Improving the memory management performance of RTSJ

M. Teresa Higuera-Toledano, Valérie Issarny

► To cite this version:

M. Teresa Higuera-Toledano, Valérie Issarny. Improving the memory management performance of RTSJ. *Concurrency and Computation: Practice and Experience*, 2005, 17 (5-6), pp.715-737. inria-00414952

HAL Id: inria-00414952

<https://hal.inria.fr/inria-00414952>

Submitted on 10 Sep 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Improving the memory management performance of RTSJ

M. Teresa Higuera-Toledano^{1,*} and Valérie Issarny²

¹*Universidad Complutense de Madrid, Ciudad Universitaria, Madrid 28040, Spain*

²*INRIA-Rocquencourt, Domaine de Voluceau, BP 105, 78153 Le Chesnay Cédex, France*

SUMMARY

From a real-time perspective, the garbage collector (GC) introduces unpredictable pauses that are not tolerated by real-time tasks. Real-time collectors eliminate this problem but introduce a high overhead. Another approach is to use memory regions (MRs) within which allocation and deallocation is customized. This facility is supported by the memory model of the Real-Time Specification for Java (RTSJ). RTSJ imposes strict access and assignment rules to avoid both the dangling inter-region references and the delays of critical tasks of the GC. A dynamic check solution can incur high overhead, which can be reduced by taking advantage of hardware features. This paper provides an in-depth analytical investigation of the overhead introduced by dynamic assignments checks in RTSJ, describing and analysing several solutions to reduce the introduced overhead. Copyright © 2005 John Wiley & Sons, Ltd.

KEY WORDS: Java; real-time; embedded systems; garbage collection; memory regions; Java microprocessor; write barriers; performance

1. INTRODUCTION

The original Java platform provides attributes that make it a powerful platform for developing embedded real-time applications; however, it also presents some important deficiencies regarding its use in real-time systems. The National Institute of Standards and Technology (NIST), has produced a basic requirements document for a standard real-time Java API extension. Solutions that comply with this document are the Real-time Specification for Java (RTSJ) [1] and the Real-time Core Extension for the Java Platform [2]. We have analysed how these solutions resolve the problems that Java presents to effectively support embedded real-time applications [3]. As a conclusion to our study, we found

*Correspondence to: M. Teresa Higuera-Toledano, DACYA, Facultad de Informática, Universidad Complutense de Madrid, Ciudad Universitaria, Madrid 28040, Spain.

†E-mail: mthiguer@dacya.ucm.es

Contract/grant sponsor: Ministry of Education of Spain (CICYT); contract/grant number: TIC2002–00334

Figure 1. The MemoryArea hierarchy in RTSJ.

that the RTSJ is the most adequate solution for a Java environment aimed at embedded systems executing multimedia applications (e.g. wireless PDAs), even if some issues are still open (e.g. memory management).

This paper focuses on how to make Java memory management real-time while accounting for relevant Java specifications: the RTSJ, the KVM [4] targeting limited-resource and network connected devices, and the microprocessor core *picoJava* [5].

1.1. Memory areas in RTSJ

The MemoryArea abstract class supports the region paradigm in the RTSJ specification [6] through three kinds of regions (see Figure 1): (i) immortal memory, supported by the ImmortalMemory and the ImmortalPhysicalMemory classes, which contain objects whose life ends only when the JVM terminates; (ii) (nested) scoped memory, supported by the ScopedMemory abstract class, which enables the grouping of objects to have well-defined lifetimes that may either offer temporal guarantees on the time taken to create objects (i.e. supported by the LMemory class) or not (i.e. supported by the VMemory class); and (iii) the conventional heap, supported by the HeapMemory class. An application can allocate memory into the system heap, the immortal system memory region (MR), as well as several scoped MRs, and several immortal regions associated with physical characteristics. Objects allocated within immortal MRs live until the end of the application and are never subject to garbage collection. Objects with limited lifetime can be allocated into the heap or into a scoped region. Garbage collection within the application heap relies on the (real-time) garbage collector (GC) of the Java virtual machine (JVM). RTSJ further defines the GarbageCollector abstract class, which can be customized through an incremental collector allowing the application to execute while the GC is being launched.

RTSJ is able to distinguish between three main kinds of tasks: (i) *low-priority tasks* that are tolerant with the GC; (ii) *high-priority tasks* that cannot tolerate unbounded preemption latencies; and (iii) *critical tasks* that cannot tolerate preemption latencies. Whereas high-priority tasks require a real-time GC, critical tasks must not be affected by the GC, and as a consequence cannot access any object within the heap. A scoped region gets collected as a whole once it is no longer used. Then, since immortal and scoped MRs are not garbage collected, they may be exploited by critical tasks, especially LMemory objects, which guarantee allocation time proportional to the object size. The lifetime of

Table I. Assignment rules in RTSJ.

	Reference to heap	Reference to immortal	Reference to scoped
Heap	Yes	Yes	No
Immortal	Yes	Yes	No
Scoped	Yes	Yes	Same or outer

objects allocated in scoped regions is governed by the control flow. Strict assignment rules placed on assignments to or from MRs prevent the creation of dangling pointers (see Table I). Several related threads, possibly real-time, can share a MR, and the region must be active until at least the last thread has exited. The strict assignment rules imposed by RTSJ to avoid dangling inter-region pointers incur high overhead, which can be reduced by taking advantage of hardware features.

1.2. Related work

The JVM must check for the above assignment rules before executing an assignment statement, and emit an `illegalAssignment()` exception if they are violated. This check includes the possibility of static analysis of the application logic [1]. The Tofte–Talpin calculus [7] uses a lexically scoped expression to delimit the lifetime of a region. Memory for the region is allocated when the control enters into the scope of the region constructor, and is de-allocated when the control leaves the scope. This mechanism is implemented by a stack of regions where regions are ordered by lifetimes. The allocation and de-allocation of regions is determined at compilation time by a type-based analysis, which consists of annotating every expression creating a value with a region variable in the source program, where region allocation and de-allocation are explicit. As this solution, our solution is based on a stack of scoped regions ordered by lifetimes [8]. However, given that in RTSJ a region can be shared among several threads, this solution requires more complex mechanisms because the region will remain active until the last thread has exited, and this fact makes it difficult to determine the de-allocation of regions at compilation time.

Our proposed solution consists of only checking the imposed assignment rules, preserving dangling pointers dynamically when executing the assignment statement. In order to do this, we introduce an extra code in all bytecodes causing an object assignment [9]. This extra code, normally called a write barrier, must be executed beforehand to update the object reference. A similar approach is given in [10], which also uses a stack-based memory management that operates dynamically. This solution proposes a contaminated GC based on the idea that each object in the heap is alive due to references that begin in the runtime stack. However, in contrast to our solution, this solution collects memory within the heap and does not treat another MR.

