

A Verified Implementation of Priority Monitors in Java

Ángel Herranz and Julio Mariño*

Babel research group
Universidad Politécnica de Madrid

Abstract. Java monitors as implemented in the `java.util.concurrent.locks` package provide *no-priority nonblocking* monitors. That is, threads signalled after blocking on a *condition* queue do not proceed immediately, but they have to wait until both the signalling thread and possibly some of the others which have requested the lock release it. This can be a source of errors (if threads that get in the middle leave the monitor in a state incompatible with the signalled thread re-entry) or inefficiency (if repeated evaluation of preconditions is used to ensure safe re-entry). A concise implementation of priority nonblocking monitors in Java is presented. Curiously, our monitors are implemented on top of the standard no-priority implementation. In order to verify the correctness of our solution, a formal transition model (that includes a formalisation of Java locks and conditions) has been defined and checked using Uppaal. This model has been adapted to PlusCal in order to obtain a formal proof in TLA independent of the number of threads.

Keywords: Monitors, Java, model checking, priority, nonblocking, TLA, *PlusCal*.

1 Introduction

A *model-driven* approach to the development of concurrent software [7] advocates the use of high-level, language-independent entities that can be subject to formal analysis (e.g., to early detect risks due to concurrency in reactive, critical systems) and that are later translated into a concrete programming language by means of safe transformations. This, rather than the unrestricted use of the concurrency mechanisms in the target language, aims at improving portability (always a must in embedded systems) and preventing hazards due to the use of error-prone or poorly documented primitives. We have used this approach for teaching concurrency for more than fifteen years at our university, first using Ada95 [3], and now Java, which is certainly lacking in many aspects when compared to the former.

* Authors partly funded by grants DESAFIOS10 (TIN2009-14599-C03-00) from the Spanish Ministry of Science and Innovation, PROMETIDOS (P2009/TIC-1465) from the Madrid Regional Government and COST Action IC0701 on Formal Verification of Object Oriented Software.

<p>CADT Counter ACTION Inc: $C_Type[io]$ ACTION Dec: $C_Type[io]$</p> <p>SEMANTICS TYPE: $C_Type = \mathbb{Z}$ INVARIANT: $\forall c \in C_Type. c \geq 0$ INITIAL(c): $c = 0$</p> <p>CPRE: $True$ Inc(c) POST: $c^{out} = c^{in} + 1$</p> <p>CPRE: $c > 0$ Dec(c) POST: $c^{out} = c^{in} - 1$</p>	<pre> class Counter { final Lock mutex = new ReentrantLock(true); final Condition strictlyPos = mutex.newCondition(); private int c = 0; public void inc() { mutex.lock(); c++; // signal some pending dec strictlyPos.signal(); mutex.unlock(); } public void dec() { mutex.lock(); // check the CPRE now if (c == 0) { try {strictlyPos.await();} catch(Exception e) {} } // if we are here, that means // that c > 0 c--; // it CAN be proved that // no more signals are needed mutex.unlock(); } } </pre>
--	---

Fig. 1. A minimal example showing the risks of Java’s no-priority locks and conditions.

Indeed, the early mechanisms for thread synchronization provided by Java were so limited that one of the pioneers of concurrent programming, Brinch Hansen, wrote a famous essay [6] alerting on the risks of their usage. As a response to these adverse reactions, more powerful concurrency mechanisms were introduced in later versions of Java, although without being explicitly presented as realizations of the *classical* ones.

In this work we will focus on the *locks & conditions* library (`java.util.concurrent.locks`), that can be seen as an attempt to implement *monitors* (actually, Brinch Hansen’s contribution) in Java. Figure 1 shows an (apparently correct) implementation of a shared resource, specified in a formal notation, using the aforementioned library. On the left side of the figure, a shared counter is specified as a kind of abstract data type with implicit mutually exclusive access and an interface that specifies the only operations (**Inc** and **Dec**) that can be used to modify its state. Moreover, these *actions* are pure and transactional. In this case, we also state the invariant that the counter cannot be negative at any time, and specify the conditional synchronization (CPRE, *concurrency precondition*) for **Dec**, that ensures the invariant after the execution of any of the two actions.

To the right, we see an implementation of the resource specification as a Java class using locks and conditions. Class `Counter` encapsulates an integer variable `c` to represent the counter. Also, a lock object `mutex` is used to ensure execution of the actions in mutual exclusion. Finally, a condition object `strictlyPos` is used to implement the conditional synchronization of threads invoking `dec()`, i.e., to put them to sleep when the action is invoked with the counter set to 0.

The code has been written following the principle that only 0 or 1 `await()` calls are executed during an action, and the same applies to `signal()`. This coding guideline really pays off when applied to concurrency problems of greater significance. Basically, it reduces the context switching points inside the code for an action to (at most) one (the execution of `await()`), dividing it into two sequential fragments, which facilitates formal reasoning. It is the responsibility of the signalling thread (in this case, any thread executing `Inc`) to ensure that the signalled thread wakes up in a safe state. In this case, there is nothing to check, as `strictlyPos.signal()` is placed right after incrementing the counter and thus, if the signalled thread resumes execution right after the signalling one executes `mutex.unlock()` it is safe to decrement the counter. That justifies the comment above the line containing `"c--;"`.

Unfortunately, this code is unsafe. Why? The reason is that Java's implementation is an example of *no-priority* signalling [2]. That means that threads signalled after blocking on a *condition* queue do not proceed immediately, but they have to wait until both the signalling thread and also other threads that might have requested the lock release it. In other words, when the decremter thread resumes execution it could be the case that other decremter thread was already queued on `mutex` and set the counter to 0 right before. Moreover, given that the execution of the `inc()` and `dec()` methods takes very little time, the probability that this scenario happens is relatively low, which makes this the kind of error that can go unnoticed for thousands of runs.

The official documentation for the `Condition` class leaves considerable room for the operational interpretation of the `await()` and `signal()` methods in future implementations and, in fact, the no-priority behaviour is not really implied by the natural language specification of the API. There is, however, a generic recommendation to enclose the calls to `await()` inside invariant-checking loops, motivated by the possibility of *spurious* wakeups. Although the use of such loops can be seen as a defensive coding technique that can avoid hazards in general, we think that forcing the application programmer to use them can be quite unsatisfactory in many situations as testing the concurrency precondition repeatedly can increase contention of certain concurrent algorithms intolerably.

The use of the 0-1 coding scheme allows for low-contention, deterministic implementations of shared resources. Also, in [2] the reader can find a classification and comparison of many monitor variations, concluding that, in general, priority implementations are less error-prone. With this in mind, we decided to reuse as much from the existing reference implementation of condition objects and turn them into a concise, priority implementation of monitors. Our proposal is described in the following section.

```

1 package es.upm.babel.cclib;
2 import java.util.concurrent.locks.Lock;
3 import java.util.concurrent.locks.ReentrantLock;
4 import java.util.concurrent.locks.Condition;
5
6 public class Monitor {
7     private Lock mutex = new ReentrantLock(true);
8     private Condition purgatory = mutex.newCondition();
9     private int inPurgatory = 0;
10    private int pendingSignals = 0;
11
12    public Monitor() {
13    }
14
15    public void enter() {
16        mutex.lock();
17        if (pendingSignals > 0 || inPurgatory > 0) {
18            inPurgatory++;
19            try { purgatory.await(); }
20            catch (InterruptedException e) { }
21            inPurgatory--;
22        }
23    }
24
25    public void leave() {
26        if (pendingSignals == 0 && inPurgatory > 0) {
27            purgatory.signal();
28        }
29        mutex.unlock();
30    }
31
32    public Cond newCond() {
33        return new Cond();
34    }
35
36    

|                             |
|-----------------------------|
| inner class Cond goes here. |
|-----------------------------|


37 }

```

Fig. 2. Java source for priority monitors (class `Monitor`).

2 Implementing Priority Monitors with Java's No-priority Monitors

Figures 2 and 3 show a stripped-down version of the source code for `Monitor.java`, our implementation of priority monitors. Due to space limitations, comments and exception managers have been omitted. A full version, including comments and a couple of extra operations can be found at <http://babel.ls.fi.upm.es/software/cclib>.

In addition to the `Monitor` class, `Monitor.java` defines the `Cond` class. Both classes are intended to be a replacement for the original `Lock` and `Condition` classes. Figure 3 contains the code for class `Cond`, and the rest of `Monitor.java` is

```

37  public class Cond {
38      private Condition condition;
39      private int waiting;
40      private Cond() {
41          condition = mutex.newCondition();
42          waiting = 0;
43      }
44
45      public void await() {
46          waiting++;
47          if (pendingSignals == 0 && inPurgatory > 0 ) {
48              purgatory.signal();
49          }
50          try { condition.await(); }
51          catch (InterruptedException e) { }
52          pendingSignals--;
53      }
54
55      public void signal() {
56          if (waiting > 0) {
57              pendingSignals++;
58              waiting--;
59              condition.signal();
60          }
61      }
62  }

```

Fig. 3. Java source for priority condition variables (class `Cond`).

shown in Figure 2. However, this separation is a mere convenience for displaying the code: the implementation of class `Cond` requires access to the state variables in class `Monitor`.

The functionality provided by the class is similar to that of locks and conditions. Method `enter()` provides exclusive access to the monitor, like method `lock()` in class `Lock`. If one thread got exclusive access to a monitor object m , subsequent calls to $m.\text{enter}()$ from other threads will force these to wait in m 's queue. The monitor is released by invoking $m.\text{leave}()$ (analogous to `unlock()`). If there are threads waiting at m 's entry queue, executing $m.\text{leave}()$ will give control to the first thread in the queue.

Condition queues associated with m are created by invoking $m.\text{newCond}()$ rather than calling the constructor of class `Cond`. This is similar to the behaviour of `newCondition()` in class `Lock`. Instances c, c' ... of class `Cond` associated with a given monitor object m are managed like condition variables. A thread that invokes $c.\text{await}()$ after $m.\text{enter}()$ is blocked unconditionally and put to wait in a queue associated with c . Also, executing $c.\text{await}()$ releases m , so that other threads can get exclusive access to m by executing $m.\text{enter}()$. As in the case of `leave()`, if there are threads waiting at m 's entry queue, executing $m.\text{await}()$ will give control to the first thread in queue.

A thread that executes `c.signal()` after getting exclusive access to `m` will continue executing (signal-and-continue policy) but it is guaranteed that if there is a thread waiting on `c`, it will be awakened *before* any process waiting on `m`'s entry queue. This is where the behavior differs from that of signals in the standard locks and conditions package, where signalled threads are requeued at the tail of the lock's entry queue.

The implementation is done on top of the existing implementation of locks and conditions, without reimplementing their functionality say, by using low-level concurrency mechanisms such as semaphores, nor accessing the Java runtime in order to "magically move" threads from one queue to another.

The technique used to simulate the effect of moving signalled threads from the head of the condition queue to the head (rather than the tail) of the monitor's entry queue is to *flush* the contents of the latter until the signalled thread becomes the first. This is achieved by letting some threads get in the monitor, but only to make them await in a special condition queue devised for that purpose, and which we call "the purgatory". Of course, these threads will have to be eventually awakened and allowed to get access to the monitor in the right sequence. As they have been moved to a condition variable, signalling them will bring them back to the lock's entry queue, so a little care is needed to ensure that the whole system progresses appropriately. Two global counters are responsible for this.

