Protection traps and alternatives for memory management of an object-oriented language.*

Antony L. Hosking   J. Eliot B. Moss

Object Systems Laboratory
Department of Computer Science
University of Massachusetts
Amherst, MA 01003

Abstract

Many operating systems allow user programs to specify the protection level (inaccessible, read-only, read-write) of pages in their virtual memory address space, and to handle any protection violations that may occur. Such page-protection techniques have been exploited by several user-level algorithms for applications including generational garbage collection and persistent stores. Unfortunately, modern hardware has made efficient handling of page protection faults more difficult. Moreover, page-sized granularity may not match the natural granularity of a given application. In light of these problems, we reevaluate the usefulness of page-protection primitives in such applications, by comparing the performance of implementations that make use of the primitives with others that do not. Our results show that for certain applications software solutions outperform solutions that rely on page-protection or other related virtual memory primitives.

1 Introduction

Paged virtual memory mechanisms perform admirably when put to their intended purpose, which is to extend the address space of user programs beyond the physical memory of the machine, and for protection from other processes in multiprogrammed systems. Hardware and operating system software have been refined to achieve this sleight of hand with performance broadly acceptable to most applications.

Many operating systems now allow user-level programs to exploit virtual memory mechanisms for their own purposes by providing primitives to manipulate page protections (inaccessible, read-only, read-write). User programs can also provide "handlers" to to be invoked in the event of an access violation. Using these primitives, user-level applications are able to monitor access to any of the pages in their virtual address space without explicit checks, by exploiting the paging hardware's ability to trap on access violations. As a result, application programmers have exercised their ingenuity in devising implementation solutions that make use of these virtual memory primitives. A number of these applications are enumerated by Appel and Li [3], where they argue that in light of programmers demands, designers of operating systems and hardware architectures must pay more attention to support for virtual memory primitives to make their implementations more efficient and robust.

Meanwhile, there is evidence [1] to indicate that the evolution of architectures towards pipelined RISC microprocessors, and operating systems towards micro-kernels, is making efficient implementation of these operating system primitives more difficult. As a result of this tension between the supposed demand from application programmers and the evolutionary trends of architectures and operating systems, we examine two applications cited by Appel and Li as benefiting from the availability of virtual memory primitives.

The contributions of this paper include a comprehensive performance evaluation of page-protection primitives for garbage collection and persistence, and direct comparison with corresponding software implementations. The nature of our experimental setup allows meaningful direct comparison. In addition, we project the optimal performance of applications that exploit page-protection techniques in their implementation. Our results indicate that alternative software implementations can approach, and in some cases outperform, the optimal performance of page-protection implementations.

The rest of the paper is organized as follows. In the next section we briefly describe the applications we examine, and how they are able to take advantage of virtual memory primitives. We then present the experimental setup used for gathering performance data, and our alternative implementations

of each of the applications along with their their relative performance. Finally, we summarize the major points of the paper and present our conclusions.

2 Applications

Appel and Li [3] describe a number of applications of virtual memory primitives, including concurrent garbage collection, shared virtual memory, concurrent checkpointing, generational garbage collection, persistent stores, extending addressability, data-compression paging, and heap overflow detection. Of these, we directly address generational garbage collection, and certain aspects of persistent stores, including object fault handling, database checkpointing, and extending addressability. Our results also have implications for the other applications, since they show that well-designed software solutions are competitive with hardware-assisted techniques.

2.1 Generational garbage collection

Generational garbage collectors [16, 24, 25] achieve short collection pause times partly because they separate heap-allocated objects into two or more generations and do not process all generations during each collection. Empirical studies have shown that in many programs most objects die young, so separating objects by age and focusing collection effort on the younger generations is a popular strategy. However, *any* collection scheme that processes only a small portion of the heap must somehow know or discover all pointers outside the collected area that refer to objects within the collected area. Since the areas *not* collected are generally assumed to be large, most generational collectors employ some sort of pointer tracking scheme, to avoid scanning the uncollected areas. Again, empirical studies show that in many programs, the older-to-younger pointers of interest to generational collection are rare, so avoiding scanning presumably improves performance. This is intuitively explained by the fact that newly allocated objects can only be immediately initialized to point to pre-existing (i.e., older) objects. Pointers from older generations to younger generations can be created only through assignment to pre-existing objects. Detecting such assignments requires special action at every pointer assignment to see whether that pointer must now be considered by the garbage collector when collecting the younger generations.

A number of schemes have been suggested for generating and maintaining the older-to-younger pointer information needed by generational collectors, including special-purpose hardware support [24, 25] and generation by compilers of the necessary inline code to perform the checks in software [2] (adding to the overhead of pointer stores). Ungar [24, 25] uses *remembered sets* to maintain the necessary information on a per-generation basis, recording the locations in older generations that may contain pointers into the generation.

The garbage collector examines all the locations recorded in the remembered sets of the younger generations being collected to determine the live (i.e., reachable) objects.

Alternatively, dirty bits can be maintained for older generations indicating whether the generation contains pointers to objects in younger generations. The heap is divided into aligned logical regions of size 2^k bytes—the address of the first byte in the region will have k low bits zero. These regions are called *cards* [23, 29]. Each card has a corresponding entry in a table indicating whether the card might contain a pointer of interest to the garbage collector. Mapping an address to an entry in the table involves shifting the address right by k bits and using the result to index the table.

The card table can be maintained explicitly by generating code to index and dirty the corresponding table entry at every store site in the program. Alternatively, by setting the card size to correspond to the virtual memory page size, updates to clean cards can be detected using the virtual memory hardware. All clean pages in the heap are protected from writes. When a write occurs to a protected page, the trap handler records the update in the card table and unprotects the page. Subsequent writes to the now dirty page incur no further overhead. Note that *all* writes to a clean page cause a protection trap, not just those that store pointers.

The time required to determine the relevant older-to-younger pointers for garbage collection varies with the granularity of the information recorded [11]. Remembered sets have the advantage of recording just those locations that can possibly contain older-to-younger pointers. In contrast, the time to scan dirty cards is proportional to the size of the cards. While software-implemented card marking schemes are free to choose any power of two for the card size, a page trapping scheme is bound by the size of a virtual memory page. Since modern operating systems and architectures typically use a relatively large virtual memory page size (on the order of thousands of bytes), scanning overheads will be proportionally higher.