As the GC coexists with MRs, objects within the heap that have references from objects outside the heap must be considered as roots by tracing-based collectors (i.e. based on copying or mark-and-sweep techniques). In our solution, to maintain the root-set of the GC, we use write barriers. As we use an incremental mark-and-sweep GC strategy based on the tri-colour algorithm given in [11], we also need write barriers to maintain the tri-colour invariant. Hence, to detect new roots of the GC, instead of

Table II. SPECjvm98 programs used.

Program	Description
JESS	Expert shell system based on NASA's CLIPS system
DB	Emulates data operations on resident memory
JAVAC	Java compiler from the JDK 1.0.2
MTRT	Multithreaded raytracer
JACK	Parser generator (early JavaCC version)

employing a technique based on the region to which the objects belong (e.g. as used by generational collectors to maintain inter-generational pointers [12]), we introduce a fourth colour indicating whether an object is outside the heap [13].

Another problem is with the critical tasks that cannot access objects within the heap. In order to detect violations of this rule and emit a `memoryAccessError()` exception, we use read barriers (i.e. we introduce an extra code in all bytecodes causing an object access). The most common approach to implementing read/write barriers is by inline code, consists of generating the instructions for executing write barrier events for every load/store operation. This solution requires compiler cooperation and presents a serious drawback because it increases the size of the application object code [14]. Our alternative solution instruments the bytecode interpreter, avoiding space problems, but this still requires a complementary solution to handle native code. A solution that minimizes the introduced overhead consists of improving the write barrier performance by using hardware support such as the picoJava-II microprocessor [15], which allows write barrier checks to be performed in parallel with the store operation. The use of hardware support for write barriers was the subject of [9], where we proposed to improve the performance of checking *illegal references* in two different ways: using existing hardware support and modifying existing hardware. The performance improvement introduced by these solutions has been compared in [16].

1.3. Organization of this paper

In the following, we propose several memory management solutions more or less compliant with the RTSJ `javax.realtime` package. We first present a software-based solution as a possible implementation of the MR abstraction presented by the RTSJ (Section 2). Two hardware-based solutions improving the performance of both MR and the GC are then given (Section 3). In Section 4, we evaluate the overhead introduced by three different solutions supporting both MRs and incremental GCs. In Section 5, we give some details in order to implement a prototype within the KVM [4], by modifying the original collector to make it incremental, and by introducing MRs. Finally, a summary of our contribution concludes this paper (Section 6).

2. THE BASIC STRATEGY

In this section, we introduce a software-based solution to check both illegal assignments and memory access from critical tasks to objects within the heap. Since the RTSJ specification allows

the implementation of real-time compliant memory management without prescribing to any specific solution, we combine an incremental GC within the heap and a stack-based algorithm supporting the scoped MRs.

2.1. Illegal assignments

A MR implementation must ensure that the following conditions are checked before the assignment is executed: (i) objects within the heap region cannot reference objects within a scoped region; (ii) objects within the immortal region cannot reference objects within a scoped region; and (iii) objects in a scoped region cannot reference objects within another scoped region that is non-outer. In order to detect *illegal assignments* every real-time thread has a region-stack associated with it containing all of the scoped MRs that the thread can hold. The basic idea to detect illegal assignments is to take actions upon those instructions that cause one object to reference another (i.e. by using write barriers):

- the `putfield` (`aputfield_quick`) bytecode causes a reference from one object *X* to another *Y*, and the `aastore` (`aastore_quick`) bytecode stores a reference (*Y*) into an array of references (*X*);
- the `putstatic` (`aputstatic_quick`) bytecode causes a reference from an object within persistent memory (i.e. an outermost region) to another object *Y*.

Then, the MR to which the object *Y* belongs must be outside of the MR to which the object *X* belongs. Figure 2 shows the write barrier pseudo-code, which we must introduce in the interpretation of the aforementioned bytecodes, where we denote the object that makes the reference as *X* and the referenced object as *Y*. The region to which an object belongs must be specified in the header of the object. Then, when an object/array is created by executing the `new` (`new_quick`) or `newarray` (`newarray_quick`) bytecode, it is associated with the scope of the active region. The `nestedRegions(X, Y)` function, throws the `IllegalAssignment()` exception when the region to which the *Y* object belongs is not outside to the region to which the *X* object belongs and returns `true` when the region to which the *Y* object belongs is further or equal to the region to which the *X* object belongs. Following this we describe the algorithm of the `nestedRegions(X, Y)` function, which requires two steps.

1. The region-stack of the active task is explored, from the top to the bottom, to find the MR to which the *X* object belongs. If it is not found, this is notified by throwing a `MemoryAccessError()` exception[‡].
2. The region-stack is again explored, but this time we take the MR found in the previous step as the top of the stack. Then, we start the search from the region to which the *X* object belongs, and the objective is to find the MR to which the *Y* object belongs (i.e. the region to which the object *Y* belongs must be outside the region to which the object *X* belongs). If the scoped region of *Y* is not found in the new region-stack, this is notified by throwing an `IllegalAssignment()` exception. If it is found, the `nestedRegions(X, Y)` returns `true`.

[‡]This exception is thrown upon any attempt to refer to an object in an inaccessible `MemoryArea`.

```
if ((region(X) ≠ scoped) and (region(Y) = scoped)) illegalAssignment();
if ((region(X) = scoped) and (region(Y) = scoped)) nestedRegions(X, Y);
```

Figure 2. Write barrier to detect illegal assignment.

```
if ((type(τ) = critical) and (region(X) = heap)) memoryAccessError();
```

Figure 3. Read barrier code for load references.

```
if ((type(τ) = critical) and ((region(X) = heap) or (region(Y) = heap))) memoryAccessError();
```

Figure 4. Read barrier code for store references.

2.2. Memory access errors

A reference of a critical task to an object allocated in the heap causes the `MemoryAccessError()` exception, which can be achieved by using *read barriers*. Note that read barriers occur upon all object accesses, which means upon executing both of the following types of bytecodes:

- those causing a *load reference* (i.e. `getfield`, `getstatic`, `agetfield_quick`, `agetstatic_quick`, or `aaload` bytecodes);
- those causing a *store reference* (i.e. those causing write barriers: `putfield`, `putstatic`, `aputfield_quick`, `aputstatic_quick`, `aastore`, or `aastore_quick` bytecodes).

For bytecodes causing a load reference, we introduce the code given in Figure 3, and for bytecodes causing a store reference, we introduce the code given in Figure 4. Where the `type()` function returns `thread`, `task` and `critical` depending on the type of the parameter task, τ is the active task.

2.3. Collecting the heap

We specifically consider a *mark-and-sweep* GC based on the tri-colour algorithm [11]. The basic algorithm is as follows: an object is coloured white when not reached by the GC, black when reached, and grey when it has been reached, but its descendants may not have been (i.e. they are white). Grey objects make a wavefront, separating the white (unreached) from the black (reached) objects (see Figure 5), and the application must preserve the invariant that no black objects have a pointer to a white object, which is achieved by using write barriers [12].

