that was occupied by the nursery becomes available again for allocation.

In the second step of the cost-benefit model we compute the cost of additional collections over the FCP's live/time pattern for each of the three scenarios. In fact, for each scenario, the cost-benefit model analytically simulates how the heap will evolve over time when going through an FCP's live/time pattern. Therefore, we compute when the (nursery) heap will be full—when the

application has allocated all memory available in the heap. In case the system would normally trigger a full collection (i.e. when the nursery size drops below the given threshold), we need to compute the cost of a full-heap collection. This is done the same way as above, by getting the amount of live data from the FCP's live/time pattern—note that we linearly interpolate the live/time pattern—and use the major-GC cost function to compute its cost. In case the nursery size is above the given threshold, we need to compute the cost of a nursery collection. Computing the cost for a nursery collection is done by reading the number of live bytes from the FCP's live/time pattern and subtracting the number of live bytes in the previous GC point; this number gives us an estimate for the amount of live data in the nursery. This estimated amount of live nursery data is used through the nursery-GC cost function to compute an estimated nursery-GC cost.

When the end of the FCP's live/time pattern is reached within the model, an end cost is added to the aggregated GC cost calculated so far. This end cost is proportional to the fraction of used nursery space and indicates that the next (expected) collection will be close by or far away. After computing all the costs for each of the three scenarios, the scenario that results in the minimal total cost is chosen.

Note that the cost-benefit model presented above is specifically developed for GenMS and GenCopy, two Appel-style generational garbage collectors with a variable nursery size. However, a cost-benefit model could also be easily constructed for other collectors. For example, modifying this model for generational collectors with a fixed nursery, or non-generational collectors should not be that difficult.

3.4 GCH Across Inputs

GCH is a profile-driven garbage collection method which implies that the input used for the offline profiling run is typically different from the input used during online execution. Getting GCH to work across inputs needs a few enhancements since the size and the height of an FCP's live/time pattern varies across inputs; the general shape of the FCP's live/time pattern however is relatively insensitive to the given input. We define the height of an FCP's live/time pattern as the difference in live data at the top of an FCP's live/time pattern and at the FCP itself. For example, for `_213_javac`, see Figure 1, the height is approximately 6MB. The size of an FCP's live/time pattern is defined as the number of allocated bytes at an FCP's live/time pattern; this is approximately 60MB for `_213_javac`, see also Figure 1. To address this cross-input issue we just scale the size and the height of the FCP live/time pattern. In practice, the amount of live data at the top of an FCP's live/time pattern can be computed at runtime when a GC is triggered at the top of an FCP's live/time pattern. The amount of allocated bytes over an FCP's live/time pattern can be computed at run-time as well by just querying the VM. These scaling factors are then to be used to rescale the data in the FCP live/time pattern.

4 Experimental Results

4.1 Garbage Collection Time

In order to evaluate the applicability of GCH, we have set up the following experiment. We used the profile information from the `-s10` run to drive the execution of the `-s100` run after cross-input rescaling as discussed in section 3.4.

Figure 6 shows the reduction through GCH in GC time over a range of heap sizes. (Reduction is defined as $100 \times (1 - \frac{time_{GCH}}{time_{old}})$. Thus, a 50% reduction means GC time is halved.) Each reduction number is an average number over three runs; numbers are shown for both GenMS and GenCopy. Figure shows that GCH improves GC time for both collectors and for nearly all heap sizes. For both collectors GCH achieves substantial speedups, up to 30X for `_213_javac` and 8X for `_209_db`.

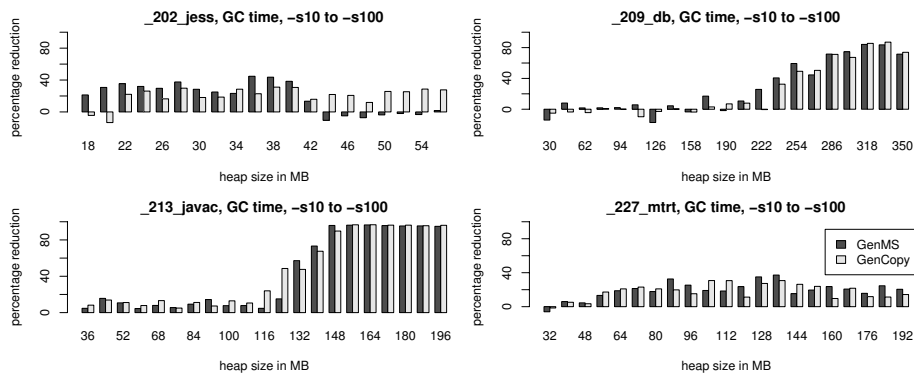


Fig. 6. Reduction in garbage collection time through GCH across inputs. The profile input is `-s10`; the reported results are for `-s100`.