2.1.1 User-level dirty bits

If operating systems were to provide user-level dirty bits (as suggested by Shaw [21], and Appel and Li [3]), the overhead to reflect page traps through to the user-level protection violation handler can be avoided. Presumably, an extra user-level dirty bit would be added to each page table entry, and a system call (`dirty`) provided to return a list of pages dirtied in a given address range since the last time it was called. The system call would clear the user-level dirty bits and enable traps on the specified pages. Traps could then be handled directly in the operating system. This can have substantial savings. As reported for a MIPS R2000 [1], the time for a user program to trap to a null C routine in the *kernel* and return to the user program is 15.4μs round trip. In contrast, Appel and Li report the corresponding overhead to handle page-fault traps in *user-mode* to be 210μs on a DECstation 3100 (MIPS R2000) running Ultrix 4.1. We have confirmed

(a) Fault blocks stand in for non-resident objects

(b) Target object faulted in

(c) Indirect block bypassed by garbage collector

○ non-resident object   □ fault block

● resident object   ▨ indirect block

Figure 1: Fault Blocks

this with our own measurement of page traps in a tight loop using the same hardware and operating system configuration, obtaining a round-trip time of $\sim250\mu s$. Note that these measurements are for a tight loop executing many repetitions, and so may tend to underestimate trap costs. Traps interspersed throughout a program's normal execution may perform less favorably, since the OS trap handling code and data structures needed to service the trap may no longer be in the hardware caches. Meanwhile, a call to dirty should be no more expensive than current primitives for manipulating page protections, except in copying out the dirty bit information, adding little if any extra overhead to applications that use the new primitive.

2.2 Persistent stores

A persistent store is a dynamic allocation heap that persists from one program invocation to the next [5, 6]. Persistent programming languages allow traversal of the data structures in a persistent store to be programmed transparently, without the need for complicated I/O or database calls to retrieve the data. Rather, the objects in the persistent store are faulted into memory on demand much as non-resident pages are automatically made resident by the virtual memory system. Moreover, a persistent program may modify the objects in the store, and commit these modifications so that their effects are permanent.

We consider three aspects in the implementation of persistence: detecting and handling object faults, extending addressability, and checkpointing of modifications.

2.2.1 Detecting and handling object faults

A persistent program may refer to both resident and non-resident persistent objects. Ideally, a memory-resident persistent object will be referred to by its virtual address, so that accessing the object can be as fast as accessing a non-persistent object. If the program traverses a reference to a non-resident object then it must be made available to the program in memory: we call this an *object fault*.

Tagging the references is one way to distinguish between references to resident and non-resident objects. An untagged reference is a direct memory pointer to the object in memory. A tagged reference contains an *object identifier* (OID) sufficient to locate the object on stable storage. By aligning resident objects on word boundaries, there are sufficient bits in a word for the tag. Every time a reference is traversed the tag is checked to make sure it points to a resident object; if it does not then an object fault is triggered.

An alternative is to use direct pointers for all object references, and to have resident proxy objects (we call them *fault blocks*) stand in for non-resident objects, as illustrated in Figure 1(a). A fault block contains the OID of the target object, and is tagged to distinguish it from an ordinary object. Whenever a pointer is followed, if it refers to a fault block, then an object fault is triggered. The target object is made resident and any pointers it contains are converted to direct pointers to resident objects or fault blocks. The fault block is changed to contained a tagged pointer to the now-resident object (see Figure 1(b)). We call the updated fault block an *indirect block*. If a traversed pointer refers to an indirect block then the target object can be located at the cost of an indirection. Occasional scanning (possibly by a garbage collector) can be used to bypass indirect blocks, as shown in Figure 1(c).

108

Object oriented programming languages can exploit the indirection implicit in the method invocation mechanism to fold residency checks into the overhead of method invocation. Method code can then directly access the fields of the object on which the method was invoked without performing further residency checks.

Rather than performing residency checks in software at each pointer traversal, fault blocks can be allocated in protected virtual memory pages so that dereferencing a pointer to a fault block is trapped. A residency check is thus implemented as a single load instruction. The protection violation handler unprotects the pages containing fault and indirect blocks, overwrites the offending fault block with an indirect block, reprotects the page, and arranges for the load instruction that caused the fault to be restarted with a direct pointer to the resident object. The fault and indirect block pages are then reprotected before resuming execution of the program.

Other approaches use page protections in a different way [15, 20, 22, 28]. When a given persistent object is to be assigned a virtual address, a page of virtual memory is reserved (although not necessarily allocated) for the page in the persistent store that contains the object. The offset of the object in the persistent page is known, allowing the virtual address of the object in the reserved virtual memory page to be calculated. Accessing the page triggers a virtual memory page trap. The trap handler reads the persistent page from the store and maps it into the previously reserved virtual page. All of the pointers in the page are then converted to direct virtual memory pointers, reserving virtual memory pages for the objects to which they refer if those objects are not already mapped into virtual memory. The faulted page is unprotected, and execution resumes. As execution proceeds, pages are reserved in a "wave-front" just ahead of the most recently faulted and swizzled pages, guaranteeing that the program will only ever see virtual memory addresses.

2.2.2 Extending addressability

Persistent stores may grow so large that they contain more objects than can be addressed directly by the available hardware.[1] Dealing with this problem involves converting persistent store OIDs into virtual memory addresses, a process which has been termed *swizzling* [19]. This technique originated in early attempts to extend the address space of Smalltalk-80 [12, 13]. In any case, it relies on an OID-to-virtual-address mapping, maintained in our case by the object store software on behalf of the application.

2.2.3 Database checkpointing

Modifications made to persistent data by a persistent program become permanent only when some sort of *checkpoint* operation is invoked, perhaps as the result of a database

---

[1] The recent arrival of 64-bit machines addresses this problem, but there are other good reasons to have different formats in the persistent store and in virtual memory.

transaction commit. Given an application that modifies only a small fraction of the resident data, writing all the data back to the stable database will be hopelessly inefficient. Instead, checkpoints can *log* just those parts of the database that have been changed, allowing programs to continue execution with minimal delay. The log records can be incorporated into the database at a later time, possibly by some process running in the background. Thus, in the face of a system crash all modifications since the last checkpoint can be recovered, and the database restored to its state at that checkpoint. Generating recovery information is an important function of any persistent store, since the reliability and resilience of the database depend on it.