Figure 5. The GC strategy.

```
if ((colour(X) = black) and (colour(Y) = white)) greyObject(Y);
```

Figure 6. Write barrier to maintain the tri-colour invariant.

For each thread, we maintain a stack of root pointers. We start the *marking* phase by colouring all objects referenced by grey root pointers. Each root stack pointer is processed root by root, and each object referenced by a root is inserted into a grey list. If during this phase, the application tries to make a reference from a black object to a white one, the colour of the referenced object is turned grey and the object is moved from the white list to the grey list. When all the descendants of a grey object are processed (i.e. the grey object has no white descendants), the grey object is turned black and moved from the grey list to the black list. The collection is completed when there are no more grey objects. During the *sweeping* phase, all the white objects can be recycled and all the black objects become white. In this process, objects that must execute the `finalize()` method are moved to the finalize list. The finalizes are executed by a specialized thread such as that in [17]. Finally, for white objects that have finalized, their memory is freed.

The code checking a reference violating the tri-colour invariant (i.e. from a black object to a white one) must be executed when updating an object reference (i.e. when executing the `putfield`, `putstatic`, `aputfield_quick`, `aputstatic_quick`, `aastore`, or `aastore_quick` bytecode). Then, we introduce the write barrier pseudo-code shown in Figure 6 into the interpretation of these bytecodes, where the `colour()` function gives the colour of the object parameter and the `greyObject(Y)` procedure unlinks the `Y` object from the white list linking it to the grey list.

Figure 7. Objects outside the heap are allocated red.

2.4. The GC and MR interaction

Since objects allocated within regions may contain references to objects within the heap, the GC must take into account these external references, adding them to its reachability graph. To facilitate this task, we introduce a fourth colour (e.g. *red*) meaning that the object is allocated outside the heap (see Figure 7). Hence, we introduce a new invariant as follows.

Definition (Fourth-colour invariant). There are no red objects within the heap, and there are no black, grey or white objects outside the heap.

The fourth colour allows us to detect when the X object must be added to the root set of the collector, where the root list is updated in a similar way to how generational collectors maintain *inter-generational pointers* (i.e. by using write barriers).

Note that when using a write-barrier-based GC, read barriers detecting memory access errors are not strictly necessary because read operations do not change the colour of the object, therefore they do not affect the GC coherence [13]. The restriction on critical tasks can be reduced to write barriers checks since reads do not interfere with the GC. We apply the same optimization as for the incremental GC, which is to use write barriers instead of read barriers. In this case, the `MemoryAccessError()` exception that is raised when a critical task attempts to access an object X within the heap, is changed

```
if (colour(Y) ≠ red)
  if (type(τ) = critical) illegalAssignment()
  else if (colour(X) = red) updateRootSet(X, Y)
      else if ((colour(X) = black) and (colour(Y) = white)) greyObject(Y);
```

Figure 8. Write barrier using the red colour.

by the `IllegalAssignment()` exception that is raised when a critical task attempts to create a reference to an object `Y` which belongs to the heap. Whereas the former solution, which is based in both write and read barriers, is 100% compliant with the abstract RTSJ specification, the introduced optimization, which consists of using write barriers in order to detect memory access errors instead of using read barriers, reduces the degree of compliance of our proposed implementation solution with RTSJ.

Whereas for non-critical tasks a reference from a red object (`X`) to another object (`Y`) allocated within the heap (i.e. non-red) causes the addition of the `X` object to the root set of the collector, for critical tasks a reference to a non-red object (i.e. white, black or grey) causes an `IllegalAssignment()` exception. We then introduce the code of Figure 8, where the `updateRootSet(X, Y)` procedure links the `X` object to the root list greying the `Y` object whether or not it is white. When the collector explores an object outside the heap (i.e. a root), that has lost its references into the heap, it is eliminated from the root set. When a scoped MR is freed, all objects within the region having references to the objects within the heap are removed from the root list of the collector.

As in our solution, objects within the heap are not reallocated, i.e. we use a non-moving collector (no-copying-based and without compactation), and read barriers can be avoided. In this case, we change write barriers detecting memory access errors for write barriers detecting illegal assignments from critical tasks to objects within the heap. This modification is not compliant with the RTSJ specification, but introduces an important improvement: about 6% of the introduced overhead (i.e. whereas 11% of executed bytecodes perform a load operation, 5% perform a store operation).

3. IMPROVING PERFORMANCE

The most common approach to implement write barriers is by inline code, which consists of generating the instructions to execute write barrier events for every store operation. This solution requires compiler cooperation (e.g. JIT), and presents a serious drawback because it nearly doubles the application's size [14]. Regarding systems with limited memory such as PDAs, this code expansion overhead is considered prohibitive. Alternatively, we can instrument the bytecode interpreter, avoiding space problems, but this still requires a complementary solution to handle native code. A solution that minimizes the write barrier overhead consists of improving the write barrier performance by using hardware support such as the picoJava-II microprocessor [15], which allows write barrier checks to be performed in parallel with the store operation. This alternative solution was the subject of [9].

3.1. Using the write barrier support of the picoJava-II microprocessor

Upon each instruction execution, the picoJava-II core checks for conditions that cause a trap. From the standpoint of hardware support for garbage collection, this microprocessor checks for the occurrence of write barriers, and notifies them using the `gc_notify` trap. This trap is triggered under certain conditions when to a reference field of some object (or an element of an array) is assigned a new reference. Our proposal includes the use and adaptation of existing hardware support. We use two types of write barriers: those typically used in generational collectors, based on the region to which the object belongs, and those typically used in incremental collectors, based on the colour of the object. A combination of both techniques, is further used.

The page-based barrier mechanism of picoJava-II was designed specifically to assist collectors based on the train algorithm. This will trap objects when, within a memory area divided into a number of fixed-size spaces that is also divided into a number of fixed-size *cars*, an object (X) references another object (Y) located in the same space but in a different car. We can use this mechanism to detect references across different MRs (i.e. inter-region references), by using only a memory space and considering that each car is a memory region. Configuring the `gc_notify` signal to trap when inter-region references, avoiding the execution of write barrier codes for intra-region assignments.

The reference-based write barriers of picoJava-II were designed to implement incremental collectors based on the tri-colour algorithm like the one we use. So we must configure the `gc_notify` signal to trap when the application attempts to make a reference from a black object (X) to a white one (Y). This mechanism allows us to improve the performance of both the collector and the application by disabling the write barriers execution in order to preserve the tri-colour invariant when disabling the collector. Similarly, we use the reference-based write barrier mechanism to detect when the root set of the collector must be updated (i.e. when the X object is red and the Y object is not red). Also, in order to avoid illegal assignments of critical tasks to objects within the heap, we must configure the `gc_notify` signal to trap whether the Y object is red. In this way, reference-based write barriers avoid the execution of write barrier code when the object assignment does not attempt to either violate the tri-colour invariant, create a new root for the GC or create a possible interaction of the GC and a critical task.