The sources for these huge speedups are twofold. First, GCH generally results in fewer (nursery) collections than without GCH, see Table 2 which shows the average number of GCs over all heap sizes; we observe fewer GCs with GCH for all benchmarks except one, namely `_209_db` for which the number of collections remains unchanged with and without GCH (we will discuss `_209_db` later on). For `_213_javac` we observe a 30% reduction in the number of (nursery) collections. The second reason for these huge GC time speedups, is the substantial reduction in the amount of time spent per GC. This was already mentioned in Figure 4.

Note that also for `_209_db`, GCH is capable of reducing the GC time substantially, especially for large heap sizes. The reason is not the reduced number of collections, but the intelligent selection of full-heap collections instead of nursery collections. The underlying reason is that `_209_db` suffers from a very large remembered set. GCH triggers more full-heap collections that do not suffer from processing the remembered set, see Table 2. A full-heap collection typically only takes 250ms whereas a nursery collection can take up to 2,000ms, see Figure 4.

Table 2. The average number of garbage collections across all heap sizes with and without GCH

Benchmark	GenMS collector				GenCopy collector			
	Full-heap		Nursery		Full-heap		Nursery	
	no GCH	GCH	no GCH	GCH	no GCH	GCH	no GCH	GCH
<code>_202_jess</code>	0	1	245	186	2	3	349	294
<code>_209_db</code>	1	3	16	14	2	4	25	25
<code>_213_javac</code>	2	2	80	62	6	5	93	61
<code>_227_mtrt</code>	0	1	45	36	2	2	81	67

Note that the remembered set increases with larger heap sizes which explains the increased speedup for larger heap sizes. The two negative speedup results for the GenCopy collector for a 110MB and 140MB heap are due to particularities in the time-varying behavior of the remembered set in `_209_db`. Our general cost model of the remembered set works fairly well across all programs but does not cover all peculiarities observed in `_209_db`. However, through an experiment we were able to tweak the cost model of the remembered set to `_209_db` which results in an overall GC time speedup (not included here due to space constraints). While the large remembered sets themselves are the consequence of the fact that JMTk uses sequential store buffers without a space cap, it shows that our analytic framework is robust in the face of extreme cases like this. This does not affect the generality of our approach since in many practical situations the cost model can be designed for a given application in a specialized environment.

Remind that because of the way Jikes RVM works, the heap contains both application data and VM data. We believe that our technique would be even more effective in a system where the collector does not have to trace VM data. In such a system, full-heap collections would generally be cheaper opening up extra opportunities to replace nursery collections by full-heap collections.

4.2 Overall Execution Time

Figure 7 depicts the impact of GCH on the overall execution time. For some benchmarks and heap sizes, the total execution time is more or less unaffected through GCH because the time spent collecting garbage is only a small fraction of the total execution time. This is the case for `_227_mtrt`. The small slowdowns or speedups observed are probably due to changed data locality behavior because of the changed GCs. However, for `_202_jess`, `_213_javac` and `_209_db` performance improved up to 5.7%, 12.5% and 10.5% respectively. For these benchmarks, the huge GC time speedups translate themselves in overall performance speedup.

4.3 Run-Time System Overhead

To explore the run-time overhead of our system, we compare the performance of a without-GCH Jikes RVM versus a with-GCH Jikes RVM. In the with-GCH version, the profile information is read, the FCPs are instrumented and at each