Detecting modifications to objects can be achieved in much the same way as for garbage collection, except that *all* updates, not just pointer stores must be recorded. Note that objects must be *unswizzled* to compare them with the object store's unmodified originals and generate log records. This is easy since we prepend OIDs to resident persistent objects. While unswizzling we may see references to new (not yet persistent) objects, which are assigned OIDs and made persistent.

3 Experiments

All of our experiments are based on a high-performance Smalltalk interpreter of our own design, using the abstract definition of Goldberg and Robson [8]. The implementation consists of two components: the *virtual machine* and the *virtual image*. The virtual machine implements a bytecode instruction set to which Smalltalk source code is compiled, as well as other primitive functionality. While we have retained the standard bytecode instruction set of Goldberg and Robson [8], our implementation of the virtual machine differs somewhat from their original definition to allow for more efficient execution. Our virtual machine running on the DECstation 3100 performs around three times faster than a microcoded implementation on the Xerox Dorado.

The virtual image is derived from an early commercial version of Smalltalk with minor modifications. It implements (in Smalltalk) all the functionality of a Smalltalk development environment, including editors, browsers, the bytecode compiler, and class libraries, all of which are first-class objects in the Smalltalk sense. Booting a Smalltalk environment involves loading the virtual image into memory for execution by the virtual machine.

In our persistent implementation of Smalltalk the virtual image resides in the database, and the Smalltalk environment is booted by loading that subset of the objects in the image sufficient to resume execution by the virtual machine. The bytecode instruction set is the same as in our non-persistent virtual machine, and changes to the virtual image have been minor. Rather, all extensions for persistence affected only the virtual machine, which has been augmented carefully to fault persistent objects into memory as they are needed by

the executing image. We have implemented the previously described schemes for detecting and handling object faults.

All benchmarks are coded directly in Smalltalk, and measured using a specially instrumented version of the interpreter. The instrumentation is kept constant across all implementation variants being considered, so that direct comparisons can be made—any differences in the results can only be due to the particular implementation variant being used.

3.1 Experimental setup

We ran our experiments on a DECstation 3100 (MIPS R2000A CPU clocked at 16.67MHz) running ULTRIX 4.1.[2] The benchmarks were run with the system in single user mode and the process's address space was locked in main memory to prevent paging. Database checkpoint operations included a call to fsync to force the log data to the local disk before completing. For the persistent store experiments the database was accessed remotely via NFS, with the client and server connected via a private EtherNet.

We measured elapsed time on the client machine using a custom timer board[3] having a resolution of 100 ns. The fine-grained accuracy of this timer allows separate measurement of each phase of a benchmark's execution.

All benchmarks involving random execution were made repeatable by presenting the same seed to the random number generator for each run. Random numbers are also generated before measurement of the benchmark execution, so that the elapsed times do not include the numerical computation overhead of random number generation.

3.2 Garbage collection

We measured three implementations of the schemes for garbage collection: remembered sets, card marking, and page traps. In contrast to our earlier performance studies [11], we have reimplemented the card and page trap schemes to avoid unnecessary scanning, by combining the precision of remembered sets with the simplicity of card marking. As the dirty cards are scanned prior to each scavenge, the older-to-younger pointers in those cards are summarized to the appropriate remembered sets, which are then used as the basis of the scavenge. The cards are thereafter treated as clean. Subsequent scavenges need only update the remembered sets by rescanning just those cards that have been dirtied since the previous scavenge.[4]

We varied the card size by multiples of four from 16 bytes up to the virtual memory page size (4K bytes). We also

measured the performance of an implementation that assumes an *oracle* to discover which pages of the heap are dirty at each garbage collection. This allows us to determine the optimal performance that could be expected if a zero-cost implementation of the dirty operating system primitive discussed in Section 2.1 were available.

3.2.1 Implementation

To avoid making the remembered sets too large we record only those stores that create pointers from older objects to younger objects. This involves extra conditional overhead at every store site to perform the check, in addition to a subroutine call to update the remembered set if the condition is true. Smalltalk object references are tagged to allow direct encoding of non-pointer immediate values such as integers. Since many object references are immediate, the first action performed by the check is to filter out non-pointer stores. This is followed by a generation test to filter out "initializing" stores to objects in the youngest generation (such stores cannot create older-to-younger pointers). Finally, if the store creates a pointer from an older object to a younger object the remembered set is updated with a subroutine call. On the MIPS R2000 non-pointers are filtered in 2 cycles. Filtering initializing stores requires another 8 cycles, while filtering the remaining uninteresting stores consumes a further 8 cycles. The size of the entire inline sequence for a store typically comes to 22 instructions, including the store itself, filtering of uninteresting stores, and the call to update the remembered set; some of these are frequently skipped because of the filtering.

For the card schemes we implement the card table as a contiguous byte array, one byte per card, so as to simplify the store check.[5] By interpreting zero bytes as dirty entries and non-zero bytes as clean, a pointer store can be recorded using just a shift, index, and byte store of zero. Since the most attractive feature of card marking is the simplicity of the store check, we omit the checks used in the pure remembered set scheme to filter uninteresting stores. On the MIPS R2000 stores are recorded with just 5 instructions: 2 to load the base of the card table, a shift to determine the index, an add to index the table, and a byte store of zero. Including the store, the entire inline sequence comes to 6 instructions. If we kept the card table base in a register this sequence would shrink to 4 instructions (registers are at a premium in the interpreter). We note that the byte store instruction on the R2000 is implemented in hardware as a read-modify-write instruction, requiring several cycles for execution.

The page trap scheme requires no inline code at store sites to detect pointer stores, relying instead on the page protection hardware to trap updates to protected pages. Thus there is no longer any advantage in using a byte table to simplify the store check. Rather, it is more important that the dirty page table consume the smallest possible space. For this reason we

---

use a bit table; setting a bit indicates that the corresponding page is dirty. When a protection trap occurs the bit in the table corresponding to the modified page is set and the page unprotected.

3.2.2 Benchmarks

We use two benchmarks to evaluate garbage collection performance. The first is a synthetic benchmark of our own devising based on tree creation. The second consists of several iterations through the standard "macro" benchmark suite that is used to compare the relative performance of Smalltalk implementations [17]. Our benchmarks have the following characteristics:

- **Destroy**—trees with destructive updates: A large initial tree (~2M bytes) is repeatedly mutated by randomly choosing a subtree to be replaced and fully recreated. The effect is to generate large amounts of garbage, since the subtree that is destroyed is no longer reachable, while retaining the rest of the tree to the next iteration. Rebuilding the subtree causes many pointer stores, some of which create older-to-younger pointers of interest to the garbage collector. Each run performs 160 garbage collections.