Since the conditions upon the `gc_notify` traps are different for critical and non-critical tasks, we must configure the reference-based write barrier each time that a task is scheduled to execute (i.e. at context-switch time).

3.2. The write barrier trap handler

When handling a write barrier we can distinguish between three main conditions depending on the MR of the referenced object (i.e. the Y object): (A) when it is within the heap; (B) when it is within the immortal region; and (C) when it is within a scoped region. Since the exception treatment and the configuration for the reference-based write barriers of the picoJava-II microprocessor are different depending on whether the task is critical or not, we analyse the above conditions for both critical tasks and non-critical tasks. Table III shows the actions that we must execute when an object Y is assigned to another object X.

Table III. Actions for inter-region references.

Category condition	X object region	Y object region	Action	
			Threads and tasks	Critical tasks
A	heap	heap	Tri-colour invariant	Illegal assignment
	immortal	heap	Update root set	Illegal assignment
	scoped	heap	Update root set	Illegal assignment
B	heap	immortal	Allowed reference	Allowed reference
	immortal	immortal	Allowed reference	Allowed reference
	scoped	immortal	Allowed reference	Allowed reference
C	heap	scoped	Illegal assignment	Illegal assignment
	immortal	scoped	Illegal assignment	Illegal assignment
	scoped	scoped	Nested regions (when regions are different)	Nested regions (when regions are different)

- Condition A (i.e. $\text{region}(Y) = \text{heap}$): this is not allowed for critical tasks. For non-critical tasks (i.e. threads and high-priority tasks), we must make the distinction depending on the colour of the object that makes the reference (i.e. the X object): when it is black and the Y object is white (A.1) we must avoid the tri-colour invariant violation by greying the Y object; when it is red (A.2), we must maintain the root set of the collector by including the X object into it and greying the Y object if it is white.
- Condition B (i.e. $\text{region}(Y) = \text{immortal}$): this is allowed and does not require any treatment.
- Condition C (i.e. $\text{region}(Y) = \text{scoped}$): we make the distinction depending on the MR of the X object. When it is within the immortal region or the heap (C.1), we must avoid an illegal assignment by throwing the `IllegalAssignment()` exception. When it is within a scoped region different than the region to which the Y object belongs (C.2), we must detect an illegal assignment by exploring the region stack associated with the active task.

At this point there are two possibilities, the first is to change the address of the exception vector value for the `gc_notify` exception at context-switch time, depending on the type of the scheduled task: critical (see Figure 9) or non-critical (see Figure 10). The other possibility consists of merging the handler for both critical and non-critical tasks. The former solution increases the context-switch overhead. However, it is more efficient than the latter, which introduces a write barrier overhead each time that a non-critical task makes an assignment to an object within the common heap. Our objective is to minimize the write barrier overhead, but also to prioritize critical tasks, therefore we chose the latter solution.

```

gc_notify_critical
  if (region(Y) = scoped)
    if (region(X) = scoped) nestedRegions(X,Y) // C.2
    else illegalAssignment(); // C.1
  if (region(Y) ≠ red) illegalAssignment(); // A
  priv_ret_from_trap // B

```

Figure 9. Write barrier handler for critical tasks.

```

gc_notify_non_critical
  if (region(Y) = scoped)
    if (region(X) = scoped) nestedRegions(X,Y) // C.2
    else illegalAssignment(); // C.1
  if (region(Y) ≠ red) if (colour(X) = red) updateRootSet(X, Y) // A.2
    else if ((colour(X) = black) and (colour(Y) = white)) greyObject(Y); // A.1
  priv_ret_from_trap // B

```

Figure 10. Write barrier handler for non-critical tasks.

3.3. Modifying the hardware support

Note that two different mechanisms detect the aforementioned conditions: (i) conditions B and C are detected by the page-based write barrier mechanism; (ii) conditions A and A.1 are detected by reference-based write barriers; and (iii) condition A.2 is detected by both mechanism reference-based and page-based write barriers. The hardware support of picoJava-II throws the `gc_notify` with priority level 14 for both reference-based and page-based write barriers. Since we must treat each condition in a different way, it is pretty interesting to make a distinction by hardware as to whether the trap is caused by a reference-based condition or a page-based condition, therefore we propose to introduce a new signal for distinguishing between them. The proposed modification further requires the addition of the three following entries in the exception vectors table of the picoJava-II microprocessor: (i) `gc_notify_0`, which traps with priority level 14 upon reference-based write barriers; (ii) `gc_notify_1`, which traps with priority level 13 upon page-based write barriers; and (iii) `gc_notify_1_0`, which traps with priority level 12 upon both conditions. The introduction of these three exceptions improves the system performance because it avoids the cause of the `gc_notify` trap being analysed by software.

4. EVALUATING THE OVERHEAD

In this section, we first review the three different write barrier implementations that we have proposed to support the RTSJ memory model. Next, we estimate the write barrier overhead introduced by both the collector and memory regions in the proposed solutions. A complementary analysis of the parameters characterizing the behaviour of the GC can be found in [16,18].

```
if (region(Y) = scoped)
    if (region(X) = scoped) nestedRegions(X,Y)           // C.2
    else illegalAssignment();                             // C.1
if (region(Y) ≠ red)
    if (r = critical) illegalAssignment();               // A
    else if (colour(X) = red) updateRootSet(X, Y)        // A.2
        else if ((colour(X) = black) and (colour(Y) = white)) greyObject(Y); // A.1
priv_ret_from_trap                                     // B
```

Figure 11. Write barrier handler for both critical and non-critical tasks.

```
gc_notify
if (region(Y) = scoped)
    if (region(X) = scoped) nestedRegions(X,Y)           // C.2
    else illegalAssignment();                             // C.1
if (region(Y) ≠ red)
    if (r = critical) illegalAssignment();               // A
    else if (colour(X) = red) updateRootSet(X, Y)        // A.2
        else if ((colour(X) = black) and (colour(Y) = white)) greyObject(Y); // A.1
priv_ret_from_trap                                     // B
```

Figure 12. Handling the `gc_notify` exception.

4.1. Write barrier implementations solutions

4.1.1. Solution 1. Modifying the Java interpreter. This solution consists of modifying the JVM by introducing the code given in Figure 11 into the interpretation of each bytecode whose function consists of assigning an object *Y* to another object *X*[§].

4.1.2. Solution 2. Using existing hardware. We improve the performance of Solution 1 by using the write barrier support of the picoJava-II microprocessor, as proposed in [9]. In this solution, write barriers must be configured at context-switch time depending on the scheduled task. Non-critical tasks throw the `gc_notify` exception when a white object is assigned to a black one, or when an object is assigned to another object allocated in a different MR. Critical tasks throw the `gc_notify` exception when the assigned object is within the heap, or a different MR to the other one. The code executed by the `gc_notify` exception handler is the same as that introduced in the interpreter in the former solution (see Figure 12).

