

# 1 A Hardware Abstraction Layer in Java

2 MARTIN SCHOEBERL, Vienna University of Technology  
3 STEPHAN KORSHOLM, Aalborg University  
4 TOMAS KALIBERA, Purdue University  
5 ANDERS P. RAVN, Aalborg University

6 Embedded systems use specialized hardware devices to interact with their environment, and since they  
7 have to be dependable, it is attractive to use a modern, type-safe programming language like Java to de-  
8 velop programs for them. Standard Java, as a platform-independent language, delegates access to devices,  
9 direct memory access, and interrupt handling to some underlying operating system or kernel, but in the  
10 embedded systems domain resources are scarce and a Java Virtual Machine (JVM) without an underlying  
11 middleware is an attractive architecture. The contribution of this article is a proposal for Java packages  
12 with hardware objects and interrupt handlers that interface to such a JVM. We provide implementations  
13 of the proposal directly in hardware, as extensions of standard interpreters, and finally with an operating  
14 system middleware. The latter solution is mainly seen as a migration path allowing Java programs to co-  
15 exist with legacy system components. An important aspect of the proposal is that it is compatible with the  
16 Real-Time Specification for Java (RTSJ).

17 Categories and Subject Descriptors: D.4.7 [Operating Systems]: Organization and Design—*Real-time*  
18 *systems and embedded systems*; D.3.3 [Programming Languages]: Language Constructs and Features—  
19 *Input/output*

20 General Terms: Languages, Design

21 Additional Key Words and Phrases: Device driver, embedded system, Java, Java virtual machine

## 22 ACM Reference Format:

23 Schoeberl, M., Korsholm, S., Kalibera, T., and Ravn, A. P. 2011. A hardware abstraction layer in Java. *ACM*  
24 *Trans. Embed. Comput. Syst.* 10, 4, Article 42 (November 2011), 40 pages.  
25 DOI = 10.1145/2043662.2043666 <http://doi.acm.org/10.1145/2043662.2043666>

## 26 1. INTRODUCTION

27 When developing software for an embedded system, for instance an instrument, it is  
28 necessary to control specialized hardware devices, for instance a heating element or an  
29 interferometer mirror. These devices are typically interfaced to the processor through  
30 device registers and may use interrupts to synchronize with the processor. In order  
31 to make the programs easier to understand, it is convenient to introduce a *Hardware*  
32 *Abstraction Layer* (HAL), where access to device registers and synchronization through  
33 interrupts are hidden from conventional program components. A HAL defines an in-  
34 terface in terms of the constructs of the programming language used to develop the

---

The research leading to these results received funding from the European Community's Seventh Framework Programme [FP7/2007-2013] under grant agreement number 216682 (JEOPARD).

Authors' addresses: M. Schoeberl (corresponding author), Vienna University of Technology, Austria; email: mschoebe@mail.tuwien.ac.at; S. Korsholm, Aalborg University, Denmark; T. Kalibera, Purdue University, U.S.A.; A. P. Ravn, Aalborg University, Denmark.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from the Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or [permissions@acm.org](mailto:permissions@acm.org).

© 2011 ACM 1539-9087/2011/11-ART42 \$10.00

DOI 10.1145/2043662.2043666 <http://doi.acm.org/10.1145/2043662.2043666>

35 application. Thus, the challenge is to develop an abstraction that gives efficient ac-  
36 cess to the hardware, while staying within the computational model provided by the  
37 programming language.

38 Our first ideas on a HAL for Java have been published in Schoeberl et al. [2008] and  
39 Korsholm et al. [2008]. This article combines the two papers, provides a much wider  
40 background of related work, gives two additional experimental implementations, and  
41 gives performance measurements that allow an assessment of the efficiency of the  
42 implementations. The remainder of this section introduces the concepts of the Java-  
43 based HAL.

#### 44 1.1 Java for Embedded Systems

45 Over the nearly 15 years of its existence Java has become a popular programming  
46 language for desktop and server applications. The concept of the Java Virtual Ma-  
47 chine (JVM) as the execution platform enables portability of Java applications. The  
48 language, its API specification, as well as JVM implementations have matured; Java  
49 is today employed in large-scale industrial applications. The automatic memory man-  
50 agement takes away a burden from the application programmers and together with  
51 type safety helps to isolate problems and, to some extent, even run untrusted code. It  
52 also enhances security; attacks like stack overflow are not possible. Java integrates  
53 threading support and dynamic loading into the language, making these features eas-  
54 ily accessible on different platforms. The Java language and JVM specifications are  
55 proven by different implementations on different platforms, making it relatively easy  
56 to write platform-independent Java programs that run on different JVM implemen-  
57 tations and underlying OS/hardware. Java has a standard API for a wide range of  
58 libraries, the use of which is thus again platform-independent. With the ubiquity of  
59 Java, it is easy to find qualified programmers who know the language, and there is  
60 strong tool support for the whole development process. According to an experimental  
61 study [Phipps 1999], Java has lower bug rates and higher productivity rates than C++.  
62 Indeed, some of these features come at a price of larger footprint (the virtual machine  
63 is a nontrivial piece of code), typically higher memory requirements, and sometimes  
64 degraded performance, but this cost is accepted in industry.

65 Recent real-time Java virtual machines based on the Real-Time Specification for  
66 Java (RTSJ) provide controlled and safe memory allocation. Also there are platforms  
67 for less critical systems with real-time garbage collectors. Thus, Java is ready to make  
68 its way into the embedded systems domain. Mobile phones, PDAs, or set-top boxes run  
69 Java Micro Edition, a Java platform with a restricted set of standard Java libraries.  
70 Real-time Java has been and is being evaluated as a potential future platform for  
71 space avionics both by NASA and ESA space agencies. Some Java features are even  
72 more important for embedded than for desktop systems because of missing features  
73 of the underlying platform. For instance, the RTEMS operating system used by ESA  
74 for space missions does not support hardware memory protection even for CPUs that  
75 do support it (like LEON3, a CPU for ESA space missions). With Java's type safety  
76 hardware protection is not needed to spatially isolate applications. Moreover, RTEMS  
77 does not support dynamic libraries, but Java can load classes dynamically.

78 Many embedded applications require very small platforms, therefore it is interest-  
79 ing to remove as much as possible of an underlying operating system or kernel, where  
80 a major part of code is dedicated to handling devices. Furthermore, Java is considered  
81 as the future language for safety-critical systems [Henties et al. 2009]. As certifica-  
82 tion of safety-critical systems is very expensive, the usual approach is to minimize the  
83 code base and supporting tools. Using two languages (e.g., C for programming device  
84 handling in the kernel and Java for implementing the processing of data) increases

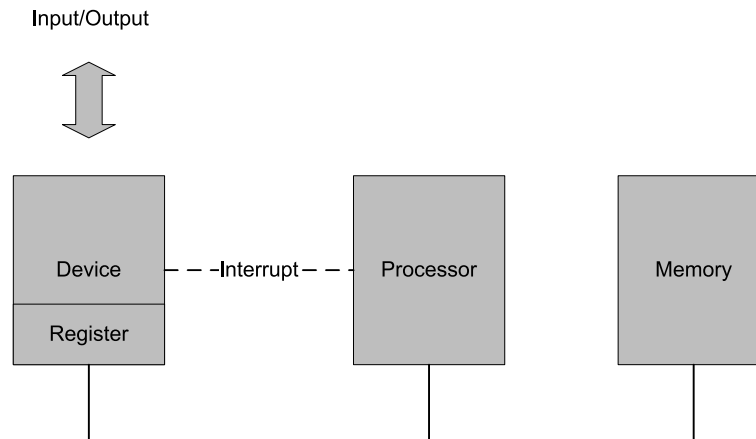


Fig. 1. The hardware: a bus connects a processor to device registers and memory, and an interrupt bus connects devices to a processor.

85 the complexity of generating a safety case. A Java-only system reduces the complexity  
 86 of the tool support and therefore the certification effort. Even in less critical systems  
 87 the same issues will show up as decreased productivity and dependability of the soft-  
 88 ware. Thus it makes sense to investigate a general solution that interfaces Java to the  
 89 hardware platform; that is the objective of the work presented here.

## 90 1.2 Hardware Assumptions

91 The hardware platform is built up along one or more buses (in small systems typically  
 92 only one) that connect the processor with memory and device controllers. Device con-  
 93 trollers have reserved some part of the address space of a bus for its device registers.  
 94 They are accessible for the processor as well, either through special I/O instructions or  
 95 by ordinary instructions when the address space is the same as the one for addressing  
 96 memory, a so called memory mapped I/O solution. In some cases the device controller  
 97 will have Direct Memory Access (DMA) as well, for instance for high-speed transfer of  
 98 blocks of data. Thus the basic communication paradigm between a controller and the  
 99 processor is shared memory through the device registers and/or through DMA. With  
 100 these facilities only, synchronization has to be done by testing and setting flags, which  
 101 means that the processor has to engage in some form of busy waiting. This is elim-  
 102 inated by extending the system with an interrupt bus, where device controllers can  
 103 generate a signal that interrupts the normal flow of execution in the processor and di-  
 104 rect it to an interrupt handling program. Since communication is through shared data  
 105 structures, the processor and the controllers need a locking mechanism; therefore in-  
 106 terrupts can be enabled or disabled by the processor through an interrupt control unit.  
 107 The typical hardware organization is summarized in Figure 1.

## 108 1.3 A Computational Model

109 In order to develop a HAL, the device registers and interrupt facilities must be mapped  
 110 to programming language constructs, such that their use corresponds to the computa-  
 111 tional model underlying the language. In the following we give simple device examples  
 112 which illustrate the solution we propose for doing it for Java.

113 *1.3.1 Hardware Objects.* Consider a simple Parallel Input/Output (PIO) device control-  
 114 ling a set of input and output pins. The PIO uses two registers: the *data register* and

```

public final class ParallelPort {
    public volatile int data;
    public volatile int control;
}

int inval, outval;
myport = JVMMechanism.getParallelPort();
...
inval = myport.data;
myport.data = outval;

```

Fig. 2. The parallel port device as a simple Java class.

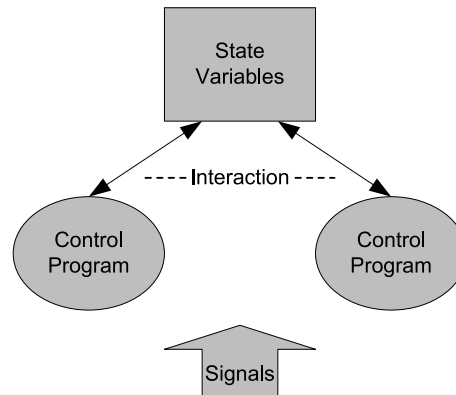


Fig. 3. Computational model: several threads of execution communicate via shared state variables and receive signals.

115 the *control register*. Writing to the data register stores the value into an internal latch  
 116 that drives the output pins. Reading from the data register returns the value that is  
 117 present on the input pins. The control register configures the direction for each PIO  
 118 pin. When bit  $n$  in the control register is set to 1, pin  $n$  drives out the value of bit  $n$  of  
 119 the data register. A 0 at bit  $n$  in the control register configures pin  $n$  as input pin. At  
 120 reset the port is usually configured as input port, a safe default configuration.

121 In an object-oriented language the most natural way to represent a device is as an  
 122 object: the *hardware object*. Figure 2 shows a class definition, object instantiation,  
 123 and use of the hardware object for the simple parallel port. An instance of the class  
 124 `ParallelPort` is the hardware object that represents the PIO. The reference `myport` points  
 125 to the hardware object. To provide this convenient representation of devices as objects,  
 126 a JVM internal mechanism is needed to access the device registers via object fields  
 127 and to *create* the device object and receive a reference to it. We elaborate on the idea  
 128 of hardware objects in Section 3.1 and present implementations in Section 4.

129 **1.3.2 Interrupts.** When we consider an interrupt, it must invoke some program code  
 130 in a method that handles it. We need to map the interruption of normal execution to  
 131 some language concept, and here the concept of an asynchronous event is useful. The  
 132 resulting computational model for the programmer is shown in Figure 3. The signals  
 133 are external, asynchronous events that map to interrupts.

134 A layered implementation of this model with a kernel close to the hardware and  
 135 applications on top has been very useful in general-purpose programming. Here one  
 136 may even extend the kernel to manage resources and provide protection mechanisms  
 137 such that applications are safe from one another, as for instance when implementing

```
public class RS232ReceiveInterruptHandler extends InterruptHandler {
    private RS232 rs232;
    private InterruptControl interruptControl;

    private byte UartRxBuffer[];
    private short UartRxWrPtr;

    ...

    protected void handle() {

        synchronized(this) {
            UartRxBuffer[UartRxWrPtr++] = rs232.P0_UART_RX_TX_REG;
            if (UartRxWrPtr >= UartRxBuffer.length) UartRxWrPtr = 0;
        }
        rs232.P0_CLEAR_RX_INT_REG = 0;
        interruptControl.RESET_INT_PENDING_REG = RS232.CLR_UART_RX_INT_PENDING;
    }
}
```

Fig. 4. An example interrupt handler for an RS232 interface. On an interrupt the method `handle()` is invoked. The private objects `rs232` and `interruptControl` are hardware objects that represent the device registers and the interrupt control unit.

138 trusted interoperable computing platforms [Group 2008]. Yet there is a price to pay  
139 which may make the solution less suitable for embedded systems: adding new device  
140 drivers is an error-prone activity [Chou et al. 2001], and protection mechanisms impose  
141 a heavy overhead on context switching when accessing devices.

142 The alternative we propose is to use Java directly since it already supports multi-  
143 threading and use methods in the special `InterruptHandler` objects to handle interrupts.  
144 The idea is illustrated in Figure 4, and the details, including synchronization and in-  
145 teraction with the interrupt control, are elaborated in Section 3.2. Implementations  
146 are found in Section 4.

#### 147 1.4 Mapping Between Java and the Hardware

148 The proposed interfacing from hardware to Java does not require language extensions.  
149 The Java concepts of packages, classes, and synchronized objects turn out to be power-  
150 ful enough to formulate the desired abstractions. The mapping is done at the level of  
151 the JVM. The JVM already provides typical OS functions handling:

- 152 — address space and memory management;
- 153 — thread management;
- 154 — interprocess communication.

155 These parts need to be modified so they cater for interfaces to the hardware.

156 Yet, the architectures of JVMs for embedded systems are more diverse than on  
157 desktop or server systems. Figure 5 shows variations of Java implementations in  
158 embedded systems and an example of the control flow for a Web server application.  
159 The standard approach with a JVM running on top of an Operating System (OS) is  
160 shown in Figure 5(a).

161 A JVM without an OS is shown in Figure 5(b). This solution is often called *running*  
162 *on the bare metal*. The JVM acts as the OS and provides thread scheduling and low-  
163 level access to the hardware. In this case the network stack can be written entirely in

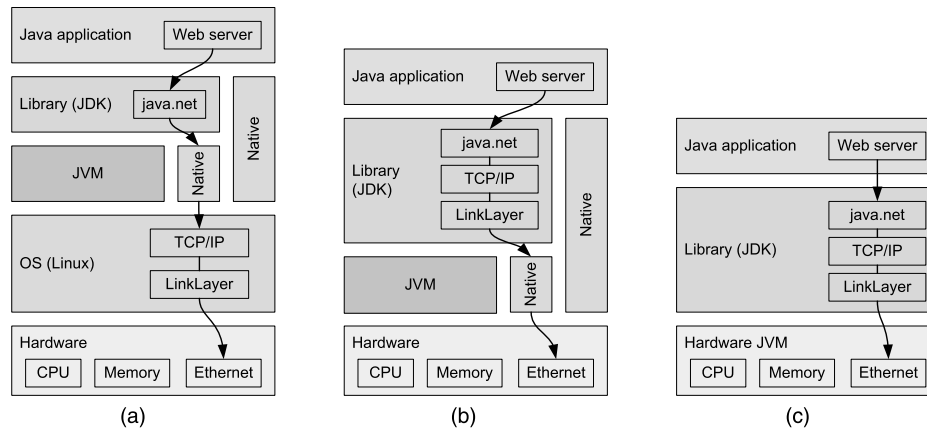


Fig. 5. Configurations for an embedded JVM: (a) standard layers for Java with an operating system, equivalent to desktop configurations; (b) a JVM on the bare metal; and (c) a JVM as a Java processor.