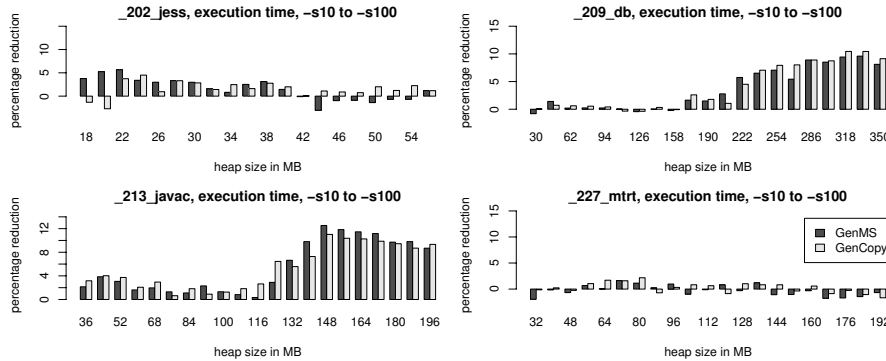


Fig. 7. Performance improvement in total execution time through GCH across inputs. The profile input is `-s10`; the reported results are for `-s100`.

invocation of an FCP the cost-benefit model is computed, however, it will never trigger a collection. For computing the overhead per benchmark, each benchmark is run multiple times and the average overhead is computed over these runs. The average overhead over all benchmarks is 0.3%; the maximum overhead is 1.3% for `_227_mtrt`.

5 Related Work

We now discuss previously proposed GC strategies that are somehow related to GCH, i.e. all these approaches implement a mechanism to decide when *or* how to collect. The work presented in this paper differs from previous work in that we combine the decision of both when *and* how to collect in a single framework.

The Boehm-Demers-Weiser (BDW) [7] garbage collector and memory allocator include a mechanism that determines whether to collect garbage or to grow the heap. The decision whether to collect or grow the heap is based on a static variable called the *free space divisor (FSD)*. If the amount of heap space allocated since the last garbage collection exceeds the heap size divided by FSD, garbage is collected. If not, the heap is grown. Brecht *et al.* [8] extended the BDW collector by taking into account the amount of physical memory available and by proposing dynamically varying thresholds for triggering collections and heap growths.

Wilson *et al.* [9] observe that (interactive) programs have phases of operation that are compute-bound. They suggest that tagging garbage collection onto the end of larger computational pauses, will not make those pauses significantly more disruptive. While the main goal of their work is to avoid or mitigate disruptive pauses, they reason that at these points, live data is likely to be relatively small since objects representing intermediate results of the previous computations have become garbage. They refer to this mechanism as *scavenge scheduling* but present no results.

More recently, Ding *et al.* [10] presented preliminary results of a garbage collection scheme called *preventive memory management* that also aims to exploit

phase behavior. They unconditionally force a garbage collection at the beginning of certain execution phases. In addition, they avoid garbage collections in the middle of a phase by growing the heap size unless the heap size reaches the hard upperbound of the available memory. They evaluated their idea using a single Lisp program and measured performance improvements up to 44%.

Detlefs *et al.* [11] present the garbage-first garbage collectors which aims at satisfying soft real-time constraints. Their goal is to spend no more than x ms during garbage collection for each y ms interval. This is done by using a collector that uses many small spaces and a *concurrent marker* that keeps track of the amount of live data per space. The regions containing most garbage are then collected first. In addition, collection can be delayed in their system if they risk violating the real-time goal.

Velasco *et al.* [12] propose a mechanism that dynamically tunes the size of the copy reserve of an Appel collector [2]. Tuning the copy reserve's size is done based on the ratio of surviving objects after garbage collection. Their technique achieves performance improvements of up to 7%.

Stefanovic *et al.* [13] evaluate the older-first generational garbage collector which only copies the oldest objects in the nursery to the mature generation. The youngest objects are not copied yet; they are given enough time to die in the nursery. This could be viewed of as a way deciding when to collect.

Recent work [14,15,16] selects the most appropriate garbage collector during program execution out of a set of available garbage collectors. As such, the garbage collector is made adaptive to the program's dynamic execution behavior. The way GCH triggers nursery or full-heap collections could be viewed as a special form of what these papers proposed.

6 Summary and Future Work

This paper presented garbage collection hints which is a profile-directed approach to guide garbage collection. The goal of GCH is to guide in terms of *when* and *how* to collect. GCH uses offline profiling to identify favorable collection points in the program code where the amount of live data is relatively small (in order to reduce the amount of work per collection) and the amount of dead bytes is relatively large (in order to increase the amount of available heap after collection). Triggering collections in these FCPs can reduce the number of collections as well as the amount of work per collection. Next to guiding when to collect, GCH also uses an analytical cost-benefit model to decide how to collect, i.e. whether to trigger a nursery or a full-heap collection. This decision is made based on the available heap size, and the cost for nursery and full-heap collections. Our experimental results using SPECjvm98 showed substantial reductions in GC time (up to 30X) and significant overall performance improvements (more than 10%).

