

Formal Concurrency

In this handout we take a more formal view of concurrency than in handout 14. Our goal is to be able to prove that a general concurrent program implements a spec.

We begin with a fairly precise account of the non-atomic semantics of Spec, though our treatment is less formally the one for atomic semantics in handout 9. Next we explain the general method for making large atomic actions out of small ones (easy concurrency) and prove its correctness. We continue with a number of examples of concurrency, both easy and hard: mutexes, condition variables, read-write locks, buffers, and non-atomic clocks. Finally, we give fairly careful proofs of correctness for some of the examples.

In this handout we will usually use ‘code’ as a synonym for ‘implementation’.

Non-atomic semantics of Spec

We have already seen that a Spec module is a way of defining an automaton, or state machine, with transitions corresponding to the invocations of external atomic procedures. This view is sufficient if we only have functions and atomic procedures, but when we consider concurrency we need to extend it to include internal transitions. To properly model crashes, we introduced the idea of atomic commands that may not be interrupted. We did this informally, however, and since a crash “kills” any active procedure, we did not have to describe the possible behaviors when two or more procedures are invoked and running concurrently. This handout describes the concurrent semantics of Spec.

The most general way to describe a concurrent system is as a collection of independent atomic actions that share a collection of variables. If the actions are A_1, \dots, A_n then the entire system is just the ‘or’ of all these actions: $A_1 \parallel \dots \parallel A_n$. In general only some of the actions will be enabled, but for each transition the system non-deterministically chooses an action to execute from all the enabled actions.

Usually, however, we find it convenient to carry over into the concurrent world as much of the framework of sequential computing as possible. To this end, we model the computation as a set of tasks, processes, or *threads*, each of which executes a sequence of atomic actions; we denote threads by variables h, h' , etc. To define its sequence, each thread has a state variable called its ‘program counter’ $\$pc$, and each of its actions has a guard of the form $(h.\$pc = \alpha) \Rightarrow c$, so

that c can only execute when h 's program counter equals α . Each action advances the program counter with an assignment of the form $h.\$pc := \beta$, thus enabling the next action in sequence. We now explain how to use this view to understand the non-atomic semantics of Spec.

Non-atomic commands and threads

Unlike an atomic command, a non-atomic command cannot be described simply as a relation between states and outcomes, that is, an atomic transition. Rather, a non-atomic command corresponds to a sequence of atomic transitions, which may be interleaved with the sequences of other commands executing concurrently. To describe this interleaving, we use *labels* and *program counters*. We also need to distinguish the various *threads* of concurrent computation.

Intuitively, threads represent sequential processes. Roughly speaking, each point in the program between two atomic commands is assigned a label. Each thread's program counter $\$pc$ takes a label as its value,¹ indicating where that thread is in the program, that is, what command it is going to execute next.

Spec threads make all the concurrency explicit at the top level of each module. A thread is syntactically much like a procedure, but instead of being invoked by a client or by another procedure, it is automatically invoked in parallel initially, for every possible value of its arguments.² When it executes a `RET` (or reaches the end of its body), a thread simply makes no more transitions. However, threads are often written to loop indefinitely.

A thread is named by the name in the declaration and the argument values. Thus, the threads declared by `THREAD Foo(bool) = ...`, for example, would be named `Foo(true)` and `Foo(false)`. The names of local variables are qualified by both the name of the thread that is the root of the call stack, and by the name of the procedure invoked.³ In other words, each procedure in each thread has its own set of local variables. So for example, the local variable `p` in the `Sieve` example below appears in the state as `Sieve(0).p`, `Sieve(1).p`, ... If there were a `PROC Foo` called by `Sieve` with a local variable `baz`, the state might be defined at `Sieve(0).Foo.baz`, `Sieve(1).Foo.baz`, ... The pseudo-names `$a`, `$x`, and `$pc` are qualified only by the thread.

Spec does not have `COBEGIN` or `FORK` constructs, as many programming languages do, since these are considerably more difficult to define precisely, being tangled up with the control structure of the program. Also, because one Spec thread starts up for every possible argument of thread declaration, they tend to be more convenient for most of the specifications and implementations in this course. To keep the thread from doing anything until a certain point in the computation, use an initial guard for the entire body.