164 Java. JNode<sup>1</sup> is an approach to implement the OS entirely in Java. This solution has  
 165 become popular even in server applications.<sup>2</sup>

166 Figure 5(c) shows an embedded solution where the JVM is part of the hardware  
 167 layer: it is implemented in a Java processor. With this solution the native layer can  
 168 be completely avoided and all code (application and system code) is written entirely in  
 169 Java.

170 Figure 5 shows also the data and control flow from the application down to the  
 171 hardware. The example consists of a Web server and an Internet connection via Ether-  
 172 net. In case (a) the application Web server talks with java.net in the Java library. The  
 173 flow goes down via a native interface to the TCP/IP implementation and the link-layer  
 174 device driver within the OS (usually written in C). The device driver talks with the  
 175 Ethernet chip. In (b) the OS layer is omitted: the TCP/IP layer and the link-layer de-  
 176 vice driver are now part of the Java library. In (c) the JVM is part of the hardware  
 177 layer and direct access from the link-layer driver to the Ethernet hardware is mandatory.

178 With our proposed HAL, as shown in Figure 6, the native interface within the JVM  
 179 in (a) and (b) disappears. Note how the network stack moves up from the OS layer  
 180 to the Java library in example (a). All three versions show a pure Java implemen-  
 181 tation of the whole network stack. The Java code is the same for all three solutions.  
 182 Version (b) and (c) benefit from hardware objects and interrupt handlers in Java as  
 183 access to the Ethernet device is required from Java source-code. In Section 5 we show  
 184 a simple Web server application implemented completely in Java as evaluation of our  
 185 approach.

## 186 1.5 Contributions

187 The key contribution of this article is a proposal for a Java HAL that can run on the  
 188 bare metal while still being *safe*. This idea is investigated in quite a number of places  
 189 which are discussed in the related work section where we comment on our initial ideas  
 190 as well. In summary, the proposal gives an interface to hardware that has the following  
 191 benefits.

<sup>1</sup><http://www.jnode.org/>

<sup>2</sup>BEA System offers the JVM LiquidVM that includes basic OS functions and does not need a guest OS.

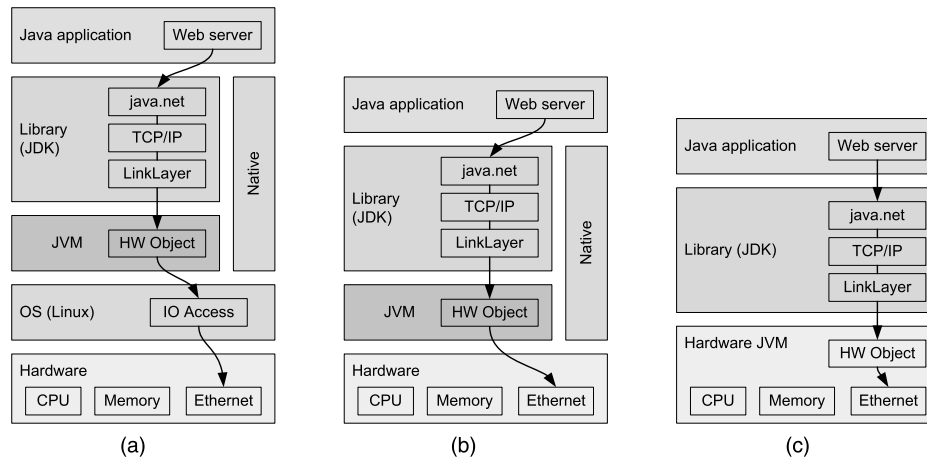


Fig. 6. Configurations for an embedded JVM with hardware objects and interrupt handlers: (a) standard layers for Java with an operating system, equivalent to desktop configurations; (b) a JVM on the bare metal; and (c) a JVM as a Java processor.

192 *Object-oriented.* An object representing a device is the most natural integration into  
 193 an object-oriented language, and a method invocation to a synchronized object is a  
 194 direct representation of an interrupt.

195 *Safe.* The safety of Java is not compromised. Hardware objects map object fields to  
 196 device registers. With a correct class that represents the device, access to it is safe.  
 197 Hardware objects can be created only by a factory residing in a special package.

198 *Generic.* The definition of a hardware object and an interrupt handler is independent  
 199 of the JVM. Therefore, a common standard for different platforms can be defined.

200 *Efficient.* Device register access is performed by single bytecodes `getfield` and `putfield`.  
 201 We avoid expensive native calls. The handlers are first-level handlers; there is no delay  
 202 through event queues.

203 The proposed Java HAL would not be useful if it had to be modified for each partic-  
 204 ular kind of JVM; thus a second contribution of this article is a number of prototype  
 205 implementations illustrating the architectures presented in Figure 6: implementa-  
 206 tions in Kaffe [Wilkinson 1996] and OVM [Armbruster et al. 2007] represent the  
 207 architecture with an OS (Figure 6(a)), the implementation in SimpleRTJ [RTJ Com-  
 208 puting 2000] represents the bare metal solution (Figure 6(b)), and the implementation  
 209 in JOP [Schoeberl 2008] represents the Java processor solution (Figure 6(c)).

210 Finally, we must not forget the claim for efficiency, and therefore the article ends  
 211 with some performance measurements that indicate that the HAL layer is generally  
 212 as efficient as native calls to C code external to the JVM.

## 213 2. RELATED WORK

214 Already in the 1970's it was recognized that an operating system might not be the  
 215 optimal solution for special-purpose applications. Device access was integrated into  
 216 high-level programming languages like Concurrent Pascal [Hansen 1977; Ravn 1980]  
 217 and Modula (Modula-2) [Wirth 1977, 1982] along with a number of similar languages,  
 218 for example, UCSD Pascal. They were meant to eliminate the need for operating sys-  
 219 tems and were successfully used in a variety of applications. The programming lan-  
 220 guage Ada, which has been dominant in defence and space applications to this day,

221 may be seen as a continuation of these developments. The advent of inexpensive mi-  
222 croprocessors, from the mid 1980's and on, lead to a regression to assembly and C  
223 programming. The hardware platforms were small with limited resources and the de-  
224 velopers were mostly electronic engineers, who viewed them as electronic controllers.  
225 Program structure was not considered a major issue in development. Nevertheless, the  
226 microcomputer has grown, and is now far more powerful than the minicomputer that  
227 it replaced. With powerful processors and an abundance of memory, the ambitions for  
228 the functionality of embedded systems grow, and programming becomes a major issue  
229 because it may turn out to be the bottleneck in development. Consequently, there is a  
230 renewed interest in this line of research.

231 An excellent overview of historical solutions to access hardware devices from and  
232 implement interrupt handlers in high-level languages, including C, is presented in  
233 Chapter 15 of Burns and Wellings [2001]. The solution to device register access in  
234 Modula-1 (Chapter 15.3) is very much like C; however, the constructs are safer because  
235 they are encapsulated in modules. Interrupt handlers are represented by threads that  
236 block to wait for the interrupt. In Ada (Chapter 15.4) the representation of individual  
237 fields in registers can be described precisely by *representation* classes, while the corre-  
238 sponding structure is bound to a location using the *Address* attribute. An interrupt is  
239 represented in the current version of Ada by a protected procedure, although initially  
240 represented (Ada 83) by task entry calls.

241 The main ideas in having device objects are thus found in the earlier safe languages,  
242 and our contribution is to align them with a Java model, and in particular, as discussed  
243 in Section 4, implementation in a JVM. From the Ada experience we learn that direct  
244 handling of interrupts is a desired feature.

## 245 2.1 The Real-Time Specification for Java

246 The Real-Time Specification for Java (RTSJ) [Bollella et al. 2000] defines a JVM exten-  
247 sion which allows better timeliness control compared to a standard JVM. The core fea-  
248 tures are: fixed priority scheduling, monitors which prevent priority inversion, scoped  
249 memory for objects with limited lifetime, immortal memory for objects that are never  
250 finalized, and asynchronous events with CPU time consumption control.

251 The RTSJ also defines an API for direct access to physical memory, including hard-  
252 ware registers. Essentially one uses `RawMemoryAccess` at the level of primitive data  
253 types. Although the solution is efficient, this representation of physical memory is not  
254 object-oriented, and there are some safety issues: When one raw memory area rep-  
255 resents an address range where several devices are mapped to, there is no protection  
256 between them. Yet, a type-safe layer with support for representing individual registers  
257 can be implemented on top of the RTSJ API.

258 The RTSJ specification suggests that asynchronous events are used for interrupt  
259 handling. Yet, it neither specifies an API for interrupt control nor semantics of the  
260 handlers. Any interrupt handling application thus relies on some proprietary API  
261 and proprietary event handler semantics. Second-level interrupt handling can be  
262 implemented within the RTSJ with an `AsyncEvent` that is bound to a *happening*.  
263 The happening is a string constant that represents an interrupt, but the meaning  
264 is implementation-dependent. An `AsyncEventHandler` or `BoundAsyncEventHandler` can  
265 be added as handler for the event. Also an `AsyncEventHandler` can be added via a  
266 `POSIXSignalHandler` to handle POSIX signals. An interrupt handler, written in C, can  
267 then use one of the two available POSIX user signals.

268 RTSJ offers facilities very much in line with Modula or Ada for encapsulating  
269 memory-mapped device registers. However, we are not aware of any RTSJ implemen-  
270 tation that implements `RawMemoryAccess` and `AsyncEvent` with support for low-level



271 device access and interrupt handling. Our solution could be used as specification of  
272 such an extension. It would still leave the first-level interrupt handling hidden in an  
273 implementation; therefore an interesting idea is to define and implement a two-level  
274 scheduler for the RTSJ. It should provide the first-level interrupt handling for asyn-  
275 chronous events bound to interrupts and delegate other asynchronous events to an  
276 underlying second-level scheduler, which could be the standard fixed priority preemptive  
277 scheduler. This would be a fully RTSJ-compliant implementation of our proposal.

## 278 2.2 Hardware Interface in JVMs

279 The aJile Java processor [aJile 2000] uses native functions to access devices. Inter-  
280 rupts are handled by registering a handler for an interrupt source (e.g., a GPIO pin).  
281 Systronix suggests<sup>3</sup> to keep the handler short, as it runs with interrupts disabled,  
282 and delegate the real handling to a thread. The thread waits on an object with ceil-  
283 ing priority set to the interrupt priority. The handler just notifies the waiting thread  
284 through this monitor. When the thread is unblocked and holds the monitor, effectively  
285 all interrupts are disabled.

286 Komodo [Kreuzinger et al. 2003] is a multithreaded Java processor targeting real-  
287 time systems. On top of the multiprocessing pipeline the concept of interrupt service  
288 threads is implemented. For each interrupt one thread slot is reserved for the interrupt  
289 service thread. It is unblocked by the signaling unit when an interrupt occurs. A  
290 dedicated thread slot on a fine-grain multithreading processor results in a very short  
291 latency for the interrupt service routine. No thread state needs to be saved. However,  
292 this comes at the cost to store the complete state for the interrupt service thread in the  
293 hardware. In the case of Komodo, the state consists of an instruction window and the  
294 on-chip stack memory. Devices are represented by Komodo-specific I/O classes.

295 Muvium [Caska 2009] is an ahead-of-time compiling JVM solution for very resource-  
296 constrained microcontrollers (Microchip PIC). Muvium uses an Abstract Peripheral  
297 Toolkit (APT) to represent devices. APT is based on an event-driven model for inter-  
298 action with the external world. Device interrupts and periodic activations are repre-  
299 sented by events. Internally, events are mapped to threads with priority dispatched by  
300 a preemptive scheduler. APT contains a large collection of classes to represent devices  
301 common in embedded systems.

302 In summary, access to device registers is handled in both aJile, Komodo, and Mu-  
303 vium by abstracting them into library classes with access methods. This leaves the im-  
304 plementation to the particular JVM and does not give the option of programming them  
305 at the Java level. It means that extension with new devices involves programming at  
306 different levels, which we aim to avoid. Interrupt handling in aJile is essentially first  
307 level, but with the twist that it may be interpreted as RTSJ event handling, although  
308 the firing mechanism is atypical. Our mechanism would free this binding and allow  
309 other forms of programmed notification, or even leaving out notification altogether.  
310 Muvium follows the line of RTSJ and has a hidden first-level interrupt handling. Ko-  
311 modo has a solution with first-level handling through a full context switch; this is very  
312 close to the solution advocated in Modula 1, but it has in general a larger overhead  
313 than we would want to incur.

## 314 2.3 Java Operating Systems

315 The JX Operating System [Felser et al. 2002] is a microkernel system written mostly  
316 in Java. The system consists of components which run in *domains*, each domain

<sup>3</sup>A template can be found at <http://practicalembeddedjava.com/tutorials/aJileISR.html>.

317 having its own garbage collector, threads, and a scheduler. There is one global preemptive  
318 scheduler that schedules the domain schedulers which can be both preemptive  
319 and nonpreemptive. Inter-domain communication is only possible through communication  
320 channels exported by services. Low-level access to the physical memory, memory-mapped  
321 device registers, and I/O ports are provided by the core (“zero”) domain services,  
322 implemented in C. At the Java level ports and memory areas are represented by objects,  
323 and registers are methods of these objects. Memory is read and written by access  
324 methods of Memory objects. Higher layers of Java interfaces provide type-safe  
325 access to the registers; the low-level access is not type safe.

326 Interrupt handlers in JX are written in Java and are run through portals; they can  
327 reside in any domain. Interrupt handlers cannot interrupt the garbage collector (the  
328 GC disables interrupts), run with CPU interrupts disabled, must not block, and can  
329 only allocate a restricted amount of memory from a reserved per-domain heap. Execution  
330 time of interrupt handlers can be monitored: on a deadline violation the handler is  
331 aborted and the interrupt source disabled. The first-level handlers can unblock a  
332 waiting second-level thread either directly or via setting a state of a AtomicVariable  
333 synchronization primitive.

334 The Java New Operating System Design Effort (JNode<sup>4</sup>) [Lohmeier 2005] is an OS  
335 written in Java where the JVM serves as the OS. Drivers are written entirely in Java.  
336 Device access is performed via native function calls. A first-level interrupt handler,  
337 written in assembler, unblocks a Java interrupt thread. From this thread the device  
338 driver-level interrupt handler is invoked with interrupts disabled. Some device drivers  
339 implement a synchronized `handleInterrupt(int irq)` and use the driver object to signal the  
340 upper layer with `notifyAll()`. During garbage collection all threads are stopped including  
341 the interrupt threads.

342 The Squawk VM [Simon et al. 2006], now available open-source,<sup>5</sup> is a platform for  
343 wireless sensors. Squawk is mostly written in Java and runs without an OS. Device  
344 drivers are written in Java and use a form of peek and poke interface to access the  
345 device registers. Interrupt handling is supported by a device driver thread that waits  
346 for an event from the JVM. The first-level handler, written in assembler, disables the  
347 interrupt and notifies the JVM. On a rescheduling point the JVM resumes the device  
348 driver thread. It has to reenables the interrupt. The interrupt latency depends on the  
349 rescheduling point and on the activity of the garbage collector. For a single device  
350 driver thread an average case latency of 0.1 ms is reported. For a realistic workload  
351 with an active garbage collector a worst-case latency of 13 ms has been observed.