Variable `inPurgatory` counts the number of threads that have been sent to the `purgatory` condition queue and have not been given access to the monitor yet – even if they have already been signalled. Variable `pendingSignals` counts how many threads that have been signalled (in regular condition variables, not `purgatory`) are yet waiting for re-entry.

Method `enter()` starts executing `lock()` on the monitor's internal lock object, called `mutex` (line 16). However, if `pendingSignals` is not zero, this thread should give way to some signalled thread that was waiting for monitor re-entry later in the queue associated with `mutex`, and thus has to execute a `purgatory.await()`, updating counter `inPurgatory` before and after (lines 17–22). The same is done if other threads have been sent to the purgatory earlier and have not re-entered yet (`inPurgatory > 0`).

The implementation of method `leave()` is quite symmetric. It ends by invoking `mutex.unlock()` (line 29), but before, a chance to re-gain access to the monitor must be given to threads moved to the purgatory (line 27), only if there are no signalled threads queued for monitor re-entry (line 26).

Method `newCond` is implemented just by invoking the standard constructor of class `Cond` (line 33). Association to a given monitor object is just implicit in the visibility of the monitor's state variables that the `Cond` object has. A `Cond` object is basically a pair of a `Condition` object associated with `mutex` and a counter `waiting` that keeps track of the number of threads waiting on it (lines 38–43).

The core of method `await()` is the corresponding call on its condition object (line 50) but, before, counter `waiting` is incremented (line 46) and the right to

execute inside the monitor is given to threads in the purgatory only if there are no pending signals (lines 47–49). Finally, if the thread that has invoked `await()` reaches line 52, the whole signalling procedure has been successful including monitor re-entry, so `pendingSignals` must be decremented accordingly.

Finally, method `signal()` checks whether the queue size is nonzero (line 56) and if so increments `pendingSignals` (line 57), as monitor re-entry for the signalled thread will be incomplete, decrements `waiting` (line 58) and executes `signal()` on the condition object (line 59). Executing `signal()` on an empty queue has no effect.

Although the implementation of the `Monitor` and `Cond` classes is quite concise, the interplay between the different threads and synchronization barriers, governed by the different counters is, admittedly, complex and hard to follow. After testing the class on a number of applications specifically devised to force errors related to incoming threads getting in the way of signalled ones, we still felt that a more convincing proof of the correctness of our implementation was necessary.

3 A Formal Model of Java Locks and Conditions

As a first step towards a formal argumentation of the correctness of our implementation of priority monitors, we decided to define a transition model for the underlying implementation of locks and condition variables.

The model has been defined as a network of (parametric) automata in Uppaal [1]: one automaton for modelling the concurrency mechanism, i.e., the lock object and its associated conditions, and one automaton per thread that makes use of it. While we give a complete formalization of the concurrency mechanism, unrelated thread code is abstracted away by just showing the interaction with the lock and the conditions.

Each thread has its own thread identifier, *pid*, and so do condition variables (*cid*). All interaction between the different automata takes place via communication channels:

- `lock[pid]` is the channel the thread *pid* uses when it needs to lock the lock.
- `await[cid][pid]` is the channel the thread *pid* uses when it needs to wait in condition *cid*.
- `signal[cid][pid]` is the channel a thread uses to signal condition *cid*.
- `lock.granted[pid]` is the channel the concurrency mechanism uses to grant a lock to the thread *pid*.
- `lock[pid]` is the channel the thread *pid* uses to unlock the lock.

Let's start with the abstract model for the threads, since, in our view, it is the most didactic way to present the whole model.

3.1 Modelling the Threads

Each thread has its own thread identifier, *pid*, and the model of a thread is the automaton (actually an Uppaal template parametrised by thread identifiers) shown in Figure 4.




Fig. 4. State model for a Java thread invoking operations on locks.

The automaton has four locations that represent the state of a thread with respect to a lock:

- Out is the location that represents that the thread neither has got the lock nor is waiting in a condition. It is an abstract representation of any program point of any arbitrary thread *outside* of the monitor.
- HasTheLock is the location that represents that the thread has the lock. It is another abstract representation, this time, of any program point of any arbitrary thread *in* the monitor.
- Locking is the location that represents that the thread is blocked while waiting for the lock to be granted to it. From the monitor perspective, it is trying to enter the monitor.
- Awaiting is the location that represents that the thread is blocked in a condition. From the monitor perspective, it is waiting for a signal that allows it to re-enter the monitor.

Edges are extremely simple. The idea behind this simplicity is trying to *contaminate* the thread models as little as possible. They do not involve access to variables, neither global nor local so adapting the model to any particular thread is immediate. Let us see the intended meaning of the actions at the edges:

- Out–Locking with action $\text{lock}[pid]!$, a send-action that indicates that the thread needs the lock. As we will see, the co-action ($\text{lock}[pid]?$) is continuously enabled in the model of the mechanism.
- Locking–HasTheLock with action $\text{lock.granted}[pid]?$, a receive-action that indicates the thread is allowed to get the lock. The co-action will occur when the thread pid has the *top* priority to lock the monitor.
- HasTheLock–Awaiting with action $\text{await}[cid][pid]!$, a send-action that indicates that the thread waits in the condition variable cid .¹
- Awaiting–HasTheLock with action $\text{lock.granted}[pid]?$, a receive-action that indicates the thread is allowed to re-take the lock (re-enter the monitor).

¹ In Uppaal, expressions such as $cid : cid.t$ non-deterministically bind the identifier cid to a value in the range $cid.t$. It can be understood as an edge *replicator*.

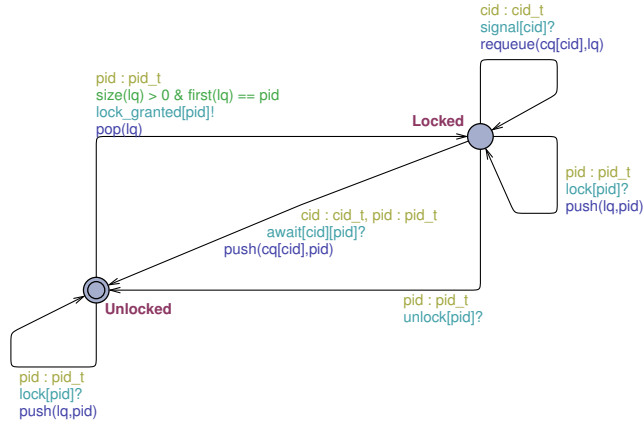


Fig. 5. State model for a Java lock.

- HasTheLock–Out with action `unlock[pid]!`, a send-action that indicates that the thread releases the lock (leave the monitor).
- HasTheLock–HasTheLock with action `signal[cid]!`, a send-action that indicates that the thread signals a thread on the condition variable `cid`.

3.2 Modelling Java Locks and Conditions

The model of a lock object and its associated condition variables is shown in Figure 5. The main task of the automaton is to keep track of threads by *listening* for their actions (on lock, await and signal channels) and responding appropriately (on channel `lock.granted`).

To keep track of the threads, the automaton defines two queues:

- A bounded queue of thread identifiers, `lq` (*lock queue*), that represent the queue for getting the lock.
- A bounded queue per condition `cid` of thread identifiers, `cq[cid]` (*condition queue of cid*), that represent the condition variable queues.

The automaton has two locations, `Locked` and `Unlocked`, that represent that the lock is locked by a thread or unlocked, respectively.

Most logic is encoded in the edges. To explain this logic we will explore the actions that fire them. We have to take into account that edges are *indexed* by thread identifiers (`pid`) and condition variable identifiers (`cid`).

- Edges with the receive-action `lock[pid]?` do not change locations, are always enabled and their assignments just push the thread identifier index `pid` in the access queue `lq`.
- The edge `Locked–Unlocked` with the receive-action `await[cid][pid]?` is always enabled and the assignment just pushes the thread identifier index `pid` in the condition variable queue `cq[cid]`.

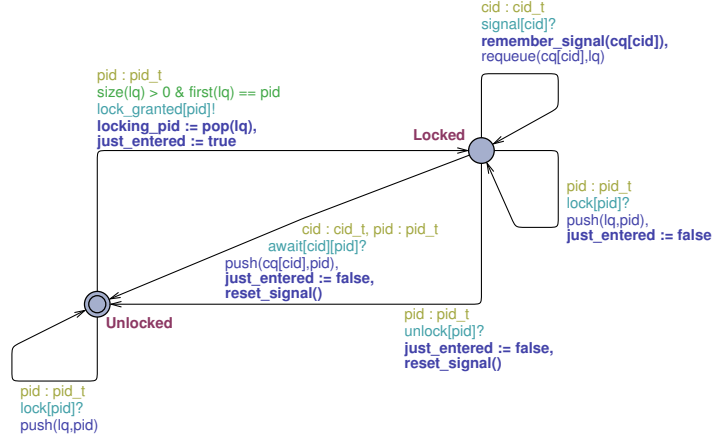


Fig. 6. An instrumented version of the Lock automaton.

- The edge Locked–Locked with the receive-action $\text{signal}[cid][pid]?$ is always enabled and the assignment just moves the first thread identifier index pid from the condition queue $cq[cid]$ to the access queue lq .
- The edge Unlocked–Locked with index pid is just enabled when the access queue lq is not empty and its first thread identifier is pid . The send-action of this edge is $\text{lock_granted}[pid]!$, granting the access to the thread with identifier pid .
- The edge Locked–Unlocked with the receive-action $\text{unlock}[pid]?$ is always enabled and does not modify any queue.

The full code for the declarations and auxiliary method definitions in this model can be found under <http://babel.ls.fi.upm.es/software/cclib>.

The natural step after defining this model is to check some properties on it. The first one is the correctness property: mutual exclusion. Encoding mutual exclusion in the Uppaal subset of CTL, the Computational Tree Logic [5], is easy:

$$A\Box (\forall (pid1 : pid.t) (\forall (pid2 : pid.t) ((Thread(pid1).HasTheLock \wedge Thread(pid2).HasTheLock) \Rightarrow pid1 = pid2)))$$

That is, whether for all paths (A path quantifier), invariantly (\Box modality), if *two* threads got the lock it is because they are the same thread. The tool finds the property correct for models of various sizes.

3.3 Instrumenting the Model

Nevertheless, for us, the most relevant property (correctness aside) is whether signalled threads resume their execution inside the monitor immediately after the signalling thread leaves. We will call this the *characteristic* property. As a

sanity check, we would like to check that this property does *not* hold for the model just presented.

However, stating this property in the query language supported by Uppaal is not possible. Basically, the query language is a subset of CTL that only allows temporal modalities as the topmost symbol of formulae, but a query representing the characteristic property for our model would require to nest modal operators.

A workaround for this limitation consists in instrumenting the model so that part of its recent history is somehow *remembered*. This way, we can encode certain temporal properties as state predicates, thus reducing the number of temporal operators required to express the characteristic property.

Again, we avoid contaminating the thread model and Figure 6 shows the resulting lock automaton (changes bold-faced). Basically, the system has new global variables to store:

- whether the n th thread that got access to the lock did execute an effective signal on some condition, and
- the identifier for the thread that received the signal (if any).