Each atomic command defines a transition, just as in the sequential semantics. However, now a transition is enabled by the program counter value. That is, a transition can only occur if the

¹ The variables declared by a program may not have labels as their values, hence there is no `Label` type.

² This differs from the threads in Java, in Modula 3, or in many C implementations. These languages start a computation with one thread and allow it to create and destroy threads dynamically using `fork` and `join` operations.

³ This works for non-recursive procedures. To accommodate recursive procedures, the state must involve something equivalent to a stack. Probably the simplest solution is to augment the state space by tacking on the nesting depth of the procedure to all the names and program counter values defined above. For example, `h + ".P.v"` becomes `h + ".P.v" + d.enc`, for every positive integer `d`. An invocation transition at depth `d` goes to depth `d+1`.

program counter of some thread equals the label before the command, and the transition sets the program counter of that thread to the label after the command. If the command at the label in the program counter fails, the thread is “stuck”, and it does not make any transitions. However, it may become unstuck later, because of the transitions of some other threads. Thus, a command failing does not necessarily (or even usually) mean that the thread fails.

We won’t give the non-atomic semantics precisely here as we did for the atomic semantics in handout 9, since it involves a number of fussy details that don’t add much insight. It’s important, however, to understand the basic ideas.

- Each atomic command in a thread or procedure defines a transition (atomic procedures and functions are taken care of by the atomic semantics).
- The program counters enable transitions: a transition can only occur if the program counter for some thread equals the label before the command, and the transition sets that program counter to the label after the command.

Thus the state of the system is the global state plus the state of all the threads. The state of a thread is its $\$pc$, $\$a$, and $\$x$ values, the local variables of the thread, and the local variables of each procedure that has been called by the thread and has not yet returned.

Suppose the label before the command c is α and the one after the command is β , and the transition function defined by $MC(c)$ is $(\ \ s, \ o \ | \ rel)$. Then if c is in thread h , its transition function is

$$(\ \ s, \ o \ | \ s(h + ".\$pc") = \alpha \ \wedge \ o(h + ".\$pc") = \beta \ \wedge \ rel')$$

If c is in procedure P , that is, c can execute for any thread whose program counter has reached α , its transition function is

$$(\ \ s, \ o \ | \ (\text{EXISTS } h: \text{Thread} \ | \ s(h + ".P.\$pc") = \alpha \ \wedge \ o(h + ".P.\$pc") = \beta \ \wedge \ rel'))$$

Here rel' is rel with each reference to a local variable v changed to $h + ".v"$ or $h + ".P.v"$.

Labels in Spec

What are the atomic transitions in a Spec program? In other words, where do we put the labels? The basic idea is to build in as little atomicity as possible (since you can always put in what you need with $\langle\langle \dots \rangle\rangle$). However, expression evaluation must be atomic, or reasoning about expressions would be a mess. To model code in which expression evaluation is not atomic, you must add temporary variables. Thus $x := a + b + c$ becomes

```
VAR t | << t := a >>; << t := t + b >>; << x := t + c >>
```

The simple commands, SKIP, HAVOC, RET, and RAISE, are atomic, as is any command in atomicity brackets $\langle\langle \dots \rangle\rangle$.

For an invocation, there is a transition to evaluate the argument and set the $\$a$ variable, and one to send control to the start of the body. The RET command’s transition sets $\$a$ and leaves control at the end of the body. The next transition leaves control after the invocation. So every procedure invocation has at least four transitions: evaluate the argument and set $\$a$, send control to the body, do the RET and set $\$a$, and return control from the body. The reason for these fussy details is to ensure that the invocation of an external procedure has start and end transitions that do not

change any other state. These are the transitions that appear in the trace and therefore must be identical in an implementation and a spec.

Minimizing atomicity means that an assignment is broken into separate transitions, one to evaluate the right hand side and one to change the left hand variable. This also has the advantage of consistency with the case where the right hand side is a non-atomic procedure invocation.