- **Interactive**—10 iterations of the "macro" benchmarks: These measure a system's support for the programming activities that constitute typical interaction with the Smalltalk programming environment, such as keyboard activity, compilation of methods to bytecodes, and browsing. Each run performs 137 garbage collections.

3.2.3 Results

We report the elapsed time of each phase of execution of the benchmark, including:

- *running*: the time spent in the interpreter executing the program, as opposed to the garbage collector (note that running includes the cost of store checks or page traps);

- *roots*: the time spent scanning through remembered sets or card/page tables and copying the immediate survivors; [6]

- *promoted*: the time spent copying any remaining survivors; and

- *other*: the time spent in any remaining GC bookkeeping activities.

Figure 2 plots the results for the remembered set (remsets), page trap (pages), and card implementations (for card sizes of 16, 64, 256, 1024, and 4096 bytes) on the Destroy benchmark. The performance that might be obtained using a zero-cost implementation of dirty is estimated by taking the *running*, *roots*, and *promoted* times for the oracle-based implementation along with the *other* overheads for the card

---

[6] In Smalltalk the stack is stored as heap objects so there is no separate stack processing. In fact, all the process stacks are copied during each scavenge. Also, Smalltalk has only a few global variables, in the interpreter.


Figure 2: Destroy

Figure 3: Interactive

marking scheme. This is plotted alongside the other measurements (dirty). The results for the Interactive benchmark are similarly shown in Figure 3.

To the extent that garbage collection overheads affect total execution time, the results are conclusive, with the page-sized granularity imposing significant overhead in scanning to determine root objects for collection. In contrast with our earlier results [11], we see that summarizing interesting pointer information into remembered sets for use in subsequent scavenges can reduce this scanning overhead, such that the card schemes are competitive with the pure remembered set scheme. Nevertheless, the pure remembered set scheme has markedly less overhead to determine the roots. Also, using a bit table versus a byte table has little effect on root processing time (the *roots* times are very similar for dirty, which scans a bit table, and cards, which scans a byte table).

The results are somewhat less conclusive for running-time overheads. The variation in running time amongst the card schemes can only be explained by hardware data cache effects (such as the specific mapping of virtual pages to physical addresses for this physically addressed cache), since the card schemes all execute the exact same code (barring differences in the shift value used to index the card table). Similarly, the fact that the oracle-based dirty scheme does not exhibit the best running time of the different implementations can only be explained as a result of such data cache effects. Neverthe-

less, dirty has running time less than pages for both benchmarks. Since these implementations use exactly the same virtual machine and garbage collection data structures, any difference is unlikely to be due to the prevously mentioned cache effects. Thus we can get some idea of the overhead to field a trap from the operating system, unprotect the appropriate page, and return to normal execution, by subtracting the running time of the oracle-based dirty scheme from that of the pages scheme, and dividing by the number of page traps. This yields a per-trap overhead of $915\mu s$ for the Destroy benchmark (864 traps), and $744\mu s$ for the Interactive benchmark (1656 traps), showing that the traps can be much more expensive than the lower bound of $250\mu s$ we obtained by measuring their cost in a tight loop. These results suggest that the frequency of traps affects their cost. Presumably, more frequent traps mean that the hardware caches are more likely to contain the operating system code and data required to service a trap, making for faster trap handling.

Given the number of store checks executed by the card schemes and the number of traps incurred by the page trap scheme for each benchmark, we can determine the trade-off between using explicit code to maintain dirty bits and a page trapping approach. Ignoring virtual-to-physical mapping cache effects, the break-even point is determined by the formula:

$$cx = fty$$

where

c = the number of store checks executed by an explicitly coded software scheme;

x = cycles per check;

f = clock frequency (16.67 MHz for DECstation 3100);

t = the number of traps incurred by a page trapping scheme;

y = μs per trap.

For these benchmarks this yields Table 1, which gives the maximum page trap overhead such that a page trapping approach will incur less running time than an alternative explicit implementation having the given overhead per store check. Let us assume that our current 5-instruction sequence for card marking executes in no more than 10 cycles. To be competitive a page trap implementation would have to incur no more than $41\mu s$ and $237\mu s$ per trap, for the Tree and Interactive benchmarks respectively. These values are significantly lower than the estimated trap overheads for these benchmarks quoted above, and lower even than the $250\mu s$ lower bound obtained for a tight loop.

We summarize the results in Tables 2 and 3, indicating the elapsed time for each of the phases as a percentage of those for dirty, and note that the total elapsed time for the 1K byte card scheme is best overall.

	Tree	Interactive
Store checks (c)	59646	654245
Page traps (t)	864	1656
Cycles per check (x)	μs per trap (y)	
1	4	24
2	8	47
4	17	95
5	21	118
10	41	237
15	62	355
20	83	474
50	207	1185
100	414	2370
150	621	3555
200	828	4740

Table 1: Break-even points for GC implementations that use page trapping vs explicit checks

3.3 Object fault handling

Persistent Smalltalk is obtained by extending the virtual machine to handle persistence. No modifications have been made to the bytecode instruction set to support persistence, and changes to the virtual image have been minor. Rather, all extensions involve the virtual machine, so that objects are faulted as they are needed by the executing image.

Permanent storage for the virtual image is provided by an underlying persistent object storage manager [18]. Since objects are too small a unit for efficient individual transfer to and from disk, the storage manager groups objects together into *physical segments* for transfer between the permanent database and its in-memory buffers. Physical segments may have arbitrary size (up to some large system-defined limit). Thus a physical segment may contain any number of objects. Objects within a physical segment are further grouped into *logical segments* (of at most 255 objects) for efficient management of the OID space (objects in a given logical segment have the same high bits in their OID). Applications can take advantage of these groupings to cluster related objects for retrieval.

3.3.1 Implementation

We have implemented two variants of the fault block approach to detecting non-residency.[7] The first uses explicit software residency checks in the virtual machine while the other exploits page-protection.