The first is done thanks to a new variable `just.signalled`, and the second with variable `thread.just.awakened`. Both of them are set by operation `remember.signal` while operation `reset.signal` resets `just.signalled` once the thread leaving the monitor coincides with `thread.just.awakened`. Also, the automata for threads are slightly constrained so that at most one signal is allowed per lock.

With these changes, the characteristic property is violated iff the $(n + 1)$ th thread in gaining access to the lock is different from the thread signalled by the n th thread. In order to avoid having generation counters in the model, we have added yet another boolean variable `just.entered` to represent that the last transition to take place in the lock automaton is precisely the one that follows a `lock.granted` message. Violation of the characteristic property can then be encoded in the Uppaal query

$$E\Diamond (\text{just.entered} \wedge \text{just.signalled} \wedge \text{locking_pid} \neq \text{thread.just.awakened}).$$

That is, whether there is some path (E path quantifier) that eventually leads (\Diamond modality) to a state in which the aforementioned proposition holds. The tool finds an example for 3 threads almost immediately, as expected.

4 Verifying our Implementation

Since our implementation of priority monitors is based on the existing lock and conditions model:

- The lock `mutex` is represented by the lock model in Figure 5.
- The condition `purgatory` is represented by the condition identifier `purgacid` (a symbolic name for referring to condition identifier 0).
- Integers `inPurgatory` and `pendingSignals` are represented by `int` Uppaal variables with the same name (`inPurgatory` and `pendingSignals`).

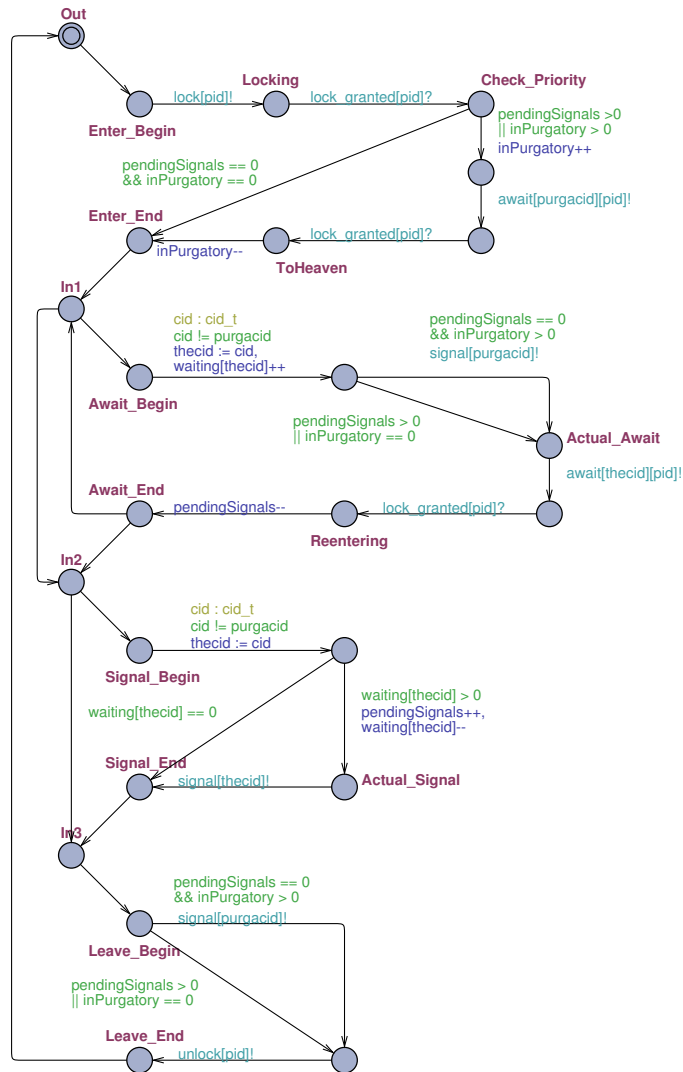


Fig. 7. State model for a thread using the Monitor class operations.

- The implementation of the methods **enter**, **leave**, **await**, and **signal** have been *inlined* in the thread model (Figure 7).

The automaton in Figure 7 models a very general thread that enters the monitor, then executes several **await** calls (possibly 0), then executes one, at the most, **signal** call, and finally leaves the monitor. To make the thread model easier to follow, most of the locations have been named according to fragments of the source code in `Monitor.java`. For example, execution of method `enter()` starts at location `Enter_Begin` and finishes at location `Enter_End` and includes the

	2	3	4	5	6
1:	0.01	0.02	0.18	6.38	274.03
2:	0.01	0.03	1.14	87.62	
3:	0.01	0.06	3.90		
4:	0.01	0.10	9.78		

locks & conditions

	2	3	4	5	6
1:	0.00	0.03	0.50	17.06	442.26
2:	0.01	0.08	3.35	214.99	
3:	0.02	0.18	12.06		
4:	0.02	0.39	35.70		

priority monitors

Table 1. Times spent in checking the mutual exclusion property.

	3	4	5	6
1:	0.00	0.02	0.08	0.40
2:	0.01	0.04	0.16	0.82
3:	0.01	0.08	0.32	1.48
4:	0.03	0.11	0.55	2.40

locks & conditions
(property violated)

	3	4	5	6
1:	0.03	0.36	11.18	450.73
2:	0.07	2.43	152.99	
3:	0.16	9.42		
4:	0.30	24.95		

priority monitors

Table 2. Times spent in checking the characteristic property.

possibility of a short excursion to the purgatory. The same scheme has been used to represent all the methods.

4.1 Experimental Results

The experimental results shown in this section have run on the instrumented models of the Java lock and conditions and of our priority monitor implementation. The model of the threads are those presented in previous sections except for the elimination of the main loop² (otherwise the state explosion makes the tool useless).

We have checked the correctness property for both models (lock and conditions and monitor) with the expected result: the models satisfy mutual exclusion. Table 1 shows execution times (in seconds) given different numbers of client threads (indicated at the top) and condition queues (indicated on the left).³

The state model for priority monitors has been instrumented following the ideas in Sec. 3.3. Table 2 shows execution times for checking the characteristic property for both instrumented models.

In addition to the characteristic property, we have also checked the property that threads sent to the purgatory eventually enter the monitor, as an implementation that would keep them blocked forever would satisfy the characteristic property trivially. The results are shown in Table 3.

² Represented by edge `Leave_End-Out` in Figure 7.

³ Figures obtained on a Dual-Core AMD Opteron(tm) Processor 2218 @ 2.6GHz, RAM 8GB, running the Academic version of Uppaal 4.1.7 (rev. 4934) running a Debian distro with kernel Linux 2.6.39-bpo.2-amd64 SMP.

	2	3	4	5
1:	0.01	0.02	0.42	14.04
2:	0.01	0.07	2.53	169.37
3:	0.01	0.15	9.49	
4:	0.01	0.30	27.52	
5:	0.02	0.58	64.18	

Table 3. Times spent in checking that the stay in the purgatory is transitory.

5 Towards a Mathematical Proof

In order to prove our implementation of monitors correct for an unbounded number of threads we have encoded it in TLA, the temporal logic of actions [8], a logic for specifying and reasoning about concurrent systems. This has been done in two steps:

- A straightforward, manual, transformation of the Java code to *PlusCal* [10], an algorithm description language specially suited for concurrency.
- An automatic transformation of the PlusCal descriptions into TLA+ [9], a complete specification language based on TLA, using The TLA Toolbox [12], which allow both model-checking finite systems⁴ but also symbolic proof.

The PlusCal code is shown in figures 8–10. Figure 8 contains the specification of the original locks and conditions library as a set of four procedures which can be invoked by concurrent processes. Analogously to the Uppaal model, locks and conditions are modelled as FIFO queues (“lq” and “cq[*cid*]”). The **await** statements (used at labels *java_lock_blocked* and *java_await_blocked*) ensure the synchronization behaviour. Labels in PlusCal specifications are not only an aid to the reader, but also help structuring the resulting TLA spec, and will be used for reasoning about the system behaviour.

The direct translation of our monitors is shown in Figure 9. The only relevant change is that part of the code (see labels *cclib_leave_resetlastsignal* and *cclib_signal_setlastsignal*) manages an extra variable (*lastSignal*) used to ease stating and proving the characteristic property, similarly to the instrumentation in the Uppaal model. Variable *lastSignal* is a tuple with two process identifiers. Initially, and when no signal is in progress, *lastSignal* = $\langle 0, 0 \rangle$ (0 is a non-valid process identifier). Otherwise, it records the pids of the signaler and signaled processes.

Finally, a system of several processes to reason about the behaviour of our algorithm is defined in Figure 10. These clients enter the monitor (*proc_enter*), execute 0 or 1 calls to await (*proc_await*), execute 0 or 1 signal (*proc_signal*), and then leave the monitor (*proc_leave*).

Theorem 1 (Characteristic property). *When a process performs a signal on a nonempty condition (label *cclib_signal_setlastsignal* in Figure 9), the next*

⁴ All theorems in this section have been model-checked with TLC, the TLA+ model-checker, prior to the symbolic proof.

```

java_lock_begin:      procedure java_lock() {
java_lock_blocked:   lq := Append(lq,self);
java_lock_end:       await self = Head(lq);
                    return;
                    }
                    procedure java_unlock() {
java_unlock_begin:   lq := Tail(lq);
java_unlock_end:     return;
                    }
                    procedure java_await(wcid) {
java_await_begin:    cq[wcid] := Append(cq[wcid],self);
java_await_unlocking: lq := Tail(lq);
java_await_blocked: await Len(lq) > 0 ∧ self = Head(lq);
java_await_end:      return;
                    }
                    procedure java_signal(scid) {
java_signal_begin:   if (Len(cq[scid]) > 0) {
java_signal_requeue: lq := Append(lq,Head(cq[scid]));
                    cq[scid] := Tail(cq[scid]);
                    };
java_signal_end:     return;
                    }

```

Fig. 8. PlusCal specification of Java locks and conditions.

process to enter the monitor (label *proc_in* in Figure 10) is the signaled one, and not an “opportunist” process:

$$\forall pid \in ProcSet : pc[pid] = \text{“proc_in”} \Rightarrow lastSignal[2] \in \{0, pid\}$$

Theorem 2 (Purgatory is transitory). *Every process sent to the purgatory ($cq[0]$) eventually enters the monitor (label *proc_in* in Figure 10):*

$$\forall pid \in ProcSet : (InPurgatory(pid) \rightsquigarrow pc[pid] = \text{“proc_in”})$$

(The temporal formula $\phi \rightsquigarrow \psi$ means that every state satisfying ϕ will eventually progress to another in which ψ holds.)