352 Our proposed constructs should be able to support the Java operating systems. For  
353 JX we observe that the concepts are very similar for interrupt handling, and actually  
354 for device registers as well. A difference is that we make device objects distinct from  
355 memory objects which should give better possibilities for porting to architectures with  
356 separate I/O buses. JNode is more traditional and hides first-level interrupt handling  
357 and device accesses in the JVM, which may - be less portable than our implementation.  
358 The Squawk solution has to have a very small footprint, but on the other hand it can  
359 probably rely on having few devices. Device objects would be at least as efficient as the  
360 peeks and pokes, and interrupt routines may eliminate the need for multithreading  
361 for simple systems, for example, with cyclic executives. Overall, we conclude that our  
362 proposed constructs will make implementation of a Java OS more efficient and perhaps  
363 more portable.

---

<sup>4</sup><http://jnode.org/>

<sup>5</sup><https://squawk.dev.java.net/>

## 364 2.4 TinyOS and Singularity

365 TinyOS [Hill et al. 2000] is an operating system designed for low-power, wireless sen-  
366 sor networks. TinyOS is not a traditional OS, but provides a framework of compo-  
367 nents that are linked with the application code. The component-based programming  
368 model is supported by nesC [Gay et al. 2003], a dialect of C. TinyOS components pro-  
369 vide following abstractions: *commands* represent requests for a service of a component;  
370 *events* signal the completion of a service; and *tasks* are functions executed nonpreemp-  
371 tive by the TinyOS scheduler. Events also represent interrupts and preempt tasks. An  
372 event handler may post a task for further processing, which is similar to a second-level  
373 interrupt handler.

374 I/O devices are encapsulated in components and the standard distribution of  
375 TinyOS includes a rich set of standard I/O devices. A Hardware Presentation Layer  
376 (HPL) abstracts the platform-specific access to the hardware (either memory or port  
377 mapped). Our proposed HAL is similar to the HPL, but represents the I/O devices  
378 as Java objects. A further abstraction into I/O components can be built above our  
379 presented Java HAL.

380 Singularity [Hunt et al. 2005] is a research OS based on a runtime managed lan-  
381 guage (an extension of C#) to build a software platform with the main goal to be de-  
382 pendable. A small HAL (IoPorts, IoDma, IoIrq, and IoMemory) provides access to PC  
383 hardware. C# style attributes (similar to Java annotations) on fields are used to define  
384 the mapping of class fields to I/O ports and memory addresses. The Singularity OS  
385 clearly uses device objects and interrupt handlers, thus demonstrating that the ideas  
386 presented here transfer to a language like C#.

## 387 2.5 Summary

388 In our analysis of related work we see that our contribution is a selection, adapta-  
389 tion, refinement, and implementation of ideas from earlier languages and platforms  
390 for Java. A crucial point, where we have spent much time, is to have a clear interface  
391 between the Java layer and the JVM. Here we have used the lessons from the Java  
392 OS and the JVM interfaces. Finally, it has been a concern to be consistent with the  
393 RTSJ because this standard and adaptations of it are the instruments for developing  
394 embedded real-time software in Java.

## 395 3. THE HARDWARE ABSTRACTION LAYER

396 In the following section the hardware abstraction layer for Java is defined. Low-level  
397 access to devices is performed via hardware objects. Synchronization with a device  
398 can be performed with interrupt handlers implemented in Java. Finally, portability  
399 of hardware objects, interrupt handlers, and device drivers is supported by a generic  
400 configuration mechanism.

### 401 3.1 Device Access

402 Hardware objects map object fields to device registers. Therefore, field access with  
403 bytecodes `putfield` and `getfield` accesses device registers. With a correct class that rep-  
404 resents a device, access to it is safe; it is not possible to read or write to an arbitrary  
405 memory address. A memory area (e.g., a video frame buffer) represented by an array  
406 is protected by Java's array bounds check.

407 In a C-based system the access to I/O devices can either be represented by a C  
408 struct (similar to the class shown in Figure 2) for memory-mapped I/O devices or needs  
409 to be accessed by function calls on systems with a separate I/O address space. With  
410 the hardware object abstraction in Java the JVM can represent an I/O device as a  
411 class independent of the underlying low-level I/O mechanism. Furthermore, the strong

```

public abstract class HardwareObject {
    HardwareObject() {};
}

```

Fig. 7. The marker class for hardware objects.

```

public final class SerialPort extends HardwareObject {

    public static final int MASK_TDRE = 1;
    public static final int MASK_RDRF = 2;

    public volatile int status;
    public volatile int data;

    public void init(int baudRate) {...}
    public boolean rxFull() {...}
    public boolean txEmpty() {...}
}

```

Fig. 8. A serial port class with device methods.

412 typing of Java avoids hard to find programming errors due to wrong pointer casts or  
 413 wrong pointer arithmetic.

414 All hardware classes have to extend the abstract class `HardwareObject` (see  
 415 Figure 7). This empty class serves as type marker. Some implementations use it to  
 416 distinguish between plain objects and hardware objects for the field access. The  
 417 package-visible-only constructor disallows creation of hardware objects by the appli-  
 418 cation code that resides in a different package. Figure 8 shows a class representing a  
 419 serial port with a status register and a data register. The status register contains flags  
 420 for receive register full and transmit register empty; the data register is the receive  
 421 and transmit buffer. Additionally, we define device-specific constants (bit masks for the  
 422 status register) in the class for the serial port. All fields represent device registers that  
 423 can change due to activity of the hardware device. Therefore, they must be declared  
 424 volatile.

425 In this example we have included some convenience methods to access the hardware  
 426 object. However, for a clear separation of concerns, the hardware object represents only  
 427 the device state (the registers). We do not add instance fields to represent additional  
 428 state, that is, mixing device registers with heap elements. We cannot implement a com-  
 429 plete device driver within a hardware object; instead a complete device driver owns a  
 430 number of private hardware objects along with data structures for buffering, and it  
 431 defines interrupt handlers and methods for accessing its state from application pro-  
 432 cesses. For device-specific operations, such as initialization of the device, methods in  
 433 hardware objects are useful.

434 Usually each device is represented by exactly one hardware object (see Sec-  
 435 tion 3.3.1). However, there might be use cases where this restriction should be re-  
 436 laxated. Consider a device where some registers should be accessed by system code only  
 437 and some other by application code. In JOP there is such a device: a system device  
 438 that contains a 1 MHz counter, a corresponding timer interrupt, and a watchdog port.  
 439 The timer interrupt is programmed relative to the counter and used by the real-time  
 440 scheduler, a JVM internal service. However, access to the counter can be useful for the  
 441 application code. Access to the watchdog register is required from the application level.  
 442 The watchdog is used for a sign-of-life from the application. If not triggered every  
 443 second the complete system is restarted. For this example it is useful to represent one

```
public final class SysCounter extends HardwareObject {  
  
    public volatile int counter;  
    public volatile int timer;  
    public volatile int wd;  
}  
  
public final class AppCounter extends HardwareObject {  
  
    public volatile int counter;  
    private volatile int timer;  
    public volatile int wd;  
}  
  
public final class AppGetterSetter extends HardwareObject {  
  
    private volatile int counter;  
    private volatile int timer;  
    private volatile int wd;  
  
    public int getCounter() {  
        return counter;  
    }  
  
    public void setWd(boolean val) {  
        wd = val ? 1 : 0;  
    }  
}
```

Fig. 9. System and application classes, one with visibility protection and one with setter and getter methods, for a single hardware device

444 hardware device by two different classes: one for system code and one for application  
445 code. We can protect system registers by private fields in the hardware object for  
446 the application. Figure 9 shows the two class definitions that represent the same  
447 hardware device for system and application code respectively. Note how we changed  
448 the access to the timer interrupt register to private for the application hardware object.  
449 Another option, shown in class `AppGetterSetter`, is to declare all fields private for  
450 the application object and use setter and getter methods. They add an abstraction on  
451 top of hardware objects and use the hardware object to implement their functionality.  
452 Thus we still do not need to invoke native functions.

453 Use of hardware objects is straightforward. After obtaining a reference to the object  
454 all that has to be done (or can be done) is to read from and write to the object fields.  
455 Figure 10 shows an example of client code. The example is a *Hello World* program  
456 using low-level access to a serial port via a hardware object. Creation of hardware  
457 objects is more complex and described in Section 3.3. Furthermore, it is JVM-specific  
458 and Section 4 describes implementations in four different JVMs.

459 For devices that use DMA (e.g., video frame buffer, disk, and network I/O buffers) we  
460 map that memory area to Java arrays. Arrays in Java provide access to raw memory  
461 in an elegant way: the access is simple and safe due to the array bounds checking  
462 done by the JVM. Hardware arrays can be used by the JVM to implement higher-level  
463 abstractions from the RTSJ such as `RawMemory` or `scoped memory`.

464 Interaction between the Garbage Collector (GC) and hardware objects needs to be  
465 crafted into the JVM. We do not want to collect hardware objects. The hardware object

```

import com.jopdesign.io.*;

public class Example {

    public static void main(String[] args) {

        BaseBoard fact = BaseBoard.getBaseFactory();
        SerialPort sp = fact.getSerialPort();

        String hello = "Hello World!";

        for (int i=0; i<hello.length(); ++i) {
            // busy wait on transmit buffer empty
            while ((sp.status & SerialPort.MASK_TDRE) == 0)
                ;
            // write a character
            sp.data = hello.charAt(i);
        }
    }
}

```

Fig. 10. A “Hello World” example with low-level device access via a hardware object.

Table I. Dispatching Properties of Different ISR Strategies

ISR	Context switches	Priorities
Handler	2	Hardware
Event	3–4	Software

466 should not be scanned for references.<sup>6</sup> This is permissible when only primitive types  
 467 are used in the class definition for hardware objects; the GC scans only reference fields.  
 468 To avoid collecting hardware objects, we mark the object to be skipped by the GC. The  
 469 type inheritance from `HardwareObject` can be used as the marker.

### 470 3.2 Interrupt Handling

471 An Interrupt Service Routine (ISR) can be integrated with Java in two different ways:  
 472 as a first-level *handler* or a second-level *event* handler.

473 *ISR handler.* The interrupt is a method call initiated by the device.

474 Usually this abstraction is supported in hardware by the processor and called a  
 475 first-level handler.

476 *ISR event.* The interrupt is represented by an asynchronous notification directed to  
 477 a thread that is unblocked from a wait-state. This is also called deferred interrupt  
 478 handling.

479 An overview of the dispatching properties of the two approaches is given in Table I.  
 480 The ISR handler approach needs only two context switches and the priority is set  
 481 by the hardware. With the ISR event approach, handlers are scheduled at software  
 482 priorities. The initial first-level handler, running at hardware priority, fires the event

<sup>6</sup>If a hardware coprocessor, represented by a hardware object, itself manipulates an object on the heap and holds the only reference to that object it has to be scanned by the GC.

483 for the event handler. Also the first-level handler will notify the scheduler. In the best  
484 case three context switches are necessary: one to the first-level handler, one to the ISR  
485 event handler, and one back to the interrupted thread. If the ISR handler has a lower  
486 priority than the interrupted thread, an additional context switch from the first-level  
487 handler back to the interrupted thread is necessary.

488 Another possibility is to represent an interrupt as a thread that is released by the  
489 interrupt. Direct support by the hardware (e.g., the interrupt service thread in Ko-  
490 modo [Kreuzinger et al. 2003]) gives fast interrupt response times. However, standard  
491 processors support only the handler model directly.

492 Direct handling of interrupts in Java requires the JVM to be prepared to be inter-  
493 rupted. In an interpreting JVM an initial handler will reenter the JVM to execute the  
494 Java handler. A compiling JVM or a Java processor can directly invoke a Java method  
495 as response to the interrupt. A compiled Java method can be registered directly in the  
496 ISR dispatch table.

497 If an internal scheduler is used (also called *green threads*) the JVM will need some  
498 refactoring in order to support asynchronous method invocation. Usually JVMs control  
499 the rescheduling at the JVM level to provide a lightweight protection of JVM internal  
500 data structures. These preemption points are called pollchecks or yield points; also  
501 some or all can be GC preemption points. In fact the preemption points resemble  
502 cooperative scheduling at the JVM level and use priority for synchronization. This  
503 approach works only for uniprocessor systems; for multiprocessors explicit synchro-  
504 nization has to be introduced.

505 In both cases there might be critical sections in the JVM where reentry cannot  
506 be allowed. To solve this problem the JVM must disable interrupts around critical  
507 nonreentrant sections. The more fine-grained this disabling of interrupts can be  
508 done, the more responsive to interrupts the system will be.

509 One could opt for second-level handlers only. An interrupt fires and releases an  
510 associated schedulable object (handler). Once released, the handler will be sched-  
511 uled by the JVM scheduler according to the release parameters. This is the RTSJ  
512 approach. The advantage is that interrupt handling is done in the context of a normal  
513 Java thread and scheduled as any other thread running on the system. The drawback  
514 is that there will be a delay from the occurrence of the interrupt until the thread gets  
515 scheduled. Additionally, the meaning of interrupt priorities, levels, and masks used by  
516 the hardware may not map directly to scheduling parameters supported by the JVM  
517 scheduler.

518 In the following we focus on the ISR handler approach, because it allows program-  
519 ming the other paradigms within Java.

520 *3.2.1 Hardware Properties.* We assume interrupt hardware as it is found in most com-  
521 puter architectures: interrupts have a fixed priority associated with them; they are  
522 set with a solder iron. Furthermore, interrupts can be globally disabled. In most  
523 systems the first-level handler is called with interrupts globally disabled. To allow  
524 nested interrupts (being able to interrupt the handler by a higher-priority interrupt  
525 as in preemptive scheduling) the handler has to enable interrupts again. However, to  
526 avoid priority inversion between handlers only interrupts with a higher priority will  
527 be enabled, either by setting the interrupt level or setting the interrupt mask. Soft-  
528 ware threads are scheduled by a timer interrupt and usually have a lower priority  
529 than interrupt handlers (the timer interrupt has the lowest priority of all interrupts).  
530 Therefore, an interrupt handler is never preempted by a software thread.

531 Mutual exclusion between an interrupt handler and a software thread is ensured by  
532 disabling interrupts: either all interrupts or selectively. Again, to avoid priority inver-  
533 sion, only interrupts of a higher priority than the interrupt that shares the data with

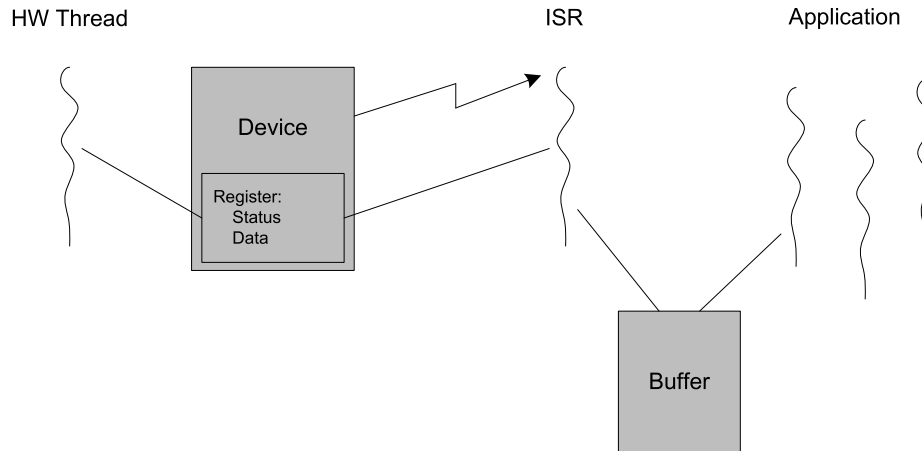


Fig. 11. Threads and shared data.