[§]The bytecodes causing write barriers are `putfield`, `putstatic`, `aputfield_quick`, `aputstatic_quick`, `aastore` and `aastore_quick`.

```

gc_notify_l0:
  if (τ ≠ critical) updateRootSet(X,Y) else illegalAssignment(); // A.2
  priv_ret_from_trap
gc_notify_l:
  if (region(Y) = scoped) nestedRegions(X,Y); // C
  priv_ret_from_trap // B
gc_notify_0:
  if (τ ≠ critical) greyObject(Y) else illegalAssignment(); // A.1
  priv_ret_from_trap

```

Figure 13. Write barrier exception handlers.

4.1.3. Solution 3. Modifying the existing hardware. This solution modifies the hardware support of picoJava-II to have three different traps (see Figure 13). In this solution, non-critical tasks cause the execution of: (i) the `gc_notify_l0` exception when a non-red object is assigned to a red one; (ii) the `gc_notify_l` exception when any object is assigned to another one allocated in a different MR; and (iii) the `gc_notify_0` exception when a white object is assigned to a black one. Critical tasks also cause the `gc_notify_0` exception when a non-red object is assigned.

4.2. Quantifying the overhead

All of the objects created in Java are allocated within the JVM heap (i.e. dynamic memory, which in RTSJ may be either the heap or another MR); only primitive types are allocated in the runtime stack [19]. In most applications of the SPECjvm98 benchmark suite[¶], less than half (e.g. 45%) of the references are to objects within the heap rather than primitive types (e.g. bytes or integers); the other half is to either the *Java* or the *native stack*. We also note that about 35% of the total executed bytecodes require an object reference, where typically 70% is for load operations and 30% for store operations [20]. Then, 15% (i.e. 0.45×0.35) of the bytecodes reference an object within the heap, where 30% make an assignment operation. As such references to objects within MRs require write barriers when assignment operations, 5% (i.e. 0.15×0.30) of the bytecodes require the execution of a write barrier code.

To obtain the write barrier overhead for the solutions given in Section 4.1, two measures are combined: (i) the number of events (*NEvents*); and (ii) the measured cost of the event (*ECost*). We also take into account the percentage of bytecodes requiring write barriers, which has been evaluated as 5%. Then, we compute the total write barrier overhead introduced by both MRs and the GC:

$$\begin{aligned}
 MR_{\text{Overhead}} &= NEvents_{\text{MR}} \times ECost_{\text{MR}} + NEvents_{\text{scoped}} \times ECost_{\text{scoped}} \\
 GC_{\text{Overhead}} &= NEvents_{\text{GC}} \times ECost_{\text{GC}} + NEvents_{\text{incGC}} \times ECost_{\text{incGC}}
 \end{aligned}$$

[¶]<http://www.spec.org/osg/jvm98>.

Table IV. Memory reference behaviour.

	Executed bytecodes	Object accesses	Object accesses (%)	Heap references (%)
JESS	1820×10^6	707×10^6	38.84	39.40
DB	3700×10^6	1464×10^6	39.56	45.61
JAVAC	1953×10^6	724×10^6	37.07	28.70
MTRT	2122×10^6	575×10^6	27.09	50.97
JACK	2996×10^6	1022×10^6	34.11	50.74

Where $ECost_{MR}$, $ECost_{scoped}$, $ECost_{GC}$, and $ECost_{incGC}$ parameters can be obtained as

$$\begin{aligned}
 ECost_{MR} &= \text{MemoryArea.getWriteBarrierOverhead}() \\
 ECost_{scoped} &= \text{ScopedMemory.getWriteBarrierOverhead}(n) \\
 ECost_{GC} &= \text{GarbageCollector.getWriteBarrierOverhead}() \\
 ECost_{incGC} &= \text{IncrementalGC.getWriteBarrierOverhead}()
 \end{aligned}$$

4.3. Event parameters

To quantify the write barrier overhead, we are interested in fixing a maximum bound for the number of events that: (i) make an inter-region assignment; (ii) explore the region stack; (iii) create an external reference for the collector; and (iv) attempt to break the tri-colour invariant.

Notation. We assume here that each object has an equal probability to being referenced. Let r , b , g and w be, respectively, the number of red, black, grey and white objects, and h , i and s be the number of objects within the heap, an immortal region or a scoped region, respectively, found in the system at a given instant. Further, let x and z denote the number of inter-region and intra-region assignments, respectively, found in m assignments made by the task τ .

Axiom 1. We have

$$\frac{h}{x} + \frac{i}{x} + \frac{s}{x} = 1$$

In x inter-region assignments of the task τ , there are h assignments from the heap, i assignments from an immortal region and s assignments from a scoped region.

Axiom 2. We have

$$\frac{b}{h} + \frac{g}{h} + \frac{w}{h} = 1$$

In h objects within the heap there are b black objects, g grey objects, and w white objects.

THEOREM. The probability that a task τ breaks the tri-colour invariant when making m assignments is bounded by $0.25 \times h$.

Table V. Maximum bound on write barrier events.

$NEvents$	Solution 1	Solution 2	Solution 3
$NEvent_{MR}$	1	$\frac{x}{m} + 0.25\frac{h}{m}$	$\frac{x}{m} - \frac{h}{m}$
$NEvent_{scoped}$	$\frac{s}{m}$	$\left(\frac{x}{m} + 0.25\frac{h}{m}\right)\left(\frac{s}{m}\right)$	$\frac{s}{m}$
$NEvent_{GC}$	1	$\frac{x}{m} + 0.25\frac{h}{m}$	$\frac{h}{m}$
$NEvent_{incGC}$	$1 - \frac{x}{m}$	$\left(\frac{x}{m} + 0.25\frac{h}{m}\right)\left(1 - \frac{x}{m}\right)$	$0.25\frac{h}{m}$

Proof. We have $h = b + g + w$. We can further express the probability to break the tri-colour invariant as

$$\frac{b \times w}{h^2} = \frac{b \times (h - (b + g))}{h^2}$$

this probability is maximum when there are no grey objects in the system (i.e. $h = b + w$). Then $b \times (h - (b + g)) \leq b \times h - b^2$, where the $b \times h - b^2$ expression takes its maximum value for $b = h/2$ (i.e. $0 = h - 2 \times b$) and $w = h/2$ (i.e. $h = h/2 + w$). $\square$

Note that for critical tasks, the overhead due to the GC is zero (i.e. $Event_{GC}$ and $Event_{incGC}$ are zero, otherwise the `IllegalAssignment()` exception increases). Then we estimate the maximum probability to execute the write barrier code when a non-critical task makes an assignment, as given in Table V.