Computation in Smalltalk proceeds by sending *messages* to objects. The effect of sending a message is to invoke a *method* on the receiver of the message. Invoking a method

---

[7]We also implemented a reference tagging scheme, but it was clearly uncompetitive

Phase	remsets	cards-0016	cards-0064	cards-0256	cards-1024	cards-4096	pages
Running	101%	98%	99%	103%	98%	98%	102%
Roots	19%	542%	153%	64%	58%	93%	138%
Promoted	97%	107%	102%	100%	99%	100%	101%
Other	94%	107%	104%	96%	96%	100%	360%
Total	98%	115%	101%	101%	97%	98%	105%

Table 2: Phases of execution for Destroy as percentage of dirty

Phase	remsets	cards-0016	cards-0064	cards-0256	cards-1024	cards-4096	pages
Running	98%	97%	97%	102%	95%	94%	103%
Roots	22%	605%	173%	71%	56%	95%	141%
Promoted	100%	110%	103%	101%	99%	100%	100%
Other	96%	106%	104%	98%	97%	100%	255%
Total	95%	111%	99%	102%	94%	94%	105%

Table 3: Phases of execution for Interactive as percentage of dirty

may be thought of as a procedure call. Precisely which method is invoked depends on the *class* of the receiver, so every Smalltalk object contains a pointer to its class, which is itself a Smalltalk object. Because computation is driven by the sending of messages, most objects will become resident only when a message is sent to them. By arranging for fault and indirect blocks to respond to messages by forwarding the message to their target object (faulting the object as necessary), message sends to resident objects typically incur no extra overhead.

Byte-compiled methods (code) and stack frames are also first-class objects in Smalltalk. By making further constraints on the residency of certain references contained in these objects we are able to restrict all residency checking to method invocation; even there the overhead is typically incurred only when a method is invoked on a non-resident object.[8]

The effect of the residency constraints is felt whenever an object is made resident. Pointer fields that are subject to a residency constraint must be swizzled to refer directly to their target. Unconstrained pointer fields are swizzled to point directly to their target only if the target object is already resident—the storage manager supports the efficient mapping of OIDs to resident objects. Otherwise, the pointer field is swizzled to refer to a fault block.

Since there may be multiple references to a given fault block dispersed through the registers and memory of the virtual machine, we arrange for object faults to bypass the indirection that would otherwise be created when a fault block is converted to an indirect block. This is not strictly necessary for the software fault detection scheme, since traversing

a pointer to an indirect block can be quickly handled at the cost of an indirection—for a fault block, the full object fault mechanism must be invoked to translate the OID contained in the fault block. However, for a page trapping scheme, the overhead of the traps is high enough to justify expending some effort to eliminate references to (ex-fault) indirect blocks, to avoid repeated loading and faulting on those references. To support this, each page of allocated fault blocks has a remembered set associated with it, recording all persistent objects whose pointer fields have been swizzled to refer to fault blocks that lies in the page. At each object fault we scan the objects in the remembered set to update any pointers and bypass the indirection. For a fair comparison with hardware-assisted variants we also apply this indirection elimination to the software scheme.

The architecture leaves open the possibility of making any number of objects resident at one time. In an earlier study [10] we considered the granularities inherent in the underlying object storage manager: individual objects, logical segments, and physical segments. Swizzling just one object at a time has the advantage of faulting just those objects needed by the program for it to continue execution. Swizzling an entire logical or physical segment at a time allows the program to take advantage of any clustering present in the physical layout of objects in the database. Also, since all the objects in a segment can be mapped before they are swizzled, any intra-segment references can be converted to direct pointers. If the static clustering is a good approximation to the dynamic locality of access by the program, then the speed of program execution will improve since fewer object faults will occur.

For these experiments we swizzle entire logical segments, and compare several versions of the virtual machine that differ only in their implementation of fault detection, and

---

[8]Primitives may need to perform additional residency checks if they access objects other than those whose residency is guaranteed by the constraints

Scheme	Description
non-persistent	Non-persistent, unadorned with residency checks
FB-resident	Non-persistent, augmented with software fault block residency checks
PF-resident	Non-persistent, augmented with the page trap handling code, plus necessary support to decode load instructions that might cause a trap
FB	Persistent, fault blocks, swizzle 1 logical segment at a time
PF	Persistent, fault blocks allocated in protected pages, swizzle 1 logical segment at a time

Table 4: Fault detection schemes measured in experiments

whether they are running against a completely resident virtual image (non-persistent Smalltalk) or against an image that is faulted on demand (persistent Smalltalk). Table 4 enumerates the variants.

3.3.2 Benchmarks

We use the Lookup and Traversal portions of the OO1 object operations benchmarks [7]. The OO1 benchmark database consists of a collection of 20,000 "part" objects, indexed by part numbers in the range 1 through 20,000, with exactly three "connections" from each part to other parts. The connections are randomly selected to produce some locality of reference: 90% of the connections are to the "closest" 1% of parts, with the remainder being made to any randomly chosen part. Closeness is defined as parts with the numerically closest part numbers. The part database and the benchmarks are implemented entirely in Smalltalk, including the B-tree used to index the parts.

The benchmarks operate as follows:

- **Lookup** fetches 1,000 randomly chosen parts from the database. For each part a null procedure is invoked, taking as its arguments the x, y, and *type* fields of the part.

- **Traversal** fetches all parts connected to a randomly chosen part, or to any part connected to it, up to seven hops (for a total of 3,280 parts, with possible duplicates). Similarly to the Lookup benchmark, a null procedure is invoked for each part, taking as its arguments the x, y, and *type* fields of the part.

Each measure is typically run 10 times, the first when the system is cold, with none of the database cached (apart from any schema or system information necessary to initialize the system). Each successive iteration fetches a *different* set of random parts. Before the first run of each series of benchmark iterations a "chill" program is executed on the client to ensure that the operating system file buffers of both client and server have been flushed of all database segments, so that the first iteration is truly cold.

In addition to the ten cold-warm iterations, we measured the elapsed time for a *hot* iteration of the Traversal benchmark, by beginning at the same initial part used in the last of the warm iterations. This hot run is guaranteed to traverse


Figure 4: Lookup

Figure 5: Traversal

only resident objects, and so will be free of any overheads due to swizzling and retrieval of non-resident objects.

3.3.3 Results

The elapsed times for the cold-warm iterations of each benchmark are plotted in Figures 4 and 5, expanding the scale to focus on the warm performance, with the non-persistent performance as a baseline. The FB and PF schemes behave very similarly, with warm performance close to optimal. However, the software-mediated FB scheme has better performance overall.