534 the software thread can be enabled. This mechanism is in effect the priority ceiling  
 535 emulation protocol [Sha et al. 1990], sometimes called immediate ceiling protocol. It  
 536 has the virtue that it eliminates the need for explicit locks (or Java monitors) on shared  
 537 objects. Note that mutual exclusion with interrupt disabling works only in a unipro-  
 538 cessor setting. A simple solution for multiprocessors is to run the interrupt handler  
 539 and associated software threads on the same processor core. A more involved scheme  
 540 would be to use spin-locks between the processors.

541 When a device asserts an interrupt request line, the interrupt controller notifies the  
 542 processor. The processor stops execution of the current thread. A partial thread con-  
 543 text (program counter and processor status register) is saved. Then the ISR is looked  
 544 up in the interrupt vector table and a jump is performed to the first instruction of the  
 545 ISR. The handler usually saves additional thread context (e.g., the register file). It is  
 546 also possible to switch to a new stack area. This is important for embedded systems  
 547 where the stack sizes for all threads need to be determined at link time.

548 **3.2.2 Synchronization.** Java supports synchronization between Java threads with the  
 549 `synchronized` keyword, either as a means of synchronizing access to a block of state-  
 550 ments or to an entire method. In the general case this existing synchronization support  
 551 is not sufficient to synchronize between interrupt handlers and threads.

552 Figure 11 shows the interacting active processes and the shared data in a scenario  
 553 involving the handling of an interrupt. Conceptually three threads interact: (1) a hard-  
 554 ware device thread representing the device activity, (2) the ISR, and (3) the application  
 555 or device driver thread. These three share two types of data.

556 **Device data.** The hardware thread and ISR share access to the device registers of the  
 557 device signaling the interrupt.

558 **Application data.** The ISR and application or device driver share access to, for exam-  
 559 ple, a buffer conveying information about the interrupt to the application.

560 Regardless of which interrupt handling approach is used in Java, synchronization  
 561 between the ISR and the device registers must be handled in an ad hoc way. In gen-  
 562 eral there is no guarantee that the device has not changed the data in its registers;  
 563 but if the ISR can be run to completion within the minimum inter-arrival time of the  
 564 interrupt the content of the device registers can be trusted.



565 For synchronization between the ISR and the application (or device driver) the fol-  
566 lowing mechanisms are available. When the ISR handler runs as a software thread,  
567 standard synchronization with object monitors can be used. When using the ISR han-  
568 dler approach, the handler is no longer scheduled by the normal Java scheduler, but  
569 by the hardware. While the handler is running, all other executable elements are sus-  
570 pended, including the scheduler. This means that the ISR cannot be suspended, must  
571 not block, or must not block via a language-level synchronization mechanism; the ISR  
572 must run to completion in order not to freeze the system. This means that when the  
573 handler runs, it is guaranteed that the application will not get scheduled. It follows  
574 that the handler can access data shared with the application without synchronizing  
575 with the application. As the access to the shared data by the interrupt handler is not  
576 explicitly protected by a synchronized method or block, the shared data needs to be  
577 declared volatile.

578 On the other hand the application must synchronize with the ISR because the ISR  
579 may be dispatched at any point. To ensure mutual exclusion we redefine the seman-  
580 tics of the monitor associated with an `InterruptHandler` object: acquisition of the monitor  
581 disables all interrupts of the same and lower priority; release of the monitor enables  
582 the interrupts again. This procedure ensures that the software thread cannot be in-  
583 terrupted by the interrupt handler when accessing shared data.

584 *3.2.3 Using the Interrupt Handler.* Figure 12 shows an example of an interrupt handler  
585 for the serial port receiver interrupt. The method `handle()` is the interrupt handler  
586 method and needs no synchronization as it cannot be interrupted by a software thread.  
587 However, the shared data needs to be declared volatile as it is changed by the device  
588 driver thread. Method `read()` is invoked by the device driver thread and the shared  
589 data is protected by the `InterruptHandler` monitor. The serial port interrupt handler  
590 uses the hardware object `SerialPort` to access the device.

591 *3.2.4 Garbage Collection.* When using the ISR handler approach it is not feasible to  
592 let interrupt handlers be paused during a lengthy stop-the-world collection. Using  
593 this GC strategy the entire heap is collected at once and it is not interleaved with  
594 execution. The collector can safely assume that data required to perform the collection  
595 will not change during the collection, and an interrupt handler shall not change data  
596 used by the GC to complete the collection. In the general case, this means that the  
597 interrupt handler is not allowed to create new objects, or change the graph of live  
598 objects.

599 With an incremental GC the heap is collected in small incremental steps. Write bar-  
600 riers in the mutator threads and nonpreemption sections in the GC thread synchronize  
601 the view of the object graph between the mutator threads and the GC thread. With  
602 incremental collection it is possible to allow object allocation and changing references  
603 inside an interrupt handler (as it is allowed in any normal thread). With a real-time  
604 GC the maximum blocking time due to GC synchronization with the mutator threads  
605 must be known.

606 Interruption of the GC during an object move can result in access to a stale copy  
607 of the object inside the handler. A solution to this problem is to allow for pinning of  
608 objects reachable by the handler (similar to immortal memory in the RTSJ). Concur-  
609 rent collectors have to solve this issue for threads anyway. The simplest approach is to  
610 disable interrupt handling during the object copy. As this operation can be quite long  
611 for large arrays, several approaches to split the array into smaller chunks have been  
612 proposed [Bacon et al. 2003; Siebert 2002]. A Java processor may support incremental  
613 array copying with redirection of the access to the correct part of the array [Schoeberl  
614 and Puffitsch 2008]. Another solution is to abort the object copy when writing to the

```

public class SerialHandler extends InterruptHandler {

    // A hardware object represents the serial device
    private SerialPort sp;

    final static int BUF_SIZE = 32;
    private volatile byte buffer[];
    private volatile int wrPtr, rdPtr;

    public SerialHandler(SerialPort sp) {
        this.sp = sp;
        buffer = new byte[BUF_SIZE];
        wrPtr = rdPtr = 0;
    }

    // This method is scheduled by the hardware
    public void handle() {
        byte val = (byte) sp.data;
        // check for buffer full
        if ((wrPtr+1)%BUF_SIZE!=rdPtr) {
            buffer[wrPtr++] = val;
        }
        if (wrPtr>=BUF_SIZE) wrPtr=0;
        // enable interrupts again
        enableInterrupt();
    }

    // This method is invoked by the driver thread
    synchronized public int read() {
        if (rdPtr!=wrPtr) {
            int val = ((int) buffer[rdPtr++]) & 0xff;
            if (rdPtr>=BUF_SIZE) rdPtr=0;
            return val;
        } else {
            return -1;        // empty buffer
        }
    }
}

```

Fig. 12. An interrupt handler for a serial port receive interrupt.

615 object. It is also possible to use replication; during an incremental copy operation,  
 616 writes are performed on both from-space and to-space object replicas, while reads are  
 617 performed on the from-space replica.

### 618 3.3 Generic Configuration

619 An important issue for a HAL is a safe abstraction of device configurations. A def-  
 620 inition of factories to create hardware and interrupt objects should be provided by  
 621 board vendors. This configuration is isolated with the help of Java packages; only the  
 622 objects and the factory methods are visible. The configuration abstraction is indepen-  
 623 dent of the JVM. A device or interrupt can be represented by an identical hardware  
 624 or interrupt object for different JVMs. Therefore, device drivers written in Java are  
 625 JVM-independent.

626 *3.3.1 Hardware Object Creation.* The idea to represent each device by a single object or  
 627 array is straightforward, the remaining question is: How are those objects created?

```
package com.board-vendor.io;

public class IOSystem {

    // some JVM mechanism to create the hardware objects
    private static ParallelPort pp = jvmPPCreate();
    private static SerialPort sp = jvmSPCreate(0);
    private static SerialPort gps = jvmSPCreate(1);

    public static ParallelPort getParallelPort() {
        return pp;
    }

    public static SerialPort getSerialPort() {...}
    public static SerialPort getGpsPort() {...}
}
```

Fig. 13. A factory with static methods for Singleton hardware objects.

628 An object that represents a device is a typical Singleton [Gamma et al. 1994]. Only  
629 a single object should map to one instance of a device. Therefore, hardware objects  
630 cannot be instantiated by a simple `new`: (1) they have to be mapped by some JVM  
631 mechanism to the device registers and (2) each device instance is represented by a  
632 single object.

633 Each device object is created by its own factory method. The collection of these  
634 methods is the board configuration, which itself is also a Singleton (the application  
635 runs on a single board). The Singleton property of the configuration is enforced by  
636 a class that contains only static methods. Figure 13 shows an example for such a  
637 class. The class `IOSystem` represents a system with three devices: a parallel port,  
638 as discussed before to interact with the environment, and two serial ports: one for  
639 program download and one which is an interface to a GPS receiver.

640 This approach is simple, but not very flexible. Consider a vendor who provides  
641 boards in slightly different configurations (e.g., with different number of serial ports).  
642 With the preceding approach each board requires a different (or additional) `IOSystem`  
643 class that lists all devices. A more elegant solution is proposed in the next section.

644 *3.3.2 Board Configurations.* Replacing the static factory methods by instance methods  
645 avoids code duplication; inheritance then gives configurations. With a factory object  
646 we represent the common subset of I/O devices by a base class and the variants as  
647 subclasses.

648 However, the factory object itself must still be a Singleton. Therefore the board-  
649 specific factory object is created at class initialization and is retrieved by a static  
650 method. Figure 14 shows an example of a base factory and a derived factory. Note  
651 how `getBaseFactory()` is used to get a single instance of the factory. We have applied  
652 the idea of a factory two times: the first factory generates an object that represents  
653 the board configuration. That object is itself a factory that generates the objects that  
654 interface to the hardware device.

655 The shown example base factory represents the minimum configuration with a sin-  
656 gle serial port for communication (mapped to `System.in` and `System.out`) represented  
657 by a `SerialPort`. The derived configuration `ExtendedBoard` contains an additional serial  
658 port for a GPS receiver and a parallel port for external control.

659 Furthermore, we show in Figure 14 a different way to incorporate the JVM mech-  
660 anism in the factory: we define well-known constants (the memory addresses of the

```

public class BaseBoard {

    private final static int SERIAL_ADDRESS = ...;
    private SerialPort serial;
    BaseBoard() {
        serial = (SerialPort) jvmHWOCreate(SERIAL_ADDRESS);
    };
    static BaseBoard single = new BaseBoard();
    public static BaseBoard getBaseFactory() {
        return single;
    }
    public SerialPort getSerialPort() { return serial; }

    // here comes the JVM internal mechanism
    Object jvmHWOCreate(int address) {...}
}

public class ExtendedBoard extends BaseBoard {

    private final static int GPS_ADDRESS = ...;
    private final static int PARALLEL_ADDRESS = ...;
    private SerialPort gps;
    private ParallelPort parallel;
    ExtendedBoard() {
        gps = (SerialPort) jvmHWOCreate(GPS_ADDRESS);
        parallel = (ParallelPort) jvmHWOCreate(PARALLEL_ADDRESS);
    };
    static ExtendedBoard single = new ExtendedBoard();
    public static ExtendedBoard getExtendedFactory() {
        return single;
    }
    public SerialPort getGpsPort() { return gps; }
    public ParallelPort getParallelPort() { return parallel; }
}

```

Fig. 14. A base class of a hardware object factory and a factory subclass.

661 devices) in the factory and let the native function `jvmHWOCreate()` return the correct  
 662 device type.

663 Figure 15 gives a summary example of hardware object classes and the correspond-  
 664 ing factory classes as an UML class diagram. The figure shows that all device classes  
 665 subclass the abstract class `HardwareObject`. Figure 15 represents the simple abstrac-  
 666 tion as it is seen by the user of hardware objects.

667 **3.3.3 Interrupt Handler Registration.** We provide a base interrupt handling API that can  
 668 be used both for non-RTSJ and RTSJ interrupt handling. The base class that is ex-  
 669 tended by an interrupt handler is shown in Figure 16. The `handle()` method contains  
 670 the device server code. Interrupt control operations that have to be invoked before  
 671 serving the device (i.e., interrupt masking and acknowledging) and after serving the  
 672 device (i.e., interrupt reenabling) are hidden in the `run()` method of the base `Inter-`  
 673 `ruptHandler`, which is invoked when the interrupt occurs.

674 The base implementation of `InterruptHandler` also provides methods for enabling and  
 675 disabling a particular interrupt or all local CPU interrupts and a special monitor im-  
 676 plementation for synchronization between an interrupt handler thread and an applica-  
 677 tion thread. Moreover, it provides methods for non-RTSJ registering and deregistering  
 678 the handler with the hardware interrupt source.

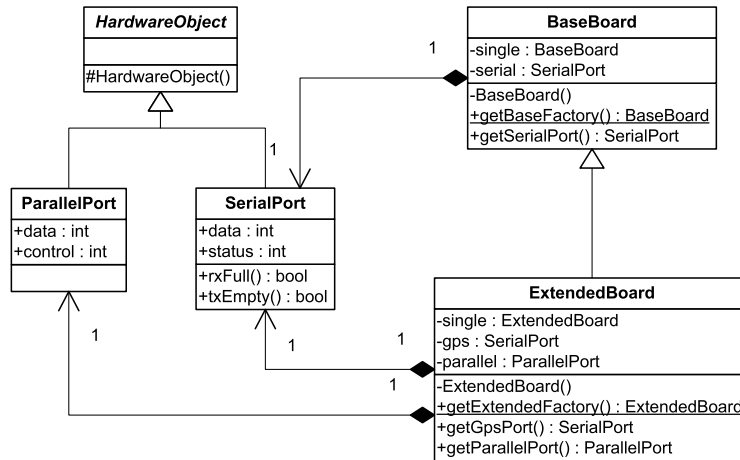


Fig. 15. Device object classes and board factory classes.

```

abstract public class InterruptHandler implements Runnable {
    ...

    public InterruptHandler(int index) { ... };

    protected void startInterrupt() { ... };
    protected void endInterrupt() { ... };

    protected void disableInterrupt() { ... };
    protected void enableInterrupt() { ... };
    protected void disableLocalCPUInterrupts() { ... };
    protected void enableLocalCPUInterrupts() { ... };

    public void register() { ... };
    public void unregister() { ... };

    abstract public void handle() { ... };

    public void run() {
        startInterrupt();
        handle();
        endInterrupt();
    }
}
  
```

Fig. 16. Base class for the interrupt handlers.

679 Registration of a RTSJ interrupt handler requires more steps (see Figure 17). The  
 680 `InterruptHandler` instance serves as the RTSJ logic for a (bound) asynchronous event  
 681 handler, which is added as handler to an asynchronous event which then is bound to  
 682 the interrupt source.

### 683 3.4 Perspective

684 An interesting topic is to define a common standard for hardware objects and interrupt  
 685 handlers for different platforms. If different device types (hardware chips) that do not  
 686 share a common register layout are used for a similar function, the hardware objects

```

ih = new SerialInterruptHandler(); // logic of new BAEH

serialFirstLevelEvent = new AsyncEvent();
serialFirstLevelEvent.addHandler(
    new BoundAsyncEventHandler( null, null, null, null, null, false, ih )
);

serialFirstLevelEvent.bindTo("INT4");

```

Fig. 17. Creation and registration of a RTSJ interrupt handler.

