

# Communicating Cooperatively Scheduled Processes

On the Unlikelihood of Implementing a Pure CSP Channel

KEVIN CHALMERS, Ravensbourne University London, United Kingdom

JAN BÆKGAARD PEDERSEN, University of Nevada Las Vegas, United States of America

This paper presents the modelling of the ProcessJ cooperative runtime environment for the implementation of a CSP-inspired communication channel. We used the CSP tool FDR to verify the correctness of the ProcessJ runtime and primitives, demonstrating how we have overcome a limitation in the existing ProcessJ runtime to improve behaviour. However, our work has demonstrated a limitation when trying to claim a cooperatively-scheduled channel implementation meets the abstract specification of a CSP channel. Our conclusion is that without sufficient hardware to execute all processes at once, a channel implementation cannot fully meet its specification when considering the execution environment in which the channel must operate. We are assured that the ProcessJ channel works correctly, such as other implementations including JCSP and CSO, but we are not assured that modelling can use a pure CSP channel in place of a ProcessJ channel or other implementation when undertaking modelling due to unintended prioritisation occurring when not enough resources are available for full concurrency. Our work has demonstrated the need to consider the execution environment as it may cause behavioural issues not detected when only modelling a system in the abstract.

CCS Concepts: • **Theory of computation** → **Program verification**; • **Software and its engineering** → **Runtime environments**.

Additional Key Words and Phrases: CSP, process-orientation, ProcessJ, formal verification, cooperative scheduler, verified runtime

## ACM Reference Format:

Kevin Chalmers and Jan Bækgaard Pedersen. 2025. Communicating Cooperatively Scheduled Processes: On the Unlikelihood of Implementing a Pure CSP Channel. 1, 1 (October 2025), 45 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 1 Introduction

Ensuring correct concurrent behaviour poses a significant challenge. Consider a synchronising channel that passes messages between processes. The conventional method of verifying a channel implementation often overlooks the hardware and execution context, such as schedulers, in which a concurrent system operates. This comes from an implicit assumption of the availability of resources or scheduling—eventually a waiting process will be serviced. In Hoare’s Communicating Sequential Processes (CSP) [Hoa78, Hoa85, Ros98, Ros10], when an event is ready, it can happen, assuming the external environment will be ready and available for code execution when the process is ready. This is an oversimplification of how code is executed on a real computer. The execution resources/environment **must** be taken into account thus also included in the formal verification of a channel-based runtime.

---

Authors’ Contact Information: [Kevin Chalmers](#), Ravensbourne University London, School of Computing, Architecture, and Emerging Technologies, London, United Kingdom, [k.chalmers@rave.ac.uk](mailto:k.chalmers@rave.ac.uk); [Jan Bækgaard Pedersen](#), University of Nevada Las Vegas, Department of Computer Science, Las Vegas, NV, United States of America, [matt.pedersen@unlv.edu](mailto:matt.pedersen@unlv.edu).

---

Permission to make digital or hard copies of all or part 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 bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM XXXX-XXXX/2025/10-ART

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

This paper extends our previously presented work [PC23]. In that paper, we followed the standard approach to verification using CSP—we modelled our synchronisation primitives for our runtime and language (ProcessJ) and demonstrated that they met their specifications. However, part of our goal was to understand what would happen when the cooperative scheduler of ProcessJ was made multithreaded. We needed to understand the execution environment that our primitives existed within. We came to the conclusion that not understanding and modelling the execution environment was a significant limitation in channel-based runtime verification, potentially leading to undiscovered issues. Indeed, the ProcessJ scheduler had issues—although stemming from a simpler implementation of the original scheduling system. To solve the issues, we were left with two options:

- (1) Change the implementation of the synchronisation primitives to work in the execution environment.
- (2) Change the execution environment so it worked correctly with the synchronisation primitives.

In this paper we focus on the second option as we were confident our synchronisation primitives were correct (we have verified them using the standard approach; see Section 2.3), but importantly we can update the ProcessJ scheduler to change the execution environment, an option not always open to implementers of concurrent runtimes.

## 1.1 ProcessJ

ProcessJ is a process-oriented programming language being developed at the University of Nevada Las Vegas in collaboration with the University of Roehampton. ProcessJ is a language that supports concurrency through parallel processes and synchronous non-shared channel communication. The syntax of ProcessJ closely resembles that of Java with extra keywords added to support process-orientation. The ProcessJ compiler produces Java source code which subsequently is compiled with the Java compiler; the resulting class files are then instrumented to support cooperative scheduling. The runtime and scheduler are written in Java, and the whole system ultimately runs on the JVM.

ProcessJ uses *processes* to specify the behaviour of its concurrently-executing components. A process is a series of steps that are undertaken sequentially. These steps can include the execution of other processes, possibly concurrently. Here, it is important to realise that a ProcessJ process is **not** an operating system process or thread. Though the code is executed by the CPU, the (operating system) thread that runs the ProcessJ code is what we later refer to as a *scheduler*. A ProcessJ process is an object on the JVM heap that contains the state of the process and has a *run()* method that contains the code to be run. The JVM thread that runs the code (the scheduler) simply calls this *run()* method every time it needs to run the ProcessJ process. This leads to jumping back to where the ProcessJ process yielded the last time it was run. The execution continues until a *yield* point is reached, at which time the *run()* method terminates and the scheduler is free to call the *run()* method of another ProcessJ process. All local state that would normally be on the operating system stack is now stored inside the object as fields.

To communicate between processes, ProcessJ provides synchronous CSP-like channels. Both writer/sender and receiver/reader must be present for a message exchange to happen. If the writer/sender arrives first, it must wait for the receiver/reader and vice-versa.

This style of concurrency requires three fundamental functions to specify system behaviour:

- The ability to execute two or more processes *concurrently* (*in parallel*).
- The ability to communicate between processes via a *channel*.
- The ability to *select* between two or more channel communications to define behaviour (*choice*).

Below is an example of a small ProcessJ program that runs two processes concurrently using the **par** keyword<sup>1</sup>. The two processes communicate an integer value 42 (passed in as a parameter) across a channel using the keywords **read** and **write**:

```

99
100
101
102 void reader(chan<int>.read in) {
103     int v;
104     v = in.read();           // read a value from the channel
105     println(v);
106 }
107
108 void writer(chan<int>.write out, int val) {
109     out.write(val);         // write the value val on the channel
110 }
111
112 void main(string args[]) {
113     chan<int> c;             // declare a channel carrying ints
114     par {                   // run in parallel the following two processes:
115         reader(c.read);     // pass reading end of channel c to reader proc
116         writer(c.write, 42); // pass writing end of channel c and the value 42 to writer proc
117     }
118 }
119

```

The ProcessJ type `chan<int>` denotes a *channel type* that carries integer values. Furthermore, `chan<int>.read` and `chan<int>.write` denote the *channel end types* of a channel carrying integer values. **read** refers to the reading end of such a channel. **write** refers to the writing end.

The ProcessJ scheduler is a cooperative scheduler that operates with a single run queue of processes. In its simplest form, it looks like this:

```

125 Queue<Process> runQueue;
126 ...
127 // enqueue 1 or more processes to run ...
128 while (!runQueue.isEmpty()) {
129     Process p = runQueue.dequeue();
130     if (p.ready())
131         p.run();
132     if (!p.terminated())
133         runQueue.enqueue(p);
134 }
135

```

Note, what we refer to as a *scheduler* (as shown above) is perhaps better understood as a runner. This is because there is no scheduling policy enforced, but processes are simply removed in order from the queue and run by having their `run()` method invoked. There may be multiple runners (schedulers) all sharing the same process queue.

In cooperative scheduling, active processes must *yield* control rather than have it *preemptively* taken by a scheduler. The ProcessJ scheduler is a thread that runs application code. It consists of a single loop that repeatedly dequeues a process, and if the process is ready, runs it. Once the process yields control back to the scheduler, or if the process is not ready to run, it is checked for termination, and if not terminated the process is placed at the back of the run queue. An important

<sup>1</sup>The process with the **par** block will itself block until all processes in the block have terminated.

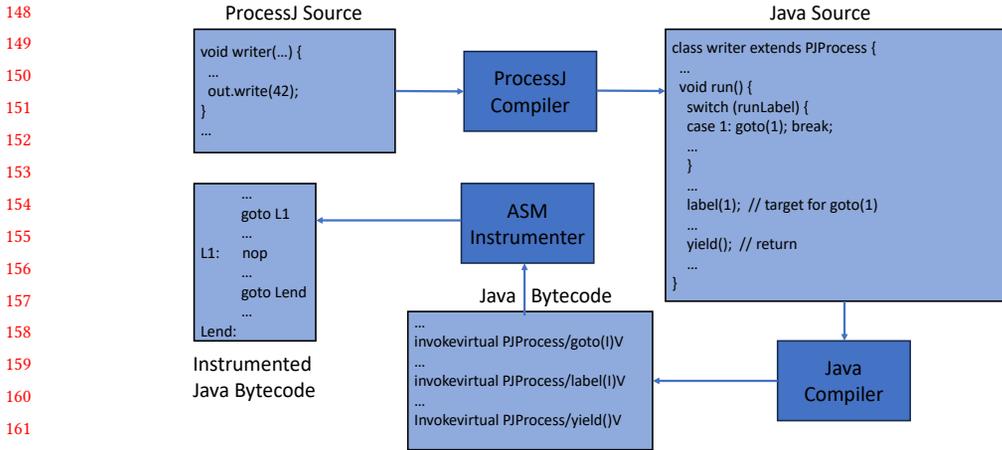


Fig. 1. The ProcessJ Toolchain.

observation about the way this single queue system works is that the run queue contains both ready and not ready processes. This is the scheduler we used in [PC23].

1.1.1 *The ProcessJ Tool Chain.* We touch upon the scheduling part of ProcessJ Processes in Section 2.2 and show the code generated for the ProcessJ example shown earlier in this section, but it is worth understanding the ProcessJ tool chain.

Figure 1 illustrates how ProcessJ source code becomes runnable Java bytecode:

- We start with a snippet of ProcessJ source code.
- The ProcessJ source is read as input by the ProcessJ compiler. The code generation module of the ProcessJ compiler produces Java source code. This Java code contains a number of “dummy” methods (*goto*, *label*, and *yield*) all declared in the class *PJProcess*. Their implementation is just an empty block. That is, if they are ever called nothing would happen, but the Java program correctly compiles to bytecode.
- Next, the Java compiler is invoked and it produces Java bytecode. Note, the calls to the dummy methods are still in this code.
- Finally, we replace the dummy method calls with appropriate bytecode:
  - *label(x)* becomes a nop instruction, but the address of this instruction is noted.
  - *goto(x)* becomes *goto Lx* where *Lx* is the the noted address of the nop instruction from the previous item.
  - *yield()* becomes *goto Lend*, where *Lend* is the address of a nop instruction at the very end of the code.
- The resulting Java bytecode file is now compatible with a cooperative scheduler that simply calls the *run()* method of the object representing the process. The process has code (the switch statement) at the top to correctly jump to where execution left off the last time it was run, and yield points have been added (by the code generation) to the appropriate places in the code.

More information about the code generation as well as the instrumentation can be found in [Cis18].

## 1.2 Limitations of the Cooperative Scheduler

Let us consider what consequences we may have once we extend our system to run multiple schedulers concurrently. First, observe that with a single scheduler, if no processes in the run queue are ready to run, a deadlock has occurred. Progress is impossible as there are no ready-to-run processes to drive the system forward by changing the *ready* flags of any of the not-ready processes in the run queue. This is a deadlock from the perspective of the application being run, as no processes will ever again become ready to run. From a verification point of view, this is exhibited as a livelock: non-ready processes keep getting removed from and put back onto the run queue. However, if there is at least one process ready to run, it will eventually move to the front of the queue and be executed, potentially altering the readiness status of other processes.

When multiple schedulers are involved, there is a scenario where a scheduler might keep removing non-ready processes from the shared queue endlessly. For instance, in a scenario where a reader and a writer are in continuous communication, it is conceivable (and can be shown in a simulated system trace) that one of the schedulers will manage both the reader and the writer alternately, while the other scheduler consistently deals with dequeuing and enqueueing non-ready processes. A possible system trace illustrating this situation is as follows:

```
<dequeue.P1, write, yield.P1, enqueue.P1, dequeue.P2, dequeue.P1, enqueue.P1, dequeue.P1, ...>
```

where the identifiers *P1* and *P2* represent the process identifiers to be executed. Until *P2* makes progress (it has only been dequeued), dequeue and enqueue of *P1* will continue as *P1* is not ready. To better illustrate this, let us split this trace across the two active scheduler threads:

```
T1 = dequeue.P1 write yield.P1 enqueue.P1 dequeue.P2 ... .. .  
T2 = ... .. . dequeue.P1 enqueue.P1 dequeue.P1 ...
```

The issue becomes more apparent in a scenario where there are more schedulers than processes in the system. In such a case, a likely sequence of events could involve a continuous cycle of dequeuing and enqueueing operations. This straightforward setup of a cooperative scheduler was demonstrated in our earlier work [PC23], leading us to the following finding: “Changing how processes are added and removed from the queue, and managing their state, should enable us to remove the spinning (livelock), but careful consideration is required to ensure that verification can still occur.”

In this paper we make these changes to the scheduler and the runtime system; we then perform the same verifications (i.e., using the same formal CSP models as the specification) to determine if we indeed managed to solve the problems discovered in [PC23]. The changes we propose in this paper include the following:

- Remove non-ready processes from the run queue. They do not need to be held in the queue as they are now held in the channel on which they are blocked.
- Make the run queue blocking (i.e., the dequeue operation will be blocking if the queue is empty).
- Add a *running* field to the process meta-data (similar to the *ready* field).
- Add a shared scheduling interface (scheduler manager).

The details about the implementation of this new scheduler are presented in Section 4. In [PC23] we started with an implementation of the runtime system of ProcessJ and then proceeded to create CSP models of the implementation. These models were then verified against specifications (from Welch and Martin [WM00]). In this paper we have no ProcessJ implementation, but start with a CSP model of the proposed design, and verify against the same Welch and Martin specification as we used previously. Once we know the design is correct, we can turn to changing the Java implementation of the scheduler and the runtime of ProcessJ.

In [PW18] a *Concurrency and Verification Workflow* diagram and approach are presented showing the typical workflow between the problem domain and a verified working implementation. In particular, the duality of a formal specification and an actual implementation is pointed out. In [PC23] we started with an implementation and developed the formal model to determine if the implementation was correct. In this paper we work in the opposite direction: here, we start with a formal model and once it has been verified, the implementation can commence.

### 1.3 Contributions

This research makes three key contributions. Firstly, this work has developed a verified cooperatively scheduled runtime for our process-oriented language, ProcessJ. This completes a verification journey that has always demonstrated that our channel implementation was correct, but we have now demonstrated that correctness when run in a multiprocessor environment. There are limitations to our verification due to the resources made available to the runtime, but this in itself is a further contribution.

Our second key contribution is the identification of limitations in CSP-inspired channel implementations. Although our context is within cooperative scheduling—where we have formally modelled the limitation—we also provide empirical evidence and discussion of this limitation for the preemptively scheduled JCSP library. Our argument is that without enough resources (i.e., hardware processors) to enable all active processes to make progress in parallel, a channel concurrency primitive mapping to CSP is still a challenge. The limitation comes from unintended prioritisation at the end of a channel communication, which can lead to potential starvation issues that would propagate through a system.

The final key contribution is our models working with the execution environment that our concurrency runtime operates within. Without our consideration of the execution environment, we would not have discovered flaws in the current implementation of the ProcessJ scheduling system which would have occurred when transitioning to a multithreaded version. We believe our approach to modelling the execution environment is unique and enables us to question implementations in operation. As we have verified our implementation (with limitations) in this context, we are more assured of its correct execution.

### 1.4 Definitions

To avoid confusion when reading the paper, we provide this definition section that explains the meaning attributed to a number of commonly used concepts that may be slightly different in the context of this paper:

- A **process**: In ProcessJ, a process is an entity that can be executed by a scheduler/runner (see the next bullet point). For ProcessJ, a process and all its state is an object that contains the state in fields and the code to be executed as a single method called *run()*. A process is a series of sequential steps. An individual step can be a concurrent execution of multiple other processes.
- **Yielding**: When a process decides to yield, its *run()* method is terminated. Yielding means giving control back to the scheduler, which will find a different process/object that is ready to run and call its *run()* method. As all state is kept in fields of the object there is no need to save any stack state. A process's *run()* method contains code that will jump to the latest yield point when called again. This is done through instrumentation of the bytecode.
- A **scheduler**: The typical definition of a scheduler is a module in the operating system that is in charge of selecting the next job that is to be run by the hardware. In ProcessJ (and in this paper) a scheduler is more of a *runner*; that is, a scheduler will interact with the

run queue and more importantly, it is the scheduler that **runs** the code. Thus, if we want multiple execution units, we add more schedulers. ProcessJ has no other scheduling policy than a FIFO queue of ready processes.