Theorem 3 (Enter call order is preserved). *No process (pid_2) entering the monitor (label *java_lock_blocked* in Figure 8) can overtake⁵ any process (pid_1) sent to the purgatory:*

$$\begin{aligned} \forall pid_1, pid_2 \in ProcSet : & (Member(cq[0], pid_1) \wedge pc[pid_2] = \text{“java_lock_blocked”}) \\ & \rightsquigarrow (\wedge Member(order, pid_1) \wedge Member(order, pid_2) \\ & \wedge Index(order, pid_1) < Index(order, pid_2)) \end{aligned}$$

The proof of the first theorem is lengthy due to the number of cases, but relatively straightforward, as is an invariant maintained by all transitions, and

⁵ Now the use of the instrumental queue *order* becomes apparent.

```

cclib_enter_begin:
cclib_enter_check_priority:
cclib_enter_toheaven:
cclib_enter_end:
cclib_leave_begin:
cclib_leave_saveit:
cclib_leave_resetlastsignal:
cclib_leave_unlocking:
cclib_leave_end:
cclib_await_begin:
cclib_await_saveit:
cclib_await_actualawait:
cclib_await_reentering:
cclib_await_end:
cclib_signal_begin:
cclib_signal_actualsignal:
cclib_signal_setlastsignal:
cclib_signal_actualsignal:
cclib_signal_end:

procedure cclib_enter() {
  call java_lock();
  if (pendingSignals > 0  $\vee$  inPurgatory > 0) {
    inPurgatory := inPurgatory + 1;
    call java_await(0);
    inPurgatory := inPurgatory - 1;
  };
  return;
}

procedure cclib_leave() {
  if (pendingSignals = 0  $\wedge$  inPurgatory > 0) {
    call java_signal(0);
  };
  if (lastSignal[1]  $\neq$  self) { lastSignal := (0, 0); };
  call java_unlock();
  return;
}

procedure cclib_await(cclwcid) {
  waiting[cclwcid] := waiting[cclwcid] + 1;
  if (pendingSignals = 0  $\wedge$  inPurgatory > 0) {
    call java_signal(0);
  };
  call java_await(cclwcid);
  pendingSignals := pendingSignals - 1;
  return;
}

procedure cclib_signal(cclscid) {
  if (waiting[cclscid] > 0) {
    pendingSignals := pendingSignals + 1;
    waiting[cclscid] := waiting[cclscid] - 1;
    lastSignal := (self, Head(cq[cclscid]));
    call java_signal(cclscid);
  };
  return;
}

```

Fig. 9. Priority monitors encoded in PlusCal.

independent of the number of processes. The second property is proved by induction on the number of awaiting processes. The third property is not strictly necessary for the safety requirements of the Monitor class, but is nice anyway. The full TLA theory and an appendix detailing the formal proof (done by hand) of the first two theorems can be found with the rest of the *cclib* software, at <http://babel.ls.fi.upm.es/software/cclib/doc>.


```

        process (procid ∈ ProcId) {
proc_enter:   call cclib_enter();
proc_in:     order := Append(order,self); \* to keep the order to gain access
proc_await:  with (cid ∈ CId) { if (cid # 0) { call cclib_await(cid); }; };
proc_signal: with (cid ∈ CId) { if (cid # 0) { call cclib_signal(cid); }; };
proc_leave:  call cclib_leave();
        }

```

Fig. 10. Process specification.

6 Conclusion

We have presented an implementation of nonblocking (signal-and-continue) priority monitors in Java, implemented on top of the existing nonblocking, no-priority implementation in the standard locks & conditions package. Moreover, we have provided a state model for our solution (extending a model of the existing Java mechanism) that gives *formal evidence* that the priority mechanism is actually implemented by our algorithm.

Our **Monitor** class encourages the use of certain coding styles that cannot be used with the standard Java implementation (i.e., the 0-1 await/signal coding idiom) and that result in cleaner, safer code, and we are actually using it in the classroom as a replacement of standard locks and conditions.

To our knowledge, there are just two other publicly available implementations of priority monitors in Java. The first one, by Chiao, Wu and Yuan [4], is really a language extension devised to overcome the limitations of synchronized methods and *per-object* waitsets – the paper is from 1999 and locks and conditions had not been added to Java yet. The implementation is based on a preprocessor which translates extended Java programs into standard Java code that invokes the methods of some **EMonitor** class in the right sequence. In a more recent work, T.S. Norvell [11] already considers the limitations of Java monitors in the light of the provided locks and conditions library. However, his approach is different from ours, reimplementing the monitor functionality on top of lower level concurrency mechanisms (semaphores) and using explicit queues of threads. The fact that our code is conciser and based on higher-level methods has facilitated the use of model-checking as a validation tool, while his implementation is not verified.⁶

Some details of the reference implementation of Java locks have been omitted in our model, namely the possibility of spurious wakeups and reentrant locking. Regarding the first, it is fairly reasonable to assume the implementation of class **ReentrantLock** to be free of spurious wakeups as, in the most liberal interpretation of the API specification, they would make most attempts at formal reasoning useless. Regarding the second, it can be seen as a benefit of the model-driven approach: we do not consider features useless for the intended idioms.

⁶ To be fair, Norvell’s implementation is longer also because he extends the functionality of Java monitors in other ways.

Methodologically speaking, the use of PlusCal/TLA as an intermediate stage between model checking/Uppaal and a deductive proof using some program verification tool seems an adequate compromise. On one hand, the PlusCal model is closer to the actual Java code while containing all the information in the Uppaal one. Also, now that we have crafted a TLA proof of the key properties of our system by hand, we have a clearer idea of the techniques that should be supported by a program verification system in order to facilitate the task that remains future work.

Acknowledgments. We are very grateful to Bart Jacobs and the anonymous reviewers for their comments on an earlier version of this paper.

References

1. Gerd Behrmann, Alexandre David, and Kim G. Larsen. A tutorial on Uppaal. In M. Bernardo and F. Corradini, editors, *International School on Formal Methods for the Design of Computer, Communication, and Software Systems, SFM-RT 2004. Revised Lectures*, volume 3185 of *Lecture Notes in Computer Science*, pages 200–237. Springer Verlag, 2004.
2. Peter A. Buhr, Michel Fortier, and Michael H. Coffin. Monitor classification. *ACM Computing Surveys*, 27:63–107, 1995.
3. Manuel Carro, Julio Mariño, Ángel Herranz, and Juan José Moreno-Navarro. Teaching how to derive correct concurrent programs (from state-based specifications and code patterns). In C.N. Dean and R.T. Boute, editors, *Teaching Formal Methods, CoLogNET/FME Symposium, TFM 2004, Ghent, Belgium*, volume 3294 of *LNCS*, pages 85–106. Springer, 2004. ISBN 3-540-23611-2.
4. Hsin-Ta Chiao, Chi-Houng Wu, and Shyan-Ming Yuan. A more expressive monitor for concurrent Java programming. In Arndt Bode, Thomas Ludwig, Wolfgang Karl, and Roland Wismüller, editors, *Euro-Par 2000 Parallel Processing*, volume 1900 of *Lecture Notes in Computer Science*, pages 1053–1060. Springer Berlin / Heidelberg, 2000.
5. E. Allen Emerson and Joseph Y. Halpern. sometimes and not never revisited: on branching versus linear time temporal logic. *J. ACM*, 33(1):151–178, January 1986.
6. Per Brinch Hansen. Java’s insecure parallelism. *ACM SIGPLAN Notices*, 34:38–45, 1999.
7. Ángel Herranz, Julio Mariño, Manuel Carro, and Juan José Moreno-Navarro. Modeling concurrent systems with shared resources. In *Formal Methods for Industrial Critical Systems, 14th International Workshop, FMICS 2009, Eindhoven, The Netherlands, November 2-3, 2009. Proceedings*, volume 5825 of *Lecture Notes in Computer Science*, pages 102–116, 2009.
8. Leslie Lamport. The temporal logic of actions. *ACM Trans. Program. Lang. Syst.*, 16(3):872–923, May 1994.
9. Leslie Lamport. *Specifying Systems*. Addison Wesley, 2004.
10. Leslie Lamport. The PlusCal algorithm language. In Martin Leucker and Carroll Morgan, editors, *Theoretical Aspects of Computing (ICTAC2009)*, number 5684 in *LNCS*, pages 36–60. Springer Verlag, 2009.
11. Theodore S. Norvell. Better monitors for Java. *Javaworld*, October 2007. <http://www.javaworld.com/javaworld/jw-10-2007/jw-10-monitors.html>.
12. TLA+. The Way to Specify. <http://www.tlaplus.net/>.




Fig. 11. State of the process queues in the PlusCal simulation.

A Proofs

A proof of the three properties stated in Section 5 follows. Section A.5 contains the whole specification of the algorithm in TLA+.

A.1 Facts of the Model

The model of our monitor implementation is based on the following facts that will be used in the proofs:

- For a process to gain access to the monitor, from *cclib_enter* or *cclib_await(cid)*, its process id must be the first in the “locking queue” *lq* (and just one process can be at the head).
- Processes ids are inserted in this queue in three situations:
 1. When the process execute *java_lock*: label *java_lock_begin* in Figure 8.
 2. When the process is signalled by other process: the signaller moves the signalled process id from a condition queue (*cq(cid)*) to *lq*: label *java_signal_requeue* in Figure 8. This requeue can be done from two different points:
 - (a) When the signalled process was in the purgatory: labels *cclib_leave_saveit* and *java_await_saveit* in Figure 9.
 - (b) When the signalled process was in a user condition: label *cclib_actual_signal* in Figure 9.
- The only process that can progress is the one at the head of *lq*, as this is precisely the guard that controls the execution of processes after performing a *java_lock* (label *java_lock_blocked* in Figure 8) or a *java_await* (label *java_await_blocked* in Figure 8).

Several lemmata are assumed for all of them, which are stated first. Also, a graphical representation of a state of the PlusCal execution is shown in Figure 11.

Queue *lq* is the “lock” or “enter” queue. Process identifiers are *always* enqueued in *lq* when trying to execute an *enter*. Moreover, a pid can be inserted in this queue in two more situations: right after receiving a *signal* and after being

sent to the *purgatory*. It is important to note that apart from the processes that try to execute *enter* – which simply enqueue their pids in *lq* and get blocked – the only process that can progress is the one at the head of *lq*, as this is precisely the guard that controls the execution of processes after performing an *enter*, thus ensuring mutual exclusion inside the monitor.

Queue $pq = cq_0$ is the purgatory queue. A process trying to enter “moves” its own pid from the head of *lq* to the rear of *pq* when there are pending signals – to give way to the signalled process – or when there are already some pid’s annotated in the purgatory – so as to prevent overtaking other processes in a similar situation.

Queues $cq_1 \dots cq_k$ are the “condition” queues. A process that executes *await* on condition *i* moves its pid from the head of *lq* to the rear of cq_i . If, later, another process executes a signal on condition *i*, it will move the pid at the front of cq_i to the rear of *lq*, and *pendingSignals* is incremented. This is the second “state” that a process annotated in *lq* can be in.

As $pq = cq_0$, i.e., it is one distinguished condition queue, processes blocked on it are eventually signalled when no better option exists. That means that those pids will also be moved from the front of *pq* to the rear of *lq*. As this happens when the process is blocked at a different code point than a normal *signal*, we regard this as a third state. The colors in Figure 11 are used to distinguish the point of the code at which the process whose pid is at a given cell of *lq* blocked: *red* for processes visiting *lq* for the first time, *blue* for processes which come from the purgatory and *yellow* for processes signalled after being blocked on a (normal) condition.⁷

The aforementioned lemmata follow. Lemmas only needed for proving one of the main properties can be found in the corresponding sections.