We summarize the benchmark results in Table 5, reporting the average elapsed time (in seconds) of the 10 iterations for the non-persistent variants (since the database is

114

	Lookup	Traversal
Scheme	Average	Average
non-persistent	0.57	0.089
FB-resident	0.57	0.089
PF-resident	0.57	0.090

	Cold	Warm	Cold	Warm	Hot
FB	28.11	0.57	14.79	0.154	0.088
PF	29.29	0.59	15.34	0.162	0.089

Table 5: Elapsed times for object faulting benchmarks (seconds)

	Lookup			Traversal		
Scheme	a	b	r	a	b	r
FB	0.568	0.000214	0.9015	0.090	0.000181	0.9885
PF	0.585	0.005967	0.9996	0.102	0.004369	0.9805

Table 6: Fault overheads (seconds)

always resident and warm), and cold (first iteration), warm (10th iteration), and hot times for the persistent variants. The non-persistent variants exhibit little difference in their performance, indicating that the overhead of the run-time residency checks is slight.

To get some sense of the cost of the page traps and software-mediated object faults we have obtained linear regression fits of the running time (time spent in the interpreter actually executing bytecodes as opposed to swizzling, bypassing indirections, or reading from disk) versus the number of faults occurring for each iteration of the benchmark. The model used is:

$$y = a + bx$$

where

y = running time (excluding swizzling and other fault handling overheads);

a = y-axis (running time) intercept;

b = seconds per fault;

x = number of faults.

The fits obtained are good and the coefficients are given in Table 6, as well as the linear correlation coefficient r. The b coefficient is a measure of the number of seconds required to get in and out of the object fault handler, either through software checks to detect faults or through a protection trap handler. The results for PF show that trap handling overhead is once again much greater than the $250\mu s$ value obtained when measured for a tightly coded loop, but the high per-trap cost is not unreasonable considering that each fault involves substantial work to eliminate references to the trapped fault block, which will significantly disturb the state of the

hardware caches. Upon resumption of normal execution the hardware instruction and data caches must be reloaded before peak execution speeds can be achieved. The results for FB reveal just how fast protection traps have to be in order to outstrip the software implementation—software-mediated fault detection overheads are less than $250\mu s$ for both benchmarks.

To summarize, we have shown that software object faulting schemes can be made to have performance close to optimal. On the other hand, page trap schemes can be significantly slowed by the cost of the traps. While we cannot vouch for the performance of a direct-mapped scheme in the style of Texas and ObjectStore, the fact that software object faulting can give performance close to optimal makes it difficult to beat.

3.4 Database checkpointing

A checkpoint operation consists of copying and unswizzling modified and newly-created objects (or modified subranges of objects) back to the storage manager's buffers for eventual return to the stable database, along with generating a log record describing the range and values of the modified region of the object. The log record is generated by comparing the old and new versions of the object as it is unswizzled—we generate a difference log indicating the changes made to the object. Unswizzling may encounter pointers to objects newly-created since the last checkpoint. These objects must be assigned an OID and unswizzled in turn, perhaps dragging further newly-created objects into the database.

115

3.4.1 Implementation

We have implemented and measured four schemes for detecting updates to persistent objects. The first uses a remembered set to record persistent objects that have been modified since the last checkpoint. To keep the remembered set from becoming too large we record only updates to persistent objects. This requires a check to see that the updated object is located in the separately managed persistent area of the volatile heap. If the updated object is persistent then a subroutine is invoked to enter the object's pointer in the remembered set. On the MIPS R2000, non-persistent objects are filtered in 9 cycles. The entire inline sequence is 12 instructions long.

Rather than noting updated objects in a remembered set, the second scheme marks objects as they are updated by setting a bit in the header of the object when it is modified. Upon checkpoint all resident persistent objects must be scanned to find those that have been updated.

Card marking can also be used to track updates for log generation. We compare card schemes that use card sizes of 16, 64, 256, 1024, and 4096 bytes, as well as a fourth approach which uses page protections to monitor updates. These schemes are implemented exactly the same as for the garbage collector store checks. The card schemes use a byte table and 5 instructions to dirty a card. The page protection scheme uses a bit table and traps updates to protected pages.

For small objects the remembered set scheme is ideal. However, updates to large objects may suffer from poor locality with respect to the object size, resulting in unnecessary unswizzling upon checkpoint. Thus checkpoint overheads are bounded solely by the size of the object. The object marking scheme suffers from the need to examine every resident persistent object to find those that have been modified. If only a few objects have been modified then it must examine many more objects than need to be unswizzled. The card and page protection schemes record updates based on fixed-size units of the address space. Similarly to garbage collection, we can expect the size of the cards to influence checkpoint costs, since large cards imply higher unswizzling overheads.

These benchmarks all use the exact same object faulting and swizzling scheme, while varying the update detection mechanism.

3.4.2 Benchmarks

Previous studies have extended the Traversal operation of the OO1 object database benchmarks to also perform some modification of part objects [27]. Each part accessed during the traversal may be updated based on some known probability fixed in advance. For example, if the probability of update is 0.5 then approximately half of all parts visited will be modified. The update consists of incrementing the x and y 4-byte integer fields of the part. A checkpoint operation is performed at the end of each traversal to commit the changes to the database.

In order to best understand the behavior of the update

Figure 6: Hot update

detection mechanisms in the absence of other effects such as object faults and swizzling, we measured the time to run 10 *hot* iterations of the update benchmarks, by beginning each hot iteration at the same initial part used in the last of the warm iterations. These hot runs are guaranteed to traverse only resident objects, and so will be free of any overheads due to swizzling and retrieval of non-resident objects.

3.4.3 Results

Figure 6 summarizes the average elapsed time for the ten hot iterations at each of the update probabilities. The results are clearer when we break the total elapsed time down into separate phases of execution, as plotted in Figures 7–10. *Running* is the time spent in the virtual machine executing the bytecodes of the benchmark program, including the overhead to note modifications. The checkpoint operation itself is decomposed into:

- *old*: time to locate and unswizzle old modified objects and generate log entries for them;

- *new*: time to unswizzle new persistent objects and generate log entries for them;

- *write*: time to flush the log records to disk;

- *other*: time to perform other bookkeeping, such as managing page protections.