- A **thread** is typically defined as a sub-execution unit of a process. For almost all purposes, whether we refer to a thread or a process (from an operating system point-of-view) there is no difference. The technical differences do not matter—all that matters is that they are both independent execution units.

## 1.5 Breakdown of the Paper

In Section 2 we consider cooperative scheduling in general, and we discuss cooperative scheduling in the ProcessJ runtime in particular. In addition, we describe the verification results that we have achieved so far. Section 3 introduces CSP and FDR. In Section 4 we present the CSP for the new scheduler and the changes made to the runtime system. In Section 5 we present the specifications for both channel and alternation that we will be using in the verification phase of the paper. In Sections 6 and 7 we combine the new scheduler with the channel and choice implementations from [PC23] and redo the verification. In Section 8 we consider a number of issues concerned with scheduling. Section 9 summarises the results, and, finally, Section 10 wraps up with future work and possible extensions.

The CSP models presented in this paper can be found at <https://github.com/kevin-chalmers/processj-csp/tree/main/non-spinning-scheduler>.

## 2 Background and Related Work

In this section we start with some background information about cooperative scheduling and then consider the results from [PC23].

### 2.1 Cooperative Scheduling and Multitasking

To support multi-processing, a runtime typically takes one of two scheduling approaches:

- *preemptive scheduling*, where processes are allocated processor time-slices managed centrally (e.g., via the operating system).
- *cooperative scheduling*, where processes decide when to give up (*yield*) the processor.

If all processes have the same priority, preemptive scheduling guarantees a ready process will eventually be allocated processor time. Cooperative scheduling does not guarantee processor time as a process may use resources indefinitely—it is for the process to give up the processor. When it does give up the processor, a method is required to determine the next process to service. Typically, and in the case of ProcessJ, a queue of processes is maintained.

On a single processor system, a simple approach to such cooperative scheduling can be used—the next (ready) process is continuously dequeued from the queue and executed until the queue is empty. In a multiprocessor system, consideration of cross-core synchronisation is required. Often, as in the case of ProcessJ, we rely on preemptive locking mechanisms provided by the operating system to ensure correct concurrent behaviour. We therefore have to account for the cooperative scheduling (via yield points) within the preemptive scheduling (via locking) when designing a system. Thus, for ProcessJ runtime verification, we **must** incorporate the implementation of the cooperative scheduling environment.

Cooperative scheduling offers some advantages over preemptive scheduling. With cooperative scheduling, as we use in ProcessJ, processes are stackless (remember, a ProcessJ process is not a thread, but an object with data that is maintained to enable scheduling). Process local data is stored as object data. Therefore, a process is incredibly lightweight. Previous ProcessJ publications [SP16]

344 have demonstrated hundreds of millions of processes in operation. This compares to preemptive  
 345 thread-based libraries, such as JCSP, managing only a few thousand processes. The limitation in  
 346 such libraries is often due to large stack sizes (often greater than a megabyte) that each thread  
 347 utilises as a stack. ProcessJ can support at least four orders of magnitude greater system scale.

348 The lightweight processes of ProcessJ also promote fast context switching [SP16]. As no pre-  
 349 emptive thread swapping is required, context switches, and thus channel communication time, is  
 350 fast. A thread provided by the operating system, such as in standard Java and used by JCSP, must  
 351 save a register record and load a new one onto the CPU with each context switch. If this switch is  
 352 cross-core, the impact is greater due to the stack swapping between caches.

353 Cooperative scheduling currently has limited popularity due to the programmer typically having  
 354 to manage scheduling. Operating systems favour preemptive scheduling (although cooperative  
 355 scheduling was dominant prior to modern 32-bit operating systems). Threading libraries in pro-  
 356 gramming languages typically use operating system provided primitives. There are some exceptions  
 357 (e.g., the Win32 fiber library), but preemptive scheduling is the *de facto* approach to threading. No  
 358 mainstream language directly supports cooperative scheduling as part of its standard library.

359 We believe that hiding cooperative scheduling (via communication), as in ProcessJ, alleviates  
 360 some of the challenges in managing such scheduling. Indeed, until Go 1.10, the scheduler was a  
 361 mix of cooperative and preemptive. The current Go scheduler is preemptive, but Go is a compiled  
 362 language that runs directly on hardware. As the aim of ProcessJ is to run within the JVM for  
 363 portability, we are limited in how fast and lightweight processes are supported. *occam- $\pi$*  also  
 364 featured a cooperative scheduler [RSB09] which was likewise hidden from the programmer.

365 In contrast, cooperative *multitasking* has a growing popularity, primarily from the *async/await*  
 366 pattern. Coroutines [MI09] have gained popularity in languages ranging from JavaScript and  
 367 Python to C++ and Rust. Indeed, this has led to the discussion of new concurrency paradigms such  
 368 as structured concurrency [CY23] which have influenced thinking in Kotlin and Swift. However,  
 369 coroutines are not scheduled as in the case of a cooperative runtime as used in ProcessJ.

370 The ProcessJ scheduler is fairly unique. Other lightweight concurrency languages, such as Go  
 371 and Erlang, typically use preemptive scheduling [CR22, CT09]. Indeed, the only similar runtime  
 372 approach—hiding cooperation in synchronisation—we have found is in Stackless Python [Tis00].  
 373

## 374 2.2 Cooperative Scheduling in ProcessJ

375 In cooperative scheduling, a process is responsible for yielding control of computational resources. In  
 376 the Java code generated from a ProcessJ program, this is achieved at various points by inserting *yield*  
 377 points in connection with synchronisation (e.g., communication, choice, and other synchronisation  
 378 primitives like barriers and timers). Yielding is achieved by a jump to the end of the run method  
 379 code. These jumps are created during an instrumentation phase (see Section 1.1.1) of the class files  
 380 after the Java compiler has run. In addition, code is inserted at the start of all *run()* methods to  
 381 correctly jump back to where the process last yielded when it is run again. A *yield* also sets the  
 382 next run label—the location where the execution resumes. It may also set the process not ready by  
 383 setting the *ready* field of the process to false.  
 384

385 A ProcessJ procedure is compiled to a Java class. The code of the ProcessJ procedure is produced  
 386 in the *run()* method of the class. In order for all the locals and parameters of a ProcessJ procedure  
 387 to survive between invocations, they are turned into fields in the generated class. Thus, a ProcessJ  
 388 process is stackless to the point that when a process yields there is no need for maintaining a stack  
 389 as all the data needed by a process is kept as fields in an object on the heap.

390 The result of the compilation of the ProcessJ *reader* procedure from the introduction is the  
 391 following Java class (which shows the general shape of the code produced for a ProcessJ procedure):  
 392

```

393 class reader extends PjProcess {
394     // params and locals are turned into fields here
395     protected PjOne2OneChannel in;
396     protected int v;
397
398     // constructor to set the fields representing the params.
399     public reader(PjOne2OneChannel in) {
400         this.in = in;
401     }
402
403     public void run() {
404         // generated code to jump to the right place in the code
405         switch (this.runLabel) {
406             case 0: break;
407             case 1: resume(1); break;           // resume(1) becomes goto L1
408             case 2: resume(2); break;       // resume(2) becomes goto L2
409             default: break;
410         }
411
412         // translated ProcessJ code with yields
413         if (!in.isReadyToRead(this)) {           // if a channel is not ready to be read
414             this.runLabel = 1;                   // resume at label(1) (L1) next time.
415             this.yield();                         // no writer is ready yet, so yield
416         }
417         label(1);                                 // L1 is the address of this invocation
418         v = in.read(this);                       // perform the actual read
419         this.runLabel = 2;                       // resume at label(2) (L2) next time.
420         this.yield();                             // courtesy yield
421
422         label(2);                                 // L2 is the address of this invocation
423
424         // Here goes the code for the println call.
425     }
426 }

```

Note, the *isReadyToRead()* method sets the process's *ready* field to false, if the channel does not have a value ready to communicate (i.e., if there is no writer present). Also, the fields *runLabel* and *ready* are found in the super class *PjProcess*.

The *label(x)* method invocations are removed but their addresses are recorded and associated with the number *x*—indicated by the labels  $L_i$  in the code above. A *resume(y)* method invocation is replaced by a *goto* instruction that jumps to the address associated with the corresponding *label(y)* invocation.

### 2.3 Related Work

In order to appreciate the approach we have taken in this paper to verify a *scheduled* system, it is important to understand the *standard* approach to verification of channel-based systems using CSP and FDR. In the standard approach, no runtime system, execution system, hardware etc.

442 is modelled. When modelling such systems with CSP, the assumption is that when an event is  
 443 ready that resources will eventually become available. However, this assumption also includes the  
 444 situation where the resource is available at exactly the point where the event has become ready.  
 445 This paper in effect is controlling for this condition of resource availability. As we have a model of  
 446 our scheduler, we can explicitly see the impact when resources are not available at a given time,  
 447 testing the assumption of eventual resource availability under scheduling conditions.

448 The standard assumption means that events in the formal model of such a system can possibly  
 449 happen whenever the system is ready to perform the event. No attention is paid to external  
 450 resources or the execution environment—the verification is in the abstract rather than in a concrete  
 451 environment. However, we argue that such an approach does not model what actually happens  
 452 when a channel-based system is brought into service. In essence, the standard way of modelling  
 453 such systems assumes that a ready event will remain available until scheduled. However, this  
 454 also assumes the extreme case—that an event may immediately occur, implying that resources  
 455 are available to execute that particular event when it immediately becomes available. Consider a  
 456 process that is at the end of a queue of ready processes. It must wait to be scheduled—its events  
 457 will not be immediately available. By modelling the execution environment as we do in this and  
 458 previous papers, we are much closer to the behaviour of a real running implementation of a channel.  
 459 In previous work, as the extreme case (no scheduling considerations taken) was not precluded,  
 460 then certain behaviours were met by implementations that enabled meeting of pure channel  
 461 specifications. Our work in ProcessJ is to consider the limitations when precluding the extreme  
 462 case through the use of scheduling to better understand the behaviour of channel primitives in  
 463 implementation.

464 Mutual-exclusion is long-solved (e.g., see [Lam87]). Model checking concurrent algorithms often  
 465 begins with the assumption that multi-processor behaviour is typically managed via a preemptive  
 466 scheduler. Preemptive scheduling is not a panacea for ensuring events happen in a fair manner or  
 467 even that they will happen in a way enabling all potential behaviours (for example, see [Reg02]).  
 468 Atomic operations introduce new challenges for mutual exclusion, but again, the preemptive  
 469 scheduler enables certain assumptions (e.g., see [MS95]). The assumption of preemptive scheduling  
 470 has been successfully applied to model check synchronous channel communication [WM00, Low19]  
 471 in JCSP [WBM<sup>+</sup>10] and Communicating Scala Objects (CSO) [Suf08]. JCSP and CSO both rely on  
 472 preemptive scheduling provided via the JVM (which in reality is normally the operating system).  
 473 The ProcessJ scheduler more closely emulates the *occam- $\pi$*  multi-processor cooperative scheduler  
 474 defined by Ritson et. al. [RSB09].

475 In all the research we have encountered in the literature where the authors set out to prove  
 476 correctness of a channel primitive, we find little consideration of the *execution environment's*  
 477 behaviours. For example, a known issue in Java concurrent code is spurious wakeups [Ora]. A  
 478 spurious wakeup happens when a thread waiting is woken up (by the operating system) but the  
 479 condition that sent the thread to sleep is still true. That is, the thread was woken up but the  
 480 condition has not changed. We will not go into why spurious wakeups happen, but note that the  
 481 standard trick to solve this problem is replacing the if statement that sent the thread to sleep with  
 482 a while statement that checks the same condition:

```
483   if ( condition )
484       wait();
```

485 becomes

```
486   while ( condition )
487       wait();
```

488

489

490

491 However, from a modelling perspective, the code will be divergent as—unlikely as it may be—the  
 492 while loop may loop forever. Spurious wakeups feature both in the core Linux pThread library as  
 493 well as in Windows threads.

494 JCSP now has code to mitigate the issue of spurious wakeups, but this has never been formally  
 495 modelled. Therefore, the JCSP verification is limited in its scope. No code runs itself, it is always  
 496 executed within an execution environment: threads, processes, schedulers, etc.

## 498 2.4 Verification So Far

499 In [PC19] we took the first step on our verification journey. We started out approaching verification  
 500 of channel communication for ProcessJ much in the same way as Welch and Martin [WM00]. We  
 501 succeeded in verifying correctness of channel communication. Like Welch and Martin, we applied  
 502 the *standard approach* of formal verification: we translated our implementation into CSP and then  
 503 used the specification for channel communication developed by Welch and Martin.

504 In [PC23] we added formal verification of choice using a slightly optimised version of a speci-  
 505 fication, again, from Welch and Martin [WM00], and obtained the same results. However, more  
 506 profoundly, we realised that none of the formal verification exercises that have been published so  
 507 far took into account the execution environment in which processes run. That is, the processes  
 508 themselves were modelled but not the environment in which they execute. This environment could  
 509 include runtime systems, schedulers, and hardware. Since ProcessJ is cooperatively scheduled, mod-  
 510 elling the interaction and behaviour for the scheduler becomes very important. It *is the scheduler*  
 511 that runs a process’s code. If we do not take into account the behaviour of the scheduler(s) and  
 512 their interaction with the processes, then how could we ever know for sure that the verification is  
 513 correct? We redid our analyses, but with the scheduler’s implementation taken into account. These  
 514 results were promising, but we did not quite manage to replicate the results when considering the  
 515 channel communication or the choice **without** the scheduler. The major problem was divergence  
 516 in the scheduler. This divergence was caused by the run queue containing processes that were  
 517 not ready to be run. This, in turn, could cause one of more schedulers to continuously engage in  
 518 enqueue and dequeue events when a not-ready process was removed from a queue that contained  
 519 no processes that were ready to run (other schedulers would be off running code independently).

520 The solution was to develop a better scheduler (modify the execution environment) to one that  
 521 does not cause divergence due to the not-ready processes in the run queue. This paper is the  
 522 description of such an approach and the results on the verification of the channel communication  
 523 as well as choice using CSP specification models developed in [PC23] but with new implementation  
 524 models.

## 526 3 Communicating Sequential Processes

527 Communicating Sequential Processes (CSP) [Hoa78, Hoa85, Ros98, Ros10] is a process algebra that  
 528 allows the specification of a concurrent system via *processes* and *events*. A process is an abstract  
 529 component defined by the events in which it wishes to engage.

530 Events are *atomic*, *synchronous*, and *instantaneous*. An event cannot be divided, and it causes all  
 531 synchronising processes to wait until the event occurs. The occurrence of the event is immediate  
 532 when executed by the system. The CSP model defines which events processes are willing to  
 533 synchronise upon at different stages of their execution.

### 536 3.1 Simple Processes

537 The simplest processes are STOP, which engages in no events and will not terminate, and SKIP,  
 538 which can engage in the *termination event*  $\checkmark$  and then terminate.

## 3.2 Prefixing

Events are added to a process using the *prefix* operator ( $\rightarrow$ ). For example, the process  $P = x \rightarrow \text{STOP}$  engages in event  $x$ , then stops. A process definition must end by executing another process definition; the general form is  $\text{Process} = \text{event} \rightarrow \text{Process}'$ . Processes can be recursive (e.g.,  $P = x \rightarrow P$ ).

## 3.3 Choice

CSP has *choice* operators. The three most frequently used choice types are *external (or deterministic) choice*, *internal (or non-deterministic) choice*, and *prefix choice*. A process can exhibit alternate behaviour when engaging in a choice.

Given two processes  $P$  and  $Q$ , the definition  $P \sqcap Q$  is a process that will behave as either  $P$  or  $Q$ , depending on the first event to be offered by the environment; it is called an *external choice*. For example, the process:

$$P = (a \rightarrow Q) \sqcap (b \rightarrow R)$$

is willing to accept  $a$ , then behave as  $Q$ , or accept  $b$ , then behave as  $R$ . If both  $a$  and  $b$  are available, then the system may choose either.

An internal choice is represented by  $\sqcap$ . A process  $P \sqcap Q$  can behave as either  $P$  or  $Q$  without considering the external environment. The choice is internal to the process.

Both external choice and internal choice can operate across a set of events. If  $E = \{e_1, \dots, e_n\}$  is a set of events, then we can write  $\sqcap a \leftarrow E \bullet a \rightarrow P$ , which expands to  $e_1 \rightarrow P \sqcap e_2 \rightarrow P \sqcap \dots \sqcap e_n \rightarrow P$ , and  $\sqcap a \leftarrow E \bullet a \rightarrow P$ , which expands to  $e_1 \rightarrow P \sqcap e_2 \rightarrow P \sqcap \dots \sqcap e_n \rightarrow P$  (where  $a$  does not appear as an event in  $P$ ).