Lemma 1 (Weak fairness). *We assume that all processes have finite progress.*

Proof. By definition, the system specification *Spec* in TLA+ directly specifies finite progress for every process: weak fairness is required for the *Next* action and the *Next* action means “every process is eligible for execution” (by an explicit enumeration of their program points):

$$\begin{aligned} \mathbf{Spec} &\triangleq \wedge \mathit{Init} \wedge \square[\mathit{Next}]_{\mathit{vars}} \wedge \mathbf{WF}_{\mathit{vars}}(\mathbf{Next}) \\ \mathbf{Next} &\triangleq (\exists \mathit{self} \in \mathit{ProcSet} : \vee \mathit{java_lock}(\mathit{self}) \vee \mathit{java_unlock}(\mathit{self}) \\ &\quad \vee \mathit{java_await}(\mathit{self}) \vee \mathit{java_signal}(\mathit{self}) \\ &\quad \vee \mathit{cclib_enter}(\mathit{self}) \vee \mathit{cclib_leave}(\mathit{self}) \\ &\quad \vee \mathit{cclib_await}(\mathit{self}) \vee \mathit{cclib_signal}(\mathit{self})) \\ &\quad \vee \mathit{proc_enter}(\mathit{self}) \vee \mathit{proc_in}(\mathit{self}) \\ &\quad \vee \mathit{proc_await}(\mathit{self}) \vee \mathit{proc_signal}(\mathit{self}) \\ &\quad \vee \mathit{proc_leave}(\mathit{self})) \end{aligned}$$

Lemma 2 (Mutual exclusion). *The only process that can progress inside the monitor is the one whose pid is annotated at the front of *lq*.*

⁷ Abusing a bit, we have displayed the cell for process 4 in blue, state that will be reached when it is moved back to *lq* and no yellow cells exist.

Proof. Only one process id can be the first element of lq .

Lemma 3. *All queues have bounded size.*

Lemma 4 (At most one pending signal). *If processes are not allowed to call signal twice before they execute leave, the value of pendingSignals is at most one.*

A.2 Proof of Theorem 1

The characteristic property was expressed as an invariant of a slightly instrumented model in which a pair of two auxiliary variables were used to recall the pids of the signaler and signaled processes after performing a signal:

$$\forall pid \in ProcSet : (InPurgatory(pid) \rightsquigarrow pc[pid] = \text{"proc_in"})$$

Lemma 5. *There can only be one yellow pid in lq at a given time.*

We will consider several cases, depending on the execution state of the process whose pid is at the front of lq , which will be named p :

Case 1: a signaled process. Then p is a *yellow* cell. Initially, when p reaches the front of lq , we know that $pc[p]$ is not in any of the points required for the antecedent of the invariant. Then, the state evolves to a state in which the antecedent is true and (due to the previous lemma) the consequent is also true – $lastSignal[1] = p$ – and later, p performs a *leave*, which resets the pair (that makes the invariant hold) and finally removes p from the front of lq (which takes us to a different case).

Case 2: a locker. There are also several subcases here. If p is the signaler, an argument similar to the previous one shows that the system evolves only through states in which the invariant holds until p performs a *leave* – in this case $lastSignal[0] = p$ right before the *leave*.

If p is a (*red*) locker, and there is no pending signal, then $lastSignal = \langle 0, 0 \rangle$, and the invariant holds when $pc[p] = \text{locker_in}$. Otherwise, if there is one signal pending, then p blocks in the purgatory before removing itself from lq so the invariant holds in all the states in between.

The reasoning for *blue* lockers is analogous.

XM Note:
Necesitamos un lema
para esto?

A.3 Proof of Theorem 2

The heart of this proof lies in showing the “liveness” of pq , i.e. that it decreases eventually. This is so because pids cannot be in the purgatory an arbitrary number of times.

Lemma 6. *A pid cannot be in lq more than three times.*

Lemma 7 (Liveness of pq). *The pid at the front of pq will be eventually removed.*

XM Note: esta es la
parte que tiene seis
casos.

A.4 Proof of Theorem 3

A.5 CCLib Monitor Specification in TLA+

```

----- MODULE Monitor -----
EXTENDS Naturals, Sequences, TLC, FiniteSets, TLAPS
-----

Auxiliary definitions

General definitions
PositiveInteger  $\triangleq$  Nat \ {0}

Operations on
Last(q)  $\triangleq$  q[Len(q)]
Front(q)  $\triangleq$  [i  $\in$  1 .. (Len(q) - 1)  $\mapsto$  q[i]]
not used: Insert_as_2nd(q, x)  $\triangleq$   $\langle$ x $\rangle$   $\circ$  q
Member(q, e)  $\triangleq$   $\exists$  i  $\in$  1 .. Len(q) : q[i] = e
Index(q, e)  $\triangleq$  CHOOSE i  $\in$  1 .. Len(q) : q[i] = e
-----

Process identifiers
CONSTANTS N_Procs
ASSUME AtLeastOne  $\triangleq$   $\wedge$  N_Procs  $\in$  PositiveInteger

First_ProcId  $\triangleq$  1
Last_ProcId  $\triangleq$  First_ProcId + N_Procs - 1
ProcId  $\triangleq$  First_ProcId .. Last_ProcId
-----

Condition identifiers (condition number 0 will be implemented as purgatory)
CONSTANT N_Conditions
ASSUME N_Conditions  $\in$  PositiveInteger
CId  $\triangleq$  0 .. N_Conditions
-----

-algorithm TestMonitor
{

  variables
  \ * Problem variables
  cq = [cid  $\in$  CId  $\mapsto$   $\langle$  $\rangle$ ];
  lq =  $\langle$  $\rangle$ ;
  pendingSignals = 0;
  inPurgatory = 0;
  waiting = [cid  $\in$  CId  $\mapsto$  0];

  \ * Instrumentation
  lastSignal =  $\langle$ 0, 0 $\rangle$ ; \ *  $\langle$ signaller, signalled $\rangle$ 
  entryOrder =  $\langle$  $\rangle$ ;

```

```

\ *****
\ * Specification of the Java implementations of Lock and Conditions
procedure java_lock()
{
  java_lock_begin :
    lq := Append(lq, self);
  java_lock_blocked :
    await self = Head(lq);
  java_lock_end :
    return;
}

procedure java_unlock()
{
  java_unlock_begin :
    lq := Tail(lq);
  java_unlock_end :
    return;
}

procedure java_await(wcid)
{
  java_await_begin :
    cq[wcid] := Append(cq[wcid], self);
  java_await_unlocking :
    lq := Tail(lq);
  java_await_blocked :
    await Len(lq) > 0  $\wedge$  self = Head(lq);
  java_await_end :
    return;
}

procedure java_signal(scid)
{
  java_signal_begin :
    if (Len(cq[scid]) > 0){
      lq := Append(lq, Head(cq[scid]));
      cq[scid] := Tail(cq[scid]);
    };
  java_signal_end :
    return;
}

\ *****
\ * Specification of the CCLib implementations of Monitors
procedure cclib_enter()
{
  cclib_enter_begin :
    call java_lock();
  cclib_enter_check_priority :
    if (pendingSignals > 0  $\vee$  inPurgatory > 0){
      inPurgatory := inPurgatory + 1;
    };
}

```

```

cclib_enter_to_purgatory :
  call java_await(0);
cclib_enter_toheaven :
  inPurgatory := inPurgatory - 1;
};
cclib_enter_end :
  return;
}

procedure cclib_leave()
{
  cclib_leave_begin :
    if (pendingSignals = 0 ∧ inPurgatory > 0){
      call java_signal(0);
    };
  cclib_leave_resetlastsignal :
    if (lastSignal[1] ≠ self){
      lastSignal := ⟨0, 0⟩;
    };
  cclib_leave_unlocking :
    call java_unlock();
  cclib_leave_end :
    return;
}

procedure cclib_await(cclwcid)
{
  cclib_await_begin :
    waiting[cclwcid] := waiting[cclwcid] + 1;
    if (pendingSignals = 0 ∧ inPurgatory > 0){
      call java_signal(0);
    };
  cclib_await_actualawait :
    call java_await(cclwcid);
  cclib_await_reentering :
    pendingSignals := pendingSignals - 1;
  cclib_await_end :
    return;
}

procedure cclib_signal(cclscid)
{
  cclib_signal_begin :
    if (waiting[cclscid] > 0){
  cclib_signal_actualsignal :
    pendingSignals := pendingSignals + 1;
    waiting[cclscid] := waiting[cclscid] - 1;
  cclib_signal_setlastsignal :
    lastSignal := ⟨self, Head(cq[cclscid])⟩;
    call java_signal(cclscid);
  };
}

```



```

    cclib_signal_end :
        return;
    }
    \ *****
    \ * Process in the system
    process(procId ∈ ProcId)
    {
        proc_enter :
            call cclib_enter();

        proc_in :
            entryOrder := Append(entryOrder, self);

        proc_await :
            with(cid ∈ CId){
                if(cid ≠ 0){
                    call cclib_await(cid);
                };
            };

        proc_signal :
            with(cid ∈ CId){
                if(cid ≠ 0){
                    call cclib_signal(cid);
                };
            };

        proc_leave :
            call cclib_leave();
    }
}

BEGIN TRANSLATION
CONSTANT defaultInitValue
VARIABLES cq, lq, pendingSignals, inPurgatory, waiting, lastSignal,
           entryOrder, pc, stack, wcid, scid, cclwcid, cclscid

vars ≜ ⟨cq, lq, pendingSignals, inPurgatory, waiting, lastSignal,
        entryOrder, pc, stack, wcid, scid, cclwcid, cclscid⟩

ProcSet ≜ (ProcId)

Init ≜ Global variables
      ∧ cq = [cid ∈ CId ↦ ⟨⟩]
      ∧ lq = ⟨⟩
      ∧ pendingSignals = 0
      ∧ inPurgatory = 0

```

$\wedge \text{waiting} = [\text{cid} \in \text{CId} \mapsto 0]$
 $\wedge \text{lastSignal} = \langle 0, 0 \rangle$
 $\wedge \text{entryOrder} = \langle \rangle$
 Procedure *java_await*
 $\wedge \text{wcid} = [\text{self} \in \text{ProcSet} \mapsto \text{defaultInitValue}]$
 Procedure *java_signal*
 $\wedge \text{scid} = [\text{self} \in \text{ProcSet} \mapsto \text{defaultInitValue}]$
 Procedure *cclib_await*
 $\wedge \text{cclwcid} = [\text{self} \in \text{ProcSet} \mapsto \text{defaultInitValue}]$
 Procedure *cclib_signal*
 $\wedge \text{cclscid} = [\text{self} \in \text{ProcSet} \mapsto \text{defaultInitValue}]$
 $\wedge \text{stack} = [\text{self} \in \text{ProcSet} \mapsto \langle \rangle]$
 $\wedge \text{pc} = [\text{self} \in \text{ProcSet} \mapsto \text{CASE } \text{self} \in \text{ProcId} \rightarrow \text{"proc_enter"}]$

$\text{java_lock_begin}(\text{self}) \triangleq \wedge \text{pc}[\text{self}] = \text{"java_lock_begin"}$
 $\wedge \text{lq}' = \text{Append}(\text{lq}, \text{self})$
 $\wedge \text{pc}' = [\text{pc EXCEPT } ![\text{self}] = \text{"java_lock_blocked"}]$
 $\wedge \text{UNCHANGED } \langle \text{cq}, \text{pendingSignals}, \text{inPurgatory},$
 $\text{waiting}, \text{lastSignal}, \text{entryOrder},$
 $\text{stack}, \text{wcid}, \text{scid}, \text{cclwcid},$
 $\text{cclscid} \rangle$