Each transition happens atomically, even if the variable is “big”. Thus `x := exp` is

```
VAR t | << t := exp >> ; << x := t >>
```

and `x := p(y)` is

```
p(y); << x := $a >>
```

Since there are no additional labels for the `c1 [] c2` command, the initial transition of the compound command is enabled exactly when the initial transition of either of the subcommands is enabled (or if they both are). Thus, the choice is made based only on the first transition. After that, the thread may get stuck in one branch (though, of course, some other thread might unstick it later). The same is true for `[*]`, except that the initial transition for `c1 [*] c2` can only be the initial transition of `c2` if the initial transition of `c1` is not enabled. And the same is also true for `VAR`. The value chosen for `id` in `VAR id | c` must allow `c` to make at least one transition; after that the thread might get stuck.

`DO` has a label, but `OD` does not introduce any additional labels. The starting and ending program counter value for `c` in `DO c OD` is the label on the `DO`. Thus, the initial transition of `c` is enabled when the program counter is the label on the `DO`, and the last transition sets the program counter back to that label. When `c` fails, the program counter is set to the label following the `OD`.

To sum up, there's a label on each `:=` (as we just discussed), `=>`, `';`', and `EXCEPT` outside of `<< . . . >>`. There is never any label inside atomicity brackets. It's convenient to write the labels as subscripts on these symbols. There's also a label at the start of a procedure, which we write on the `=` of the declaration, and a label at the end. There is one label for a procedure invocation, after the argument is evaluated; we write it just before the closing `'`'. After the invocation is complete, the PC goes to the next label after the invocation.

As a consequence of this labeling, as we said before, a procedure invocation has one transition to evaluate the argument expression and a second to set the program counter to the label immediately before the procedure body. There is also a transition for every transition of the procedure body, using the labels of the procedure body. The `RET` command in the invoked procedure sets the program counter after the body, and the final transition sets it to the label immediately following the invocation; if the invocation is the right side of an assignment, this is the label on the `:=`.

Here is a meaningless example, just to show where the labels go.

```
PROC P() = [P1] VAR x, y |  
  IF x > 5 => [P2] x := [P4] Q(x + 1, 2 [P3]); [P5] y := [P6] 3  
  [] << y := 4 >>  
  FI; [P7]  
  VAR z | DO [P8] << P() >> OD [P9]
```

External actions

In sequential Spec a module has only external actions; each invocation of a function or atomic procedure is an external action. In concurrent Spec there are two differences.

1. There are internal actions. These can be actions of an externally invoked `PROC` or actions of a thread declared and executing in the module.
2. There are two external actions in the external invocation of a (non-atomic) procedure: the call, which sends control from after evaluation of the argument to the entry point of the procedure, and the return, which sends control from after the `RET` command in the procedure to just after the invocation in the caller. These external transitions do not affect the `$a` variable that communicates the argument and result values. That variable is set by the internal transitions that compute the argument and do the `RET` command.

There's another style of defining external interfaces in which every external action is an `APROC`. If you want to get the effect of a non-atomic procedure, you have to break it into two `APROC`'s, one that delivers the arguments and sets the internal state so that internal actions will do the work of the procedure body, and a second that retrieves the result. We will not use this style.

Examples

Here are two Spec programs that search for prime numbers and collect the result in a set `primes`; both omit the even numbers, initializing `primes` to `{2}`. Both are based on the *sieve of Erasthsthenes*, testing each prime less than $n^{1/2}$ to see whether it divides n . Since the threads may not be synchronized, we must ensure that all the numbers $\leq n^{1/2}$ have been checked before we check n .