With prefix choice we can define an event and a parameter. If we define a set of events as  $\{c.v \mid v \in \text{Values}\}$ , we can consider  $c$  as a channel willing to communicate a value  $v$ . We can then define input and output operations ( $?$  and  $!$ , respectively) to allow communication and binding of variables. This is simply a shorthand due to the following identities:

$$c!v \rightarrow P \equiv c.v \rightarrow P$$

$$c?x \rightarrow P \equiv \sqcap x \leftarrow \text{Values} \bullet c.x \rightarrow P$$

where  $\text{Values}$  is a set.

In addition, CSP supports *if* statements. *if* follows the standard functional model with each branch having to end in another process definition. It is therefore common to write:

if ( $b$ ) then

$P$

else

$Q$

The semantics of  $c?v \rightarrow (\text{if } (v == x) \text{ then } P \text{ else } Q)$  in CSP is equivalent to  $(c.x \rightarrow P) \sqcap (\sqcap v \leftarrow X \setminus \{x\} \bullet c.v \rightarrow Q)$ <sup>2</sup>.

**3.3.1 Guards.** It is possible to control whether an alternation branch is enabled or disabled (i.e., whether the branch will even be considered). This is done by placing a Boolean expression (a guard) and a  $\&$  before the branch. For example:

$$e_1 \& c?x \rightarrow \dots \sqcap e_2 \& d?x \rightarrow \dots$$

If the Boolean expression  $e_1$  is true and the Boolean expression  $e_2$  is false, only the  $c?x$  branch will be considered. Similar, if only  $e_2$  is true, only  $d?x$  will be considered; only if both  $e_1$  and  $e_2$  are true will both branches be considered in the alternation.

<sup>2</sup>Assuming that  $v$  does not appear in  $P$ .

### 3.4 Process Composition

Processes can be combined via *parallel* and *sequential* composition. There are three forms of parallel composition.

$P[A \parallel] Q$  defines two processes with a *synchronisation set*  $A$ . For the events in  $A$ , they can only occur if offered by both  $P$  and  $Q$  at the same time (i.e., synchronously). Any event not in  $A$  will not synchronise between  $P$  and  $Q$ . For example,  $(a \rightarrow b \rightarrow P)[\{b\}](b \rightarrow c \rightarrow Q)$  has two processes synchronising on  $b$ . Thus,  $a$  must happen, before both processes execute  $b$  and then  $c$  is performed. If the synchronisation set also contained  $c$ , then the right-hand side would not be able to perform this event unless the left-hand side also offered it; that is,  $P$  must offer  $c$ .

$P[A \parallel B] Q$  defines two processes that can only perform events in their respective *alphabets*;  $P$  with alphabet  $A$  and  $Q$  with alphabet  $B$ .  $P$  and  $Q$  must also synchronise on any events in the set  $A \cap B$ . For example,  $(a \rightarrow b \rightarrow P)[\{a, b\} \parallel \{b, c\}](b \rightarrow c \rightarrow Q)$  has two processes synchronising on the intersection alphabet  $\{b\}$ . Thus,  $a$  must happen, before both processes execute  $b$  and then  $c$  is performed. If the alphabet on the left-hand side did not contain  $a$ , then system would deadlock as  $a$  could not be performed.

Two processes can also *interleave*:  $P \parallel\parallel Q$ . This means that  $P$  and  $Q$  execute in parallel but *do not* synchronise on shared events. Returning to our previous example,  $(a \rightarrow b \rightarrow P) \parallel\parallel (b \rightarrow c \rightarrow Q)$  may accept  $a$  or  $b$  first. The event  $b$  is only ever performed by one process at a time.

Sequential composition is denoted as  $P ; Q$ . This means after  $P$  has terminated,  $Q$  is performed. Sequential composition operates on processes, not events. Therefore, it allows the definition of behaviour as a sequential composition of discrete process definitions.

### 3.5 Traces and Hiding

The set of all externally observed event sequences of a process is referred to as the *trace set*. For example, for a process  $P = a \rightarrow P$  the shortest trace observable is the empty trace  $\langle \rangle$ . If and when  $P$  engages on the event  $a$  for the first time we can observe this external event and we now have the trace  $\langle a \rangle$ . The second time  $P$  engages on  $a$  we can observe the trace  $\langle a, a \rangle$ , and we say that the traces of  $P$  are  $traces(P) = \{\langle \rangle, \langle a \rangle, \langle a, a \rangle, \dots\}$ . For a process  $Q = a \rightarrow b \rightarrow Q$  has a trace-set  $\{\langle \rangle, \langle a \rangle, \langle a, b \rangle, \langle a, b, a \rangle, \dots\}$ .

Observable behaviour of a process can be concealed via the hiding operator  $\backslash$ . Hidden events are replaced by  $\tau$  and are ignored when comparing observable traces. For example,  $(a \rightarrow b \rightarrow a \rightarrow SKIP) \backslash \{a\}$  has traces  $\{\langle \rangle, \langle b \rangle, \langle b, \checkmark \rangle\}$  as  $\tau$ s are not observable, and  $\checkmark$  represents termination. Similarly,  $traces(P \backslash \{a\}) = \{\langle \rangle\}$  and  $traces(Q \backslash \{a\}) = \{\langle \rangle, \langle b \rangle, \langle b, b \rangle, \dots\}$ .

### 3.6 Models

CSP defines three models of behaviour that can be used to analyse systems: the *traces* model, the *failures* model, and the *failures/divergence* model. Each model subsequently analyses further behaviours of a system to establish refinement.

**3.6.1 The Traces Model.** The **traces model** comes from the externally observed behaviour of a system. If the set of traces of a system *implementation*  $Q$  is a subset of the set of traces of a system *specification*  $P$  (i.e.,  $traces(Q) \subseteq traces(P)$ ), we state that  $P \sqsubseteq_T Q$  ( $Q$  *trace refines*  $P$ ).

A different way of thinking about this is as follows: in a refinement test *Specification*  $\sqsubseteq_T$  *Implementation*, *Specification* can be thought of as a specification that defines what traces are allowable; therefore, the refinement test asks whether *Implementation* has only allowable traces.

Hiding events allows us to analyse the external behaviour of a process, thus allowing assertions such as  $P \sqsubseteq_T Q \wedge Q \sqsubseteq_T P$ , where  $P$  is a simple specification and  $Q$  a complex implementation. An implementation may contain events that are not part of the specification. These are events

required to model the implementation of the system rather than just a specification of behaviour. We therefore hide internal events of an implementation to enable refinement checking.

**3.6.2 The Stable Failures Model.** The **stable failures model** [Ros98, Ros10] deals with events that a process may refuse to engage in after a sequence of events occur. A stable state is one where a process  $P$  cannot make internal progress (including no  $\tau$  events) and must engage externally. A refusal is a set of events a process cannot engage with when in a stable state.

The stable failures model overcomes the limitation of comparing traces. For example,  $(P = a \rightarrow P) \sqsubseteq_T ((P = a \rightarrow P) \sqcap \text{STOP})$ , the right-hand definition may non-deterministically refuse to accept any events. A *failure* is a pair  $(s, X)$  where  $s$  is a trace of a process  $P$ , and  $X$  is the set of events that can be refused after  $P$  has performed the trace  $s$ . Stating that  $P \sqsubseteq_F Q$  means that whenever  $Q$  refuses to perform a set of events,  $P$  can do likewise. More formally,  $P \sqsubseteq_F Q \Leftrightarrow \text{failures}(Q) \subseteq \text{failures}(P)$ . Therefore, a specification in the stable failures model defines which failures an implementation is allowed to have.

**3.6.3 The Failures-Divergences Model.** The final model is the **failures-divergences model**. Divergences deal with potential livelock scenarios where a process can continually perform internal events (denoted by  $\tau$ ) and not progress in its externally observed behaviour. For example, consider the following two processes:

$$P = a \rightarrow \text{STOP}$$

$$Q = (a \rightarrow \text{STOP}) \sqcap \text{DIV}$$

where DIV is a process that immediately diverges. Although  $P \sqsubseteq_T Q$  and  $P \sqsubseteq_F Q$  (and vice-versa),  $Q$  can continuously not accept  $a$  and just perform  $\tau$ . The final refinement check  $\sqsubseteq_{FD}$  allows processes to be compared in this circumstance. Formally, we consider both the failures and the divergences of a process  $P$  as a pair:  $(\text{failures}_\perp(P), \text{divergences}(P))$  and now define refinement in the failures/divergence model as follows:  $P \sqsubseteq_{FD} Q \Leftrightarrow \text{failures}_\perp(Q) \subseteq \text{failures}_\perp(P) \wedge \text{divergence}(Q) \subseteq \text{divergences}(P)$ .  $\text{failures}_\perp(P)$  is defined as  $\text{failures}(P) \cup \{(s, X) \mid s \in \text{divergences}(P)\}$ . That is, the set of traces leading to a divergent set are added to the existing set of stable failures to create an extended failures set.

It is worth noting that, if both the specification and implementation are divergence free, then analysis only needs to demonstrate equivalence in the stable failures model. Indeed, for our verification purposes, both the specification and implementation are indeed divergence free. See Section 5.

### 3.7 FDR

FDR [GRABR14] is a refinement checker that can check various assertions about CSP processes written in  $\text{CSP}_M$  (a machine readable version of CSP). FDR is used to analyse CSP models of systems and to test them against specifications (also written in CSP). FDR allows processes to be refinement checked via the three models shown in the previous section.

FDR can also check a system for determinism and divergence freedom. In the stable failures model, FDR “considers a process  $P$  to be deterministic providing no witness to non-determinism exists, where a witness consists of a trace  $tr$  and an event  $a$  such that  $P$  can both accept and refuse  $a$  after  $tr$ .” [FDR]. In the failures-divergences model, the process must also be divergent free.

### 3.8 Modelling State

All the ProcessJ runtime classes have member state variables which are accessible (read and write) by other classes. It is therefore necessary to model such shared state variables in CSP. In this section we present CSP that can be used to model general state which can be *set* and *read*. A state variable can be modelled in CSP via a process. The process maintains the current value of the state variable, providing events to either *load* the current value of the variable, or *store* a new value of the variable.

687 We can define a generic process to model a state variable (*VARIABLE*) with alphabet  $\alpha$ *VARIABLE* as  
 688 follows:

689 datatype *Operations* = *load* | *store*

690 *VARIABLE*(*var*, *val*) =

691 *var.load!val* → *VARIABLE*(*var*, *val*) – *load* the value

692 □

693 *var.store?val* → *VARIABLE*(*var*, *val*) – *store* the value

694  
 695  $\alpha$ *VARIABLE*(*var*, *T*) = { *var.o.v* | *o* ← *Operations*, *v* ← *T* }

696  
 697 *VARIABLE* takes two parameters: *var* is the channel used to communicate between the variable  
 698 and the environment. *val* is the current value of the variable. In the alphabet  $\alpha$ *VARIABLE*, *val* has  
 699 type *T*, which is provided for the given variable by the system specifier.

700 For example, each ProcessJ process has a Boolean *ready* flag (stored as a field in an object), and  
 701 we can define a channel for communicating with a variable representing such a flag as follows:

702 channel *ready* : *Processes.Operations.Bool*

703 The channel is called *ready* and communicates a triple: a process identifier (*pid*), an operation (*load*  
 704 or *store*), and a Boolean value (*true* or *false*). We then define a *VARIABLE* process for each *p* in  
 705 *Processes*.

## 706 4 Modelling a Non-Spinning Scheduler

707  
 708 In our previous work [PC23], we modelled what would happen when moving from a single-threaded  
 709 scheduler to a multi-threaded scheduler, determining the problems that this would introduce while  
 710 ensuring the system itself did not deadlock. The key issue faced was how the process queue was  
 711 managed. In our previous work, processes were immediately put back into the process queue and  
 712 the scheduling processes would check each process in the queue, looking for a ready one. This  
 713 introduced divergence—the scheduling processes would loop continuously searching for a ready  
 714 process. To solve this problem, we add an additional flag to the process meta-data—*running*—that  
 715 allows distinguishing between processes that are running and processes that are ready via the  
 716 *ready* flag. We also introduce scheduling and descheduling algorithms to ensure correct behaviour.

717 Furthermore, in our previous work, our models used a fixed schedule of process identifiers used  
 718 to initialise the run queue. The new model lets each process schedule itself by sending its process  
 719 identifier to the scheduling manager, which will add it to the run queue.

### 720 4.1 Scheduler Architecture

721  
 722 We are now ready to develop the CSP for the new *implementation*. We start with the CSP for the  
 723 new scheduler. The overall scheduler architecture can be seen in Figure 2. We will describe each of  
 724 these components in turn over the following subsections. In Figure 2 we use the following alphabets  
 725 and synchronisation sets:

- 726 •  $\alpha$ *RUNQUEUE* is the natural alphabet of the *QUEUE* process (see Section 4.1.1).
- 727 •  $\alpha$ *PROCESSES* is the natural alphabet of the *VARIABLE* processes used for the the processes  
 728 meta-data (see Section 4.1.3).
- 729 •  $\alpha$ *SCHEDULING* is the difference of the union of the natural alphabets of the *SCHEDULER*  
 730 processes and the *SCHEDULE\_MANAGER* (see Section 4.1.2) and the natural alphabet of  
 731 the *QUEUE*.

732  
 733 <sup>3</sup>Note, the *PROCESSES* part represents the *meta-data* for the processes needed by the scheduler, not the actual application  
 734 processes. They must be run in parallel to the scheduler system.

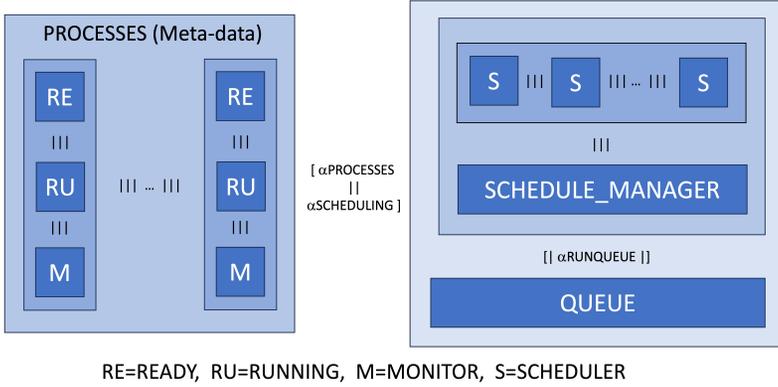


Fig. 2. Scheduler Architecture Diagram ( $N\_SCHEDULER\_SYSTEM$ ).<sup>3</sup>

4.1.1 *Queue*. The scheduler requires a queue of ready to run processes. This can be modelled in CSP as follows:

```

754  QUEUE(enqueue, dequeue, q, CAP) =
755    length(q) < CAP & enqueue?v → QUEUE(enqueue, dequeue, q^<v>, CAP)
756    □
757    length(q) > 0 & dequeue!head(q) → QUEUE(enqueue, dequeue, tail(q), CAP)

```

The queue supports enqueueing and dequeueing.  $CAP$  is the maximum capacity of the queue, required for verification. Note the use of *guarded conditions* on each of the choice branches. If the queue is not full it will offer an *enqueue* event, and if it is not empty it will offer a *dequeue* event.

4.1.2 *Scheduler(s) and Schedule Manager*. In Figure 2, schedulers are labelled  $S$ . A scheduler is responsible for running a process by dequeuing a  $pid$  (Process Identifier) from the run queue, sending a *run* event to the relevant process, and waiting for a *yield* event to occur before recursing. Schedulers interleave with a new schedule manager process which accepts *schedule* events—a request to add a  $pid$  to the run queue—then enqueueing the  $pid$  onto the queue before recursing.<sup>4</sup> These two processes can be modelled in CSP as follows:

```

767  SCHEDULER =
768    rqdequeue?p → - dequeue a process
769    run.p → - run it
770    yield.p → - wait for yield
771    SCHEDULER - recurse
772
773  SCHEDULE_MANAGER =
774    schedule?pid → - wait for a schedule event
775    rqenqueue!pid → - place the process in the run queue
776    SCHEDULE_MANAGER - recurse

```

where  $rqenqueue$  and  $rqdequeue$  are channels connected to the queue for enqueueing and dequeueing processes.

4.1.3 *Process*. In our previous work, a process was not defined as a separate entity. The only information about a process that a scheduler needed was its current ready state (represented by a *ready* flag). Our new scheduling algorithms require two additional pieces of meta-data:

<sup>4</sup>Note, only processes that are ready are ever scheduled.

- A *running* flag to determine if a process is currently being executed by a scheduler.
- A monitor to control access to the process data (*ready* flag and *running* flag) to avoid race conditions during the scheduling and descheduling stages.

*PROCESS* is therefore an interleaving of its two variables (*RE* for *ready* and *RU* for *running*) and the monitor (*M*) grouped together in the *PROCESSES* process in Figure 2. The CSP for the meta-data of a single process is as follows:

```

791 PROCESS(p) =
792     VARIABLE(ready.p, true)           - ready = true
793     |||
794     VARIABLE(running.p, false)       - running = false
795     |||
796     MONITOR(claim_process.p, release_process.p)

```

A collection of *PROCESS* processes are defined as follows:

```

798 PROCESSES = ||| p : Processes • PROCESS(p)

```

where *Processes* is the set of *pids*. *PROCESSES* is now all the meta-data the scheduler requires for all the processes in the system (but not the actual application processes).