4.4. Cost parameters

We consider the event cost as the execution time expended by the introduced write barrier code per assignment, i.e. *writeBarrierCost/assignmentCost* where the *writeBarrierCost* is the execution time of the introduced write barriers and the *assignmentCost* is the execution time of an object assignment. The write barrier cost is proportional to the number of evaluated conditions. The execution time taken by both the `greyObject(Y)` and the `updateRootSet(X, Y)` functions is considered as part of the GC overhead rather than as part of the write barrier overhead. For scoped regions, we must consider further the cost of having nested scoped levels, i.e. the cost of executing the `nestedRegions(X, Y)` function, which is proportional to the number of inner levels of the region to which the *Y* object belongs in the region-stack. Recall that in the first step of the algorithm, the region-stack is explored from the top to the bottom to find the region of the *X* object. Suppose that the number of explored levels is x and the region stack has n levels. In the second step of the algorithm, the region stack is explored from the region found in the previous step (i.e. the $n - x$ inner level) to find the region of the *Y* object. Suppose that the number of explored levels is y (i.e. it is found at the $n - x - y$ inner level). Since it is

Table VI. Evaluated conditions for write barrier.

$maxConditions$	Solutions 1 and 2	Solution 3
$maxConditions_{MR}$	2	1
$maxConditions_{scoped}$	n	n
$maxConditions_{GC}$	3	1
$maxConditions_{incGC}$	2	1

evident that $n \geq 0$, then $x + y \geq n$. We conclude by taking n as the maximum bound of the executed evaluations to check whether a region is further out than another. Then we bound the cost parameters as

$$writeBarrierCost = maxConditions \times \frac{conditionCost}{assignmentCost}$$

where the $maxConditions$ parameter is the maximum number of evaluated conditions to check whether the following actions should be executed: (i) call `nestedRegions(X, Y)`; (ii) execute `nestedRegions(X, Y)`; (iii) call `updateRootSet(X, Y)`; and (iv) call `greyObject(Y)`. The $conditionCost$ parameter is the execution time to evaluate a condition. Table VI gives the maximum value for the number of evaluated conditions, where n is the maximum number of nested scoped levels.

Let $MR_{Overhead_i}$ and $GC_{Overhead_i}$ ($1 \leq i \leq 3$) be the $MR_{Overhead}$ and $GC_{Overhead}$ parameters of each given solution, then:

$$\begin{aligned}
 MR_{Overhead_1} &< \left(2 + n \frac{s}{m}\right) \times \frac{conditionCost}{assignmentCost} \\
 MR_{Overhead_2} &< \left(\frac{x}{m} + 0.25 \frac{h}{m}\right) \times MR_{Overhead_1} \\
 MR_{Overhead_3} &< \left(\frac{i}{m} + (n+1) \frac{s}{m}\right) \times \frac{conditionCost}{assignmentCost} \\
 GC_{Overhead_1} &< \left(3 + 2 \left(1 - \frac{x}{m}\right)\right) \times \frac{conditionCost}{assignmentCost} \\
 GC_{Overhead_2} &< \left(\frac{x}{m} + 0.25 \frac{h}{m}\right) \times GC_{Overhead_1} \\
 GC_{Overhead_3} &< 1.2 \frac{h}{m} \times \frac{conditionCost}{assignmentCost}
 \end{aligned}$$

4.5. Comparison

In Solution 1, the write barrier code is executed for all references. Solution 2 reduces the cost of write barriers for intra-region references to the cost of maintaining the tri-colour invariant (i.e. by a factor of $x/m + 0.25(h/m)$). This is because the `gc_notify` exception only traps when a task makes an inter-region reference or attempts to violate the tri-colour invariant. Solution 3 also reduces the cost

for inter-region references to the cost of detecting both of the illegal assignments, when the assigned object is outside the heap and the root set updates when the referenced object is within the heap.

Note that for hardware-based solutions (i.e. Solutions 2 and 3) we must take into account the time that the picoJava-II microprocessor spends to catch a trap. Recall also that the write barrier overhead introduced by scoped regions is the execution time of the `nestedRegions(X,Y)` function. Then, to bind it, we must bind the number of nested region levels.

5. IMPLEMENTATION DETAILS

We have modified the KVM GC (Version 1.0.1) making it a stack-based tri-colour algorithm. We have implemented the `IncrementalGC` class within the KVM by modifying some files (i.e. the `garbage.c` file to implement the collector algorithm and the `interpreter.c` file to implement the write barriers, as well as the `native.h` and the `nativeCore.c` files, which support the interface for the native methods). This class supports the method related with parameters characterizing the collector behaviour (i.e. `getMinimumReclamationRate()`, `setReclamationRate()`, `getWriteBarrierOverhead()` and `getPreemptionLatency()`). We have only implemented three types of memory regions: the heap that is collected by an incremental GC; immortal that are never collected and cannot be nested; and scoped that have limited lifetime and can be nested. These regions are supported by the `HeapMemory`, the `ImmortalMemory` and the `ScopedMemory` classes. Unlike RTSJ, in our prototype the `ScopedMemory` class is a non-abstract class. RTSJ does not consider the write barrier overhead for MRs, then we add the `getWriteBarrierOverhead()` method to the `MemoryArea` class, which gives the cost of detecting illegal assignments between different types of MRs. In the same way, we add the `getWriteBarrierOverhead(int n)` method to the `ScopedMemory` class, which identifies the write barrier cost to have n nested levels for scoped regions.

Instead of using the SPECjvm98 benchmark, which is not compatible with the KVM, we use an artificial collector benchmark. This is an adaptation made by Hans Boehm from the Ellis and Kovac benchmark that we can found at http://www.hpl.hp.com/personal/Hans_Boehm/gc. Two data structures of the same size are kept around during the entire process: (i) a tree containing many pointers; and (ii) a large array containing double precision floating point numbers, which we have modified to contain integers to make it compatible with the KVM. This benchmark executes 262×10^6 bytecodes and allocates 408 MBytes. The number of executed bytecodes performing the write barrier test is 15×10^6 (i.e. `aastore`: 1×10^6 , `putfield`: 6×10^6 , `putfield_fast`: 7×10^6 , `putstatic`: 1 and `putstatic_fast`: 0) for a total of 262×10^6 executed bytecodes. This means that 5% of the executed bytecodes perform a write barrier test.

5.1. Memory footprint

We have limited the number of regions to 256. Then, we need to add a word to the object header to include the following fields: `REGION_ID` $\langle 11:4 \rangle$, `REGION_TYPE` $\langle 3:2 \rangle$ and `COLOUR` $\langle 1:0 \rangle$. Where the `REGION_ID` field specifies the MR to which the object belongs, the `REGION_TYPE` specifies the MR type (e.g. 00 for the heap, 01 for immortal and 10 for scoped), and the `COLOUR` field specifies the colour of the object (e.g. 00 for black, 01 for grey, 10 for white and 11 for red).