$\text{java_lock_blocked}(\text{self}) \triangleq \wedge \text{pc}[\text{self}] = \text{"java_lock_blocked"}$
 $\wedge \text{self} = \text{Head}(\text{lq})$
 $\wedge \text{pc}' = [\text{pc EXCEPT } ![\text{self}] = \text{"java_lock_end"}]$
 $\wedge \text{UNCHANGED } \langle \text{cq}, \text{lq}, \text{pendingSignals},$
 $\text{inPurgatory}, \text{waiting},$
 $\text{lastSignal}, \text{entryOrder}, \text{stack},$
 $\text{wcid}, \text{scid}, \text{cclwcid}, \text{cclscid} \rangle$

$\text{java_lock_end}(\text{self}) \triangleq \wedge \text{pc}[\text{self}] = \text{"java_lock_end"}$
 $\wedge \text{pc}' = [\text{pc EXCEPT } ![\text{self}] = \text{Head}(\text{stack}[\text{self}]).\text{pc}]$
 $\wedge \text{stack}' = [\text{stack EXCEPT } ![\text{self}] = \text{Tail}(\text{stack}[\text{self}])]$
 $\wedge \text{UNCHANGED } \langle \text{cq}, \text{lq}, \text{pendingSignals},$
 $\text{inPurgatory}, \text{waiting}, \text{lastSignal},$
 $\text{entryOrder}, \text{wcid}, \text{scid}, \text{cclwcid},$
 $\text{cclscid} \rangle$

$\text{java_lock}(\text{self}) \triangleq \text{java_lock_begin}(\text{self}) \vee \text{java_lock_blocked}(\text{self})$
 $\vee \text{java_lock_end}(\text{self})$

$\text{java_unlock_begin}(\text{self}) \triangleq \wedge \text{pc}[\text{self}] = \text{"java_unlock_begin"}$
 $\wedge \text{lq}' = \text{Tail}(\text{lq})$
 $\wedge \text{pc}' = [\text{pc EXCEPT } ![\text{self}] = \text{"java_unlock_end"}]$
 $\wedge \text{UNCHANGED } \langle \text{cq}, \text{pendingSignals},$
 $\text{inPurgatory}, \text{waiting},$
 $\text{lastSignal}, \text{entryOrder}, \text{stack},$

$wcid, scid, cclwcid, cclscid$

$java_unlock_end(self) \triangleq \wedge pc[self] = \text{"java_unlock_end"}$
 $\wedge pc' = [pc \text{ EXCEPT } ![self] = Head(stack[self]).pc]$
 $\wedge stack' = [stack \text{ EXCEPT } ![self] = Tail(stack[self])]$
 $\wedge \text{UNCHANGED } \langle cq, lq, pendingSignals,$
 $inPurgatory, waiting,$
 $lastSignal, entryOrder, wcid,$
 $scid, cclwcid, cclscid \rangle$

$java_unlock(self) \triangleq java_unlock_begin(self)$
 $\vee java_unlock_end(self)$

$java_await_begin(self) \triangleq \wedge pc[self] = \text{"java_await_begin"}$
 $\wedge cq' = [cq \text{ EXCEPT } ![wcid[self]] = Append(cq[wcid[self]], self)]$
 $\wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"java_await_unlocking"}]$
 $\wedge \text{UNCHANGED } \langle lq, pendingSignals,$
 $inPurgatory, waiting,$
 $lastSignal, entryOrder, stack,$
 $wcid, scid, cclwcid, cclscid \rangle$

$java_await_unlocking(self) \triangleq \wedge pc[self] = \text{"java_await_unlocking"}$
 $\wedge lq' = Tail(lq)$
 $\wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"java_await_blocked"}]$
 $\wedge \text{UNCHANGED } \langle cq, pendingSignals,$
 $inPurgatory, waiting,$
 $lastSignal, entryOrder,$
 $stack, wcid, scid, cclwcid,$
 $cclscid \rangle$

$java_await_blocked(self) \triangleq \wedge pc[self] = \text{"java_await_blocked"}$
 $\wedge Len(lq) > 0 \wedge self = Head(lq)$
 $\wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"java_await_end"}]$
 $\wedge \text{UNCHANGED } \langle cq, lq, pendingSignals,$
 $inPurgatory, waiting,$
 $lastSignal, entryOrder,$
 $stack, wcid, scid, cclwcid,$
 $cclscid \rangle$

$java_await_end(self) \triangleq \wedge pc[self] = \text{"java_await_end"}$
 $\wedge pc' = [pc \text{ EXCEPT } ![self] = Head(stack[self]).pc]$
 $\wedge wcid' = [wcid \text{ EXCEPT } ![self] = Head(stack[self]).wcid]$
 $\wedge stack' = [stack \text{ EXCEPT } ![self] = Tail(stack[self])]$
 $\wedge \text{UNCHANGED } \langle cq, lq, pendingSignals,$
 $inPurgatory, waiting, lastSignal,$
 $entryOrder, scid, cclwcid,$
 $cclscid \rangle$

$$\begin{aligned}
\text{java_await}(self) &\triangleq \text{java_await_begin}(self) \\
&\quad \vee \text{java_await_unlocking}(self) \\
&\quad \vee \text{java_await_blocked}(self) \vee \text{java_await_end}(self) \\
\\
\text{java_signal_begin}(self) &\triangleq \wedge pc[self] = \text{"java_signal_begin"} \\
&\quad \wedge \text{IF } \text{Len}(cq[scid[self]]) > 0 \\
&\quad \quad \text{THEN } \wedge lq' = \text{Append}(lq, \text{Head}(cq[scid[self]])) \\
&\quad \quad \quad \wedge cq' = [cq \text{ EXCEPT } ![scid[self]] = \text{Tail}(cq[scid[self]])] \\
&\quad \quad \text{ELSE } \wedge \text{TRUE} \\
&\quad \quad \quad \wedge \text{UNCHANGED } \langle cq, lq \rangle \\
&\quad \wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"java_signal_end"}] \\
&\quad \wedge \text{UNCHANGED } \langle \text{pendingSignals}, \text{inPurgatory}, \\
&\quad \quad \text{waiting}, \text{lastSignal}, \\
&\quad \quad \text{entryOrder}, \text{stack}, \text{wcid}, \text{scid}, \\
&\quad \quad \text{cclwcid}, \text{cclscid} \rangle \\
\\
\text{java_signal_end}(self) &\triangleq \wedge pc[self] = \text{"java_signal_end"} \\
&\quad \wedge pc' = [pc \text{ EXCEPT } ![self] = \text{Head}(\text{stack}[self]).pc] \\
&\quad \wedge scid' = [scid \text{ EXCEPT } ![self] = \text{Head}(\text{stack}[self]).scid] \\
&\quad \wedge stack' = [stack \text{ EXCEPT } ![self] = \text{Tail}(\text{stack}[self])] \\
&\quad \wedge \text{UNCHANGED } \langle cq, lq, \text{pendingSignals}, \\
&\quad \quad \text{inPurgatory}, \text{waiting}, \\
&\quad \quad \text{lastSignal}, \text{entryOrder}, \text{wcid}, \\
&\quad \quad \text{cclwcid}, \text{cclscid} \rangle \\
\\
\text{java_signal}(self) &\triangleq \text{java_signal_begin}(self) \\
&\quad \vee \text{java_signal_end}(self) \\
\\
\text{cclib_enter_begin}(self) &\triangleq \wedge pc[self] = \text{"cclib_enter_begin"} \\
&\quad \wedge stack' = [stack \text{ EXCEPT } ![self] = \langle [procedure \mapsto \text{"java_lock"}, \\
&\quad \quad \quad pc \quad \quad \mapsto \text{"cclib_enter_check_priority"}], \\
&\quad \quad \quad \circ stack[self] \rangle] \\
&\quad \wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"java_lock_begin"}] \\
&\quad \wedge \text{UNCHANGED } \langle cq, lq, \text{pendingSignals}, \\
&\quad \quad \text{inPurgatory}, \text{waiting}, \\
&\quad \quad \text{lastSignal}, \text{entryOrder}, \text{wcid}, \\
&\quad \quad \text{scid}, \text{cclwcid}, \text{cclscid} \rangle \\
\\
\text{cclib_enter_check_priority}(self) &\triangleq \wedge pc[self] = \text{"cclib_enter_check_priority"} \\
&\quad \wedge \text{IF } \text{pendingSignals} > 0 \vee \text{inPurgatory} > 0 \\
&\quad \quad \text{THEN } \wedge \text{inPurgatory}' = \text{inPurgatory} + 1 \\
&\quad \quad \quad \wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"cclib_enter_to_purgatory"}] \\
&\quad \quad \text{ELSE } \wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"cclib_enter_end"}] \\
&\quad \quad \quad \wedge \text{UNCHANGED } \text{inPurgatory} \\
&\quad \wedge \text{UNCHANGED } \langle cq, lq, \\
&\quad \quad \text{pendingSignals}, \\
&\quad \quad \text{waiting}, \text{lastSignal}, \\
&\quad \quad \text{entryOrder}, \text{wcid}, \text{scid}, \text{cclwcid}, \text{cclscid} \rangle
\end{aligned}$$

entryOrder, stack,
wcid, scid, cclwcid,
cclscid)

$$\begin{aligned}
cclib_enter_to_purgatory(self) \triangleq & \wedge pc[self] = \text{"cclib_enter_to_purgatory"} \\
& \wedge \wedge stack' = [stack \text{ EXCEPT } ![self] = \langle [procedure \mapsto \text{"java_await"}, \\
& \qquad \qquad \qquad pc \qquad \mapsto \text{"cclib_enter_toheaven"}, \\
& \qquad \qquad \qquad wcid \qquad \mapsto wcid[self]] \rangle \\
& \qquad \qquad \qquad \circ stack[self]] \\
& \wedge wcid' = [wcid \text{ EXCEPT } ![self] = 0] \\
& \wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"java_await_begin"}] \\
& \wedge \text{UNCHANGED } \langle cq, lq, pendingSignals, \\
& \qquad \qquad \qquad inPurgatory, waiting, \\
& \qquad \qquad \qquad lastSignal, entryOrder, \\
& \qquad \qquad \qquad scid, cclwcid, cclscid \rangle
\end{aligned}$$

$$\begin{aligned}
cclib_enter_toheaven(self) \triangleq & \wedge pc[self] = \text{"cclib_enter_toheaven"} \\
& \wedge inPurgatory' = inPurgatory - 1 \\
& \wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"cclib_enter_end"}] \\
& \wedge \text{UNCHANGED } \langle cq, lq, pendingSignals, \\
& \qquad \qquad \qquad waiting, lastSignal, \\
& \qquad \qquad \qquad entryOrder, stack, wcid, \\
& \qquad \qquad \qquad scid, cclwcid, cclscid \rangle
\end{aligned}$$

$$\begin{aligned}
cclib_enter_end(self) \triangleq & \wedge pc[self] = \text{"cclib_enter_end"} \\
& \wedge pc' = [pc \text{ EXCEPT } ![self] = Head(stack[self]).pc] \\
& \wedge stack' = [stack \text{ EXCEPT } ![self] = Tail(stack[self])] \\
& \wedge \text{UNCHANGED } \langle cq, lq, pendingSignals, \\
& \qquad \qquad \qquad inPurgatory, waiting, \\
& \qquad \qquad \qquad lastSignal, entryOrder, wcid, \\
& \qquad \qquad \qquad scid, cclwcid, cclscid \rangle
\end{aligned}$$