**4.1.4 Monitor.** A monitor is used to control access to shared data—in this instance the channel and process objects. Our previous work only required a monitor for channels. As we now require a monitor for each process, we have made our monitor model more generic. The CSP for the *MONITOR* is as follows:

```

806 MONITOR(claim, release) =
807     claim?pid →                       - wait for a claim
808     release.pid →                     - wait for a release on the same pid
809     MONITOR(claim, release)           - recurse

```

**4.1.5 Scheduler System.** With everything in place, we can define an *N\_SCHEDULER\_SYSTEM* process representing our new scheduling system to run in parallel with whatever processes we wish to run. The scheduler side of this process is simply the queue (*QUEUE*) run in parallel with *N* schedulers (*SCHEDULER*) and a single *SCHEDULE\_MANAGER*. The schedulers and the schedule manager are interleaved as they do not interact:

```

815 N_SCHEDULER_SYSTEM(N) =
816     (
817         (
818             QUEUE(rqenqueue, rqdequeue, <>, card(Processes))
819             [| αRUNQUEUE |]
820             (||| n : {1..N} • SCHEDULER) ||| SCHEDULE_MANAGER)
821         )
822     [αSCHEDULING || αPROCESSES]
823     PROCESSES
824 ) \ αRUNQUEUE

```

The various alphabets on which the processes synchronise can be found in the CSP code available online.

## 4.2 Scheduler Interaction

To work with the new scheduler, we require new algorithms for scheduling, descheduling, and yielding. These algorithms are used by executing processes and are not directly part of the scheduler architecture. Rather, when an executing process requires to schedule another process, deschedule itself, or yield to the scheduler, it uses these helper definitions to do so.

834 4.2.1 *Schedule*. *SCHEDULE* is a procedure that manages the interaction when an attempt is made  
 835 to add a process to the run queue. In our previous work, to schedule a process it was sufficient to  
 836 set its *ready* field to true. Our new model requires a process to determine if it must set a process  
 837 ready and if so, whether to add it to the run queue. The *SCHEDULE* procedure will handle the  
 838 scheduling of a process. There are three possible states a process can be in when it is scheduled:

- 839 (1) The process is ready, and therefore already exists in the run queue; therefore, do nothing.
- 840 (2) The process is not ready, but is still running (i.e., it is scheduled); therefore, set the process  
 841 ready but do nothing else.
- 842 (3) The process is not ready, and is not running; therefore, set the process ready, and schedule  
 843 it via the schedule manager.

844 The CSP code for the *SCHEDULE* procedure is:

```

845 SCHEDULE(me, pid) =
846   claim_process.pid.me →                               - lock the process object
847   ready.pid.load?r →                                   - load the ready field
848   if(r) then                                           - if (the process is ready) then
849     release_process.pid.me → SKIP                       -   release the lock
850   else                                                 - else
851     ready.pid.store!true →                               -   set the process ready
852     running.pid.load?r2 →                               -   get the running status
853     if(r2) then                                         -   if (the process is running) then
854       release_process.pid.me → SKIP                       -     release the lock
855     else                                               -   else
856       schedule.pid →                                    -     schedule the process
857       release_process.pid.me → SKIP                       -     release the lock

```

858 where *me* is the process identifier of the process calling *SCHEDULE* and *pid* the process identifier  
 859 of the process being scheduled. *claim\_process.pid.me* locks the process object (indicated by *pid*) to  
 860 ensure only one schedule or deschedule procedure is executing at a time. We then check the *ready*  
 861 flag, and if it is true, we release the lock and SKIP (State 1). If the *ready* flag is false, we set it to true  
 862 (State 2 and State 3), and check the *running* flag. If the *running* flag is true, we release the process  
 863 and SKIP (State 2). If not running, we add the process to the run queue via the *schedule* event (to  
 864 the schedule manager), release the process and SKIP (State 3).  
 865

866 4.2.2 *Deschedule*. The *DESCHEDULE* procedure executes prior to a process yielding control back  
 867 to the scheduler process. It must therefore set the *running* flag to false. There are two possibilities  
 868 at this point:

- 869 (1) The process is ready—either because it is courtesy yielding or another process has set it  
 870 ready prior to descheduling.
- 871 (2) The process is not ready.

872 The CSP code for *DESCHEDULE* is as follows:

```

873 DESCHEDULE(pid) =
874   claim_process.pid.pid →                               - lock the process
875   running.pid.store.false →                             - set the running field to false
876   ready.pid.load?r →                                   - get the ready field
877   if(r) then                                           - if (the process is ready) then
878     schedule.pid →                                    -   schedule the process
879     release_process.pid.pid → SKIP                       -   release the lock
880   else                                                 - else
881     release_process.pid.pid → SKIP                       -   release the lock

```

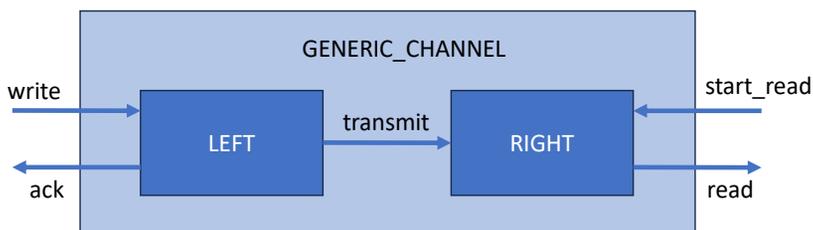


Fig. 3. The *GENERIC\_CHANNEL* specification.

Note, only a process can deschedule itself, therefore *pid* is now the process identifier of both the caller and the process being descheduled. After claiming the process and setting *running* to false, we check the *ready* flag. If *ready* is true, we schedule the process again, release the process lock, and SKIP. Otherwise, we only release the process lock and skip.

**4.2.3 Yield.** Yielding occurs when a process releases the scheduler. That is, when a process is ready to give back control to the scheduler that was running it. Before a process yields, it must call *DESCHEDULE* and, after it is resumed, it must set itself ‘as running’ (set *running* to true; this is *running.pid.set!true* in the code below). In between the *DESCHEDULE* call and the setting itself ready again, a process engages in the actual *yield* event, and then waits for another scheduler to engage it in a *run* event. The CSP code for *YIELD* is as follows:

```

905 YIELD(pid) =
906     DESCHEDULE(pid);           - deschedule the process
907     (
908         yield.pid →           - yield
909         run.pid →             - run the process again
910         running.pid.store!true → SKIP - set the running field to true
911     )

```

## 5 Specifications

To verify our new scheduler implementation, we return to our previous specifications of channel and choice. Our goal is to demonstrate that the ProcessJ channel and choice mechanisms behave as specified when executed in coordination with the scheduler system. Two specifications are required—a *generic channel* and a *generic alternation*.

### 5.1 A Generic Channel Specification

For channels, we use the specification presented in [WM00], because interleaving read and write processes with a simple transmit channel does not provide a specification behaviour to verify against. Welch and Martin’s specification has also been proven to be equivalent (in the *failures-divergence model*) to a CSP channel (See page 289 of [WM00]). Figure 3 shows the *GENERIC\_CHANNEL* specification that we are about to build. First, we define channels to represent a channel object interface:

```

926 channel read, write : Channels.Processes.Values
927 channel start_read, ack : Channels.Processes
928 transmit : Channels.Values

```

The *LEFT* (writing) process accepts a message on the *write* channel, transmits the message to the reader, and performs an *ack* event. This is the equivalent of calling a *write* method on the

932 channel, with *ack* similar to a *return*. The *RIGHT* (reading) process engages in an event on the  
 933 *start\_read* channel before receiving a message via *transmit*. The transmitted message is passed to  
 934 the environment on the *read* channel. This is equivalent to calling a *read* method on the channel,  
 935 with *read* returning the read value to the caller.

```

936 LEFT(pid, chan) = write.chan.pid?mess →           - take a message
937                  transmit.chan!mess →           - send it
938                  ack.chan.pid →                 - acknowledge the send
939                  LEFT(pid, chan)                - recurse
940
941 RIGHT(pid, chan) = start_read.chan.pid →         - Start the read
942                  transmit.chan?mess →         - received the message
943                  read.chan.pid!mess →         - deliver the message
944                  RIGHT(pid, chan)             - recurse
  
```

945 We define a *GENERIC\_CHANNEL* by executing *LEFT* and *RIGHT* in parallel:

```

946 GENERIC_CHANNEL(pid, pid', chan) =
947   (LEFT(pid, chan)[αLEFT(pid, chan) || αRIGHT(pid', chan)]RIGHT(pid', chan))
948   \ {transmit.chan.v | v ← Values}
  
```

949 Note that *write* and *ack* events represent the call and return from a write operation (a channel  
 950 write); similarly, the *start\_read* and *read* events represent the call and return from a reading operation  
 951 (a channel read). Between these pairs the processes synchronise on *transmit*, representing that the  
 952 communication actually happened.

953 We *hide* the *transmit* events as they are internal to the process and not part of the external  
 954 specification behaviour. The *GENERIC\_CHANNEL* process has been verified to be *divergence free*.

955  
 956 *5.1.1 Algebraic Proof.* The generic channel is identical to Welch and Martin's [WM00]. They also  
 957 provide an algebraic proof to demonstrate their specification is equivalent to a regular CSP channel.  
 958 As we use the same specification, we are likewise attempting to prove for ProcessJ. Our previous  
 959 paper [PC23] describes this equivalence in more detail.

## 960 5.2 A Generic Alternation Specification

961 A simple example of **alternation** or *alt* in ProcessJ where we reuse the *writer* procedure from the  
 962 introduction could look like this:

```

964 void selector(chan<int>.read c1, chan<int>.read c2) {
965   int v;
966   alt {
967     v = c1.read() : { System.out.println(v); }
968     v = c2.read() : { System.out.println(v); }
969   }
970 }
971
972 void main(string args[ ]) {
973   chan<int> c1, c2;
974   par {
975     writer(c1.write, 42);
976     writer(c2.write, 43);
977     selector(c1.read, c2.read);
978   }
979 }
980
  
```

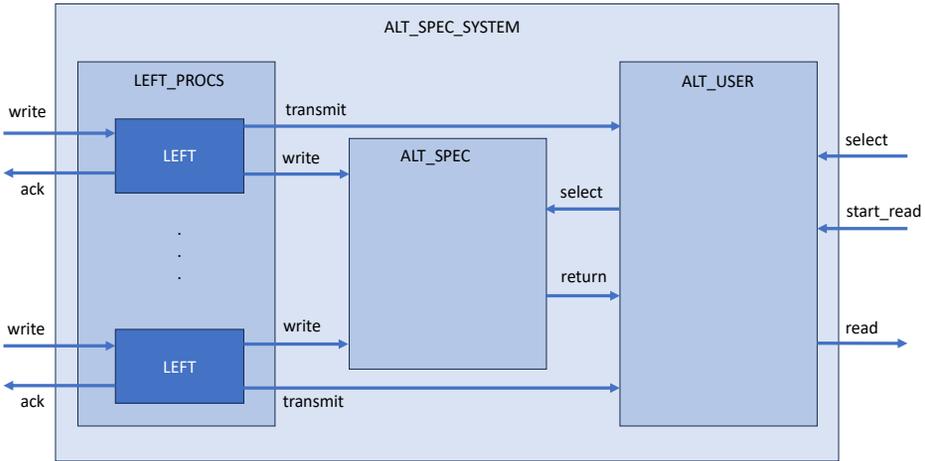


Fig. 4. The *ALT\_SPEC\_SYSTEM* specification.

An alternation (or *alt* for short) chooses one of its *ready* branches and executes it along with the statement that follows. A channel-read (as used here) is ready if there is a committed writer on the channel. The output of running this code is either 42 or 43. In addition, the code will deadlock as only one of the writers will get to go. Only alternation of channel reading on non-shared channels is supported at present in ProcessJ (apart from timer timeouts and skip guard). For this analysis, each branch is of the form  $v = c.read(): statement$ . Though in ProcessJ,  $skip: statement$  (skip guard) and  $t.timeout(v): statement$  (timeout guard) are also allowed, but not considered in this paper.

Figure 4 illustrates the *ALT\_SPEC\_SYSTEM* specification we are about to construct. The specification uses the same channels as *GENERIC\_CHANNEL* (e.g., *write*, *ack*, etc.) and also introduces a *select* event (used to signify the start of a choice being made), and an internal *return* channel to provide the channel identifier picked during the choice process.

To verify correctness of the ProcessJ *alt* (choice), we again turn to Welch and Martin [WM00]. However, contrary to the specification we used in [PC23], this time we have expanded the specification slightly. In particular, we have generalised it such that it works for any number of writers. There are three parts to this specification. The first part we need is an external choice that picks a writer. Generally, for  $n$  channels and  $n$  writing processes, we need to generate the following external choice:

```

write.C1.W1 → ALT_SPEC(...)
□
write.C2.W2 → ALT_SPEC(...)
□
...
□
write.Cn.Wn → ALT_SPEC(...)

```

We need generate this external choice ‘on the fly’, as  $n$  is not fixed. We can make this simple using a *zip* style function (e.g., from Haskell) which we can apply to such a builder function:

```

zip(<>, <>) = <>
zip(< x > ^ xs, < y > ^ ys) = < (x, y) > ^ zip(xs, ys)

```

We also need two channels: one to kick off the *alt* and one to obtain the returned channel chosen. They are as follows:

1030 channel *select* : *Processes*  
 1031 channel *return* : *Channels*

1032

1033 We can generate the first part, which is an external choice taking messages (i.e., writes) as  
 1034 follows:

1035  $\square (c, p) : \text{pairs} \bullet \text{write}.c.p?mess \rightarrow \text{ALT\_SPEC}(\text{selector}, \text{pairs}, \text{union}(\text{writing}, \{c\}), \text{waiting})$

1036 The second requirement (containing the second and third part) for an *alt* is to service the selecting  
 1037 process. There are two choices: either the select has not started and the *alt* must offer the select  
 1038 event to begin the process; or the select has started and the *alt* must offer the currently ready  
 1039 channels. These ready channels are kept in the parameter *writing*. In an external choice of writing  
 1040 processes, we offer up a *return* event:

1041  $\square c : \text{writing} \bullet \text{return}.c \rightarrow \text{ALT\_SPEC}(\text{selector}, \text{pairs}, \text{diff}(\text{writing}, \{c\}), \text{false})$

1042 where any writing process *c* is added to the *writing* set. *waiting* is a flag that will capture whether or  
 1043 not the selection process has started. Putting these two choices together with an event on the *select*  
 1044 channel indicating that the *alt* has begun, we get the following code for the *ALT\_SPEC* process:

1045  $\text{ALT\_SPEC}(\text{selector}, \text{pairs}, \text{writing}, \text{waiting}) =$   
 1046  $(\square (c, p) : \text{pairs} \bullet \text{write}.c.p?mess \rightarrow \text{ALT\_SPEC}(\text{selector}, \text{pairs}, \text{union}(\text{writing}, \{c\}), \text{waiting}))$   
 1047  $\square$   
 1048 if *waiting* then  
 1049  $(\square c : \text{writing} \bullet \text{return}.c \rightarrow \text{ALT\_SPEC}(\text{selector}, \text{pairs}, \text{diff}(\text{writing}, \{c\}), \text{false}))$   
 1050 else  
 1051  $\text{select.selector} \rightarrow \text{ALT\_SPEC}(\text{selector}, \text{pairs}, \text{writing}, \text{true})$

1052 Note, an event on *select* sets the *waiting* flag to true, indicating that the *alt* has started. Additionally,  
 1053 we need a process to use the *alt* and perform the actual read; we call this process *ALT\_USER*, and it  
 1054 is defined as follows:

1055  $\text{ALT\_USER}(pid) =$   
 1056  $\text{select}.pid \rightarrow$  - kick off the alt  
 1057  $\text{return}?idx \rightarrow$  - get a returned index of a ready channel back  
 1058  $\text{start\_read}.idx.pid \rightarrow$  - Initiate the read  
 1059  $\text{transmit}.idx?mess \rightarrow$  - transmit the message internally  
 1060  $\text{read}.idx.pid!mess \rightarrow$  - do the actual read from the chosen channel  
 1061  $\text{ALT\_USER}(pid)$  - recurse

1062 Before we can wire up the complete specification, we need a few helper functions. They are  
 1063 defined as follows:

1064  $\text{PRIVATE\_ALT\_SPEC\_SYSTEM}(pid) = \{\text{transmit}.c.v, \text{return}.c \mid c \leftarrow \text{Channels}, v \leftarrow \text{Values}\}$

1065  $\text{LEFT\_PROCS}(\text{pairs}) =$   
 1066  $\|\| (c, p) : \text{pairs} \bullet \text{LEFT}(p, c)$

1067 where *LEFT\_PROCS* creates *n* interleaved *LEFT* processes (the writers). *PRIVATE\_ALT\_SPEC\_SYSTEM*  
 1068 contains the events that the alt specification will not engage with externally. We can now finally  
 1069 wire up the *ALT\_SPEC\_SYSTEM* specification as follows:

1070  $\text{ALT\_SPEC\_SYSTEM}(\text{selector}, \text{chans}, \text{writing\_procs}) =$   
 1071  $($   
 1072  $($   
 1073  $\text{LEFT\_PROCS}(\text{set}(\text{zip}(\text{chans}, \text{writing\_procs})))$   
 1074  $[\alpha \text{LEFT\_PROCS}(\text{writing\_procs}, \text{chans}) \|\| \alpha \text{ALT\_SPEC}(\text{selector})]$   
 1075  $\text{ALT\_SPEC}(\text{selector}, \text{set}(\text{zip}(\text{chans}, \text{writing\_procs}))), \{\}, \text{false}$   
 1076  $)$   
 1077  $[\text{union}(\alpha \text{LEFT\_PROCS}(\text{writing\_procs}, \text{chans}), \alpha \text{ALT\_SPEC}(\text{selector})) \|\| \alpha \text{ALT\_USER}(\text{selector})]$

1078

```

1079     ALT_USER(selector)
1080     ) \ union(PRIVATE_ALT_SPEC_SYSTEM(selector), {| transmit |})

```

1081

1082 where *chans* represents the channels, *writing\_procs* represent the writers and *selector* is the  
 1083 selecting process. Initially, all flags and *waiting* are false, signifying that no writers are ready and  
 1084 the reader (the selector) is not waiting. When a writer becomes ready, the appropriate *ready* flag is  
 1085 set to true. The remaining events not hidden are equivalent to the events that appear outside the  
 1086 box in Figure 4.