Table II. Embedded Java Architectures

Direct (no OS)		Indirect (OS)
Interpreted	SimpleRTJ	Kaffe VM
Native	JOP	OVM

687 will be different. However, if the structure of the devices is similar, as is the case  
688 for the serial port on the three different platforms used for the implementation (see  
689 Section 4), the driver code that *uses* the hardware object is identical.

690 If the same chip (e.g., the 8250 type and compatible 16x50 devices found in all PCs  
691 for the serial port) is used in different platforms, the hardware object and the device  
692 driver, which also implements the interrupt handler, can be shared. The hardware ob-  
693 ject, the interrupt handler, and the visible API of the factory classes are independent of  
694 the JVM and the OS. Only the implementation of the factory methods is JVM-specific.  
695 Therefore, the JVM-independent HAL can be used to start the development of drivers  
696 for a Java OS on any JVM that supports the proposed HAL.

### 697 3.5 Summary

698 Hardware objects are easy to use for a programmer, and the corresponding definitions  
699 are comparatively easy to define for a hardware designer or manufacturer. For a stan-  
700 dardized HAL architecture we proposed factory patterns. As shown, interrupt han-  
701 dlers are easy to use for a programmer who knows the underlying hardware paradigm,  
702 and the definitions are comparatively easy to develop for a hardware designer or man-  
703 ufacturer, for instance using the patterns outlined in this section. Hardware objects  
704 and interrupt handler infrastructure have a few subtle implementation points which  
705 are discussed in the next section.

## 706 4. IMPLEMENTATION

707 We have implemented the core concepts on four different JVMs<sup>7</sup> to validate the pro-  
708 posed Java HAL. Table II classifies the four execution environments according to two  
709 important properties: (1) whether they run on bare metal or on top of an OS and (2)  
710 whether Java code is interpreted or executed natively. Thereby we cover the whole  
711 implementation spectrum with our four implementations. Even though the suggested  
712 Java HAL is intended for systems running on bare metal, we include systems run-  
713 ning on top of an OS because most existing JVMs still require an OS, and in order for  
714 them to migrate incrementally to run directly on the hardware they can benefit from  
715 supporting a Java HAL.

716 In the direct implementation a JVM without an OS is extended with I/O function-  
717 ality. The indirect implementation represents an abstraction mismatch; we actually

<sup>7</sup>On JOP the implementation of the Java HAL is already in use in production code.

```
public final class SerialPort extends HardwareObject {
    // LSR (Line Status Register)
    public volatile int status;
    // Data register
    public volatile int data;
    ...
}
```

Fig. 18. A simple hardware object.

718 remap the concepts. Related to Figure 6 in the Introduction, OVM and Kaffe represent  
719 configuration (a), SimpleRTJ configuration (b), and JOP configuration (c).

720 The SimpleRTJ JVM [RTJ Computing 2000] is a small, interpreting JVM that does  
721 not require an OS. JOP [Schoeberl 2005, 2008] is a Java processor executing Java byte-  
722 codes directly in hardware. Kaffe JVM [Wilkinson 1996] is a complete, full-featured  
723 JVM supporting both interpretation and JIT compilation; in our experiments with  
724 Kaffe we have used interpretative execution only. The OVM JVM [Armbruster et al.  
725 2007] is an execution environment for Java that supports compilation of Java byte-  
726 codes into the C language, and via a C compiler into native machine instructions for  
727 the target hardware. Hardware objects have also been implemented in the research  
728 JVM, CACAO [Krall and Graf 1997; Schoeberl et al. 2008].

729 In the following we provide the different implementation approaches that are neces-  
730 sary for the very different JVMs. Implementing hardware objects was straightforward  
731 for most JVMs; it took about one day to implement them in JOP. In Kaffe, after famil-  
732 iarizing us with the structure of the JVM, it took about half a day of pair programming.

733 Interrupt handling in Java is straightforward in a JVM not running on top of an  
734 OS (JOP and SimpleRTJ). Kaffe and OVM both run under vanilla Linux or the real-  
735 time version Xenomai Linux [Xenomai Developers 2008]. Both versions use a distinct  
736 user/kernel mode and it is not possible to register a user-level method as interrupt  
737 handler. Therefore, we used threads at different levels to simulate the Java handler  
738 approach. The result is that the actual Java handler is the third- or even fourth-level  
739 handler. This solution introduces quite a lot of overheads due to the many context  
740 switches. However, it is intended to provide a stepping stone to allow device drivers in  
741 Java; the goal is a real-time JVM that runs on the bare hardware.

742 In this section we provide more implementation details than usual to help other  
743 JVM developers to add a HAL to their JVM. The techniques used for the JVMs can  
744 probably not be used directly. However, the solutions (or sometimes work-arounds)  
745 presented here should give enough insight to guide other JVM developers.

#### 746 4.1 SimpleRTJ

747 The SimpleRTJ JVM is a small, simple, and portable JVM. We have ported it to run  
748 on the bare metal of a small 16-bit microcontroller. We have successfully implemented  
749 the support for hardware objects in the SimpleRTJ JVM. For interrupt handling we  
750 use the ISR handler approach described in Section 3.2. Adding support for hardware  
751 objects was straightforward, but adding support for interrupt handling required more  
752 work.

753 *4.1.1 Hardware Objects.* Given an instance of a hardware object as shown in Figure 18  
754 one must calculate the base address of the I/O port range, the offset to the actual I/O  
755 port, and the width of the port at runtime. We have chosen to store the base address  
756 of the I/O port range in a field in the common superclass for all hardware objects  
757 (`HardwareObject`). The hardware object factory passes the platform- and device-specific

```

SerialPort createSerialPort(int baseAddress ) {
    SerialPort sp = new SerialPort(baseAddress);
    return sp;
}

```

Fig. 19. Creating a simple hardware object.

758 base address to the constructor when creating instances of hardware objects (see  
759 Figure 19).

760 In the `put/getfield` bytecodes the base address is retrieved from the object instance.  
761 The I/O port offset is calculated from the offset of the field being accessed: in the  
762 example in Figure 18 status has an offset of 0 whereas data has an offset of 4. The  
763 width of the field being accessed is the same as the width of the field type. Using these  
764 values the SimpleRTJ JVM is able to access the device register for either read or write.

765 **4.1.2 Interrupt Handler.** The SimpleRTJ JVM uses a simple stop-the-world garbage  
766 collection scheme. This means that within handlers, we prohibit use of the `new` key-  
767 word and writing references to the heap. These restrictions can be enforced at runtime  
768 by throwing a preallocated exception or at class loading by an analysis of the handler  
769 method. Additionally we have turned off the compaction phase of the GC to avoid the  
770 problems with moving objects mentioned in Section 3.2.4.

771 The SimpleRTJ JVM implements thread scheduling within the JVM. This means  
772 that it had to be refactored to allow for reentering the JVM from inside the first-level  
773 interrupt handler. We got rid of all global state (all global variables) used by the JVM  
774 and instead allocate shared data on the C stack. For all parts of the JVM to still be  
775 able to access the shared data we pass around a single pointer to that data. In fact  
776 we start a new JVM for the interrupt handler with a temporary (small) Java heap and  
777 a temporary (small) Java stack. Currently we use 512 bytes for each of these items,  
778 which have proven sufficient for running nontrivial interrupt handlers so far.

779 The major part of the work was making the JVM reentrant. The effort will vary  
780 from one JVM implementation to another, but since global state is a bad idea in any  
781 case JVMs of high quality use very little global state. Using these changes we have  
782 experimented with handling the serial port receive interrupt.

## 783 4.2 JOP

784 JOP is a Java processor intended for hard real-time systems [Schoeberl 2005, 2008].  
785 All architectural features have been carefully designed to be time-predictable with  
786 minimal impact on average case performance. We have implemented the proposed  
787 HAL in the JVM for JOP. No changes inside the JVM (the microcode in JOP) were  
788 necessary. Only the creation of the hardware objects needs a JOP-specific factory.

789 **4.2.1 Hardware Objects.** In JOP, objects and arrays are referenced through an indirec-  
790 tion called *handle*. This indirection is a lightweight read barrier for the compacting  
791 real-time GC [Schoeberl 2006; Schoeberl and Vitek 2007]. All handles for objects in the  
792 heap are located in a distinct memory region, the handle area. Besides the indirection  
793 to the *real* object the handle contains auxiliary data, such as a reference to the class  
794 information, the array length, and GC-related data. Figure 20 shows an example with  
795 a small object that contains two fields and an integer array of length 4. The object and  
796 the array on the heap just contain the data and no additional hidden fields. This object  
797 layout greatly simplifies our object to device mapping. We just need a handle where  
798 the indirection points to the memory-mapped device registers instead of into the heap.  
799 This configuration is shown in the upper part of Figure 20. Note that we do not need



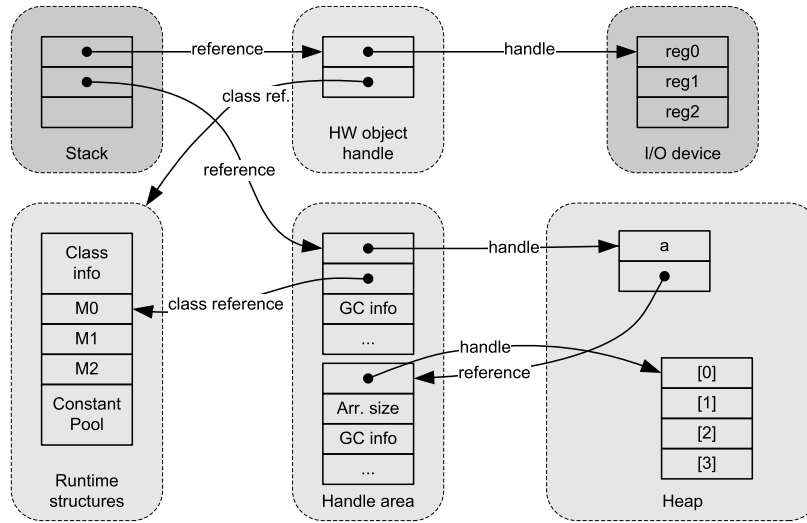


Fig. 20. Memory layout of the JOP JVM.

800 the GC information for the hardware object handles. The factory, which creates the  
 801 hardware objects, implements this indirection.

802 As described in Section 3.3.1 we do not allow applications to create hardware objects;  
 803 the constructor is private (or package visible). Figure 21 shows part of the hardware  
 804 object factory that creates the hardware object `SerialPort`. Two static fields (`SP_PTR`  
 805 and `SP_MTAB`) are used to store the handle to the serial port object. The first field is  
 806 initialized with the base address of the I/O device; the second field contains a pointer  
 807 to the class information.<sup>8</sup> The address of the static field `SP_PTR` is returned as the  
 808 reference to the serial port object.

809 The class reference for the hardware object is obtained by creating a normal instance  
 810 of `SerialPort` with `new` on the heap and copying the pointer to the class information. To  
 811 avoid using native methods in the factory class we delegate JVM internal work to a  
 812 helper class in the JVM system package as shown in Figure 21. That helper method  
 813 returns the address of the static field `SP_PTR` as reference to the hardware object.  
 814 All methods in class `Native`, a JOP system class, are *native*<sup>9</sup> methods for low-level  
 815 functions, the code we want to avoid in application code. Method `toInt(Object o)` defeats  
 816 Java's type safety and returns a reference as an `int`. Method `ToObject(int addr)` is the  
 817 inverse function to map an address to a Java reference. Low-level memory access  
 818 methods are used to manipulate the JVM data structures.

819 To disallow the creation with `new` in normal application code, the visibility is set  
 820 to package. However, the package visibility of the hardware object constructor is a  
 821 minor issue. To access private static fields of an arbitrary class from the system class  
 822 we had to change the runtime class information: we added a pointer to the first static  
 823 primitive field of that class. As addresses of static fields get resolved at class linking,  
 824 no such reference was needed so far.

<sup>8</sup>In JOP's JVM the class reference is a pointer to the method table to speed-up the `invoke` instruction. Therefore, the name is `XX_MTAB`.

<sup>9</sup>There are no real native functions in JOP; bytecode is the native instruction set. The very few native methods in class `Native` are replaced by special, unused bytecodes during class linking.

```

package com.jopdesign.io;

public class BaseFactory {

    // static fields for the handle of the hardware object
    private static int SP_PTR;
    private static int SP_MTAB;

    private SerialPort sp;

    IOFactory() {
        sp = (SerialPort) makeHWObject(new SerialPort(), Const.IO_UART1_BASE, 0);
    };

    ...

    // That's the JOP version of the JVM mechanism
    private static Object makeHWObject(Object o, int address, int idx) {
        int cp = Native.rdIntMem(Const.RAM_CP);
        return JVMHelp.makeHWObject(o, address, idx, cp);
    }
}

package com.jopdesign.sys;

public class JVMHelp {

    public static Object makeHWObject(Object o, int address, int idx, int cp) {
        // usage of native methods is allowed here as
        // we are in the JVM system package
        int ref = Native.toInt(o);
        // fill in the handle in the two static fields
        // and return the address of the handle as a
        // Java object
        return Native.toObject(address);
    }
}

```

Fig. 21. Part of a factory and the helper method for the hardware object creation in the factory.

825 **4.2.2 Interrupt Handler.** The original JOP [Schoeberl 2005, 2008] was a very puristic  
826 hard real-time processor. There existed only one interrupt: the programmable timer  
827 interrupt as time is the primary source for hard real-time events. All I/O requests were  
828 handled by periodic threads that polled for pending input data or free output buffers.  
829 During the course of this research we have added an interrupt controller to JOP and  
830 the necessary software layers.

831 Interrupts and exact exceptions are considered the hard part in the implementa-  
832 tion of a processor pipeline [Hennessy and Patterson 2002]. The pipeline has to be  
833 drained and the complete processor state saved. In JOP there is a translation stage  
834 between Java bytecodes and the JOP internal microcode [Schoeberl 2008]. On a pend-  
835 ing interrupt (or exception generated by the hardware) we use this translation stage  
836 to insert a special bytecode in the instruction stream. This approach keeps the inter-  
837 rupt completely transparent to the core pipeline. The special bytecode that is unused  
838 by the JVM specification [Lindholm and Yellin 1999] is handled in JOP as any other  
839 bytecode: execute microcode, invoke a special method from a helper class, or execute

```

static Runnable ih[] = new Runnable[Const.NUM_INTERRUPTS];
static SysDevice sys = IOFactory.getFactory().getSysDevice();

static void interrupt() {

    ih[sys.intNr].run();
}

```

Fig. 22. Interrupt dispatch with the static `interrupt()` method in the JVM helper class.

```

public class InterruptHandler implements Runnable {

    public static void main(String[] args) {

        InterruptHandler ih = new InterruptHandler();
        IOFactory fact = IOFactory.getFactory();
        // register the handler
        fact.registerInterruptHandler(1, ih);
        // enable interrupt 1
        fact.enableInterrupt(1);
        .....
    }

    public void run() {
        System.out.println("Interrupt fired!");
    }
}

```

Fig. 23. An example Java interrupt handler as `Runnable`.

840 Java bytecode from `JVM.java`. In our implementation we invoke the special method  
 841 `interrupt()` from a JVM helper class.

842 The implemented Interrupt Controller (IC) is priority-based. The number of inter-  
 843 rupt sources can be configured. Each interrupt can be triggered in software by a IC  
 844 register write as well. There is one global interrupt enable and each interrupt line can  
 845 be enabled or disabled locally. The interrupt is forwarded to the bytecode/microcode  
 846 translation stage with the interrupt number. When accepted by this stage, the inter-  
 847 rupt is acknowledged and the global enable flag cleared. This feature avoids immediate  
 848 handling of an arriving higher-priority interrupt during the first part of the handler.  
 849 The interrupts have to be enabled again by the handler at a *convenient* time. All inter-  
 850 rupts are mapped to the same special bytecode. Therefore, we perform the dispatch of  
 851 the correct handler in Java. On an interrupt the static method `interrupt()` from a system  
 852 internal class gets invoked. The method reads the interrupt number and performs the  
 853 dispatch to the registered `Runnable` as illustrated in Figure 22. Note how a hardware  
 854 object of type `SysDevice` is used to read the interrupt number.