The first example is more like a spec, using an infinite number of threads, one for every odd number.

```
TYPE Odds          = Int SUCHTHAT (\i: Int | i // 2 = 1 /\ i > 1 )
VAR primes         : SET Int := {2}
    done           : SET Int := {}

THREAD Sieve1(n: Odds) =
  (ALL i: Odds | i <= Sqrt(n) ==> i IN done) =>      % Wait for possible factors
  IF (ALL p :IN primes | p <= Sqrt(n) ==> n // p # 0) =>
    << primes := primes + {n} >>
  [*] SKIP
  FI;
  << done := done + {n} >>

FUNC Sqrt(n: Int) -> Int = RET { i: Int | i*i <= n }.max
```

The second example, on the other hand, is closer to code, running ten parallel searches. Although there is one thread for every integer, only threads `Sieve(0)`, `Sieve(1)`, ..., `Sieve(9)` are “active”. Differences from `seive1` are boxed.

```
VAR primes         : SET Int := {2}
    nThreads       := 10
    next           := 0 .. nThreads-1
```

```

THREAD Sieve2(i: Int) = next!i =>
  next(i) := 2*i + 3;
  DO VAR n: Int := next(i) |
    (ALL j :IN next.rng | j >= Sqrt(n)) =>
      IF (ALL p :IN primes | p <= Sqrt(n) ==> n // p # 0) =>
        << primes := primes + {n} >>
        [*] SKIP
      FI;
  next(i) := n + 2*nThreads

```

□

Big atomic actions

As we saw earlier, we need atomic actions for practical, easy concurrency. Spec lets you specify any grain of atomicity in the program just by writing `<< ... >>` brackets.⁴ It doesn't tell you where to write the brackets. If the environment in which the program has to run doesn't impose any constraints, it's usually best to make the atomic actions as big as possible, because big atomic actions are easier to reason about. But often big atomic actions are too hard or too expensive to implement, and we have to make do with small ones. For example, in a shared-memory multiprocessor typically only the individual instructions are atomic. So we are faced with the problem of showing that code with small atomic actions satisfies a spec with bigger ones.

The idea

The standard way of doing this is by some kind of 'non-interference'. The idea is based on the following observation. Suppose we have a program with a thread *h* that contains the sequence

$$A; B \tag{1}$$

as well as an arbitrary collection of other commands. We denote the program counter value at the semi-colon by β . We are thinking of the program as

$$A; h.\$pc = \beta \Rightarrow B [] C_1 [] C_2 [] \dots$$

where each command has an appropriate guard that enables it only when the program counter for its thread has the right value. We have written the guard for *B* explicitly.

Suppose *B* denotes an arbitrary atomic command, and *A* denotes an atomic command that commutes with every command in the program (other than *B*) that is enabled when *h* is at the semicolon, that is, when $h.\$pc = \beta$. (We give a precise meaning for 'commutes' below.) In addition, both *A* and *B* have only internal actions. Then it's intuitively clear that the program with (1) simulates a program with the same commands except that instead of (1) it has

$$\langle\langle A; B \rangle\rangle \tag{2}$$

Informally this is true because any *C*'s that happen between *A* and *B* have the same effect on the state that they would have if they happened before *A*, since they all commute with *A*.

Thus we have achieved the goal of making a bigger atomic command `<< A; B >>` out of two small ones *A* and *B*. We can call the big command *D* and repeat the process on *E*; *D* to get a still

⁴ As we have seen, Spec does treat expression evaluation as atomic. If you are dealing with an environment in which an expression like $x(i) + f(y)$ can't be evaluated atomically, you should write it as `VAR t1, t2 | t1 := x(i); t2 := f(y); ... t1 + t2 ...`

bigger command $\langle\langle E; A; B \rangle\rangle$. The same thing also works in the other order, for $B; A$ and $\langle\langle B; A \rangle\rangle$.

How do we ensure that only a command C that commutes with A can execute while $h.pc = \beta$? The simplest way is to ensure that the variables that A touches are disjoint from the variables that C writes, and vice versa; then they will surely commute. Two such commands are called ‘non-interfering’. There are two standard ways to show that commands are non-interfering. One is that A touches only local variables of h . Only actions of h touch local variables of h , and the only action of h that is enabled when $h.pc = \beta$ is B . So any sequence of commands that touch only local variables is atomic, and if it is preceded or followed by a single arbitrary atomic command the whole thing is still atomic.