1087 The writing processes are *LEFT* processes from Section 5.1 build using *LEFT\_PROCS*. *ALT\_USER*,  
 1088 replaces the *RIGHT* process in regular channel communication. *transmit* and *return* events are  
 1089 hidden by defining *PRIVATE\_ALT\_SPEC\_SYSTEM*. FDR verifies that the specification of the generic  
 1090 *alt* is deadlock-free both in the *failures* and in the *failures-divergence* model. Similar to the channel  
 1091 specification, we have also verified that the *ALT\_SPEC\_SYSTEM* is divergence free. The *alt* is also  
 1092 nondeterministic which occurs if multiple writing events are ready.

1093

## 1094 6 Channel Verification

1095 In the previous two sections we presented:

1096

- 1097 • the CSP specifications for a the scheduler architecture: queue; the scheduler and the sched-  
 1098 uler manager; a monitor; a process's meta-data (the part of a process required for the  
 1099 scheduler) and the scheduler system.
- 1100 • the code a process uses to interact with the scheduling system: schedule, deschedule, and  
 1101 yield.
- 1102 • the specifications for correct behaviour of a channel and alternation.

1103 In this section and the following one, we provide the CSP model of the proposed ProcessJ trans-  
 1104 lated code to meet the specification. In this section, the CSP for a channel communication system  
 1105 (a writer and a reader along with a channel) is presented. These processes running concurrently  
 1106 with the scheduler system allow verification of the channel implementation against the defined  
 1107 specification. In this section, we also highlight changes from our previous work [PC23] through  
 1108 underlined CSP statements.

1109 An example of a channel communication in ProcessJ can be found in the introduction section of  
 1110 this paper. Recall that channel communication is one-to-one and synchronous.

1111 A channel consists of a single monitor and three variables (fields): the reference (pid) of the  
 1112 writer process, the reference of the reader process, and the reference to the data.

1113 First we define some constants and channels required for communicating with the *VARIABLE*  
 1114 processes and the *MONITOR* used in the *CHANNEL* process:

```

1115 datatype Nullable_Processes = NULL | P1 | P2           - set of ProcessJ processes.

```

```

1116 subtype Processes = P1 | P2

```

```

1117 datatype Channels = C1                               - set of ProcessJ channels.

```

```

1118 datatype Values = A | B                             - data values.

```

```

1119 channel data : Channels.Operations.Values           - channels

```

```

1120 channel writer, reader : Channels.Operations.Nullable_Processes

```

```

1121 channel channel_claim, channel_release : Channels.Processes

```

1122 The *reader.chan* and *writer.chan* represent references to the reader and the writer of the channel.  
 1123 The *data.chan* is the field representing the value being communicated across the channel. The  
 1124 monitor process uses the *channel\_claim.chan* and *channel\_release.chan* channels for claiming and  
 1125 releasing a lock on the channel.

1126

1127

1128 A channel is the three *VARIABLE* processes and a single *MONITOR* process interleaved (along  
1129 with the code for *READ* and *WRITE*):

```
1130 CHANNEL(chan) =
1131     VARIABLE(writer.chan, NULL)           - writer = null
1132     |||
1133     VARIABLE(reader.chan, NULL)         - writer = null
1134     |||
1135     VARIABLE(data.chan, A)              - data = A
1136     |||
1137     MONITOR(channel_claim.chan, channel_release.chan)
```

1138 To verify our channel implementation working with our scheduler, we need to build a system  
1139 consisting of a single channel (as defined above) and a reader and a writer. We will consider the  
1140 reader and the writer in the following section.

## 1141 6.1 Channel Communication

1142 What does the CSP that corresponds to a ProcessJ channel-write statement (e.g. *out.write(42)*) look  
1143 like? Here, we incorporate it into a process called *RESTRICTED\_PROCESS\_WRITER*<sup>5</sup>.

1145 6.1.1 A Writing Process. A writing process behaves as follows:

- 1146 • It must be scheduled (i.e., it needs to be ready to run and placed in the run queue).
- 1147 • It then must be run by the scheduler.
- 1148 • The write commences (and loops):
  - 1149 – A message (via *write*) is taken from the environment.
  - 1150 – The channel is claimed to ensure exclusive access by the writer.
  - 1151 – The data field is written to with the message.
  - 1152 – The writer field is written to with the process identifier (pid) of the writing process.
  - 1153 – The writing process sets itself not ready.
  - 1154 – If the reading process is present (the *reader* field is not null), the reading process is  
1155 scheduled.
  - 1156 – The channel is released from exclusive access.
  - 1157 – The writing process yields.
  - 1158 – Once the writer is rescheduled, an *ack* event is raised to the environment.

1159 Here is the associated CSP code:

```
1160 RESTRICTED_PROCESS_WRITER(pid, chan) =
1161     schedule!pid →           - schedule the process
1162     run.pid →                - wait to be run
1163     running.pid.set!true →   - this.running = true
1164     RESTRICTED_PROCESS_WRITER'(pid, chan)
1165
1166 RESTRICTED_PROCESS_WRITER'(pid, chan) =
1167     write.chan.pid?message →   - take message from the environment
1168     channel_claim.chan!pid →   - claim the channel
1169     WRITE(pid, chan, message); - perform the write
1170     channel_release.chan.pid → - release the channel
1171     YIELD(pid);                - yield
1172     ack.chan.pid →           - acknowledge to the environment
1173     RESTRICTED_PROCESS_WRITER'(pid, chan) - recurse
```

1174 <sup>5</sup>Restricted because it is controlled by the scheduler as opposed to an unrestricted process that is not controlled by a  
1175 scheduler.

```

1177
1178 WRITE(pid, chan, item) =
1179     data.chan.store!item →           - set the data field of the channel
1180     writer.chan.store!pid →         - register the writer
1181     ready.pid.store!false →        - set the process not ready to run
1182     reader.chan.load?v →           - get the reader field
1183     if (v != NULL) then             - if (reader ≠ null) then
1184         SCHEDULE(pid, v)         - schedule the reader
1185     else                             - else
1186         SKIP                       - do nothing

```

The differences between the new version above and the version found in [PC23] are as follows (the lines are underlined in the code above):

- schedule!pid: the first time a process runs, it must schedule itself. When a process is created, its *ready* field is set to true and its *running* field to false. In the previous version, a fixed schedule was provided to the system.
- running.pid.store!true: when a process runs (either the first time or after resuming from a yield) it must set its *running* field to true. Note that this happens also in the *YIELD* procedure. The *running* flag is a new addition.
- SCHEDULE(pid, v): rather than simply setting the reader process's *ready* field to true, we use the new *SCHEDULE* procedure to schedule the reader again. This is due to the new scheduling algorithm in use.

The yielding of the writer is necessary to ensure that the communication is indeed synchronous. It is worth noting that a one-to-one channel is always ready for a writer to write on when it gets to the write statement. If there were data still on the channel, the reader would not yet have read it, and the writer would have been not-ready to run, but if the writer is ready to run, the data must have been read because the writer could only have been scheduled by the reader as the channel is one-to-one and synchronous.

6.1.2 *A Reading Process*. The reading process is more complicated than the writer as there is no guarantee that a reader can continue directly to reading a message from the channel, as the writer may not have written any data yet. It behaves as follows:

- The reading process is scheduled.
- The reading process is run.
- The running flag is set to true.
- The read commences and loops:
  - A *start\_read* event is accepted from the environment indicating the read has begun.
  - The channel is locked for exclusive access.
  - If the writer is not present (the *writer field is null*):
    - \* The *reader* field is updated with the reading process ID.
    - \* The reading process is set not ready.
    - \* The channel is unlocked for exclusive access.
    - \* The reading process yields.
  - Else, if the writer is present (the *writer field is not null*):
    - \* The channel is unlocked for exclusive access.
  - The channel is locked for exclusive access (necessary as the reading process might have yielded prior to this point).

- 1226           – The writing process is scheduled<sup>6</sup>.
- 1227           – The reader and writer fields are set to null.
- 1228           – The data is retrieved from the channel.
- 1229           – The channel is unlocked for exclusive access.
- 1230           – The reading process yields (a courtesy yield).
- 1231           – The data is returned to the caller (via the *read* event).

1232 The actual CSP looks like this:

```

1233 RESTRICTED_PROCESS_READER(pid, chan) =
1234   schedule.pid →                – schedule the process
1235   run.pid →                      – wait to be run
1236   running.pid.set!true →       – this.running = true
1237   RESTRICTED_PROCESS_READER'(pid, chan)
1238
1239 RESTRICTED_PROCESS_READER'(pid, chan) =
1240   start_read.chan.pid →          – external event to start the read
1241   channel_claim.chan!pid →       – claim the channel
1242   writer.chan.load?p →          – get the writer field
1243   (
1244     if (p == NULL) then          – if (writer == null) then
1245       reader.chan.store.pid →    – reader = this
1246       ready.pid.store!false →     – this.ready = false
1247       channel_release.chan.pid → – release the monitor
1248       YIELD(pid)                 – yield
1249     else                          – else
1250       channel_release.chan.pid → – release the channel
1251       SKIP
1252   );
1253   channel_claim.chan!pid →       – claim the channel
1254   READ(pid, chan);              – see READ below
1255   data.chan.load?message →       – message = data
1256   channel_release.chan.pid →    – release the monitor
1257   (
1258     YIELD(pid);                  – yield
1259     read.chan.pid!message →      – deliver the message to the environment
1260     RESTRICTED_PROCESS_READER'(pid, chan) – recurse to start over
1261   )
1262
1263 READ(pid, chan) =
1264   writer.chan.load?w →          – get the writer field.
1265   if (w == NULL) then
1266     DIV                            – never happens.
1267   else
1268     SCHEDULE(pid, w);          – schedule the writer
1269     (
1270       writer.chan.store!NULL →    – writer = null
1271       reader.chan.store!NULL →    – reader = null
1272       SKIP
1273     )
1274 
```

<sup>6</sup>Note that this writer cannot write a new value to the channel because the channel is still locked by the reader.

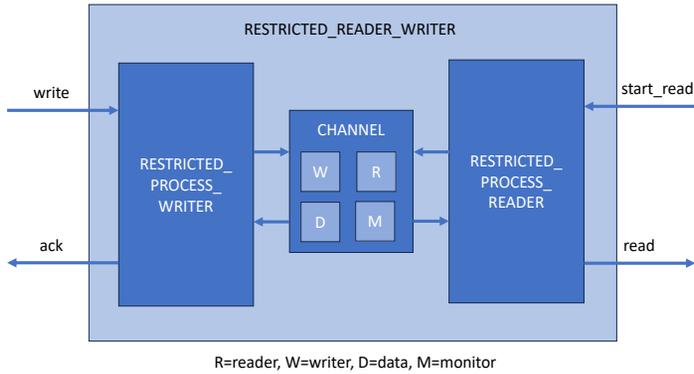


Fig. 5. A writer and a reader along with a channel.

We now have both the writing and the reading side of our channel communication system (the left and right parts of Figure 5). We can now wire them up with with a channel and run them with a scheduler.

6.1.3 *Scheduled Reader/Writer System.* We can now create a process that runs the reader and the writer interleaved and both in parallel with a channel; see Figure 5. That CSP looks as follows:

```

RESTRICTED_READER_WRITER(writer_proc, reader_proc, chan) =
  (
    (
      RESTRICTED_PROCESS_WRITER(writer_proc, chan)      - a writing process
      |||
      RESTRICTED_PROCESS_READER(reader_proc, chan)      - a reading process
    )
    [| αCHANNELS |]
    CHANNELS
  ) \ αCHANNELS

```

where *CHANNELS* is defined as follows:

$$CHANNELS = ||| c : Channels \bullet CHANNEL(c)$$

which, because *Channels* is just  $\{C1\}$ , evaluates to a single channel, namely,  $CHANNEL(C1)$ .

All we now have to do is run this process in parallel with the  $N\_SCHEDULER\_SYSTEM$  from Section 4 to have a complete CSP specification of the *implementation* of a channel communication in ProcessJ. Here,  $n$  is the number of schedulers that we want the system to run with.

```

RESTRICTED_PJ_CHAN_SYSTEM(reader_proc, writer_proc, chan, n) =
  (
    RESTRICTED_READER_WRITER(reader_proc, writer_proc, chan)
    [| αN_SCHEDULER_SYSTEM |]
    N_SCHEDULER_SYSTEM(n)
  ) \ αN_SCHEDULER_SYSTEM

```

We can now perform the verification using FDR by asserting deadlock freedom, livelock freedom, determinism, and the following refinement checks:

$$GENERIC\_CHANNEL(P1, P2, C1) \sqsubseteq_M RESTRICTED\_PJ\_CHAN\_SYSTEM(P1, P2, C1, n)$$

$$RESTRICTED\_PJ\_CHAN\_SYSTEM(P1, P2, C1, n) \sqsubseteq_M GENERIC\_CHANNEL(P1, P2, C1)$$

Table 1. Results of FDR Verification for Channels Under New Scheduling System

	Restricted (Number of Schedulers)								Unrestricted
	1		2		3		4		
	Old	New	Old	New	Old	New	Old	New	
Deadlock free	FD	FD	F	FD	F	FD	F	FD	FD
Livelock free	FD	FD	✗	FD	✗	FD	✗	FD	FD
Spec $\sqsubseteq$ Impl	T	T	F	FD	F	FD	F	FD	FD
Impl $\sqsubseteq$ Spec	✗	✗	T	FD	T	FD	T	FD	FD

*Spec* = *GENERIC\_CHANNEL*, *Impl* = *RESTRICTED\_Pj\_CHAN\_SYSTEM*.

where  $M$  is either  $T$  (trace refinement),  $F$  (failures refinement), or  $FD$  (failures-divergence refinement). We use  $n=1, 2, 3, 4$ , and  $n \geq 2$  is enough schedulers for each process to get one.

## 6.2 Verification Results for ProcessJ Channel Implementation Under New Scheduling System

In Table 1 the right-most column labelled **Unrestricted**<sup>7</sup> is the verification results we obtained in [PC23] for channel communication *without* a scheduler. That is, the results we obtained by applying the *standard* approach to verification where no execution environment or hardware was taking into account. This column represents our verification goal. If we can achieve the same results with the implementations in this paper, we will have demonstrated that the channel behaves as specified within our execution environment. As Table 1 shows, there is limited success based on the number of running schedulers available. The implementations that include a scheduler are referred to as **Restricted**. In Table 1, the results from [PC23] are listed in the **Old** columns, and the new results from this paper are listed in the **New** columns.

With the newly developed scheduler system, we can now repeat all the CSP assertions developed in [PC23] to verify if we indeed have managed to achieve success in the *failures-divergence* model, the strictest of the three FDR models<sup>8</sup>. The entries which differ from the old to the new results have been highlighted with a grey box.

As we can see, most of the new entries are indeed  $FD$ , which means that the checks succeeded in the *failures-divergence* semantic model and yielded the same as the unrestricted verification (i.e., the ‘standard’ verification approach without taking the scheduler into account). However, a number of entries are not.

- 1 Scheduler:

- *Spec*  $\sqsubseteq$  *Impl*: It is clear that the implementation’s traces are indeed a (true) subset of the traces allowed by the specification. With a single scheduler the order in which processes appear in the run queue will make some of the traces allowed by the specification impossible. The specification requires both the *start\_read* and *write* events are available. A single scheduler means only one of these two events will be available, non-deterministically based on the initial queue. The specification also requires *read* and *ack* are available at the same time. However, with a single scheduler, the order of events will become sequential –  $\langle ack, write, read \rangle$ . Note that the next *write* will occur before *read* due to the yielding and scheduling within the system. Channel behaviour

<sup>7</sup>Unrestricted because there was no scheduler to restrict the execution.