855 The timer interrupt, used for the real-time scheduler, is located at index 0. The  
 856 scheduler is just a plain interrupt handler that gets registered at mission start at index  
 857 0. At system startup, the table of `Runnables` is initialized with dummy handlers. The  
 858 application code provides the handler via a class that implements `Runnable` and regis-  
 859 ters that class for an interrupt number. We reuse the factory presented in Section 3.3.1.  
 860 Figure 23 shows a simple example of an interrupt handler implemented in Java.

861 For interrupts that should be handled by an event handler under the control of the  
 862 scheduler, the following steps need to be performed on JOP.

863 (1) Create a `SwEvent` with the correct priority that performs the second-level interrupt  
 864 handler work.

865 (2) Create a short first-level interrupt handler as Runnable that invokes fire() of the  
866 corresponding software event handler.  
867 (3) Register the first-level interrupt handler as shown in Figure 23 and start the real-  
868 time scheduler.

869 In Section 5 we evaluate the different latencies of first- and second-level interrupt  
870 handlers on JOP.

### 871 4.3 Kaffe

872 Kaffe is an open-source<sup>10</sup> implementation of the JVM which makes it possible to add  
873 support for hardware objects and interrupt handlers. Kaffe requires a fully fledged OS  
874 such as Linux to compile and run. Although ports of Kaffe exist on uCLinux we have  
875 not been able to find a bare metal version of Kaffe. Thus even though we managed  
876 to add support of hardware objects and interrupt handling to Kaffe, it still cannot be  
877 used without an OS.

878 4.3.1 *Hardware Objects.* Hardware objects have been implemented in the same man-  
879 ner as in the SimpleRTJ, described in Section 4.1.

880 4.3.2 *Interrupt Handler.* Since Kaffe runs under Linux we cannot directly support the  
881 ISR handler approach. Instead we used the ISR event approach in which a thread  
882 blocks waiting for the interrupt to occur. It turned out that the main implementation  
883 effort was spent in the signaling of an interrupt occurrence from the kernel space to  
884 the user space.

885 We wrote a special Linux kernel module in the form of a character device. Through  
886 proper invocations of ioctl() it is possible to let the module install a handler for an  
887 interrupt (e.g., the serial interrupt, normally on IRQ 7). Then the Kaffe VM can make  
888 a blocking call to read() on the proper device. Finally the installed kernel handler will  
889 release the user space application from the blocked call when an interrupt occurs.

890 Using this strategy we have performed nontrivial experiments implementing a full  
891 interrupt handler for the serial interrupt in Java. Still, the elaborate setup requiring  
892 a special-purpose kernel device is far from our ultimate goal of running a JVM on the  
893 bare metal. Nevertheless the experiment has given valuable experience with interrupt  
894 handlers and hardware objects at the Java language level.

### 895 4.4 OVM

896 OVM [Armbruster et al. 2007] is a research JVM allowing many configurations; it is  
897 primarily targeted at implementing a large subset of RTSJ while maintaining rea-  
898 sonable performance. OVM uses ahead of time compilation via the C language: it  
899 translates both application and VM bytecodes to C, including all classes that might be  
900 later loaded dynamically at runtime. The C code is then compiled by GCC.

901 4.4.1 *Hardware Objects.* To compile Java bytecode into a C program, the OVM's Java-  
902 to-C compiler internally converts the bytecode into an Intermediate Representation  
903 (IR) which is similar to the bytecode, but includes more codes. Transformations at the  
904 IR level are both optimizations and operations necessary for correct execution, such as  
905 insertion of null-pointer checks. The produced IR is then translated into C, allowing  
906 the C compiler to perform additional optimizations. Transformations at the IR level,  
907 which is similar to the bytecode, are also typical in other JVM implementations, such  
908 as Sun's HotSpot.

<sup>10</sup><http://www.kaffe.org/>

Reading from the device register `serial.data`, saving the result to the stack

<i>original bytecode</i>	<i>stack content</i>		<i>modified bytecode</i>	<i>stack content</i>
GETFIELD data	{serial} {io port address}	⇒	GETFIELD data INB	{serial} {io port address} {inval}

Writing a value on the stack into the device register `serial.data`

<i>original bytecode</i>	<i>stack content</i>		<i>modified bytecode</i>	<i>stack content</i>
PUTFIELD data	{serial}, {outval} empty	⇒	SWAP	{serial}, {outval}
		⇒	GETFIELD data	{outval}, {serial}
		⇒	OUTB	{outval}, {io port address}
				empty

Fig. 24. Translation of bytecode for access to regular fields into bytecode for access to I/O port registers.

909 We base our access to hardware objects on IR instruction transformations. We in-  
 910 troduce two new instructions `outb` and `inb` for byte-wide access to I/O ports. Then we  
 911 employ OVM's instruction rewriting framework to translate accesses to hardware ob-  
 912 ject fields, `putfield` and `getfield` instructions, into sequences centered around `outb` and  
 913 `inb` where appropriate. We did not implement word-wide or double-word-wide access  
 914 modes supported by a x86 CPU. We discuss how this could be done at the end of this  
 915 section.

916 To minimize changes to the OVM code we keep the memory layout of hardware  
 917 objects as if they were ordinary objects, and store port addresses into the fields repre-  
 918 senting the respective hardware I/O ports. Explained with the example from Figure 18,  
 919 the instruction rewriting algorithm proceeds as follows: `SerialPort` is a subclass of `Hard-`  
 920 `wareObject`; hence it is a hardware object, and thus accesses to all its public volatile int  
 921 fields, `status` and `data`, are translated to port accesses to I/O addresses stored in those  
 922 fields.

923 The translation (Figure 24) is very simple. In case of reads we append our new `inb`  
 924 instruction after the corresponding `getfield` instruction in the IR: `getfield` will store the  
 925 I/O address on the stack and `inb` will replace it by a value read from this I/O address.  
 926 In case of writes we replace the corresponding `putfield` instruction by a sequence of  
 927 `swap`, `getfield`, and `outb`. The `swap` rotates the two top elements on stack, leaving the  
 928 hardware object reference on top of the stack and the value to store to the I/O port  
 929 below it, The `getfield` replaces the object reference by the corresponding I/O address,  
 930 and `outb` writes the value to the I/O port.

931 The critical part of hardware object creation is to set I/O addresses into hardware  
 932 object fields. Our approach allows a method to turn off the special handling of hard-  
 933 ware objects. In a hardware object factory method accesses to hardware object fields  
 934 are handled as if they were fields of regular objects; we simply store I/O addresses to  
 935 the fields.

936 A method can turn off the special handling of hardware objects with a marker  
 937 exception mechanism which is a natural solution within OVM. The method declares to  
 938 throw a `PragmaNoHWIORegistersAccess` exception. This exception is neither thrown  
 939 nor caught, but the OVM IR-level rewriter detects the declaration and disables

940 rewriting accordingly. As the exception extends `RuntimeException`, it does not need  
941 to be declared in interfaces or in code calling factory methods. In Java 1.5, not  
942 supported by OVM, a standard substitute to the marker exception would be method  
943 annotation.

944 Our solution depends on the representation of byte-wide registers by 16-bit fields  
945 to hold the I/O address. However, it could still be extended to support multiple-width  
946 accesses to I/O ports (byte, 16-bit, and 32-bit) as follows: 32-bit I/O registers are rep-  
947 resented by Java long fields, 16-bit I/O registers by Java int fields, and byte-wide I/O  
948 registers by Java short fields. The correct access width will be chosen by the IR rewriter  
949 based on the field type.

950 **4.4.2 Interrupt Handler.** Low-level support depends heavily on scheduling and preemp-  
951 tion. For our experiments we chose the uniprocessor x86 OVM configuration with  
952 green threads running as a single Linux process. The green threads, delayed I/O opera-  
953 tions, and handlers of asynchronous events, such as POSIX signals, are only scheduled  
954 at well-defined points (*pollchecks*) which are by default at back-branches at bytecode  
955 level and indirectly at Java-level blocking calls (I/O operations, synchronization calls,  
956 etc). When no thread is ready to run, the OVM scheduler waits for events using the  
957 POSIX select call.

958 As OS we use Xenomai RT Linux [Gerum 2004; Xenomai Developers 2008]. Xeno-  
959 mai tasks, which are in fact user-space Linux threads, can run either in the Xenomai  
960 primary domain or in the Xenomai secondary domain. In the primary domain they are  
961 scheduled by the Xenomai scheduler, isolated from the Linux kernel. In the secondary  
962 domain Xenomai tasks behave as regular real-time Linux threads. Tasks can switch to  
963 the primary domain at any time, but are automatically switched back to the secondary  
964 domain whenever they invoke a Linux system call. A single Linux process can have  
965 threads of different types: regular Linux threads, Xenomai primary domain tasks, and  
966 Xenomai secondary domain tasks. Primary domain tasks can wait on hardware inter-  
967 rupts with a higher priority than the Linux kernel. The Xenomai API provides the  
968 interrupts using the ISR event handler approach and supports *virtualization* of basic  
969 interrupt operations – disabling and enabling a particular interrupt or all local CPU  
970 interrupts. These operations have the same semantics as real interrupts, and dis-  
971 abling/enabling a particular one leads to the corresponding operation being performed  
972 at the hardware level.

973 Before our extension, OVM ran as a single Linux process with a single (native  
974 Linux) thread, a *main OVM thread*. This native thread implemented Java green  
975 threads. To support interrupts we add additional threads to the OVM process: for each  
976 interrupt source handled in OVM we dynamically add an interrupt listener thread run-  
977 ning in the Xenomai primary domain. The mechanism that leads to invocation of the  
978 Java interrupt handler thread is illustrated in Figure 25.

979 Upon receiving an interrupt, the listener thread marks the pending interrupt in a  
980 data structure shared with the main OVM thread. When it reaches a pollcheck, it  
981 discovers that an interrupt is pending. The scheduler then immediately wakes-up and  
982 schedules the Java green thread that is waiting for the interrupt (IRQ server thread  
983 in the figure). To simulate the first-level ISR handler approach, this green thread  
984 invokes some handler method. In a non-RTSJ scenario the green thread invokes the  
985 `run()` method of the associated `InterruptHandler` (see Figure 16). In an RTSJ scenario  
986 (not shown in Figure 25), a specialized thread fires an asynchronous event bound to  
987 the particular interrupt source. It invokes the `fire()` method of the respective RTSJ's  
988 `AsyncEvent`. As mentioned in Section 3.3.3 the RTSJ logic of `AsyncEventHandler` (AEH)  
989 registered to this event should be an instance of `InterruptHandler` in order to allow the  
990 interrupt handling code to access basic interrupt handling operations.

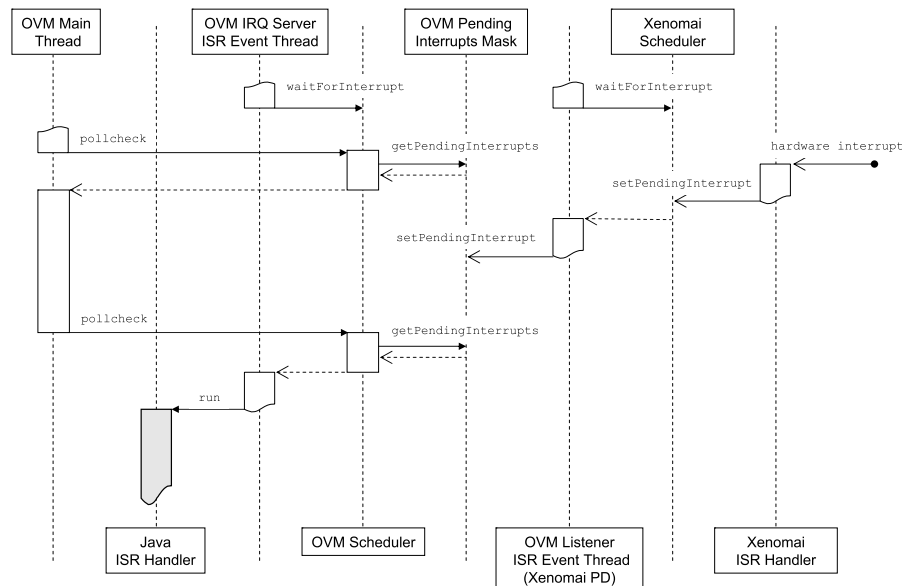


Fig. 25. Invocation of a Java interrupt handler under OVM/Xenomai.

991 As just explained, our first-level InterruptHandlers virtualize the interrupt handling  
 992 operations for interrupt enabling, disabling, etc. Therefore, we have two levels of inter-  
 993 rupt virtualization, one is provided by Xenomai to our listener thread, and the other  
 994 one, on top of the first one, is provided by the OVM runtime to the InterruptHandler  
 995 instance. In particular, disabling/enabling of local CPU interrupts is emulated, hard-  
 996 ware interrupts are disabled/enabled, and interrupt completion is performed at the  
 997 interrupt controller level (via the Xenomai API), and interrupt start is emulated; it  
 998 only tells the listener thread that the interrupt was received.

999 The RTSJ scheduling features (deadline checking, interarrival time checking, de-  
 1000 laying of sporadic events) related to release of the AEH should not require any further  
 1001 adaptations for interrupt handling. We could not test these features as OVM does not  
 1002 implement them.

1003 OVM uses *thin monitors* which means that a monitor is only instantiated (*inflated*)  
 1004 when a thread has to block on acquiring it. This semantic does not match to what  
 1005 we need: disable the interrupt when the monitor is acquired to prevent the handler  
 1006 from interrupting. Our solution provides a special implementation of a monitor for  
 1007 interrupt handlers and inflate it in the constructor of InterruptHandler. This way we  
 1008 do not have to modify the monitorenter and monitorexit instructions and we do not slow  
 1009 down regular thin monitors (noninterrupt-based synchronization).

#### 1010 4.5 Summary

1011 Support for hardware objects (see Section 3.1) and interrupt handling (see Section 3.2)  
 1012 to all four JVMs relies on common techniques. Accessing device registers through  
 1013 hardware objects extends the interpretation of the bytecodes putfield and getfield or  
 1014 redirects the pointer to the object. If these bytecodes are extended to identify the field  
 1015 being accessed as inside a hardware object, the implementation can use this infor-  
 1016 mation. Similarly, the implementation of interrupt handling requires changes to the  
 1017 bytecodes monitorenter and monitorexit or preinflating a specialized implementation of  
 1018 a Java monitor. In case of the bytecode extension, the extended codes specify if the

1019 monitor being acquired belongs to an interrupt handler object. If so, the implemen-  
1020 tation of the actual monitor acquisition must be changed to disable/enable interrupts.  
1021 Whether dealing with hardware or interrupt objects, we used the same approach of let-  
1022 ting the hardware object and interrupt handler classes inherit from the superclasses  
1023 `HardwareObject` and `InterruptHandler` respectively.

1024 For JVMs that need a special treatment of bytecodes `putfield` and `getfield` (Sim-  
1025 pleRTJ, Kaffe, and OVM) bytecode rewriting at runtime can be used to avoid the addi-  
1026 tional check of the object type. This is a standard approach (called *quick* bytecodes in  
1027 the first JVM specification) in JVMs to speed-up field access of resolved classes.