Table VII. GC register (GC_CONFIG).

Bits	Field	Description
31:21	SPACE_MASK	Allows knowledge of whether both the X and Y objects belong to the same space
20:16	CAR_MASK	Allows knowledge of whether both the X and Y objects belong to the same <i>car</i>
15:0	WB_VECTOR (Write Barrier Vector)	If the corresponding bit is set, the bytecodes <i>putstatic</i> , <i>aputfield_quick</i> , <i>aputstatic_quick</i> , <i>aastore</i> and <i>aastore_quick</i> signal a <i>gc_notify</i> trap

This increases the memory consumption by a word per object. Alternatively, we can modify the original header format of KVM objects (i.e. *SIZE* <31:8>, *TYPE* <7:2>, *MARK_BIT* <1> and *STATIC_BIT* <0>) to support the colour and region of the object (i.e. *SIZE* <31:17>, *REGION_ID* <16:10>, *REGION_TYPE* <9:8>, *TYPE* <7:2> and *COLOUR* <1:0>). The old *MARK_BIT* that is used by the original mark-and-sweep collector of the KVM to mark the object is no longer used because objects are marked by colour. The old *STATIC_BIT* is not used either because it came from an old collector based on the copying algorithm that has been changed in order to make the KVM suitable for small devices. Note that the maximum size of the object has been reduced from 16 Mbytes to 32 Kbytes; given the small average object size that the specJVM applications present (i.e. about 32 bytes), we optimize for small objects. We also maintain a region structure of four words for each MR object in the system with the following format: *REGION_ID* <63:56>, *OUTER_REGION* <55:48>, *REFERENCE_COUNTER* <47:42>, *REGION_TYPE* <41:40>, *INITIAL_SIZE* <39:25> and *MAXIMUM_SIZE* <24:0>, which increases the memory footprint by a maximum of 2 Kbytes. Note that these region structures form a scope-three where the heap is the root and immortal regions are not included.

5.2. Configuration

In picoJava-II, the conditions under which the write barrier trap is generated are governed by the values of the GC register (GC_CONFIG), the configuration of which is summarized in Table VII. The object reference in picoJava-II has four fields (i.e. the GC_TAG field determines whether to signal a write barrier GC trap, the ADDRESS field contains the address of the header object, the X field indicates whether the object is an array and the H indicates whether the object is referenced directly or indirectly). The GC_TAG, X and H fields are masked before the reference is used as an address, and also in comparison instructions (i.e. *if_acmpeq* and *if_acmpne*). For direct references, the object instance variables start one word after the header. For indirect references, we must go through the handle to access the object. Then, the <31-30> and the <1-0> bits of the address are not used to access memory; the 3 Gbytes with greater addresses are reserved for the system and cannot be used by the application, and object headers are mapped into addresses that are multiples of four.

Figure 14. Memory map for MRs (CAR_MASK = 11111).

Since we map each RTSJ MR object to a car of picoJava-II and the CAR field of the object reference indicates the MR object within which the object is allocated, the maximum number of MRs is 32, and each region is composed of several pages, where the maximum number of pages is 2048. The page size is between 16 and 512 KBytes, depending on the number of MRs (i.e. if we have 32 regions, the page size is 16 KBytes, whereas if we have only one region, the page size is 512 KBytes). For each MR, we maintain a list of pages or a map of bits indicating which pages are assigned to the region.

To have a single space, we configure the SPACE_MASK field of the GC_CONFIG register with the value 0000.0000.000, which masks the $\langle 29 : 19 \rangle$ bits in the ADDRESS field of the object reference. Considering the CAR_MASK field of the object reference configured as 1.1111, the same car is repeated each 512 KBytes in the memory map, which means every 32 pages, each belonging to a different MR (see Figure 14). More specifically, page_0 (address \$000000000-\$00003FFF) belongs to the heap, page_1 (address \$00004000-\$00007FFF) belongs to the immortal memory, and pages from page_2 to page_31 belong to 30 different scoped MRs. This mapping is repeated along the overall address space, i.e. page_32 belongs to the heap, page_33 belongs to the immortal memory, and pages from page_34 to page_63 belong to 30 different (immortal physical or scoped) MRs.

Consider the following colour codes: 11, 10, 01 and 00 that denote black, grey, red and white objects, respectively. We obtain the following configuration values.

- Reference-based write barriers for critical tasks avoid accesses to objects within the heap, detecting assignments to objects within the heap (i.e. to black, grey or white objects) which give the following combinations: 1111 (black to black), 1110, 1100, 1011, 1010, 1000, 0111, 0110, 0100 0011, 0010, 0000 (i.e. bits 15, 14, 12, 11, 10, 8, 7, 6, 4, 3, 2 and 0). Then, we must configure the WR_VECTOR with the 1101.1101.0000.1101 value (i.e. \$BB0B).

```
configureWriteBarrierCriticalTasks:      configureWriteBarrierNonCriticalTasks:
    spush 0xB BBBB // reference-based    spush 0x10C0 // reference-based
    seti 0x001F // page-based            seti 0x001F // page-based
    priv_write_gc_config                 priv_write_gc_config
    priv_ret_from_trap                   priv_ret_from_trap
```

Figure 15. Configuring write barriers.

- Reference-based write barriers for non-critical tasks enable us to preserve the tri-colour invariant (i.e. the 1100 combination) and to maintain the root-set (i.e. 0111, 0110, 0100 combinations). Then, we must configure the WR_VECTOR field with the value \$10C0 (i.e. 0001.0000.1101.0000).

In order to use the write barrier hardware aid of picoJava-II, as we have described, we introduce two routines to configure write barriers for both critical and non-critical tasks (see Figure 15). Where the `priv_read_gc_config` and `priv_write_gc_config` extended bytecodes allow access to the GC_CONFIG register.

6. CONCLUSION

A real-time GC avoids the user's need to recycle memory, but introduces high overhead and unpredictable behaviour. Memory regions that can be supported in a stack discipline offer a high level of predictability. The memory regions model of RTSJ combines the advantages of both techniques. This specification imposes restricted assignments rules that keep longer-lived objects from referencing objects in scoped memory, which possibly have a shorter life. This requires runtime checks for each assignment, which introduces a high overhead. In this paper, we have proposed a solution to the realization of the abstract memory model introduced by the RTSJ specification. In particular, garbage collection in the heap complies with real-time constraints by using write barriers to maintain both the root-set and the tri-colour invariant.