<sup>8</sup>We do realise that since both the specification and the implementation are divergence-free, verifying in the *stable failures* model is technically sufficient.

then becomes a known fixed sequence of events. Failures refinement is not possible as the refusal sets will differ: there is only one scheduler, so the process that can run is limited to the next ready process in the run queue.

- $Impl \sqsubseteq Spec$ : With a single scheduler and an ordering imposed by the order in which the processes are placed in the run queue, the implementation will never be able to mimic the specification completely. If the reader is in the queue before the writer, the reader will be run first. Similarly, if the writer is before the reader in the queue, it will be run first. The specification does not make that distinction. It should be obvious that the implementation's set of traces is indeed a true subset of that of the specification—something FDR verifies as well.<sup>9</sup>

With enough schedulers, which, in this case is just two, all results will be the same as the unrestricted version. As more processes are added to the system, further schedulers would be required to meet the specification. For example, with a three channel system featuring six processes, if five or fewer schedulers are available then likewise the implementation will only refine within the traces model of the specification, not the failures model.

## 7 Choice Verification

With channel communication handled, we now turn to choice (also known as *alternation*, or simply *alt*).

For choice we have two writers ( $W1$  and  $W2$ ) and a single reader ( $R$ ) that performs an alternation between the two writers. Thus, we also need two channels ( $C1$  and  $C2$ ). Therefore, we have the following slightly changed initial definitions:

datatype  $Nullable\_Processes = NULL \mid R \mid W1 \mid W2$       – set of ProcessJ processes

subtype  $Processes = R \mid W1 \mid W2$

subtype  $Writers = W1 \mid W2$

datatype  $Channels = NONE \mid C1 \mid C2$       – set of channels ( $NONE$  is 'NULL')

subtype  $Active\_Channels = C1 \mid C2$

datatype  $Values = A \mid B$       – data values

This time, we get two channels as follows:

$CHANNELS = \parallel c : Active\_Channels \bullet CHANNEL(c)$

The selector/reader needs two additional *VARIABLE* processes for variables called  $i$  and  $selected$ . We define them as follows:

$ALT\_VARIABLES = VARIABLE(i, NONE) \parallel VARIABLE(selected, NONE)$

where  $NONE$  serves as the *null* value for channels. We have added a  $NONE$  value as during the enable/disable process of the select, we might have no channel selected.

We now can add channels for communicating with the variables:

channel  $i, selected : Operations.Channels$

### 7.1 The Writing Processes

The implementation of the writing end of the two channels part-taking in the choice is exactly the same as we saw in Section 6.1.1. The writing processes are not aware that the reading ends of the channels on which they communicate are used in a choice operation in the reading process.

<sup>9</sup>In general, if there are not enough schedulers for each process to have its own, the implementation will *never* more than trace-refine the specification as some traces in the implementation are simply impossible. Naturally as long as the system does not deadlock (which it never does) we are fine.

## 7.2 The Reading Choice Process

The `alt` operator chooses between a number of ready channel communications (here just two) by picking one of them based on a rule set. Here this rule set is simply picking the first one that is ready (by the order in which they appear in the program)—this is called a *prioritized choice* or sometimes a *pri alt*. This choice is performed by an enable/disable sequence which first traverses all the channel communications from top to bottom and determines the top-most ready one (enable), then traverses them backwards and updates the top-most ready one if one further up is found (disable). The enable sequence maintains a variable *i* (holding a channel) and the disable sequence produces the channel to be read in a variable called *selected* (stored in the variable *c* in the CSP below). See [PC23] for more details about the working of the *ENABLE* and *DISABLE* procedures; also note that in this paper both procedures are now more general, recursive versions, that work with any number of branches like the rest of the *alt* system. The steps of the reading choice are:

- The selecting process is scheduled.
- The selecting process is run.
- The running flag is set to true.
- The select commences and loops:
  - The ready flag is set to false.
  - The select operation is started.
  - The branch enable sequence is executed.
  - The selecting process yields—it may be ready to run if a ready branch was found, or it might now wait until one of the channels signals it is ready.
  - The branch disable sequence is executed.
  - The selected channel is retrieved.
  - The read operation is started on the selected channel.
  - The selected channel is locked for exclusive access.
  - The writing process is scheduled.
  - The reader and writer fields are set to null.
  - The data is retrieved from the channel.
  - The channel is unlocked for exclusive access.
  - The selecting process yields (a courtesy yield).
  - The data is returned to the caller (via the *read* event).

The equivalent CSP is:

```

RESTRICTED_PROCESS_SELECTOR(pid) =
  schedule!pid →           - schedule the process
  run.pid →                 - wait to be run
  running.pid.set!true →   - this.running = true
  RESTRICTED_PROCESS_SELECTOR'(pid)

RESTRICTED_PROCESS_SELECTOR'(pid) =
  ready.pid.store>false →   - this.ready = false
  select.pid →               - external event to start the alt
  ENABLE(pid);               - execute the enable sequence
  YIELD(pid);                 - yield
  DISABLE(pid);               - execute the disable sequence
  selected.load?c →         - read the selected variable
  (
    start_read.c.pid →       - start the read on the selected channel
    channel_claim.c.pid →    - claim the channel
  )

```

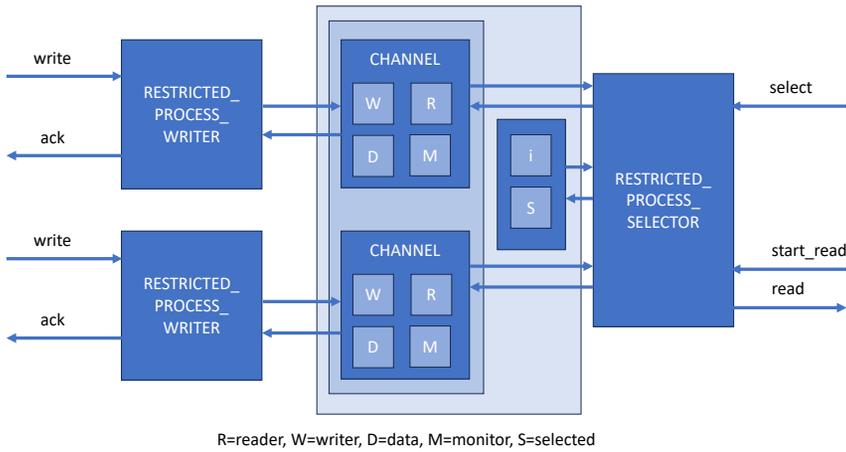


Fig. 6. Two writers and a selector/reader along with two channels and the extra alt variables.

```

1488 READ(pid, c);
1489 data.c.load?message → - perform the actual read
1490 channel_release.c!pid → - message = data
1491 YIELD(pid); - release the channel
1492 read.c.pid!message → - courtesy yield
1493 RESTRICTED_PROCESS_SELECTOR'(pid) - deliver the message to the environment
1494 ) - recurse to start over

```

where `READ` is as defined in Section 6.1.2. Note, both `ENABLE` and `DISABLE` may set the reading process's ready field back to `true` if a ready channel is encountered. We now can wire two writers and a single selector together in a system along with the channels and the required variables (`i` and `selected` in the `ALT_VARIABLES` process) and get the following CSP process (which is depicted in Figure 6).

That CSP code looks like this:

```

1502 RESTRICTED_SELECTOR_WRITERS(writing_procs, selector, chans) =
1503 (
1504   (CHANNELS ||| ALT_VARIABLES)
1505   [| union(αCHANNELS, αALT_VARIABLES) |]
1506   ALT_NETWORK(writing_procs, selector, chans)
1507 ) \ union(αCHANNELS, αALT_VARIABLES)

```

where `ALT_NETWORK` consists of a single `RESTRICTED_PROCESS_SELECTOR` and the appropriate number of `RESTRICTED_PROCESS_WRITER` processes:

```

1510 ALT_NETWORK(writing_procs, selector, chans) =
1511 RESTRICTED_PROCESS_SELECTOR(selector)
1512 |||
1513 PROCESS_WRITER_NETWORK(writing_procs, chans)
1514
1515 PROCESS_WRITER_NETWORK(<proc> ^ writing_procs, <chan> ^ chans) =
1516 RESTRICTED_PROCESS_WRITER(proc, chan)
1517 |||
1518 PROCESS_WRITER_NETWORK(writing_procs, chans)

```

Table 2. Results of FDR Verification for Choice Under New Scheduling System

	Restricted (Number of Schedulers)								Unrestricted
	1		2		3		4		
	Old	New	Old	New	Old	New	Old	New	
Deadlock free	FD	FD	F	FD	F	FD	F	FD	FD
Livelock free	FD	FD	✗	FD	✗	FD	✗	FD	FD
Spec $\sqsubseteq$ Impl	T	T	T	T	F	FD	F	FD	FD
Impl $\sqsubseteq$ Spec	✗	✗	✗	✗	T	FD	T	FD	FD

$$Spec = ALT\_SPEC\_SYSTEM, Impl = RESTRICTED\_Pj\_ALT\_SYSTEM$$

where *PROCESS\_WRITE\_NETWORK* builds an interleaving of the correct number of *RESTRICTED\_PROCESS\_WRITER* processes. Finally, just like with the channel communication, we run in parallel with a scheduling system and obtain the final process, *RESTRICTED\_Pj\_ALT\_SYSTEM*, which looks like this:

$$\begin{aligned}
 & RESTRICTED\_Pj\_ALT\_SYSTEM(writing\_procs, selector, chans, n) = \\
 & ( \\
 & \quad RESTRICTED\_SELECTOR\_WRITERS(writing\_procs, selector, chans) \\
 & \quad [ \alpha N\_SCHEDULER\_SYSTEM ] \\
 & \quad N\_SCHEDULER\_SYSTEM(n) \\
 & ) \setminus \alpha N\_SCHEDULER\_SYSTEM
 \end{aligned}$$

This time, the *RESTRICTED\_Pj\_ALT\_SYSTEM* can be checked for deadlock, livelock, and determinism, as well as verified against the *ALT\_SPEC\_SYSTEM* specification from Section 5.2:

$$\begin{aligned}
 & ALT\_SPEC\_SYSTEM(R, \langle C1, C2 \rangle, \langle W1, W2 \rangle) \sqsubseteq_M \\
 & \quad RESTRICTED\_Pj\_ALT\_SYSTEM(\langle W1, W2 \rangle, R, \langle C1, C2 \rangle, n)
 \end{aligned}$$

$$\begin{aligned}
 & RESTRICTED\_Pj\_ALT\_SYSTEM(\langle W1, W2 \rangle, R, \langle C1, C2 \rangle, n) \sqsubseteq_M \\
 & \quad ALT\_SPEC\_SYSTEM(R, \langle C1, C2 \rangle, \langle W1, W2 \rangle)
 \end{aligned}$$

Again, *M* is one of *T*, *F*, or *FD*. *R* is the reader/selector process, and *W1* and *W2* are the two writing processes, and *C1* and *C2* are the channels. We use  $n = 1, 2, 3, 4$ , and  $n \geq 3$  is enough schedulers for each process to get one.

### 7.3 Verification Results for ProcessJ Choice Implementation Under New Scheduling System

In Table 2 we present the choice verification results using the same mapping between restricted, old and new results. We have again repeated all the same CSP assertions from [PC23] to verify choice in the *failures-divergence* model. As with channel, most of the new entries are FD. In the following section we shall explain the differences in the results.

#### 7.3.1 Analysis of Choice Verification Results.

- 1 and 2 Schedulers:

- Spec  $\sqsubseteq$  Impl: The implementation traces are a subset of the specification. However, as with the single-scheduler channel, the system cannot undertake all possible event interleavings, meaning the implementation refusal set is different from the specification refusal set.
- Impl  $\sqsubseteq$  Spec: As the implementation does not contain all possible traces of the specification (due to the different refusal set), the specification traces are not a proper subset

of the implementation. The specification can allow progress of each writer and the selector concurrently, whereas only one process in the single scheduler system can continue at a time.

The limitations here are the same as with channels. With enough schedulers, which, in this case is three, all results will be the same as the unrestricted version. Without enough *SCHEDULER* processes, some traces can no longer happen due to the order in which processes sit in the run queue; if there are not enough schedulers to run them all, then the run queue order imposes certain restrictions on the traces that are possible. This will not happen if there are at least as many schedulers as there are processes.

#### 7.4 Three-Branch Choice

We have also verified a three-branch *alt* using the same CSP as for the two-branch *alt*. The results are like those in Table 2 except for three schedulers the refinement check  $Spec \sqsubseteq Impl$  which now only succeeds in the traces model and the refinement  $Impl \sqsubseteq Spec$  now fails in all models. This is exactly as the two-branch *alt* behaves with two schedulers. The assertions pass for four schedulers. In other words, we have also shown that for four active processes (three writers and one selector) we require at minimum four scheduling processes to enable meeting the specification in the stable-failures model.

It should be noted that although we have demonstrated a two-branch *alt* and a three-branch *alt* verification, the analysis is not general enough for us to prove that the *alt* would work with any arbitrary number of branches with the right number of supporting scheduling processes. However, we believe the demonstration of two-branch and three-branch *alt* is enough for us to conclude that this is likely the case.

## 8 Analysis of Scheduling Behaviours

In this section, we use different analysis and explanation techniques to illustrate the unintended prioritisation (and thus increasing likelihood of traces not occurring) that can occur during channel communication. We first consider trace-limitations that can be caused by insufficient resources, before presenting empirical evidence of the prioritisation occurring within JCSP. Finally, we present an analysis of the interactions occurring during channel communication for both ProcessJ and JCSP to demonstrate the availability to continue at points of interaction.

### 8.1 The Trace-Limitations Caused by Insufficient Resources

We have already seen how a lack of resources (not enough schedulers) can impact the possible traces of a system. In this section we illustrate this issue with a small general example. In a single-threaded execution environment, consider the following simple CSP:

$$A = a \rightarrow \text{SKIP}$$

$$B = b \rightarrow \text{SKIP}$$

$$AB = A \parallel B$$

We can easily determine the set of traces (again ignoring the termination event  $\surd$ ) of  $AB$ , and it is:

$$\text{traces}(AB) = \{\langle \rangle, \langle a \rangle, \langle b \rangle, \langle a, b \rangle, \langle b, a \rangle\}$$

After the empty trace  $\langle \rangle$ , FDR computes the set of events the process is willing to engage in as the set  $\{a, b\}$ —this set is called the *acceptance set*. Similarly, the *stable refusal set*—the set of events that the process can refuse to engage in—is the empty set  $\{\}$ . That is, either of  $A$  and  $B$  can progress (as  $A$  and  $B$  are interleaved). Now consider an execution context with a single scheduler and a single run queue. This run queue will contain the equivalent processes of  $A$  and  $B$ , say,  $P_A$  and  $P_B$ .

1618 With two processes, there are two different ways such a run queue can look (denoted  $R_1$  and  $R_2$ ),  
 1619 and they are as follows:

$$1620 \quad R_1 = [P_A, P_B] \text{ or } R_2 = [P_B, P_A]$$

1621 If the run queue has the form  $R_1$ , then  $AB$  effectively becomes  $A ; B$  (ignoring the SKIPs). If  $R_2$  is  
 1622 the run queue, then  $AB$  is  $B ; A$ . Then possible traces when using  $R_1$  as the run queue is<sup>10</sup>:

$$1623 \quad \text{traces}((P_A ||| P_B)_{R_1}) = \text{traces}(A ; B) = \{\langle \rangle, \langle a \rangle, \langle a, b \rangle\}$$

1624 that is, the set of traces that can be produced by running the single scheduler with the particular  
 1625 run queue configuration found in  $R_1$ . Similarly, if we use the  $R_2$  run queue configuration, we get:

$$1626 \quad \text{traces}((P_A ||| P_B)_{R_2}) = \text{traces}(B ; A) = \{\langle \rangle, \langle b \rangle, \langle b, a \rangle\}$$

1627 It should be clear that for any run queue schedule  $R_i$ , the following equation will therefore hold:  
 1630

$$1631 \quad \text{traces}((P_A ||| P_B)_{R_i}) \subseteq \text{traces}(AB)$$

1632 while at the same time we have the following:  
 1633

$$1634 \quad \text{traces}(AB) \not\subseteq \text{traces}((P_A ||| P_B)_{R_i})$$

1635 If  $AB$  is a *specification*, and  $(P_A ||| P_B)_{R_i}$  an *implementation*, then the implementation (when running  
 1636 under a single ProcessJ scheduler) will always *behave* and never produce a trace disallowed by  
 1637 the specification, because the scheduler is deterministic. However, there will be traces of the  
 1638 specification that a particular run queue configuration cannot produce (e.g.,  $\langle b, a \rangle$  cannot be  
 1639 produced by the  $R_1$  configuration). This will be the reason for the trace refinement failing in this  
 1640 direction when running with just a single scheduler in ProcessJ.  
 1641

## 1642 8.2 Empirical Evidence of JCSP Prioritisation

1643 We believe that preemptive implementations of a CSP-inspired runtime will likewise suffer from  
 1644 unintended priorities. Indeed, we have empirical evidence demonstrating that JCSP has such an  
 1645 unintended prioritisation. Consider the following system specification (presented here in CSP):  
 1646

```
1647 PROC_1 = a!0 → b!1 → PROC_1
1648 PROC_2 = a?0 → b!2 → PROC_2
1649 RECEIVER = b?x → b?y → RECEIVER
1650 SYSTEM = (PROC_1 [| {a} |] PROC_2) [| {b} |] RECEIVER
```

1651 We have implemented this system in both JCSP and ProcessJ. The code is not presented here for  
 1652 simplicity as the CSP definition is enough to explain the general behaviour under consideration.