1028 Historically, registers of most x86 I/O devices are mapped to a dedicated I/O ad-  
1029 dress space, which is accessed using dedicated instructions: `port read` and `port writes`.  
1030 Fortunately, both the processor and Linux allow user-space applications running with  
1031 administrator privileges to use these instructions and access the ports directly via  
1032 `ioopl`, `inb`, and `outb` calls. For both the Kaffe and OVM implementations we have imple-  
1033 mented bytecode instructions `putfield` and `getfield` accessing hardware object fields by  
1034 calls to `ioopl`, `inb`, and `outb`.

1035 Linux does not allow user-space applications to handle hardware interrupts. Only  
1036 kernel space functionality is allowed to register interrupt handlers. We have overcome  
1037 this issue in two different ways.

1038 — For Kaffe we have written a special-purpose kernel module through which the user-  
1039 space application (the Kaffe VM) can register interest in interrupts and get notified  
1040 about interrupt occurrence.

1041 — For OVM we have used the Xenomai real-time extension to Linux. Xenomai extends  
1042 the Linux kernel to allow for the creation of real-time threads and allows user-space  
1043 code to wait for interrupt occurrences.

1044 Both these work-arounds allow an incremental transition of the JVMs and the re-  
1045 lated development libraries into a direct (bare metal) execution environment. In that  
1046 case the work-arounds would no longer be needed.

1047 If a compiling JVM is used (either as JIT or ahead-of-time) the compiler needs to be  
1048 aware of the special treatment of hardware objects and monitors on interrupt handlers.  
1049 One issue which we did not face in our implementations was the alignment of object  
1050 fields. When device registers are represented by differently sized integer fields, the  
1051 compiler needs to pack the data structure.

1052 The restrictions within an interrupt handler are JVM-dependent. If an interrupt-  
1053 ible, real-time GC is used (as in OVM and JOP) objects can be allocated in the handler  
1054 and the object graph may be changed. For a JVM with a stop-the-world GC (SimpleRTJ  
1055 and Kaffe) allocations are not allowed because the handler can interrupt the GC.

## 1056 5. EVALUATION AND CONCLUSION

1057 Having implemented the Java HAL on four different JVMs we evaluate it on a sev-  
1058 eral test applications, including a tiny Web server, and measure the performance of  
1059 hardware accesses via hardware objects and the latency of Java interrupt handlers.

### 1060 5.1 Qualitative Observations

1061 For first tests we implemented a serial port driver with hardware objects and interrupt  
1062 handlers. As the structure of the device registers is exactly the same on a PC, the  
1063 platform for SimpleRTJ, and JOP, we were able to use the exact same definition of the  
1064 hardware object `SerialPort` and the test programs on all four systems.

1065 Using the serial device we run an embedded TCP/IP stack, implemented completely  
1066 in Java, over a SLIP connection. The TCP/IP stack contains a tiny Web server and we



1067 serve Web pages with a Java-only solution similar to the one shown in the Introduction  
1068 in Figure 6. The TCP/IP stack, the tiny Web server, and the hardware object for the  
1069 serial port are the same for all platforms. The only difference is in the hardware object  
1070 creation with the platform-dependent factory implementations. The Web server uses  
1071 hardware objects and polling to access the serial device.

1072 *5.1.1 A Serial Driver in Java.* For testing the interrupt handling infrastructure in OVM  
1073 we implemented a serial interrupt-based driver in Java and a demo application that  
1074 sends back the data received through a serial interface. The driver part of the applica-  
1075 tion is a full-duplex driver with support for hardware flow control and with detection  
1076 of various error states reported by the hardware. The driver uses two circular buffers,  
1077 one for receiving and the other for sending. The user part of the driver implements  
1078 blocking `getChar` and `putChar` calls, which have (short) critical sections protected by  
1079 the interrupt-disabling monitor. To reduce latencies the `getChar` call sets the DSR flag  
1080 to immediately allow receiving more data and the `putChar`, after putting the charac-  
1081 ter into the sending buffer, initiates immediately the sending, if this is not currently  
1082 being done already by the interrupt machinery. The driver supports serial ports with  
1083 a FIFO buffer. The user part of the demo application implements the loop-back using  
1084 `getChar` and `putChar`. The user part is a RTSJ `AsyncEventHandler` which is fired when  
1085 a new character is received. From a Java perspective this is a second-level interrupt  
1086 handler, invoked after the corresponding serial event is fired from the first-level han-  
1087 dler. To test the API described in the article we implemented two versions that differ  
1088 in how the first-level handler is bound to the interrupt: (a) a RTSJ-style version where  
1089 the first-level handler is also a RTSJ event handler bound using `bindTo` to the JVM  
1090 provided first-level serial event, and (b) a non-RTSJ-style version where the first-level  
1091 handler is registered using a `InterruptHandler.register` call. We have stress-tested the  
1092 demo application and the underlying modified OVM infrastructure by sending large  
1093 files to it through the serial interface and checked that they were returned intact.

1094 *5.1.2 The HAL in Daily Use.* The original idea for hardware objects evolved during de-  
1095 velopment of low-level software on the JOP platform. The abstraction with `read` and  
1096 `write` functions and using constants to represent I/O addresses just *felt* wrong with  
1097 Java. Currently hardware objects are used all over in different projects with JOP.  
1098 Old code has been refactored to some extent, but new low-level code uses only hard-  
1099 ware objects. By now low-level I/O is integrated into the language, for example, auto-  
1100 completion in the Eclipse IDE makes it easy to access the factory methods and fields  
1101 in the hardware object.

1102 For experiments with an on-chip memory for thread-local scope caching [Wellings  
1103 and Schoeberl 2009] in the context of a chip-multiprocessor version of JOP, the hard-  
1104 ware array abstraction greatly simplified the task. The on-chip memory is mapped to  
1105 a hardware array and the RTSJ-based scoped memory uses it. Creation of an object  
1106 within this special scope is implemented in Java and is safe because the array bounds  
1107 checks are performed by the JVM.

1108 *5.1.3 JNI vs. Hardware Objects.* JNI provides a way to access the hardware without  
1109 changing the code of the JVM. Nevertheless, with a lack of commonly agreed API,  
1110 using it for each application would be redundant and error prone. It would also add  
1111 dependencies to the application: hardware platform and the operating system (the C  
1112 API for accessing the hardware is not standardized). The build process is complicated  
1113 by adding C code to it as well. Moreover, the system needs to support shared libraries,  
1114 which is not always the case for embedded operating systems (an example is RTEMS,  
1115 used by ESA).

1116 In addition, JNI is typically too heavyweight to implement trivial calls such as  
 1117 port or memory access efficiently (no GC interaction, no pointers, no threads inter-  
 1118 action, no blocking). Even JVMs that implement JNI usually have some other inter-  
 1119 nal lightweight native interface which is the natural choice for hardware access. This  
 1120 leads us back to a Java HAL as illustrated here.

1121 *5.1.4 OVM-Specific Experience.* Before the addition of hardware objects, OVM did not  
 1122 allow hardware access because it did not and does not have JNI or any other native  
 1123 interface for user Java code. OVM has a simplified native interface for the virtual  
 1124 machine code which indeed we used when implementing the hardware objects. This  
 1125 native interface can as well be used to modify OVM to implement user-level access to  
 1126 hardware via regular method calls. We have done this to implement a benchmark to  
 1127 measure HWO/native overheads (later in this section). As far as simple port access  
 1128 is concerned, none of the solutions is strictly better from the point of the JVM: the  
 1129 bytecode manipulation to implement hardware objects was easy, as well as adding code  
 1130 to propagate native port I/O calls to user code. Thanks to ahead-of-time compilation  
 1131 and the simplicity of the native interface, the access overhead is the same.

1132 The OVM compiler is fortunately not “too smart” so it does not get in the way of  
 1133 supporting hardware objects: if a field is declared volatile side-effects of reading of  
 1134 that field are not a problem for any part of the system.

1135 The API for interrupt handling added to OVM allows full control over interrupts,  
 1136 typically available only to the operating system. The serial port test application has  
 1137 shown that, at least for a simple device; it really allows us to write a driver. An in-  
 1138 teresting feature of this configuration is that OVM runs in user space and therefore it  
 1139 greatly simplifies development and debugging of Java-only device drivers for embed-  
 1140 ded platforms.

## 1141 5.2 Performance

1142 Our main objective for hardware objects is a clean object-oriented interface to hard-  
 1143 ware devices. Performance of device register access is an important goal for relatively  
 1144 slow embedded processors; thus we focus on that in the following. It matters less on  
 1145 general-purpose processors where the slow I/O bus essentially limits the access time.

1146 *5.2.1 Measurement Methodology.* Execution time measurement of single instructions  
 1147 is only possible on simple in-order pipelines when a cycle counter is available. On  
 1148 a modern superscalar architecture, where hundreds of instructions are in flight each  
 1149 clock cycle, direct execution time measurement becomes impossible. Therefore, we  
 1150 performed a bandwidth-based measurement. We measure how many I/O instructions  
 1151 per second can be executed in a tight loop. The benchmark program is self-adapting  
 1152 and increases the loop count exponentially until the measurement runs for more than  
 1153 one second and the iterations per second are reported. To compensate for the loop  
 1154 overhead we perform an overhead measurement of the loop and subtract that overhead  
 1155 from the I/O measurement. The I/O bandwidth  $b$  is obtained as follows.

$$b = \frac{cnt}{t_{test} - t_{ovhd}}$$

1156 Figure 26 shows the measurement loop for the read operation in method `test()` and  
 1157 the overhead loop in method `overhead()`. In the comment above the method the byte-  
 1158 codes of the loop kernel are shown. We can see that the difference between the two  
 1159 loops is the single bytecode `getField` that performs the read request.

1160 *5.2.2 Execution Time.* In Table III we compare the access time with native functions  
 1161 to the access via hardware objects. The execution time is given in clock cycles. We

```

public class HwoRead extends BenchMark {

    SysDevice sys = IOFactory.getFactory().getSysDevice();

    /* Bytecodes in the loop kernel
    ILOAD 3
    ILOAD 4
    IADD
    ALOAD 2
    GETFIELD com/jopdesign/io/SysDevice.uscntTimer : I
    IADD
    ISTORE 3
    */
    public int test(int cnt) {

        SysDevice s = sys;
        int a = 0;
        int b = 123;
        int i;

        for (i=0; i<cnt; ++i) {
            a = a+b+s.uscntTimer;
        }
        return a;
    }

    /* Bytecodes in the loop kernel
    ILOAD 3
    ILOAD 4
    IADD
    ILOAD 2
    IADD
    ISTORE 3
    */
    public int overhead(int cnt) {

        int xxx = 456;
        int a = 0;
        int b = 123;
        int i;

        for (i=0; i<cnt; ++i) {
            a = a+b+xxx;
        }
        return a;
    }
}

```

Fig. 26. Benchmark for the read operation measurement.

1162 scale the measured I/O bandwidth  $b$  with the clock frequency  $f$  of the system under  
 1163 test by  $n = \frac{f}{b}$ .

1164 We have run the measurements on a 100 MHz version of JOP. As JOP is a sim-  
 1165 ple pipeline, we can also measure short bytecode instruction sequences with the cycle  
 1166 counter. Those measurements provided the exact same values as the ones given by our  
 1167 benchmark, such that they validated our approach.

Table III. Access Time to a Device Register in Clock Cycles

	JOP		OVM		SimpleRTJ		Kaffe	
	read	write	read	write	read	write	read	write
native	5	6	5517	5393	2588	1123	11841	11511
HW Object	13	15	5506	5335	3956	3418	9571	9394

1168 On JOP the native access is faster than using hardware objects because a native ac-  
 1169 cess is a special bytecode and not a native function call. The special bytecode accesses  
 1170 memory directly where the bytecodes `putfield` and `getfield` perform a null pointer check  
 1171 and indirection through the handle for the field access. Despite the slower I/O ac-  
 1172 cess via hardware objects on JOP, the access is fast enough for all currently available  
 1173 devices. Therefore, we will change all device drivers to use hardware objects. The  
 1174 measurement for OVM was run on a Dell Precision 380 (Intel Pentium 4, 3.8 GHz,  
 1175 3G RAM, 2M 8-way set associative L2 cache) with Linux (Ubuntu 7.10, Linux 2.6.24.3  
 1176 with Xenomai-RT patch). OVM was compiled without Xenomai support and the gen-  
 1177 erated virtual machine was compiled with all optimizations enabled. As I/O port we  
 1178 used the printer port. Access to the I/O port via a hardware object is just slightly faster  
 1179 than access via native methods. This was expected as the slow I/O bus dominates the  
 1180 access time. On the SimpleRTJ JVM the native access is faster than access to hard-  
 1181 ware objects. The reason is that the JVM does not implement JNI, but has its own  
 1182 proprietary, more efficient way to invoke native methods. It is done in a prelinking  
 1183 phase where the `invokestatic` bytecode is instrumented with information to allow an  
 1184 immediate invocation of the target native function. On the other hand, using hard-  
 1185 ware objects needs a field lookup that is more time consuming than invoking a static  
 1186 method. With bytecode-level optimization at class load time it would be possible to  
 1187 avoid the expensive field lookup.

1188 We measured the I/O performance with Kaffe on an Intel Core 2 Duo T7300,  
 1189 2.00 GHz with Linux 2.6.24 (Fedora Core 8). We used access to the serial port for  
 1190 the measurement. On the interpreting Kaffe JVM we notice a difference between the  
 1191 native access and hardware object access. Hardware objects are around 20% faster.

1192 **5.2.3 Summary.** For practical purposes the overhead on using hardware objects is  
 1193 insignificant. In some cases there may even be an improvement in performance. The  
 1194 benefits in terms of safe and structured code should make this a very attractive option  
 1195 for Java developers.

### 1196 5.3 Interrupt Handler Latency

1197 **5.3.1 Latency on JOP.** To measure interrupt latency on JOP we use a periodic thread  
 1198 and an interrupt handler. The periodic thread records the value of the cycle counter  
 1199 and triggers the interrupt. In the handler the counter is read again and the differ-  
 1200 ence between the two is the measured interrupt latency. A plain interrupt handler  
 1201 as `Runnable` takes a constant 234 clock cycles (or 2.3  $\mu$ s for a 100 MHz JOP system)  
 1202 between the interrupt occurrence and the execution of the first bytecode in the han-  
 1203 dler. This quite large time is the result of two method invocations for the interrupt  
 1204 handling: (1) invocation of the system method `interrupt()` and (2) invocation of the ac-  
 1205 tual handler. For more time-critical interrupts the handler code can be integrated in  
 1206 the system method. In that case the latency drops down to 0.78  $\mu$ s. For very low-  
 1207 latency interrupts the interrupt controller can be changed to emit different bytecodes  
 1208 depending on the interrupt number, then we avoid the dispatch in software and can  
 1209 implement the interrupt handler in microcode.

Table IV. Interrupt (and polling) Latencies in Microseconds

	Median ( $\mu\text{s}$ )	3rd Quartile ( $\mu\text{s}$ )	95% Quantile ( $\mu\text{s}$ )	Maximum ( $\mu\text{s}$ )
Polling	3	3	3	8
Kernel	14	16	16	21
Hard	14	16	16	21
User	17	19	19	24
Ovm	59	59	61	203

1210 We have integrated the two-level interrupt handling at the application level. We  
 1211 set up two threads: one periodic thread, that triggers the interrupt, and a higher-  
 1212 priority event thread that acts as second-level interrupt handler and performs the  
 1213 handler work. The first-level handler just invokes `fire()` for this second-level handler  
 1214 and returns. The second-level handler gets scheduled according to the priority. With  
 1215 this setup the interrupt handling latency is 33  $\mu\text{s}$ . We verified this time by measuring  
 1216 the time between fire of the software event and the execution of the first instruction  
 1217 in the handler directly from the periodic thread. This took 29  $\mu\text{s}$  and is the overhead  
 1218 due to the scheduler. The value is consistent with the measurements in Schoeberl and  
 1219 Vitek [2007]. There we measured a minimum useful period of 50  $\mu\text{s}$  for a high-priority  
 1220 periodic task.