$$\begin{aligned}
cclib_enter(self) \triangleq & cclib_enter_begin(self) \\
& \vee cclib_enter_check_priority(self) \\
& \vee cclib_enter_to_purgatory(self) \\
& \vee cclib_enter_toheaven(self) \\
& \vee cclib_enter_end(self)
\end{aligned}$$

$$\begin{aligned}
cclib_leave_begin(self) \triangleq & \wedge pc[self] = \text{"cclib_leave_begin"} \\
& \wedge \text{IF } pendingSignals = 0 \wedge inPurgatory > 0 \\
& \quad \text{THEN } \wedge \wedge scid' = [scid \text{ EXCEPT } ![self] = 0] \\
& \qquad \wedge stack' = [stack \text{ EXCEPT } ![self] = \langle [procedure \mapsto \text{"java_signal"}, \\
& \qquad \qquad \qquad pc \qquad \mapsto \text{"cclib_leave_resetlastsignal"}, \\
& \qquad \qquad \qquad scid \qquad \mapsto scid[self]] \rangle \\
& \qquad \qquad \qquad \circ stack[self]] \\
& \quad \wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"java_signal_begin"}] \\
& \quad \text{ELSE } \wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"cclib_leave_resetlastsignal"}]
\end{aligned}$$

$$\begin{aligned}
& \wedge \text{UNCHANGED } \langle \text{stack}, \text{scid} \rangle \\
& \wedge \text{UNCHANGED } \langle \text{cq}, \text{lq}, \text{pendingSignals}, \\
& \quad \text{inPurgatory}, \text{waiting}, \\
& \quad \text{lastSignal}, \text{entryOrder}, \text{wcid}, \\
& \quad \text{cclwcid}, \text{cclscid} \rangle \\
\text{cclib_leave_resetlastsignal}(\text{self}) \triangleq & \wedge \text{pc}[\text{self}] = \text{"cclib_leave_resetlastsignal"} \\
& \wedge \text{IF } \text{lastSignal}[1] \neq \text{self} \\
& \quad \text{THEN } \wedge \text{lastSignal}' = \langle 0, 0 \rangle \\
& \quad \text{ELSE } \wedge \text{TRUE} \\
& \quad \wedge \text{UNCHANGED } \text{lastSignal} \\
& \wedge \text{pc}' = [\text{pc EXCEPT } ![\text{self}] = \text{"cclib_leave_unlocking"}] \\
& \wedge \text{UNCHANGED } \langle \text{cq}, \text{lq}, \\
& \quad \text{pendingSignals}, \\
& \quad \text{inPurgatory}, \\
& \quad \text{waiting}, \text{entryOrder}, \\
& \quad \text{stack}, \text{wcid}, \text{scid}, \\
& \quad \text{cclwcid}, \text{cclscid} \rangle \\
\text{cclib_leave_unlocking}(\text{self}) \triangleq & \wedge \text{pc}[\text{self}] = \text{"cclib_leave_unlocking"} \\
& \wedge \text{stack}' = [\text{stack EXCEPT } ![\text{self}] = \langle [\text{procedure} \mapsto \text{"java_unlock"}, \\
& \quad \text{pc} \mapsto \text{"cclib_leave_end"}] \rangle \\
& \quad \circ \text{stack}[\text{self}]] \\
& \wedge \text{pc}' = [\text{pc EXCEPT } ![\text{self}] = \text{"java_unlock_begin"}] \\
& \wedge \text{UNCHANGED } \langle \text{cq}, \text{lq}, \text{pendingSignals}, \\
& \quad \text{inPurgatory}, \text{waiting}, \\
& \quad \text{lastSignal}, \text{entryOrder}, \\
& \quad \text{wcid}, \text{scid}, \text{cclwcid}, \\
& \quad \text{cclscid} \rangle \\
\text{cclib_leave_end}(\text{self}) \triangleq & \wedge \text{pc}[\text{self}] = \text{"cclib_leave_end"} \\
& \wedge \text{pc}' = [\text{pc EXCEPT } ![\text{self}] = \text{Head}(\text{stack}[\text{self}]).\text{pc}] \\
& \wedge \text{stack}' = [\text{stack EXCEPT } ![\text{self}] = \text{Tail}(\text{stack}[\text{self}])] \\
& \wedge \text{UNCHANGED } \langle \text{cq}, \text{lq}, \text{pendingSignals}, \\
& \quad \text{inPurgatory}, \text{waiting}, \\
& \quad \text{lastSignal}, \text{entryOrder}, \text{wcid}, \\
& \quad \text{scid}, \text{cclwcid}, \text{cclscid} \rangle \\
\text{cclib_leave}(\text{self}) \triangleq & \text{cclib_leave_begin}(\text{self}) \\
& \vee \text{cclib_leave_resetlastsignal}(\text{self}) \\
& \vee \text{cclib_leave_unlocking}(\text{self}) \\
& \vee \text{cclib_leave_end}(\text{self}) \\
\text{cclib_await_begin}(\text{self}) \triangleq & \wedge \text{pc}[\text{self}] = \text{"cclib_await_begin"} \\
& \wedge \text{waiting}' = [\text{waiting EXCEPT } ![\text{cclwcid}[\text{self}]] = \text{waiting}[\text{cclwcid}[\text{self}]] + 1] \\
& \wedge \text{IF } \text{pendingSignals} = 0 \wedge \text{inPurgatory} > 0 \\
& \quad \text{THEN } \wedge \wedge \text{scid}' = [\text{scid EXCEPT } ![\text{self}] = 0]
\end{aligned}$$

$$\begin{aligned}
& \wedge stack' = [stack \text{ EXCEPT } ![self] = \langle [procedure \mapsto \text{"java_signal"}, \\
& \quad pc \quad \mapsto \text{"cclib_await_actualawait"}, \\
& \quad scid \quad \mapsto scid[self]], \\
& \quad \circ stack[self] \rangle \\
& \wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"java_signal_begin"}] \\
\text{ELSE } & \wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"cclib_await_actualawait"}] \\
& \wedge \text{UNCHANGED } \langle stack, scid \rangle \\
& \wedge \text{UNCHANGED } \langle cq, lq, pendingSignals, \\
& \quad inPurgatory, lastSignal, \\
& \quad entryOrder, wcid, cclwcid, \\
& \quad cclscid \rangle \\
cclib_await_actualawait(self) & \triangleq \wedge pc[self] = \text{"cclib_await_actualawait"} \\
& \wedge \wedge stack' = [stack \text{ EXCEPT } ![self] = \langle [procedure \mapsto \text{"java_await"}, \\
& \quad pc \quad \mapsto \text{"cclib_await_reentering"}, \\
& \quad wcid \quad \mapsto wcid[self]], \\
& \quad \circ stack[self] \rangle \\
& \wedge wcid' = [wcid \text{ EXCEPT } ![self] = cclwcid[self]] \\
& \wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"java_await_begin"}] \\
& \wedge \text{UNCHANGED } \langle cq, lq, pendingSignals, \\
& \quad inPurgatory, waiting, \\
& \quad lastSignal, entryOrder, \\
& \quad scid, cclwcid, cclscid \rangle \\
cclib_await_reentering(self) & \triangleq \wedge pc[self] = \text{"cclib_await_reentering"} \\
& \wedge pendingSignals' = pendingSignals - 1 \\
& \wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"cclib_await_end"}] \\
& \wedge \text{UNCHANGED } \langle cq, lq, inPurgatory, \\
& \quad waiting, lastSignal, \\
& \quad entryOrder, stack, wcid, \\
& \quad scid, cclwcid, cclscid \rangle \\
cclib_await_end(self) & \triangleq \wedge pc[self] = \text{"cclib_await_end"} \\
& \wedge pc' = [pc \text{ EXCEPT } ![self] = Head(stack[self]).pc] \\
& \wedge cclwcid' = [cclwcid \text{ EXCEPT } ![self] = Head(stack[self]).cclwcid] \\
& \wedge stack' = [stack \text{ EXCEPT } ![self] = Tail(stack[self])] \\
& \wedge \text{UNCHANGED } \langle cq, lq, pendingSignals, \\
& \quad inPurgatory, waiting, \\
& \quad lastSignal, entryOrder, wcid, \\
& \quad scid, cclscid \rangle \\
cclib_await(self) & \triangleq cclib_await_begin(self) \\
& \quad \vee cclib_await_actualawait(self) \\
& \quad \vee cclib_await_reentering(self) \\
& \quad \vee cclib_await_end(self) \\
cclib_signal_begin(self) & \triangleq \wedge pc[self] = \text{"cclib_signal_begin"}
\end{aligned}$$

$$\begin{aligned}
& \wedge \text{IF } \text{waiting}[\text{cclscid}[\text{self}]] > 0 \\
& \quad \text{THEN } \wedge \text{pc}' = [\text{pc EXCEPT } ![\text{self}] = \text{"cclib_signal_actualsignal"}] \\
& \quad \text{ELSE } \wedge \text{pc}' = [\text{pc EXCEPT } ![\text{self}] = \text{"cclib_signal_end"}] \\
& \wedge \text{UNCHANGED } \langle \text{cq}, \text{lq}, \text{pendingSignals}, \\
& \quad \text{inPurgatory}, \text{waiting}, \\
& \quad \text{lastSignal}, \text{entryOrder}, \\
& \quad \text{stack}, \text{wcid}, \text{scid}, \text{cclwcid}, \\
& \quad \text{cclscid} \rangle
\end{aligned}$$

$$\begin{aligned}
\text{cclib_signal_actualsignal}(\text{self}) \triangleq & \wedge \text{pc}[\text{self}] = \text{"cclib_signal_actualsignal"} \\
& \wedge \text{pendingSignals}' = \text{pendingSignals} + 1 \\
& \wedge \text{waiting}' = [\text{waiting EXCEPT } ![\text{cclscid}[\text{self}]] = \text{waiting}[\text{cclscid}[\text{self}]] - 1] \\
& \wedge \text{pc}' = [\text{pc EXCEPT } ![\text{self}] = \text{"cclib_signal_setlastsignal"}] \\
& \wedge \text{UNCHANGED } \langle \text{cq}, \text{lq}, \text{inPurgatory}, \\
& \quad \text{lastSignal}, \\
& \quad \text{entryOrder}, \text{stack}, \\
& \quad \text{wcid}, \text{scid}, \text{cclwcid}, \\
& \quad \text{cclscid} \rangle
\end{aligned}$$