In future work, we plan to extend and evaluate GCH for other collectors than the ones considered here. We also plan to study dynamically inserted garbage collection hints in which profiling is done online during program execution.

Acknowledgments

Dries Buytaert is supported by a grant from the Institute for the Promotion of Innovation by Science and Technology in Flanders (IWT). Kris Venstermans is supported by a BOF grant from Ghent University, Lieven Eeckhout is a Postdoctoral Fellow of the Fund for Scientific Research—Flanders (Belgium) (FWO—Vlaanderen). We thank the reviewers for their insightful comments.

References

1. Blackburn, S.M., Cheng, P., McKinley, K.S.: Myths and realities: the performance impact of garbage collection. In: Proceedings of SIGMETRICS'04, ACM (2004)
2. Appel, A.W.: Simple generational garbage collection and fast allocation. *Software practices and experience* **19** (1989) 171–183
3. Blackburn, S.M., Cheng, P., McKinley, K.S.: Oil and water? High performance garbage collection in Java with JMTk. In: Proceedings of ICSE'04. (2004) 137–146
4. Alpern, B., Attanasio, C.R., Barton, J.J., Burke, M.G., Cheng, P., Choi, J.D., Cocchi, A., Fink, S.J., Grove, D., Hind, M., Hummel, S.F., Lieber, D., Litvinov, V., Mergen, M.F., Ngo, T., Russell, J.R., Sarkar, V., Serrano, M.J., Shepherd, J.C., Smith, S.E., Sreedhar, V.C., Srinivasan, H., Whaley, J.: The Jalapeño Virtual Machine. *IBM Systems Journal* **39** (2000) 211–238
5. Browne, S., Dongarra, J., Garner, N., Ho, G., Mucci, P.: A portable programming interface for performance evaluation on modern processors. *The international journal of high performance computing applications* **14** (2000) 189–204
6. Hertz, M., Blackburn, S.M., Moss, J.E.B., McKinley, K.S., Stefanovic, D.: Error free garbage collection traces: how to cheat and not get caught. In: Proceedings of SIGMETRICS'02, ACM (2002) 140–151
7. Boehm, H., Weiser, M.: Garbage collection in an uncooperative environment. *Software practices and experience* **18** (1988) 807–820
8. Brecht, T., Arjomandi, E., Li, C., Pham, H.: Controlling garbage collection and heap growth to reduce the execution time of Java applications. In: Proceedings of OOPSLA'01, ACM (2001) 353–366
9. Wilson, P.R., Moher, T.G.: Design of the opportunistic garbage collector. In: Proceedings of OOPSLA'89, ACM (1989) 23–35
10. Ding, C., Zhang, C., Shen, X., Ogihara, M.: Gated memory control for memory monitoring, leak detection and garbage collection. In: Proceedings of MSP'05, ACM (2005) 62–67
11. Detlefs, D., Flood, C., Heller, S., Printezis, T.: Garbage-first garbage collection. In: Proceedings of ISMM'04, ACM (2004) 37–48
12. J. M. Velasco, K. Olcoz, F.T.: Adaptive tuning of reserved space in an Appel collector. In: Proceedings of ECOOP'04, ACM (2004) 543–559
13. Stefanovic, D., Hertz, M., Blackburn, S.M., McKinley, K.S., Moss, J.E.B.: Older-first garbage collection in practice: evaluation in a Java virtual machine. In: Proceedings of MSP'02, ACM (2002) 25–36
14. Andreasson, E., Hoffmann, F., Lindholm, O.: Memory management through machine learning: to collect or not to collect? In: Proceedings of JVM'02, USENIX (2002)
15. Printezis, T.: Hot-swapping between a mark&sweep and a mark&compact garbage collector in a generational environment. In: Proceedings of JVM'01, USENIX (2001)
16. Soman, S., Krintz, C., Bacon, D.F.: Dynamic selection of application-specific garbage collectors. In: Proceedings of ISMM'04, ACM (2004) 49–60