For update probability $p = 0$, there are very few old modified objects to be unswizzled, so the remembered set scheme shows little overhead for this phase. The card and page schemes incur overhead to scan the card table. Since the card table is larger for smaller cards, the overheads are correspondingly higher, while the object marking scheme must scan all the objects for very little gain. Note that every checkpoint also generates a small number of new persistent objects—recall that stack frames are objects and may persist, so the checkpoint must log any newly-allocated active stack frames, to allow resumption of execution from the checkpoint in the case of a crash.

116

Figure 7: Update ($p = 0.0$), 10 hot iterations

Figure 8: Update ($p = 0.05$), 10 hot iterations


Figure 9: Update ($p = 0.5$), 10 hot iterations

Figure 10: Update ($p = 1.0$), 10 hot iterations

Scheme	μs per byte	μs per skip	r
objects	1.17	2.67	0.9998
remsets	1.27		1.0000
cards	1.01		0.9954
pages	1.07		0.9998

Table 7: Checkpoint scanning overheads

For higher update probabilities, the effects of the competing schemes are more clearly discerned. Unswizzling of old objects consumes more of the total checkpoint time, so that large-granularity schemes face higher unswizzling costs. For example, the large-granularity card schemes must unswizzle all objects that lie within a dirty card, including those that have not actually been modified. In contrast, although object marking scans all the resident objects, it need only unswizzle those that have actually been modified. Remembered sets are even better, except at probability $p = 1.0$, since they eliminate the need to scan. Note that all schemes write the same size log at each checkpoint, since we only write out the part of an object that has actually been modified. The cost of the page traps can also be seen at the higher update probabilities. Still, the recurring theme is that the cost of the page traps is a small fraction of total costs—granularity issues are much more important.

We have also managed to obtain a measure of the per-byte scanning costs associated with each scheme in unswizzling old modified persistent objects. These results were obtained by running regression fits for the *old* times versus the number of bytes unswizzled and compared, in order to generate the difference log. For the object marking scheme we actually had to use a multi-linear fit to account for the additional overhead of sifting through all the persistent objects to find those that were modified and hence needed to be unswizzled, taking as an additional independent variable the number of objects skipped as unmodified during the scan. The fits were all excellent and the results are summarized in Table 7.

The per-byte overheads are much the same for all the schemes. The variation can be accounted for by considering that there are other effects that may be correlated with the per-byte scanning costs and cannot be linearly separated by the regression. There are also slight differences among the schemes in the way the scanning loops are implemented. The skip overhead is the cost to decode each unmodified object's header and to skip to the next object in the scan. These results indicate the correlation between the time required to unswizzle objects, and the number of bytes that must be considered when unswizzling. Clearly, the best schemes are those that minimize the amount of scanning and the number of bytes that must be unswizzled.

117

4 Interpreters versus Compilers

The experiments and results we have presented are for an interpreted language. An obvious question is how they would change for a compiled language. The results for garbage collection and database checkpointing will likely stand since compilation will tend to shrink the time actually spent in the program code, but garbage collection and checkpointing costs will be unaffected since they are part of the language run-time system. Thus, collection time and checkpoint overheads become more important, strengthening our arguments for software rather than hardware approaches in many circumstances. Reducing the incremental costs of noting updates for garbage collection and checkpointing will not improve the situation. Rather, it is better to invest in improving the garbage collection and checkpoint overheads, which are heavily impacted by the granularity of the noted updates. Nevertheless, inline checks do affect program size and possibly instruction cache behavior, so we cannot predict the situation with certainty. The size and frequency of inline checks will affect overhead with respect to a *non-garbage collected* or *non-persistent* implementation, but in the absence of contrary evidence, it is reasonable to assume that the relative ranking of the various schemes will remain the same.

With respect to object faulting we observe that the nature of the programming language and programmer's style may be important. In particular, our faulting implementation exploits Smalltalk's heavy use of dynamic method dispatch. In effect, non-resident objects are a hidden subclass, all of whose methods fault the object in. This means that method dispatch does not explicitly check residence, substantially reducing the number of explicit checks. For example, in the Lookup benchmark 93% of the checks are absorbed by the dispatch mechanism, while in Traversal virtually all (99–100%) are. In languages such as C++ (or Modula-3), the same technique can be applied by having non-resident objects point to different, special, virtual function tables (method suites). Again, there is no performance impact on dispatch for resident objects. However, the more one uses direct access and non-virtual function calls, the more one must rely on explicit checks. In the case of change noting, one can apply the same sort of trick (one function table for clean objects and a different one for dirty objects). We further note that there are certainly circumstances in which an earlier check subsumes later ones, so that suitable data flow analysis can further reduce the number of explicit checks required.

Given these observations, we believe that our results are likely to carry over to compiled languages, at least for the most part, but certainly admit the necessity of further experimentation. We do note that the somewhat contrary results of White and DeWitt [27], for the E persistent programming language, are for a compiler that does not incorporate any of the techniques for eliminating explicit checks that we describe in this section.

5 Conclusions

We compared the performance of application implementations that exploit virtual memory primitives against implementations using software techniques. In summary, generational garbage collection performs best when the roots of a collection can be accurately located, without having to scan to find them; remembered sets do this admirably. Pure card-based schemes, particularly at the size of virtual memory pages, exhibit poor performance because pages are too large a granule; this is improved by a hybrid scheme that marks cards but after scanning them preserves the significant information in remembered sets. Object fault handling in persistent programming languages can be implemented efficiently without virtual memory primitives, especially given an object-oriented execution paradigm, yielding performance close to optimal. Finally, checkpointing persistent data using virtual memory primitives suffers in performance because pages are too large a granule when update locality is small.

Our software solutions require modification of the programming language implementation. One nice feature of hardware-assisted techniques for garbage collection and persistence is that they are transparent and do not need support from the language implementation. While page trapping might be indispensable for completely transparent solutions to these applications, we have shown that software solutions can be more efficient.

Our results show that for these applications the costs of the virtual memory primitives themselves are not the primary problem. Rather, the unduly large granularity of virtual memory pages in modern operating systems and architectures significantly affects overall performance for some applications (notably checkpointing). Moreover, as processor speeds improve and physical memories grow, page sizes are likely to become much larger, further degrading the performance of virtual memory solutions in applications that have a naturally smaller granularity. It is worthwhile noting that similar results have been obtained in other very different application areas, such as efficient implementation of data breakpoints for debuggers [26]. We speculate that sub-page protection and dirty bits, along with appropriate operating systems interfaces, might somewhat overcome the performance disadvantages we observed. However, it is clear that while user level virtual memory primitives are transparent solutions to various memory management problems, they do not necessarily offer the best performance.