The other easy case is a critical section protected by a mutex. Recall that a critical section for v is a command with the property that if some thread’s PC is in the command, then no other thread’s PC can be in any critical section for v . If the only commands that touch v are in critical sections for v , then we know that only B and commands that don’t touch v can execute while $h.pc = \beta$. So if every command in any critical section for v only touches v (and perhaps local variables), then the program simulates another program in which every critical section is an atomic command. A critical section is usually implemented by acquiring a *lock* or *mutex* and holding it for the duration of the section. The property of a lock is that it’s not possible to acquire it when it is already held, and this ensures the mutual exclusion property for critical sections.

In fact, reader/writer locks are sufficient for non-interference, because read operations all commute with each other. See handout 14 for more details on different kinds of locks.

Another important case is mutex acquire and release operations. These operations only touch the mutex, so they commute with everything else. What about these operations on the same mutex in different threads? If both can execute, they certainly don’t yield the same result in either order; that is, they don’t commute. When can both operations execute? We have the following cases (writing the executing thread as an explicit argument of each operation):

A	C	Possible sequence?
$m.acq(h)$	$m.acq(h')$	No: C is blocked by h holding m
$m.acq(h)$	$m.rel(h')$	No: C won’t be reached because h' doesn’t hold m
$m.rel(h)$	$m.acq(h')$	OK
$m.rel(h)$	$m.rel(h')$	No: one thread doesn’t hold m , hence won’t do rel

So $m.acq$ commutes with everything that’s enabled at β , since neither mutex operation is enabled at β in a program that avoids havoc. But $m.rel(h)$ doesn’t commute with $m.acq(h')$. The reason is that the $A; C$ sequence can happen, but the $C; A$ sequence $m.acq(h'); m.rel(h)$ cannot, because in this case h doesn’t hold m and therefore can’t be doing a rel . Hence it’s not possible to flip every C in front of $m.rel(h)$ in order to make $A; B$ atomic.

What does this mean? You can acquire more locks and still keep things atomic, but as soon as you release one, you no longer have atomicity.

Proofs

How can we make this precise and *prove* that a program containing (1) implements the same program with (2) replacing (1), using our trusty method of abstraction relations? For easy reference, we repeat (1) and (2).

A; B (1)

<< A; B >> (2)

As usual, we call the complete program containing (2) the spec S and the complete program containing (1) the code T . We need an abstraction relation AR between the states of T and the states of S under which every transition of T simulates a (possibly empty) trace of S . Note that the state spaces of T and S are the same, except that $h.\$pc$ can never be β in S . We use s and u for states of S and T , to avoid confusion with various other uses of t .


First we need a precise definition of “ C is enabled at β and commutes with A ”. For any command X , we write $u \boxtimes u'$ for $\text{MC}(X)(u, u')$, that is, if X relates u to u' . The idea of ‘commutes’ is that $\langle\langle A; C \rangle\rangle$ is the same as $\langle\langle C; A \rangle\rangle$, and the definition follows from the meaning of semicolon:

$$(\text{ALL } u_1, u_2 \mid (\text{EXISTS } u \mid u_1 \overset{A}{\boxtimes} u \wedge u \overset{C}{\boxtimes} u_2 \wedge u("h.\$pc") = \beta) \\ \implies (\text{EXISTS } u' \mid u_1 \overset{C}{\boxtimes} u' \wedge u' \overset{A}{\boxtimes} u_2))$$

It seems reasonable to do the proof by making A simulate the empty trace and B simulate $\langle\langle A; B \rangle\rangle$, since we know more about A than about B ; every other command simulates itself. So we make AR the identity everywhere except at β , where it relates any state that can be reached from s by A to s . This expresses the intention that at β we haven’t yet done A in S , but we have done A in T . (Since A may take many states to s , this can’t just be an abstraction function.) We write $s \sim u$ for “ AF maps u to s ”. Precisely, we say that $s \sim u$ if either $u("h.\$pc") \neq \beta$ and $s = u$, or $u("h.\$pc") = \beta$ and $s \overset{A}{\boxtimes} u$.