$$\begin{aligned}
\text{cclib_signal_setlastsignal}(\text{self}) \triangleq & \wedge \text{pc}[\text{self}] = \text{"cclib_signal_setlastsignal"} \\
& \wedge \text{lastSignal}' = \langle \text{self}, \text{Head}(\text{cq}[\text{cclscid}[\text{self}]]) \rangle \\
& \wedge \wedge \text{scid}' = [\text{scid EXCEPT } ![\text{self}] = \text{cclscid}[\text{self}]] \\
& \quad \wedge \text{stack}' = [\text{stack EXCEPT } ![\text{self}] = \langle [\text{procedure} \mapsto \text{"java_signal"}, \\
& \quad \quad \quad \text{pc} \quad \quad \mapsto \text{"cclib_signal_end"} \\
& \quad \quad \quad \text{scid} \quad \quad \mapsto \text{scid}[\text{self}] \rangle \\
& \quad \quad \quad \circ \text{stack}[\text{self}]] \\
& \wedge \text{pc}' = [\text{pc EXCEPT } ![\text{self}] = \text{"java_signal_begin"}] \\
& \wedge \text{UNCHANGED } \langle \text{cq}, \text{lq}, \\
& \quad \text{pendingSignals}, \\
& \quad \text{inPurgatory}, \text{waiting}, \\
& \quad \text{entryOrder}, \text{wcid}, \\
& \quad \text{cclwcid}, \text{cclscid} \rangle
\end{aligned}$$

$$\begin{aligned}
\text{cclib_signal_end}(\text{self}) \triangleq & \wedge \text{pc}[\text{self}] = \text{"cclib_signal_end"} \\
& \wedge \text{pc}' = [\text{pc EXCEPT } ![\text{self}] = \text{Head}(\text{stack}[\text{self}]).\text{pc}] \\
& \wedge \text{cclscid}' = [\text{cclscid EXCEPT } ![\text{self}] = \text{Head}(\text{stack}[\text{self}]).\text{cclscid}] \\
& \wedge \text{stack}' = [\text{stack EXCEPT } ![\text{self}] = \text{Tail}(\text{stack}[\text{self}])] \\
& \wedge \text{UNCHANGED } \langle \text{cq}, \text{lq}, \text{pendingSignals}, \\
& \quad \text{inPurgatory}, \text{waiting}, \\
& \quad \text{lastSignal}, \text{entryOrder}, \text{wcid}, \\
& \quad \text{scid}, \text{cclwcid} \rangle
\end{aligned}$$

$$\begin{aligned}
\text{cclib_signal}(\text{self}) \triangleq & \text{cclib_signal_begin}(\text{self}) \\
& \vee \text{cclib_signal_actualsignal}(\text{self}) \\
& \vee \text{cclib_signal_setlastsignal}(\text{self}) \\
& \vee \text{cclib_signal_end}(\text{self})
\end{aligned}$$

$$\begin{aligned}
proc_enter(self) &\triangleq \wedge pc[self] = \text{"proc_enter"} \\
&\wedge stack' = [stack \text{ EXCEPT } ![self] = \langle [procedure \mapsto \text{"cclib_enter"}, \\
&\quad pc \quad \mapsto \text{"proc_in"} \rangle \\
&\quad \circ stack[self]] \\
&\wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"cclib_enter_begin"}] \\
&\wedge \text{UNCHANGED } \langle cq, lq, pendingSignals, inPurgatory, \\
&\quad waiting, lastSignal, entryOrder, \\
&\quad wcid, scid, cclwcid, cclscid \rangle \\
\\
proc_in(self) &\triangleq \wedge pc[self] = \text{"proc_in"} \\
&\wedge entryOrder' = Append(entryOrder, self) \\
&\wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"proc_await"}] \\
&\wedge \text{UNCHANGED } \langle cq, lq, pendingSignals, inPurgatory, \\
&\quad waiting, lastSignal, stack, wcid, scid, \\
&\quad cclwcid, cclscid \rangle \\
\\
proc_await(self) &\triangleq \wedge pc[self] = \text{"proc_await"} \\
&\wedge \exists cid \in CId : \\
&\quad \text{IF } cid \neq 0 \\
&\quad \quad \text{THEN } \wedge \wedge cclwcid' = [cclwcid \text{ EXCEPT } ![self] = cid] \\
&\quad \quad \wedge stack' = [stack \text{ EXCEPT } ![self] = \langle [procedure \mapsto \text{"cclib_await"}, \\
&\quad \quad pc \quad \mapsto \text{"proc_signal"}, \\
&\quad \quad cclwcid \mapsto cclwcid[self]] \rangle \\
&\quad \quad \quad \circ stack[self]] \\
&\quad \quad \wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"cclib_await_begin"}] \\
&\quad \quad \text{ELSE } \wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"proc_signal"}] \\
&\quad \quad \wedge \text{UNCHANGED } \langle stack, cclwcid \rangle \\
&\wedge \text{UNCHANGED } \langle cq, lq, pendingSignals, inPurgatory, \\
&\quad waiting, lastSignal, entryOrder, \\
&\quad wcid, scid, cclscid \rangle \\
\\
proc_signal(self) &\triangleq \wedge pc[self] = \text{"proc_signal"} \\
&\wedge \exists cid \in CId : \\
&\quad \text{IF } cid \neq 0 \\
&\quad \quad \text{THEN } \wedge \wedge cclscid' = [cclscid \text{ EXCEPT } ![self] = cid] \\
&\quad \quad \wedge stack' = [stack \text{ EXCEPT } ![self] = \langle [procedure \mapsto \text{"cclib_signal"}, \\
&\quad \quad pc \quad \mapsto \text{"proc_leave"}, \\
&\quad \quad cclscid \mapsto cclscid[self]] \rangle \\
&\quad \quad \quad \circ stack[self]] \\
&\quad \quad \wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"cclib_signal_begin"}] \\
&\quad \quad \text{ELSE } \wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"proc_leave"}] \\
&\quad \quad \wedge \text{UNCHANGED } \langle stack, cclscid \rangle \\
&\wedge \text{UNCHANGED } \langle cq, lq, pendingSignals, inPurgatory, \\
&\quad waiting, lastSignal, entryOrder, \\
&\quad wcid, scid, cclwcid \rangle \\
\\
proc_leave(self) &\triangleq \wedge pc[self] = \text{"proc_leave"}
\end{aligned}$$

$$\begin{aligned}
& \wedge stack' = [stack \text{ EXCEPT } ![self] = \langle [procedure \mapsto \text{"cclib_leave"}, \\
& \quad \quad \quad pc \quad \quad \quad \mapsto \text{"Done"}] \rangle \\
& \quad \quad \quad \circ stack[self]] \\
& \wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"cclib_leave_begin"}] \\
& \wedge \text{UNCHANGED } \langle cq, lq, pendingSignals, inPurgatory, \\
& \quad \quad \quad waiting, lastSignal, entryOrder, \\
& \quad \quad \quad wcid, scid, cclwcid, cclscid \rangle
\end{aligned}$$

$$\begin{aligned}
procid(self) &\triangleq proc_enter(self) \vee proc_in(self) \\
&\vee proc_await(self) \vee proc_signal(self) \\
&\vee proc_leave(self)
\end{aligned}$$

$$\begin{aligned}
Next &\triangleq (\exists self \in ProcSet : \vee java_lock(self) \vee java_unlock(self) \\
&\quad \vee java_await(self) \\
&\quad \vee java_signal(self) \\
&\quad \vee cclib_enter(self) \\
&\quad \vee cclib_leave(self) \\
&\quad \vee cclib_await(self) \\
&\quad \vee cclib_signal(self)) \\
&\vee (\exists self \in ProcId : procid(self)) \\
&\vee \text{Disjunct to prevent deadlock on termination} \\
&\quad ((\forall self \in ProcSet : pc[self] = \text{"Done"}) \wedge \text{UNCHANGED } vars)
\end{aligned}$$

$$\begin{aligned}
Spec &\triangleq \wedge Init \wedge \square [Next]_{vars} \\
&\wedge WF_{vars}(Next)
\end{aligned}$$

$$Termination \triangleq \diamond (\forall self \in ProcSet : pc[self] = \text{"Done"})$$

END TRANSLATION



Invariants and temporal properties (theorems)

$$\begin{aligned}
Types &\triangleq \\
&\text{Global variables} \\
&\wedge cq \in [CId \rightarrow Seq(ProcSet)] \\
&\wedge lq \in Seq(ProcSet) \\
&\wedge pendingSignals \in Nat \\
&\wedge inPurgatory \in Nat \\
&\wedge waiting \in [CId \rightarrow Nat] \\
&\wedge lastSignal \in (ProcSet \cup \{0\}) \times (ProcSet \cup \{0\}) \\
&\text{Procedure } java_await \\
&\wedge wcid \in [ProcSet \rightarrow CId \cup \{defaultInitValue\}] \\
&\text{Procedure } java_signal \\
&\wedge scid \in [ProcSet \rightarrow CId \cup \{defaultInitValue\}] \\
&\text{Procedure } cclib_await \\
&\wedge cclwcid \in [ProcSet \rightarrow CId \cup \{defaultInitValue\}]
\end{aligned}$$

Procedure *cclib_signal*
 $\wedge cclscid \in [ProcSet \rightarrow CId \cup \{defaultInitValue\}]$

THEOREM *TypePreservation* $\triangleq \square Types$

0. – *Mutex*

Mutex \triangleq

$\forall pid1 \in ProcSet :$

$\forall pid2 \in ProcSet :$

$pc[pid1] = \text{"proc_in"} \wedge pc[pid2] = \text{"proc_in"} \Rightarrow pid1 = pid2$

1. – Characteristic property: *La* que hemos chequeado con *UPPAAL*, que si se hace un signal sobre una condition no vacía, el siguiente thread en tomar acceso del monitor es el primero de dicha condition, y no alguno de los que estaban en la cola del lock mutex.

As an invariant with the instrumentation:

CharacteristicAsInv \triangleq

$\forall pid \in ProcSet : pc[pid] = \text{"proc_in"} \Rightarrow lastSignal[2] \in \{0, pid\}$

2. – *ToHeaven*: se acaba saliendo del purgatorio

InPurgatory(pid) $\triangleq Member(cq[0], pid)$

ToHeaven $\triangleq \forall pid \in ProcSet : (InPurgatory(pid) \rightsquigarrow pc[pid] = \text{"proc_in"})$

BY Auto

THEOREM *ToHeavenThm* $\triangleq InPurgatory(1) \rightsquigarrow pc[1] = \text{"proc_in"}$

(1)1. ASSUME $InPurgatory(1) \wedge [Next]_{vars} \Rightarrow (InPurgatory(1)' \vee pc'[1] = \text{"proc_in"})$,

$InPurgatory(1) \wedge \langle Next \rangle_{vars} \Rightarrow pc'[1] = \text{"proc_in"}$,

$InPurgatory(1) \Rightarrow ENABLED \langle Next \rangle_{vars}$

PROVE $\square [Next]_{vars} \wedge WF_{vars}(Next) \Rightarrow (InPurgatory(1) \rightsquigarrow pc[1] = \text{"proc_in"})$

BY *RuleWF1*

(1).QED

3. – *OrderPreservation*:

NoOrderPreservation \triangleq

$\forall pid1 \in ProcSet :$

$\forall pid2 \in ProcSet :$

$(\wedge Member(cq[0], pid1)$

$\wedge pc[pid2] = \text{"java_lock_blocked"})$

\rightsquigarrow

$(\wedge Member(entryOrder, pid1)$

$\wedge Member(entryOrder, pid2)$

$\wedge \text{Index}(\text{entryOrder}, \text{pid2}) < \text{Index}(\text{entryOrder}, \text{pid1})$

$\text{OrderPreservation} \stackrel{\Delta}{=} \neg \text{NoOrderPreservation}$