1221 The runtime environment of JOP contains a concurrent real-time GC [Schoeberl  
 1222 and Vitek 2007]. The GC can be interrupted at a very fine granularity. During sections  
 1223 that are not preemptive (data structure manipulation for a new and write-barriers on  
 1224 a reference field write) interrupts are simply turned off. The copy of objects and ar-  
 1225 rays during the compaction phase can be interrupted by a thread or interrupt handler  
 1226 [Schoeberl and Puffitsch 2008]. Therefore, the maximum blocking time is in the atomic  
 1227 section of the thread scheduler and not in the GC.

1228 *5.3.2 Latency on OVM/Xenomai.* For measuring OVM/Xenomai interrupt latencies, we  
 1229 have extended an existing interrupt latency benchmark, written by Jan Kiszka from  
 1230 the Xenomai team [Xenomai Developers 2008]. The benchmark uses two machines  
 1231 connected over a serial line. The *log* machine, running a regular Linux kernel, toggles  
 1232 the RTS state of the serial line and measures the time it takes for the *target* machine  
 1233 to toggle it back.

1234 To minimize measuring overhead the *log* machine uses only polling and disables lo-  
 1235 cal CPU interrupts while measuring. Individual measurements are stored in memory  
 1236 and dumped at shutdown so that they can be analyzed offline. We have made 400,000  
 1237 measurements in each experiment, reporting only the last 100,000 (this was to warm-  
 1238 up the benchmark, including memory storage for the results). The *log* machine toggles  
 1239 the RTS state regularly with a given period.

1240 We have tested 5 versions of the benchmark on the *target* machine: a polling version  
 1241 written in C (*polling*), a kernel-space interrupt handler in C/Xenomai running out of  
 1242 control of the Linux scheduler (*kernel*), a hard-realtime kernel-space interrupt handler  
 1243 running out of control of both the Xenomai scheduler and the Linux scheduler (*hard*),  
 1244 a user-space interrupt handler written in C/Xenomai (*user*), and finally an interrupt  
 1245 handler written in Java/OVM/Xenomai (*ovm*).

1246 The results are shown in Table IV. The median latency is 3  $\mu\text{s}$  for polling, 14  $\mu\text{s}$  for  
 1247 both kernel-space handlers (hard and kernel), 17  $\mu\text{s}$  for user-space C handler (*user*),  
 1248 and 59  $\mu\text{s}$  for Java handler in OVM (*ovm*). Note that the table shows that the overhead  
 1249 of using interrupts over polling is larger than the overhead of handling interrupts  
 1250 in user space over kernel space. The maximum latency of OVM was 203  $\mu\text{s}$ , due to

1251 infrequent pauses. Their frequency is so low that the measured 95% quantile is only  
1252 61  $\mu$ s.

1253 The experiment was run on Dell Precision 380 (Intel Pentium 4 3.8 GHz, 3G RAM,  
1254 2M 8-way set associative L2 cache) with Linux (Ubuntu 7.10, Linux 2.6.24.3 with  
1255 Xenomai-RT patch). As Xenomai is still under active development we had to use  
1256 Xenomai work-arounds and bugfixes, mostly provided by Xenomai developers, to make  
1257 OVM on Xenomai work.

1258 **5.3.3 Summary.** The overhead for implementing interrupt handlers is very accept-  
1259 able since interrupts are used to signal relatively infrequently occurring events like  
1260 end of transmission, loss of carrier, etc. With a reasonable work division between  
1261 first-level and second-level handlers, the proposal does not introduce dramatic block-  
1262 ing terms in a real-time schedulability analysis, and thus it is suitable for embedded  
1263 systems.

## 1264 **5.4 Discussion**

1265 **5.4.1 Safety Aspects.** Hardware objects map object fields to the device registers. When  
1266 the class that represents a device is correct, access to it is safe; it is not possible to read  
1267 from or write to an arbitrary memory address. A memory area represented by an array  
1268 is protected by Java's array bounds check.

1269 **5.4.2 Portability.** It is obvious that hardware objects are platform-dependent; after all  
1270 the idea is to have an interface to the bare metal. Nevertheless, hardware objects give  
1271 device manufacturers an opportunity to supply supporting factory implementations  
1272 that fit into Java's object-oriented framework and thus cater for developers of embed-  
1273 ded software. If the same device is used on different platforms, the hardware object is  
1274 portable. Therefore, standard hardware objects can evolve.

1275 **5.4.3 Compatibility with the RTSJ Standard.** As shown for the OVM implementation, the  
1276 proposed HAL is compatible with the RTSJ standard. We consider it to be a very im-  
1277 portant point since many existing systems have been developed using such platforms  
1278 or subsets thereof. In further development of such applications existing and future in-  
1279 terfacing to devices may be refactored using the proposed HAL. It will make the code  
1280 safer and more structured and may assist in possible ports to new platforms.

## 1281 **5.5 Perspective**

1282 The many examples in the text show that we achieved a representation of the hard-  
1283 ware close to being platform-independent. Also, they show that it is possible to imple-  
1284 ment system-level functionality in Java. As future work we consider to add devices  
1285 drivers for common devices such as network interfaces<sup>11</sup> and hard disc controllers.  
1286 On top of these drivers we will implement a file system and other typical OS-related  
1287 services towards our final goal of a Java-only system.

1288 An interesting question is whether a common set of *standard* hardware objects is  
1289 definable. The SerialPort was a lucky example. Although the internals of the JVMs and  
1290 the hardware were different one compatible hardware object worked on all platforms.  
1291 It should be feasible that a chip manufacturer provides, beside the data sheet that  
1292 describes the registers, a Java class for the register definitions of that chip. This defi-  
1293 nition can be reused in all systems that use that chip, independent of the JVM or OS.

1294 Another interesting idea is to define the interaction between the GC and hardware  
1295 objects. We stated that the GC should not collect hardware objects. If we relax this

---

<sup>11</sup>A device driver for a CS8900-based network chip is already part of the Java TCP/IP stack.

1296 restriction we can redefine the semantics of collecting an object: on running the  
1297 finalizer for a hardware object the device can be put into sleep mode.

## 1298 ACKNOWLEDGMENTS

1299 We wish to thank Andy Wellings for his insightful comments on an earlier version of the article. We also  
1300 thank the reviewers for their detailed comments that helped to enhance the original submission.

## 1301 REFERENCES

- 1302 AJILE. 2000. aj-100 real-time low power Java processor. Preliminary data sheet.
- 1303 ARMBRUSTER, A., BAKER, J., CUNEI, A., FLACK, C., HOLMES, D., PIZLO, F., PLA, E., PROCHAZKA, M.,  
1304 AND VITEK, J. 2007. A real-time Java virtual machine with applications in avionics. *Trans. Embed.*  
1305 *Comput. Sys.* 7, 1, 1–49.
- 1306 BACON, D. F., CHENG, P., AND RAJAN, V. T. 2003. A real-time garbage collector with low overhead and  
1307 consistent utilization. In *Proceedings of the 30th ACM SIGPLAN-SIGACT Symposium on Principles of*  
1308 *Programming Languages (POPL03)*. ACM Press, New York, 285–298.
- 1309 BOLLELLA, G., GOSLING, J., BROSGOL, B., DIBBLE, P., FURR, S., AND TURNBULL, M. 2000. *The Real-Time*  
1310 *Specification for Java*. Java Series. Addison-Wesley.
- 1311 BURNS, A. AND WELLINGS, A. J. 2001. *Real-Time Systems and Programming Languages: ADA 95, Real-*  
1312 *Time Java, and Real-Time POSIX* 3rd Ed. Addison-Wesley Longman Publishing.
- 1313 CASKA, J. accessed 2009. micro [ $\mu$ ] virtual-machine. <http://mvium.com/>.
- 1314 CHOU, A., YANG, J., CHELF, B., HALLEM, S., AND ENGLER, D. 2001. An empirical study of operating  
1315 systems errors. *SIGOPS Oper. Syst. Rev.* 35, 5, 73–88.
- 1316 FELSER, M., GOLM, M., WAWERSICH, C., AND KLEINÖDER, J. 2002. The JX operating system. In *Proceed-*  
1317 *ings of the USENIX Annual Technical Conference*. 45–58.
- 1318 GAMMA, E., HELM, R., JOHNSON, R., AND VLISSIDES, J. M. 1994. *Design Patterns: Elements of Reusable*  
1319 *Object-Oriented Software*. Addison Wesley Professional.
- 1320 GAY, D., LEVIS, P., VON BEHREN, R., WELSH, M., BREWER, E., AND CULLER, D. 2003. The nesC language:  
1321 A holistic approach to networked embedded systems. In *Proceedings of the ACM SIGPLAN Conference*  
1322 *on Programming Language Design and Implementation (PLDI'03)*. ACM Press, New York, 1–11.
- 1323 GERUM, P. 2004. Xenomai - Implementing a RTOS emulation framework on GNU/Linux.  
1324 <http://www.xenomai.org/documentation/branches/v2.4.x/pdf/xenomai.pdf>.
- 1325 GROUP, T. C. 2008. Trusted computing. <https://www.trustedcomputinggroup.org/>.
- 1326 HANSEN, P. B. 1977. *The Architecture of Concurrent Programs*. Prentice-Hall Series in Automatic Comput-  
1327 ing. Prentice-Hall.
- 1328 HENNESSY, J. AND PATTERSON, D. 2002. *Computer Architecture: A Quantitative Approach* 3rd Ed. Morgan  
1329 Kaufmann Publishers, Palo Alto, CA.
- 1330 HENTIES, T., HUNT, J. J., LOCKE, D., NILSEN, K., SCHOEBERL, M., AND VITEK, J. 2009. Java for safety-  
1331 critical applications. In *Proceedings of the 2nd International Workshop on the Certification of Safety-*  
1332 *Critical Software Controlled Systems (SafeCert'09)*.
- 1333 HILL, J., SZEWCZYK, R., WOO, A., HOLLAR, S., CULLER, D. E., AND PISTER, K. S. J. 2000. System archi-  
1334 tecture directions for networked sensors. In *Proceedings of the 9th International Conference on Archi-*  
1335 *tectural Support for Programming Languages and Operating Systems (ASPLOS IX)*. 93–104.
- 1336 HUNT, G., LARUS, J. R., ABADI, M., AIKEN, M., BARHAM, P., FAHNDRICH, M., HAWBLITZEL, C., HODSON,  
1337 O., LEVI, S., MURPHY, N., STEENSGAARD, B., TARDITI, D., WOBBER, T., AND ZILL, B. D. 2005. An  
1338 overview of the singularity project. Tech. rep. MSR-TR-2005-135, Microsoft Research (MSR).
- 1339 KORSHOLM, S., SCHOEBERL, M., AND RAVN, A. P. 2008. Interrupt handlers in Java. In *Proceedings of*  
1340 *the 11th IEEE International Symposium on Object/Component/Service-Oriented Real-Time Distributed*  
1341 *Computing (ISORC'08)*. IEEE Computer Society.
- 1342 KRALL, A. AND GRAFL, R. 1997. CACAO – A 64 bit JavaVM just-in-time compiler. In *Proceedings of the*  
1343 *Workshop on Java for Science and Engineering Computation (PPoPP'97)*, G. C. Fox and W. Li Eds. ACM.
- 1344 KREUZINGER, J., BRINKSCHULTE, U., PFEFFER, M., UHRIG, S., AND UNGERER, T. 2003. Real-Time event-  
1345 handling and scheduling on a multithreaded Java microcontroller. *Microprocess. Microsyst.* 27, 1, 19–31.
- 1346 LINDHOLM, T. AND YELLIN, F. 1999. *The Java Virtual Machine Specification* 2nd Ed. Addison-Wesley,  
1347 Reading, MA.
- 1348 LOHMEIER, S. 2005. Jini on the Jnode Java os. <http://monochromata.de/jnodejini.html>.

- 1349 PHIPPS, G. 1999. Comparing observed bug and productivity rates for java and c++. *Softw. Pract. Exper.* 29, 4,  
1350 345–358.
- 1351 RAVN, A. P. 1980. Device monitors. *IEEE Trans. Softw. Engin.* 6, 1, 49–53.
- 1352 RTJ COMPUTING. 2000. SimpleRTJ a small footprint Java VM for embedded and consumer devices.  
1353 <http://www.rtjcom.com/>.
- 1354 SCHOEBERL, M. 2005. Jop: A java optimized processor for embedded real-time systems. Ph.D. thesis, Vienna  
1355 University of Technology.
- 1356 SCHOEBERL, M. 2006. Real-Time garbage collection for Java. In *Proceedings of the 9th IEEE International*  
1357 *Symposium on Object and Component-Oriented Real-Time Distributed Computing (ISORC'06)*. IEEE,  
1358 424–432.
- 1359 SCHOEBERL, M. 2008. A Java processor architecture for embedded real-time systems. *J. Syst. Archit.* 54/1–2,  
1360 265–286.
- 1361 SCHOEBERL, M. AND VITEK, J. 2007. Garbage collection for safety critical Java. In *Proceedings of the 5th*  
1362 *International Workshop on Java Technologies for Real-Time and Embedded Systems (JTRES'07)*. ACM  
1363 Press, 85–93.
- 1364 SCHOEBERL, M. AND PUFFITSCH, W. 2008. Non-Blocking object copy for real-time garbage collection. In  
1365 *Proceedings of the 6th International Workshop on Java Technologies for Real-Time and Embedded Sys-*  
1366 *tems (JTRES'08)*. ACM Press.
- 1367 SCHOEBERL, M., KORSHOLM, S., THALINGER, C., AND RAVN, A. P. 2008. Hardware objects for Java. In  
1368 *Proceedings of the 11th IEEE International Symposium on Object/Component/Service-Oriented Real-*  
1369 *Time Distributed Computing (ISORC'08)*. IEEE Computer Society.
- 1370 SHA, L., RAJKUMAR, R., AND LEHOCZKY, J. P. 1990. Priority inheritance protocols: An approach to real-  
1371 time synchronization. *IEEE Trans. Comput.* 39, 9, 1175–1185.
- 1372 SIEBERT, F. 2002. *Hard Realtime Garbage Collection in Modern Object Oriented Programming Languages*.  
1373 aicas Books.
- 1374 SIMON, D., CIFUENTES, C., CLEAL, D., DANIELS, J., AND WHITE, D. 2006. Java on the bare metal of  
1375 wireless sensor devices: The squawk Java virtual machine. In *Proceedings of the 2nd International*  
1376 *Conference on Virtual Execution Environments (VEE'06)*. ACM Press, New York, 78–88.
- 1377 WELLINGS, A. AND SCHOEBERL, M. 2009. Thread-local scope caching for real-time Java. In *Proceedings of*  
1378 *the 12th IEEE International Symposium on Object/Component/Service-Oriented Real-Time Distributed*  
1379 *Computing (ISORC'09)*. IEEE Computer Society.
- 1380 WILKINSON, T. 1996. Kaffe – A virtual machine to run java code. <http://www.kaffe.org>.
- 1381 WIRTH, N. 1977. Design and implementation of modula. *Softw. Pract. Exper.* 7, 3–84.
- 1382 WIRTH, N. 1982. *Programming in Modula-2*. Springer.
- 1383 XENOMAI DEVELOPERS. 2008. Xenomai: Real-Time framework for Linux. <http://www.xenomai.org>.
- 1384 Received August 2008; revised July 2009; accepted February 2010