Why is this an abstraction relation? It certainly relates an initial state to an initial state, and it certainly works for any transition $u \rightarrow u'$ that stays away from β , that is, in which $u("h.\$pc") \neq \beta$ and $u'("h.\$pc") \neq \beta$, since the abstract and concrete states are the same. What about transitions that do involve β ?

1. If $h.\$pc$ changes to β then we must have executed A . The picture is




The abstract trace is empty, so the abstract state doesn’t change: $s = s' = u$. But we executed A , so $u \overset{A}{\boxtimes} u'$, so $s' \sim u'$ because of the equalities.

2. If $h.\$pc$ starts at β then the command must be either B or some C that commutes with A . If the command is B , then the picture is



To show the top relation, we have to show that there exists an s_0 such that $s \stackrel{A}{\sim} s_0$ and $s_0 \stackrel{B}{\sim} s'$, by the meaning of semicolon. But u has exactly this property, since $s' = u'$.

3. If the command is c , then the picture is



But this follows from the definition of ‘commutes’: we are given $s, u,$ and u' related as shown, and we need s' related as shown, which is just what the definition gives us, with $u_1 = s, u_2 = u',$ and $u' = s'$.

Examples of concurrency

This section contains a number of example specs and codes that illustrate various aspects of concurrency. The specs have large atomic actions that keep them simple. The codes have smaller atomic actions that reflect the realities of the machines on which they have to run. Some of the examples of code illustrate easy concurrency (that is, that use locks): `RWLockImpl` and `BufferImpl`. Others illustrate hard concurrency: `SpionLock`, `Mutex2Impl`, and `ClockImpl`.

Incrementing a register

The first example involves incrementing a register that has `Read` and `Write` operations. Here is the unsurprising spec of the register, which makes both operations atomic:

```

MODULE Register EXPORT Read, Write =
VAR x          : Int := 0
APROC Read() -> Int = << RET x >>
APROC Write(i: Int) = << x := i >>
END Register

```

To increment the register, we could use the following procedure:

```

PROC Increment() = VAR t: Int | t := Register.Read(); t := t + 1; Register.Write(t)

```

Suppose that, starting from the initial state where $x = 0$, n threads execute `Increase` in parallel. Then, depending on the interleaving of the low-level steps, the final value of the register could be anything from 1 to n . This is unlikely to be what was intended. Certainly this code doesn't implement the spec

```
PROC Increase() = << x := x + 1 >>
```

Suppose that we weaken our atomicity assumptions to say that the value of a register is represented as a sequence of bits, and that the only atomic operations are reading and writing individual bits. Now what are the possible final states if n threads execute `Increase` in parallel?

Alternatively, consider a new module `RWInc` that explicitly supports `Increase` operations in addition to `Read` and `Write`. This might add the following (exportable) procedure to the `Register` module:

```
PROC Increase() = x := x+1
```

Or, more explicitly:

```
PROC Increase() = VAR t: Int | << t := x >>; << x := t+1 >>
```

Because of the fine grain of atomicity, it is still true that if n threads execute `Increase` in parallel then, depending on the interleaving of the low-level steps, the final value of the register could be anything from 1 to n . Putting the procedure inside the `Register` module doesn't help.

Mutexes

Here is a spec of a simple `Mutex` module, which can be used to ensure mutually exclusive execution of critical sections; it is copied from handout 14. The state of a mutex is `nil` if the mutex is available, and otherwise is the thread that holds the mutex. The module maintains lots of mutexes in the state variable `s`, and the type `M` selects one of them.

```
MODULE Mutex EXPORT M, New =
TYPE M          = Int WITH {acq:=Acquire, rel:=Release}
VAR s           : M -> (Thread + Nil) := {}
% Each mutex is either nil or the thread holding the mutex.
% The variable SELF is defined to be the thread currently making a transition.
APROC New() -> M = << VAR m | ~s!m => s(m) := nil; RET m >>
APROC Acquire(m) = << s(m) = nil => s(m) := SELF; RET >>
APROC Release(m) = << s(m) = SELF => s(m) := nil ; RET [*] HAVOC >>
END Mutex
```

If a thread invokes `Acquire` when `m = nil`, then the body fails. This means that there's no possible transition for that thread, and the thread is blocked, waiting at this point until the guard becomes true. If many threads are blocked at this point, then when `m` is set to `nil`, one is scheduled first, and it sets `m` to itself atomically; the other threads are still blocked.