References

[1] T. E. Anderson, H. M. Levy, B. N. Bershad, and E. D. Lazowska. The interaction of architecture and operating system design. In ASPLOS-IV [4], pages 108–120.

[2] A. Appel. Simple generational garbage collection and fast allocation. *Software: Practice and Experience*, 19(2):171–183, Feb. 1989.

[3] A. W. Appel and K. Li. Virtual memory primitives for user programs. In ASPLOS-IV [4], pages 96–107.

[4] *Proceedings of the Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, Santa Clara, California, Apr. 1991. *ACM SIGPLAN Not. 26*, 4 (Apr. 1991).

[5] M. P. Atkinson, P. J. Bailey, K. J. Chisholm, P. W. Cockshott, and R. Morrison. An approach to persistent programming. *The Computer Journal*, 26(4):360–365, Nov. 1983.

[6] M. P. Atkinson, K. J. Chisholm, W. P. Cockshott, and R. M. Marshall. Algorithms for a persistent heap. *Software: Practice and Experience*, 13(7):259–271, Mar. 1983.

[7] R. G. G. Cattell and J. Skeen. Object operations benchmark. *ACM Trans. Database Syst.*, 17(1):1–31, Mar. 1992.

[8] A. Goldberg and D. Robson. *Smalltalk-80: The Language and its Implementation*. Addison-Wesley, 1983.

[9] A. L. Hosking, E. Brown, and J. E. B. Moss. Update logging for persistent programming languages: A comparative performance evaluation. In *Proceedings of the Nineteenth International Conference on Very Large Data Bases*, Dublin, Ireland, Aug. 1993. Morgan Kaufmann.

[10] A. L. Hosking and J. E. B. Moss. Object fault handling for persistent programming languages: A performance evaluation. In *Proceedings of the Conference on Object-Oriented Programming Systems, Languages, and Applications*, Washington, DC, Oct. 1993.

[11] A. L. Hosking, J. E. B. Moss, and D. Stefanović. A comparative performance evaluation of write barrier implementations. In *Proceedings of the Conference on Object-Oriented Programming Systems, Languages, and Applications*, pages 92–109, Vancouver, Canada, Oct. 1992. *ACM SIGPLAN Not. 27*, 10 (Oct. 1992).

[12] T. Kaehler. Virtual memory on a narrow machine for an object-oriented language. In *Proceedings of the Conference on Object-Oriented Programming Systems, Languages, and Applications*, pages 87–106, Portland, Oregon, Sept. 1986. *ACM SIGPLAN Not. 21*, 11 (Nov. 1986).

[13] T. Kaehler and G. Krasner. LOOM—large object-oriented memory for Smalltalk-80 systems. In Krasner [14], chapter 14, pages 251–270.

[14] G. Krasner, editor. *Smalltalk-80: Bits of History, Words of Advice*. Addison-Wesley. 1983.

[15] C. Lamb, G. Landis, J. Orenstein, and D. Weinreb. The ObjectStore database system. *Commun. ACM*, 34(10):50–63, Oct. 1991.

[16] H. Lieberman and C. Hewitt. A real-time garbage collector based on the lifetimes of objects. *Commun. ACM*, 26(6):419–429, June 1983.

[17] K. McCall. The Smalltalk-80 benchmarks. In Krasner [14], chapter 9, pages 153–173.

[18] J. E. B. Moss. Design of the Mneme persistent object store. *ACM Trans. Inf. Syst.*, 8(2):103–139, Apr. 1990.

[19] J. E. B. Moss. Working with persistent objects: To swizzle or not to swizzle. *IEEE Trans. Softw. Eng.*, 18(8):657–673, Aug. 1992.

[20] Object Design, Inc. *ObjectStore User Guide*, Oct. 1990. Release 1.0.

[21] R. A. Shaw. Improving garbage collector performance in virtual memory. Technical Report CSL-TR-87-323, Stanford University, Mar. 1987.

[22] V. Singhal, S. V. Kakkad, and P. R. Wilson. Texas, an efficient, portable persistent store. In *Proceedings of the Fifth International Workshop on Persistent Object Systems*, pages 11–33, San Miniato, Italy, Sept. 1992.

[23] P. G. Sobalvarro. A lifetime-based garbage collector for LISP systems on general-purpose computers, 1988. B.S. Thesis, Dept. of EECS, Massachusetts Institute of Technology, Cambridge.

[24] D. Ungar. Generation scavenging: A non-disruptive high performance storage reclamation algorithm. In *Proceedings of the ACM SIGSOFT/SIGPLAN Software Engineering Symposium on Practical Software Development Environments*, pages 157–167, Pittsburgh, Pennsylvania, Apr. 1984. *ACM SIGPLAN Not. 19*, 5 (May 1984).

[25] D. M. Ungar. *The Design and Evaluation of a High Performance Smalltalk System*. ACM Distinguished Dissertations. The MIT Press, Cambridge, MA, 1987. Ph.D. Dissertation, University of California at Berkeley, February 1986.

[26] R. Wahbe. Efficient data breakpoints. In *Proceedings of the Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 200–212, Boston, Massachusetts, Sept. 1992. *ACM SIGPLAN Not. 27*, 9 (Sept. 1991).

[27] S. J. White and D. J. DeWitt. A performance study of alternative object faulting and pointer swizzling strategies. In *Proceedings of the Eighteenth International Conference on Very Large Data Bases*, pages 419–431, Vancouver, Canada, Aug. 1992. Morgan Kaufmann.

[28] P. R. Wilson and S. V. Kakkad. Pointer swizzling at page fault time: Efficiently and compatibly supporting huge address spaces on standard hardware. In *Proceedings of the 1992 International Workshop on Object Orientation in Operating Systems*, pages 364–377, Paris, France, Sept. 1992. IEEE Press.

[29] P. R. Wilson and T. G. Moher. Design of the Opportunistic Garbage Collector. In *Proceedings of the Conference on Object-Oriented Programming Systems, Languages, and Applications*, pages 23–35, New Orleans, Louisiana, Oct. 1989. *ACM SIGPLAN Not. 24*, 10 (Oct. 1989).