1653 In this system,  $PROC_1$  and  $PROC_2$  first communicate via channel  $a$  ( $PROC_1$  writing and  
 1654  $PROC_2$  reading). They then both communicate to  $RECEIVER$  via channel  $b$ .  $PROC_1$  and  $PROC_2$   
 1655 do not synchronise on channel  $b$ , and so will proceed in some order. For  $RECEIVER$ , the values of  $x$   
 1656 and  $y$  will either be 1 and 2 ( $PROC_1$  communicates on  $b$  first) or 2 and 1 ( $PROC_2$  communicates  
 1657 on  $b$  first).

1658 In a fair system, with no hidden priorities,  $RECEIVER$  will have equal probability to see either  
 1659 option for the values of  $x$  and  $y$ . On an unfair system, the probability will favour one choice over  
 1660 another. We are fully aware that CSP does not consider probabilities in this manner, and probabilistic  
 1661 models exist. However, from an implementation point-of-view, it is important to recognise the  
 1662 outcomes of limited resources, which monitoring the likelihood of an event occurring provides  
 1663 empirical evidence of.

1664  
 1665 <sup>10</sup>We will use the  $a$  and  $b$  events from the CSP here as well  
 1666

In JCSP, channel  $b$  is implemented as an *Any2OneChannel*. This channel has a shared writing end, protected by a monitor. With this implementation of the specified system, we have created a race condition between  $PROC\_1$  and  $PROC\_2$  to claim the shared channel end to use channel  $b$ .

When we implement this system in JCSP and run it on a single processor machine (such as a Raspberry Pi Zero) or otherwise constrain the available processors for the JVM to one, over 99.999% of the time,  $RECEIVER$  will observe 2 then 1 on channel  $b$ , demonstrating priority to the reading process,  $PROC\_2$ , during channel communication. Thus, we know that JCSP, when constrained of resources will demonstrate unexpected priority in its channel communication based on the fact that reading leaves the communication first. In ProcessJ, the results we see are writing leaving the communication first, giving priority to  $PROC\_1$ .

The system (execution environment) has prioritised a particular trace. As no execution environment analysis was taken for JCSP, the existing verification of JCSP would not signal an issue that certain behaviours may not be possible under certain resource and scheduling conditions. It can be surmised that the JCSP channel does not fully behave as a CSP channel under constrained conditions, and a CSP channel cannot be naïvely used as a drop-in replacement for a JCSP channel between implementation and modelling.

We believe this outcome is likely across all standard (e.g., deterministically scheduled) preemptive runtimes implementing CSP channels when under constrained resources. One side of the communication will complete first, leading to unexpected priorities. More work is required to build a suitable execution environment for a preemptively scheduled runtime to fully test these beliefs. In CSP, all processes participating in a communication event continue together (the event is atomic to all participants). In an implementation, without hardware support, we find this unlikely as only certain participants will have hardware executing them to enable those processes to continue in parallel. A process within a deterministic preemptively scheduled runtime will not yield the hardware until its time slice has ended – prioritising certain parts of the system accordingly. In JCSP terms, this would be the reading end of a channel communication being prioritised.

### 8.3 Examination of Channel Interactions

Let us now examine the interaction between a reading and writing process during the communication handshake undertaken by a channel. Both ProcessJ and JCSP will be presented<sup>11</sup>. We consider each step of the interaction, illustrating when the writing process (WRITE) and the reading process (READ) are blocked or runnable. As these interactions take place during locked sections with notifications, the reordering of operations is limited and execution ordering (through the channel communication ping-pong) is known. We present both when the writing process arrives at the channel first and when the reading process arrives at the channel first. For ProcessJ, we present the interactions as if there is a single scheduling (runner) thread in operation due to the potential for more concurrency due to reduced locking in comparison to JCSP. In Table 3 and Table 4, ■ indicates when a process is blocked from execution, whereas ▷ indicates where a process is able to continue execution.

Table 3 presents the interactions during channel communication with ProcessJ. Note that for the situation where the WRITE process accesses the channel first, at step 11 WRITE is free to continue as it has been scheduled by READ (step 10) and will now wait trying to lock the channel. READ will eventually yield (step 15), meaning with only a single processor WRITE can continue without restriction, prioritising its behaviour. Similarly, when READ arrives first, WRITE will be free to continue at step 15, while READ will yield at step 19. With two or more runner threads, write and

<sup>11</sup>See Appendix B for the JCSP code-snippets.

Table 3. Singularity Scheduled ProcessJ Channel Communication

Write first.			Read first.		
Step	WRITE	READ	Step	WRITE	READ
1	LOCK CHANNEL	█	1	█	LOCK CHANNEL
2	SET HOLD	█	2	█	SET READER
3	SET WRITER	█	3	█	SET READER NOT READY
4	SET WRITER NOT READY	█	4	█	UNLOCK CHANNEL
5	UNLOCK CHANNEL	█	5	▷	YIELD
6	YIELD	▷	6	LOCK CHANNEL	█
7	█	LOCK CHANNEL	7	SET HOLD	█
8	█	UNLOCK CHANNEL	8	SET WRITER	█
9	█	LOCK CHANNEL	9	SET WRITER NOT READY	█
10	█	SCHEDULE WRITER	10	SCHEDULE READER	█
11	▷	SET WRITER TO NULL	11	UNLOCK CHANNEL	█
12	▷	SET READER TO NULL	12	YIELD	▷
13	▷	TEMP STORE HOLD	13	█	LOCK CHANNEL
14	▷	UNLOCK CHANNEL	14	█	SCHEDULE WRITER
15	▷	YIELD	15	▷	SET WRITER TO NULL
16	RETURN	▷	16	▷	SET READER TO NULL
...	▷	RETURN TEMP STORE	17	▷	TEMP STORE HOLD
			18	▷	UNLOCK CHANNEL
			19	▷	YIELD
			20	RETURN	▷
			...	▷	RETURN TEMP STORE

Table 4. JCSP Channel Communication

Write first.			Read first.		
Step	WRITE	READ	Step	WRITE	READ
1	LOCK CHANNEL	█	1	█	LOCK CHANNEL
2	SET HOLD	█	2	█	SET EMPTY TO FALSE
3	SET EMPTY TO FALSE	█	3	█	WAIT (UNLOCK CHANNEL)
4	WAIT (UNLOCK CHANNEL)	█	4	LOCK CHANNEL	█
5	█	LOCK CHANNEL	5	SET EMPTY TO TRUE	█
6	█	SET EMPTY TO TRUE	6	NOTIFY	█
7	█	NOTIFY	7	WAIT (UNLOCK CHANNEL)	█
8	█	RETURN HOLD	8	█	LOCK CHANNEL
9	█	UNLOCK CHANNEL	9	█	NOTIFY
10	LOCK CHANNEL	▷	10	█	RETURN HOLD
11	UNLOCK CHANNEL	▷	11	█	UNLOCK CHANNEL
12	RETURN	▷	12	LOCK CHANNEL	▷
			13	UNLOCK CHANNEL	▷
			14	RETURN	▷

read can make progress together, as the final yields by READ are a courtesy with the process still able to continue if resources are available.

Table 4 presents the interactions in JCSP when the channel communication occurs. Note that the opportunity for concurrency is reduced due to locking of the channel through synchronised methods, leading to WRITE being blocked from continuing until READ has completed (step 10 when write arrives first, and step 12 when read arrives first). READ is able to continue freely after it releases the lock.

With enough resources, it can be seen that both ProcessJ and JCSP have the opportunity for concurrent continuation. However, if resources are limited, then a more rigid ordering will be followed. For example, with ProcessJ we know that the writing process will always be prioritised leaving the channel. In JCSP, the reading process is prioritised. In JCSP, READ will notify WRITE that it can continue, but WRITE is blocked until READ unlocks the channel. As such, READ can continue freely until it is preempted. As READ was only recently provided resources (i.e., step 5 when WRITE comes first, step 8 when READ comes first), and only four operations occur (two for locks, one to notify), a preemptive scheduler would have to have aggressively short time slices to preempt READ at this point. As such, these priorities on write and read will occur under resource

1765 constraints, and as such certain traces are not possible. This prioritisation will likewise propagate  
 1766 through a system, as illustrated in the empirical evidence presented in Section 8.2.

## 1767 9 Discussion and Conclusion

1769 Although our results demonstrate ProcessJ’s channel and choice working correctly within certain  
 1770 operating conditions (or fully when not considering restricted operating conditions as demonstrated  
 1771 in our previous work [PC23], much like other papers on channel system modelling), this was not the  
 1772 ultimate goal of our work. Indeed, we had previously verified our primitives as correct outside their  
 1773 execution environment. Our primary aim was to ensure ProcessJ would work when multithreading  
 1774 was introduced to the scheduling system. Our previous paper [PC23] did not fully succeed, although,  
 1775 we were confident the divergence was not an issue for a system to make progress. In this paper, we  
 1776 have proposed a change to the ProcessJ scheduling runtime to remove divergence, and therefore  
 1777 verified we can model (and soon implement) an improved ProcessJ scheduling system. The work  
 1778 we have presented is on the verification of the scheduling environment that the channel and choice  
 1779 mechanisms must work within.

1780 However, our work has demonstrated limitations with scheduling ProcessJ processes, which  
 1781 we believe will be evident in other channel-based implementations. Full specification verification  
 1782 is only possible if there are enough resources (i.e., processors) to run all active processes. This is  
 1783 demonstrated in our channel results (see Table 1), where full specification verification of a channel  
 1784 is only possible with two or more schedulers. Likewise, for choice, three or more schedulers are  
 1785 required for full system verification of a two-branch alt (with three active processes), and four  
 1786 schedulers for a three-branch alt (with four active processes). In both cases, the amount of hardware  
 1787 available **must** be equal to or greater than the number of active processes for full verification. We  
 1788 consider this a general problem—no real-world implementation of a channel-based runtime can  
 1789 meet the abstract specification of a channel unless there are enough processors available.

1790 The problem we are presenting is illustrated in Figure 7. Welch and Martin [WM00] have  
 1791 algebraically proven their channel specification is the same as a pure CSP channel. We do not argue  
 1792 with this proof, and indeed use this specification ourselves. However, we do not believe a claim  
 1793 that a channel implementation fully meets Welch and Martin’s channel specification is sound in all  
 1794 circumstances as demonstrated in our work. There is an additional level of indirection that has been  
 1795 introduced. The indirection occurs due to the claim that the channel implementation will refine the  
 1796 channel specification under failures-divergences, demonstrating a behavioural equivalence to a CSP  
 1797 channel. Welch and Martin argued that as their model of the JCSP channel implementation refined  
 1798 their channel specification, the JCSP channel could be replaced by a CSP channel when modelling.  
 1799 Thus, a system specification that used CSP channels could check a system implementation that  
 1800 used JCSP channels without modelling the full JCSP channel implementation. Only CSP channels  
 1801 were necessary.

1802 However, if the channel implementation does not meet the specification in failures-divergences,  
 1803 then likewise any system implementation will not meet the system specification if only CSP channels  
 1804 are used to model actual channel implementations. As we have argued, meeting Welch and Martin’s  
 1805 channel specification when considering the execution environment of the implementation is  
 1806 problematic. Thus, a broad claim of using a CSP channel to model an actual channel implementation  
 1807 to verify a system will not account for the unexpected prioritisation that occurs at the end of the  
 1808 channel implementation’s ping-pong of synchronisation.

1809 The additional level of indirection for verification has introduced potential unintended be-  
 1810 havioural limitations. Problems will occur if a system is modelled and verified in CSP using the  
 1811 argument that the channel implementation can be replaced with a standard CSP channel. Our work  
 1812 demonstrates that this is only the case *if* enough processors are available to run all active processes.  
 1813

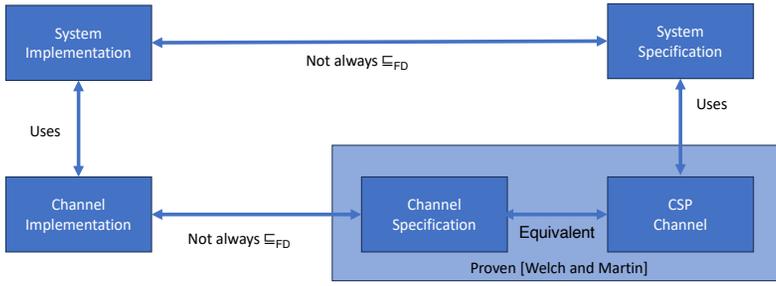


Fig. 7. Indirection of Verification with a Channel Implementation.

The same argument can be made for choice. This is a limitation of implementations without enough processors or the assumption that scheduling will provide all possible outcomes.

When taking a component-based approach to specification and verification (e.g., see [GRBC<sup>+</sup>17], Section 6), we can replace an implementation with its specification when the implementation has been verified correct within the failures-divergences model. Taking this approach with a channel implementation while not considering the execution environment is a concern. Claiming part of a system meets its specification through using a CSP channel in place of a channel implementation can lead to certain behaviours not being possible. For example, in ProcessJ, we know that writing is prioritised during a channel communication, and likewise JCSP prioritises reading. Without sufficient resources, this will mean that certain traces are not possible in the implementation in a live environment (and thus the failures-divergence verification does not hold in all circumstances). Such unintended priorities will lead to certain behaviours not occurring, thus creating unintended starvation in the system. This starvation can possibly lead to deadlock for parts of the system which are unserved as some event traces are prioritised over the general set of all event traces.

With ProcessJ, we know explicitly the order that events will occur in a singularly scheduled system. Behaviour will be completely determined when the system is started—there will be only one trace of possible behaviour. A read-write handshake will always ensure the writing process will continue first (a fuller explanation is provided in our previous work [PC23]). Such assurance will mean any follow-on behaviour will be based on the writing process making progress before the reading process. Such behaviours will propagate through the system. As certain behaviours will thus be prioritised unintentionally, unexpected starvation can occur due to parts of the system not making fair progress as assumed in the specification.

We likewise believe such prioritisation will occur in a preemptively scheduled runtime. For example, in JCSP the reading process always exits the communication handshake first. Although preemptive scheduling means the writer will be allowed to make progress without having to wait for the reader to yield, the read process will have priority to progress immediately after the handshake. As the handshake involves a ping-pong of waits, notifies and locking, the reader leaves the handshake having only recently been allocated processor time, and is unlikely to be descheduled for the writing process. Likewise this priority will propagate through the system, and may lead to potential starvation scenarios. Preemptive scheduling is obviously less sensitive to such occurrences, and therefore the limitations of cooperative scheduling become clearer in a limited resource situation.

The potential issues we are highlighting do not lead directly to deadlock in a system. A system will still make progress, but only in a subset of any behaviour defined in a specification. Other implementations of CSP primitives have likewise been demonstrated to lack certain behaviours (e.g.,

1863 see Section 8.4 of [Cha09] demonstrating the limitations of JCSP's multiway sync [WBM<sup>+</sup>10]). We  
 1864 have implementations that trace refine specifications, but do not provide equivalent behaviours in a  
 1865 manner to ensure confidence in using pure CSP modelling for a system that relies on channel-based  
 1866 communication.

1867 Although our work has demonstrated a behaviour limitation in implementations of concur-  
 1868 rency primitives, limiting the assurance that equivalent behaviour is possible between model and  
 1869 implementation, it does not mean such runtimes cannot be used. Our modelling demonstrates  
 1870 that the number of schedulers must equal the number of processes in a system to achieve full  
 1871 failures/divergence refinement in both directions. However, this will not be the case in a real system.  
 1872 A real system will only have a subset of its processes ready to run at any one time, and if all these  
 1873 processes can be serviced, we will reach specification equivalence.

1874 On a single processor/threaded system, such as ProcessJ's current implementation, we will have  
 1875 limited behaviour to the point of potential singular trace behaviour. However, as the number of  
 1876 processors is increased, the more specification traces are possible. As the number of processors  
 1877 approaches the number of potentially ready processes, the specification and implementation trace  
 1878 sets will become closer to equivalent. When the number of processors equals the number of  
 1879 potentially ready processes, all system traces will be available. An interesting question we plan to  
 1880 explore is the potential relationship between the amount of resources available and the completeness  
 1881 of the implementation traces compared to the specification.