The spec says that if a thread that doesn't hold `m` does `m.rel`, the result is `HAVOC`. As usual, this means that the code is free to do anything when this happens. As we shall see in the `SpinLock` code below, one possible 'anything' is to free the mutex anyway.

Here is a simple use of a mutex `m` to make the `Increment` procedure atomic:

```
PROC Increment() = VAR t: Int |
    m.acq; t := Register.Read(); t := t + 1; Register.Write(t); m.rel
```

This keeps concurrent calls of `Increment` from interfering with each other. If there are other accesses to the register, they must also use the mutex to avoid interfering with threads executing `Increment`.

Spin locks

A simple way to implement a mutex is to use a *spin lock*. The name is derived from the behavior of a thread waiting to acquire the lock—it “spins”, repeatedly attempting to acquire the lock until it is finally successful.

Here is incorrect code:

```
MODULE BrokenSpinLock EXPORT M, New =
TYPE AH          = ENUM[available, held]
    M            = Int WITH {acq:=Acquire, rel:=Release}
VAR s           : M -> AH := {}

APROC New() -> M = << VAR m | ~ s!m => s(m) := available; RET m >>
PROC Acquire(m) = DO << s(m) = held => SKIP >> OD; << s(m) := held >>
PROC Release(m) = << s(m) := available >>

END BrokenSpinLock
```

This is wrong because two concurrent invocations of `Acquire` could both find `s(m) = available` and subsequently both set `s(m) := held` and return.

Here is correct code. It uses a more complex atomic command in the `Acquire` procedure. This command corresponds to the atomic “test-and-set” instruction provided by many real machines to implement locks. It records the initial value of the lock, and then sets it to `held`. Then it tests the initial value; if it was `available`, then this thread was successful in atomically changing the state of the lock from `available` to `held`. Otherwise some other thread must hold the lock, so we “spin”, repeatedly trying to acquire it until we succeed. The important difference in `SpinLock` is that the guard now involves only the local variable `t`, instead of the global variable `s(m)` in `Mutex`. A thread acquires the lock when it is the one that changes it from `available` to `held`, which it checks by testing the value returned by the test-and-set.

```

MODULE SpinLock EXPORT M, New =

TYPE AH          = ENUM[available, held]
  M              = Int WITH {acq:=Acquire, rel:=Release}
VAR s            : M -> AH := {}

APROC New() -> M = << VAR m | ~ s!m => s(m) := available; RET m >>

PROC Acquire(m) = VAR t: AH |
  DO << t := s(m); s(m) := held >>; IF t # held => RET [*] SKIP FI OD

PROC Release(m) = << s(m) := available >>

END SpinLock

```

Of course this implementation is not practical in general unless each thread has its own processor. Later, in `MutexImpl`, we present a practical implementation that queues a waiting thread.

The `SpinLock` code differs from the `Mutex` spec in another important way. It “forgets” which thread owns the mutex. The following `ForgetfulMutex` module is useful in understanding the `SpinLock` code—in `ForgetfulMutex`, the threads get forgotten, but the atomicity is the same as in `Mutex`.

```

MODULE ForgetfulMutex EXPORT M, New =

TYPE AH          = ENUM[available, held]
  M              = Int WITH {acq:=Acquire, rel:=Release}
VAR s            : M -> AH := {}

APROC New() -> M = << VAR m | ~ s!m => s(m) := available, RET m >>

PROC Acquire() = << s(m) = available => s(m) := held >>
PROC Release() = << s(m) := available >>

END ForgetfulMutex

```

Note that `ForgetfulMutex` releases a mutex regardless of which thread acquired it, and it does a `SKIP` if the mutex is already free. This is one of the behaviors permitted by the spec, which allows anything under these conditions.