In our solution, the detection of illegal assignments related with MRs and illegal accesses related with critical tasks are made dynamically by introducing a write barrier mechanism based on a region stack associated with the active task. We improve the performance of our solution by using the write barrier support of the picoJava-II microprocessor, as proposed in [9]. In this solution, an exception trap for inter-region references must be configured at context-switch time depending on the scheduled task, because non-critical tasks trap when a white object is assigned to a black one, whereas critical tasks trap when the assigned object is not red. We also propose a modification to the hardware support of picoJava-II to have three different traps: (i) to preserve the root set of the collector; (ii) to detect illegal assignments; and (iii) to preserve the tri-colour invariant.

Hardware-based implementations are efficient, but quite inflexible. We must configure the system to determine the virtual region memory map. In addition, our solution requires the size of a region to be a multiple of the car size, which may possibly introduce internal fragmentation. Finally, for a VTMemory scoped region that can change its size up to its *maximumSize*, the additional memory must

be assigned in terms of cars. This problem can be unpractical for classes dealing with I/O mapped memory (e.g. `ScopedPhysicalMemory`), which specify not only the *size* of the region, but also the *base* address in their constructor.

Another problem with this is that we omit write barriers in native code, which may be addressed using either of the following solutions: forcing the native code to register their writes explicitly, or using virtual memory protection to detect and register changes. The latter solution needs further investigation because it is not trivial to combine real-time bounded collection with barriers supported in the MMU.

To support critical applications in RTSJ, the GC of the heap must be disabled and all MRs (i.e. scoped and immortal physical) must be created at initialization time. In this way, the application runs with static memory. The prototype given in [21] avoids both collectors, those collecting objects within the heap and those collecting unused scoped regions, which facilitates an accurate pre-runtime analysis of the memory management behaviour. In general, static analysis is used to determine the worst case and performance measures are used to determine the average case. It is also advisory to use a combination of both analytical and empirical techniques to evaluate resource usage and to optimize resource consumption.

ACKNOWLEDGEMENTS

This work has been partially funded by Texas Instruments. The authors gratefully acknowledge the comments and guidance of Miguel Angel de Miguel.

REFERENCES

1. The Real-Time for Java Expert Group (RTJEG). Real-Time Specification for Java. <http://www.rtg.org> [June 2000].
2. J Consortium Inc. Core Real-Time Extensions for the Java Platform. <http://www.j-consortium.org/rtjwg.html> [August 1999].
3. Higuera-Toledano MT, Issarny V, Banâtre M, Cabillic G, Lesot JP, Parain F. Java embedded real-time systems: An overview of existing solutions. *Proceedings 3th International Symposium on Object-Oriented Real-Time Distributed Computing (ISORC)*. IEEE Computer Society Press: Los Alamitos, CA, 2000.
4. Sun Microsystems. KVM Technical Specification. <http://java.sun.com> [May 2000].
5. McGhan H, O'Connor M. picoJava: A direct execution engine for Java bytecode. *IEEE Computer* 1998; **31**(2):22–30.
6. Bollella G, Gosling J. The real-time specification for Java. *IEEE Computer* 2000.
7. Tofte M. Implementing the call-by-value lambda-calculus using a stack of regions. *Proceedings of the Conference of Programming Language Design and Implementation (PLDI) (ACM SIGPLAN)*, January 1994.
8. Higuera-Toledano MT, de Miguel MA. Dynamic detection of access errors and illegal references in RTSJ. *Proceedings 8th IEEE Real-time and Embedded Technology and Applications Symposium (RTAS)*. IEEE Computer Society Press: Los Alamitos, CA, 2002.
9. Higuera-Toledano MT, Issarny V, Banâtre M, Cabillic G, Lesot JP, Parain F. Region-based memory management for real-time Java. *Proceedings 4th International Symposium on Object-Oriented Real-Time Distributed Computing (ISORC)*. IEEE Computer Society Press: Los Alamitos, CA, 2000.
10. Cannarozzi DJ, Plezbert MP, Cytron RK. Contaminated garbage collection. *Proceedings of the Conference of Programming Language Design and Implementation (PLDI) (ACM SIGPLAN)*, May 2000.
11. Dijkstra EW, Lamport L, Martin AJ, Scholtenand CS, Steffens EFM. On-the-fly garbage collection: An exercise in cooperation. *Communications of the ACM* 1978; **21**(11).
12. Wilson PR, Johnstone MS. Real-time non-copying garbage collection. *Workshop on Garbage Collection and Memory Management*. ACM Press: New York, 1993.
13. Higuera MT, Issarny V, Banâtre M, Cabillic G, Lesot JP, Parain F. Memory management for real-time Java: An efficient solution using hardware support. *Real-Time Systems Journal* (to appear).
14. Zorn B. Barrier Methods for Garbage Collection. <http://www.cs.colourado> [1990].

15. Sun Microsystems. picoJava-II Programmer's Reference Manual. <http://www.sun.com/microelectronics/picoJava> [March 1999].
16. Higuera-Toledano MT, Issarny V. Analyzing the performance of memory management in RTSJ. *Proceedings 5th International Symposium on Object-Oriented Real-Time Distributed Computing (ISORC)*. IEEE Computer Society Press: Los Alamitos, CA, 2000.
17. Petit-Bianco A, Tromeu T. Garbage collection for Java in embedded systems. *Proceedings IEEE Workshop on Programming Languages for Real-Time Industrial Applications*, 1998.
18. Higuera-Toledano MT. *Memory Management Solutions for Real-time Java*. INRIA-Rocquencourt, 2002.
19. Gay D, Steensgaard B. Stack Allocating Objects in Java. <http://www.research.microsoft.com/apl> [1998].
20. Kim JS, Hsu Y. Memory system behavior of Java programs: Methodology and analysis. *Proceedings of the ACM Java Grande 2000 Conference*, June 2000.
21. Puschner P, Wellings A. A profile for high-integrity real-time Java programs. *Proceedings 4th International Symposium on Object-Oriented Real-Time Distributed Computing (ISORC)*. IEEE Computer Society Press: Los Alamitos, CA, 2001.
22. Fresko N, Foote W. Proposed Java Real Time Extension Specification. <http://www.sdct.itl.nist.gov/~carnahan/real-time/sun/index.html> [1998].
23. Standard Performance Evaluation Council. SPEC JVM98 Benchmarks. <http://www.spec.org/osg/jvm98> [1998].
24. Gay D, Aiken A. Memory management with explicit regions. *Proceedings of the Conference of Programming Language Design and Implementation (PLDI) (ACM SIGPLAN)*, June 1998.
25. Wilson PR. Uniprocessor Garbage Collection Techniques. <ftp://ftp.cs.utexas.edu/pub/garbage/bigsur.ps>.
26. Nilsen K. Adding real-time capabilities to Java. *Communications of the ACM* 1998; **41**(6).
27. McDowell CE. Reducing Garbage in Java. <http://www.cs.Berkeley.edu> [1997].
28. Henriksson R. Predictable automatic memory management for embedded systems. *Proceedings of the Workshop on Garbage Collection and Memory Management (ACM SIGPLAN-SIGACT)*, October 1997.