1882 Given our discussion, we have demonstrated why we needed to create a suitable multithreaded  
 1883 scheduling system for ProcessJ. The current implementation is singular in its behaviour, thus  
 1884 negating many of the potential behaviours we may expect when designing concurrent systems. By  
 1885 introducing a multithreaded scheduler, we will enable a larger set of potential trace behaviours  
 1886 to occur. This work has clearly demonstrated this need which was not apparent in the standard  
 1887 verification approach of the ProcessJ concurrency primitives.

## 1888 9.1 Conclusion

1889 We now have verified the implementation of both channel communication and choice in ProcessJ  
 1890 and we have shown that with enough schedulers ( $\geq 2$  for channel communication,  $\geq 3$  for a two-  
 1891 branch choice, and  $\geq 4$  for a three-branch choice), a cooperatively scheduled system (a *restricted*  
 1892 system) behaves **exactly** like an *unrestricted* non-scheduled one. We have *failures/divergence* results  
 1893 for all checks.  
 1894

1895 In the situation where the number of schedulers is less than the number of processes, we have  
 1896 shown that the system is deadlock free and will behave as specified. The lack of F/FD results in these  
 1897 cases have no impact on the correctness (i.e., the ability of the implementation to meet part of the  
 1898 specification behaviour) but are solely caused by execution behaviours imposed by the order of the  
 1899 processes in the run queue. It may slow down the execution compared to a system with an ample  
 1900 amount of schedulers as some processes need to cycle the run queue as they may be ahead of other  
 1901 processes whose prior execution is required in order to make progress. For example, with a single  
 1902 scheduler and a channel communication in which the reader is ahead of the writer in the run queue,  
 1903 the dequeuing sequence for a complete communication to happen will be: reader, writer, reader,  
 1904 writer, reader. For the same system but with the writer first in the queue, the dequeuing sequence is:  
 1905 writer, reader, writer, reader. As we can see, the first sequence has an additional dequeue of a reader  
 1906 at the beginning which eventually enqueued the reader to the back of the queue, and execution can  
 1907 progress as in the second example where the writer was first.

1908 The execution environment is important. This has been demonstrated across two papers: in the  
 1909 first paper we did not get the results that we expected, but it was not because the implementation of  
 1910 the channel communication or choice was incorrect (we verified them using the standard approach),  
 1911

1912 but simply that the execution environment was not suited to support the expected behaviour.  
 1913 Therefore, in this paper, we made changes to the execution environment in order to bring the  
 1914 results of the verification of the channel communication and choice into line with what we expected  
 1915 from the standard approach. This strongly supports our assertion that verification should never be  
 1916 done without taking the execution environment into account.

1917 This is also demonstrated by the results when there are not enough resources to achieve maximal  
 1918 concurrency. This results in the specification not being fully met. Whether or not this is a problem  
 1919 must be further investigated, but it may leave a concern that certain behaviours in the abstract  
 1920 system may not happen in the concrete system. When considering the concrete system interacting  
 1921 with other components of a larger system, we may have issues that verification against the abstract  
 1922 system may not identify. When creating concurrency primitives, such as a channel, and arguing  
 1923 behaviours are equivalent to a formal specification, significant limitations may result.

## 1924 10 Future Work

1926 The first step is to implement the results from this paper in the actual ProcessJ runtime/scheduler.  
 1927 The only challenging issue is the queue which must be blocking (without spinning) on the dequeue  
 1928 operation. However, the *java.util.concurrent* package contains a *LinkedBlockingQueue* class that  
 1929 serves exactly the purpose we need. See Appendix A for a prototype Java outline for how we plan  
 1930 to implement the new scheduling system in ProcessJ.

1931 Clearly, one of the design goals of a concurrent system is speed and efficiency, not only correctness.  
 1932 However, a speedy system that is not correct is of no use to anyone. If we study our solution closely,  
 1933 we see that a number of locks are required to ensure safe and correct operation. There are locks on:

- 1934 • An entire process; this is necessary to correctly handle setting the *ready* flag and placing
- 1935 processes in the run queue.
- 1936 • An entire channel; this is necessary in order to avoid race conditions when reading and
- 1937 writing to the channel.

1938 Inevitably, locking forces less concurrency, thus potentially making a system slower. The next  
 1939 step in developing the ProcessJ scheduler and runtime will therefore be to reduce the need and  
 1940 dependence on locking objects and data structures as far as possible. We believe that we can remove  
 1941 most of the locking we currently have by making use of the following two techniques:

- 1942 • Atomic variables: atomic variables (fields) support compare-and-swap. That is, a compare-
- 1943 and-swap operation allows for comparing the variable ( $v$ ) to a value ( $m$ ) and, if the compar-
- 1944 ison succeeds, set  $v$  to a new value ( $n$ ).
- 1945 • Concurrent queue: a concurrent queue allows multiple processes to access it safely at the
- 1946 same time, but which also must be blocking when the queue is empty.

1948 We believe that all the primitive fields (e.g., channel *data* store, process *ready* field etc.) can be  
 1949 correctly implemented using atomics.

1950 Naturally, to support these claims, we need to develop CSP models for both atomic variables and  
 1951 a concurrent queue—this will be the next step in this story.

1952 In addition, we also want to verify the correctness of shared channels in ProcessJ.

1953 Our main argument is that to truly verify a concurrency runtime, the execution environment  
 1954 must be modelled. In the introduction, we stated two choices were possible when an issue in the  
 1955 model was discovered:

- 1956 (1) Change the implementation of the synchronisation primitives to work better in the execution
- 1957 environment.
- 1958 (2) Change the execution environment so it worked correctly with the synchronisation primi-
- 1959 tives.

To fully explore these issues, a preemptive scheduling approach in CSP would need to be developed, and then used to verify JCSP and other concurrency runtimes. We believe the spurious wakeup problem would lead to further consideration on how to ensure correctness of the JCSP runtime, which would likely mean updates to the synchronisation primitives as the environment would need to be changed. Or, it could mean the spurious wakeup problem could be solved using the execution environment model to discover a solution.

## Acknowledgement

We would like to thank the reviewers for their helpful feedback, especially for the comments focusing on improving and optimizing the CSP.

## References

- [Cha09] Kevin Chalmers. *Investigating Communicating Sequential Processes for Java to Support Ubiquitous Computing*. PhD thesis, Edinburgh Napier University, 2009.
- [Cis18] Benjamin Cisneros. *ProcessJ: The JVM CSP Code Generator*. Master’s thesis, University of Nevada Las Vegas, 2018.
- [CR22] Milind Chabbi and Murali Krishna Ramanathan. A Study of Real-World Data Races in Golang. In *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation, PLDI 2022*, page 474–489, New York, NY, USA, 2022. Association for Computing Machinery.
- [CT09] Francesco Cesarini and Simon Thompson. *Erlang programming: a concurrent approach to software development*. O’Reilly Media, Inc., 2009.
- [CY23] Yi-An Chen and Yi-Ping You. Structured concurrency: A review. In *Workshop Proceedings of the 51st International Conference on Parallel Processing, ICPP Workshops ’22*, New York, NY, USA, 2023. Association for Computing Machinery.
- [FDR] FDR 4.2.7 Documentation. <https://cocotec.io/fdr/manual>. Accessed: 2023-08-23.
- [GRABR14] Thomas Gibson-Robinson, Philip Armstrong, Alexandre Boulgakov, and A. W. Roscoe. FDR3 — A Modern Refinement Checker for CSP. In Erika Ábrahám and Klaus Havelund, editors, *Tools and Algorithms for the Construction and Analysis of Systems*, volume 8413 of *Lecture Notes in Computer Science*, pages 187–201, 2014.
- [GRBC+17] Thomas Gibson-Robinson, Guy Broadfoot, Gustavo Carvalho, Philippa Hopcroft, Gavin Lowe, Sidney Nogueira, Colin O’Halloran, and Augusto Sampaio. *FDR: From Theory to Industrial Application*. Springer International Publishing, 2017.
- [Hoa78] C. A. R. Hoare. Communicating Sequential Processes. *CACM*, 21(8):666–677, August 1978.
- [Hoa85] C. A. R. Hoare. *Communicating Sequential Processes*. Prentice-Hall, 1985.
- [Lam87] Leslie Lamport. A fast mutual exclusion algorithm. *ACM Transactions on Computer Systems*, 5(1):1–11, January 1987.
- [Low19] Gavin Lowe. Discovering and correcting a deadlock in a channel implementation. *Formal Aspects of Computing*, 31(4):411–419, August 2019.
- [MI09] Ana Lúcia De Moura and Roberto Ierusalimsky. Revisiting coroutines. *ACM Trans. Program. Lang. Syst.*, 31(2), February 2009.
- [MS95] Maged M. Michael and Michael L. Scott. Simple, fast, and practical non-blocking and blocking concurrent queue algorithms. Technical report, Rochester University, Department of Computer Science, 1995.
- [Ora] Oracle. Oracle Help Center. <https://docs.oracle.com/javase/10/docs/api/java/util/concurrent/locks/Condition.html>. Accessed: 2022-12-09.
- [PC19] Jan B. Pedersen and Kevin Chalmers. Verifying Channel Communication Correctness for a Multi-core Cooperatively Scheduled Runtime Using CSP. In *Proceedings of the 7th International Workshop on Formal Methods in Software Engineering, FormaliSE ’19*, pages 65–74, Piscataway, NJ, USA, 2019. IEEE Press.
- [PC23] Jan B. Pedersen and Kevin Chalmers. Towards Verifying Cooperatively-Scheduled Runtimes using CSP. *Formal Aspects of Computing*, 1(1), January 2023.
- [PW18] Jan B. Pedersen and Peter H. Welch. The Symbiosis of Concurrency and Verification: Teaching and Case Studies. *Formal Aspects of Computing*, 30(2):239–277, March 2018.
- [Reg02] J. Regehr. Scheduling tasks with mixed preemption relations for robustness to timing faults. In *23rd IEEE Real-Time Systems Symposium, 2002. RTSS 2002.*, pages 315–326, 2002.
- [Ros98] A. W. Roscoe. *The theory and practice of concurrency*. Prentice Hall, 1998. The text book teaching material can be found at <http://www.comlab.ox.ac.uk/publications/books/concurrency/>.
- [Ros10] A. W. Roscoe. *Understanding Concurrent Systems*. Springer, 2010.

- 2010 [RSB09] Carl G. Ritson, Adam T. Sampson, and Frederick R. M. Barnes. Multicore scheduling for lightweight communi-  
 2011 cating processes. In John Field and Vasco T. Vasconcelos, editors, *Coordination Models and Languages*, pages  
 2012 163–183, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg.
- 2013 [SP16] Cable Shrestha and Jan B. Pedersen. JVMCSP—Approaching Billions of Processes on a Single-Core JVM. In  
 2014 Kevin Chalmers, Jan B. Pedersen, Jan F. Broenink, Brian Vinter, and Peter H. Welch, editors, *Proceedings of*  
 2015 *Communicating Process Architecture 2016*, Amsterdam, The Netherlands, August 2016. IOS Press.
- 2016 [Suf08] Bernard Sufrin. Communicating Scala Objects. In *CPA*, volume 66, pages 35–54, 2008.
- 2017 [Tis00] Christian Tismer. Continuations and stackless Python. In *Proceedings of the 8th international python conference*,  
 2018 volume 1, 2000.
- 2019 [WBM<sup>+</sup>10] Peter Welch, Neil Brown, James Moores, Kevin Chalmers, and Bernhard Spth. Alting barriers: synchronisation  
 2020 with choice in Java using JCSP. *Concurrency and Computation: Practice and Experience*, 22(8):1049–1062, 2010.
- 2021 [WM00] Peter H. Welch and Jeremy M. R. Martin. Formal Analysis of Concurrent Java Systems. In Peter H. Welch  
 2022 and André W. P. Bakkers, editors, *Communicating Process Architectures 2000*, volume 66, WoTUG-31, pages  
 2023 275–301, Amsterdam, The Netherlands, sep 2000. IOS Press. ISBN: 978-1-58603-907-3.

## 2023 A Java Code

2024 In this section we present real Java code implemented using the verified CSP. The *Process* class is  
 2025 basically the *PjProcess* class from the *ProcessJ* runtime. Note, the use of monitors in the CSP code  
 2026 is mirrored by declaring methods *synchronized* and by using *synchronized* blocks. Additionally, the  
 2027 *LinkedBlockingQueue* is thread safe.

### 2028 A.1 SchedulerManager

```

2029 import java.util.concurrent.*;
2030 import java.util.*;
2031
2032 class SchedulerManager {
2033
2034     // This is the process queue - only processes that are ready are in this queue.
2035     static BlockingQueue <Process> queue = new LinkedBlockingQueue <>();
2036
2037     // This is a hash set of all the schedulers
2038     static HashSet <Thread> schedulers = new HashSet <Thread>();
2039
2040     // Number of processJ processes on the heap.
2041     static int processCount = 0;
2042
2043     // This method adds a process to the queue.
2044     public static synchronized void schedule(Process p) {
2045         try {
2046             queue.put(p);
2047         } catch (Exception e) {}
2048     }
2049
2050     // Increment the number of processes.
2051     public static synchronized void processCountInc() {
2052         processCount++;
2053     }
2054
2055     // Decrement the number of processes.
2056     // If all processes have terminated, terminate all the schedulers
  
```

```

2059 public static synchronized void processCountDec() {
2060     processCount--;
2061     if (processCount == 0)
2062         new Thread() {
2063             public void run() {
2064                 SchedulerManager.shutdown();
2065             }
2066         }.start();
2067 }
2068
2069 // Create a new Scheduler
2070 public static void newScheduler() {
2071     Scheduler s = new Scheduler();
2072     schedulers.add(s);
2073     s.start();
2074 }
2075
2076 // Shut down all the active scheduler.
2077 public static void shutdown() {
2078     for (Thread t : schedulers)
2079         t.interrupt();
2080 }
2081 }

```

## 2083 A.2 Scheduler

```

2084 class Scheduler extends Thread {
2085     public void run() {
2086         while (true) {
2087             Process p = null;
2088             try {
2089                 // get a new process to run
2090                 p = Manager.queue.take();
2091             } catch (InterruptedException ie) {
2092                 return;
2093             }
2094             // run the process
2095             p.run();
2096         }
2097     }
2098 }
2099 }
2100
2101

```

## 2102 A.3 Process (PjProcess)

```

2103 class Process {
2104     // Ready field
2105     private boolean ready = true;
2106
2107

```

```

2108 // Running field
2109 private boolean running = false;
2110
2111 // Method for scheduling a process
2112 public synchronized void schedule() {
2113     if (!ready) {
2114         ready = true;
2115         if (!running)
2116             ScheduleManager.schedule(this);
2117     }
2118 }
2119
2120 // Method for descheduling a process
2121 public synchronized void deschedule() {
2122     running = false;
2123     if (ready)
2124         ScheduleManager.schedule(this);
2125 }
2126
2127 // Constructor
2128 Process(int pid) {
2129     this.ready = true;
2130     ScheduleManager.processCountInc();
2131     ScheduleManager.schedule(this);
2132 }
2133 }
2134
2135

```

#### A.4 An Application Process

The class *Foo* below is a (somewhat abstract) example of a compiled *ProcessJ* procedure.

```

2138 class Foo extends Process {
2139     public void run() {
2140         // Processes always set themselves to running at the start and after each yield point.
2141         // Since every re-activation starts here setting running to true here is enough.
2142         synchronized (this) {
2143             running = true;
2144         }
2145
2146         // This is the switch that jumps to the appropriate place in the application code below.
2147         switch(runLabel) {
2148             ...
2149         }
2150
2151         // Application code
2152         ...
2153
2154         // At every yield point:
2155         this.deschedule();
2156

```

```

2157
2158     // At termination:
2159     ScheduleManager.processCountDec();
2160     return;
2161
2162     }
2163 }
2164
2165

```

## 2166 B JCSP Read and Write Code

2167 In this section we present the relevant JCSP code for reading and writing. The complete code can  
 2168 be found at:

2169 [github.com/JonKerridge/jcsp/blob/master/src/main/java/jcsp/lang/One2OneChannelImpl.java](https://github.com/JonKerridge/jcsp/blob/master/src/main/java/jcsp/lang/One2OneChannelImpl.java)

2170 First the code for writing:

```

2171
2172     synchronized (rwMonitor) {
2173         hold = value;
2174         if (empty)
2175             empty = false;
2176         else {
2177             empty = true;
2178             rwMonitor.notify();
2179         }
2180         rwMonitor.wait();
2181     }
2182

```

2183 and now the code for reading:

```

2184
2185     synchronized (rwMonitor) {
2186         if (empty) {
2187             empty = false;
2188             rwMonitor.wait();
2189         } else
2190             empty = true;
2191         rwMonitor.notify();
2192         return hold;
2193     }
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205

```