Later we will show that `SpinLock` implements `ForgetfulMutex` and that `ForgetfulMutex` implements `Mutex`, from which it follows that `SpinLock` implements `Mutex`.

Read/write locks

Here is a spec of a module that provides locks with two modes, read and write, rather than the single mode of a mutex. Several threads can hold a lock in read mode, but only one thread can hold a lock in write mode, and no thread can hold a lock in read mode if some thread holds it in write mode. In other words, read locks can be shared, but write locks are exclusive; hence the locks are also known as ‘shared’ and ‘exclusive’.

```

MODULE RWLock EXPORT L, New =

TYPE L          = Int WITH { rAcq:=ReadAcquire, rRel:=ReadRelease,
                             wAcq:=WriteAcquire, wRel:=WriteRelease }
  ST            = SET Thread

VAR r           : L -> SET Thread := {}
  w             : L -> SET Thread := {}

APROC New() -> L = << VAR l | ~ r!l => r(l) := {}; w(l) := {}; RET l >>

APROC ReadAcquire(l) =
% Acquires read lock if there are no current write locks.
  << SELF IN r(l) + w(l) => HAVOC
  [*] w(l) = {} => r(l) := r(l) + {SELF} >>

APROC WriteAcquire(l) =
% Acquires write lock if there are no current locks.
  << SELF IN r(l) + w(l) => HAVOC
  [*] r(l) + w(l) = {} => w(l) := {SELF} >>

APROC ReadRelease(l) =
% Releases lock if the thread has it; otherwise anything can happen.
  << ~ (SELF IN r(l)) => HAVOC [*] r(l) := r(l) - {SELF} >>

APROC WriteRelease(l) =
  << ~ (SELF IN w(l)) => HAVOC [*] w(l) := {} >>

END RWLock

```

The following simple code is similar to `ForgetfulMutex`. It has the same atomicity as `RWLock`, but uses a different data structure to represent possession of the lock. Specifically, it uses a single integer variable `rw` to keep track of the number of readers (positive) or the existence of a writer (-1).

```

MODULE ForgetfulRWL EXPORT L, New =

TYPE L          = Int WITH { rAcq:=ReadAcquire, rRel:=ReadRelease,
                             wAcq:=WriteAcquire, wRel:=WriteRelease }

VAR rw          : L -> Int := {}
% >0 gives number of readers, 0 means free, -1 means one writer

APROC New() -> L = << VAR l | ~ rw(l) => rw(l) := {}, RET l >>

APROC ReadAcquire (l) = << rw(l) >= 0 => rw(l) := rw(l) + 1 >>
APROC WriteAcquire(l) = << rw(l) = 0 => rw(l) := -1 >>

APROC ReadRelease (l) = << rw(l) := rw(l) - 1 >>
APROC WriteRelease(l) = << rw(l) := 0 >>

END ForgetfulRWL

```

We will see later how to implement `ForgetfulRWL` using a mutex.

Condition variables

Mutexes are used to protect shared variables. Often a thread `h` cannot proceed until some condition is true of the shared variables, a condition produced by some other thread. Since the variables are protected by a lock, and can be changed only by the thread holding the lock, `h` has to

release the lock. It is not efficient to repeatedly release the lock and then re-acquire it to check the condition. Instead, it's better for `h` to *wait* on a *condition variable*. Whenever any thread changes the shared variables in such a way that the condition might become true, it *signals* the threads waiting on that variable. Sometimes we say that the waiting threads 'wake up' when they are signaled. Depending on the application, a thread may signal one or several of the waiting threads.

Here is the spec for condition variables, copied from handout 14.

```

MODULE Condition EXPORT C, New =

TYPE C          = Int WITH {wait:=Wait, signal:=Signal, broadcast:=Broadcast}
M              = Mutex.M
VAR s          : C -> SET Thread := {}
% Each condition variable is the set of waiting threads.

APROC New() -> C = << VAR c | ~s!c => s(c) := {}; RET c >>

PROC Wait(c, m) =
  << s(c) := s(c) + {SELF}; m.rel >>;           % HAVOC if SELF doesn't hold m
  << ~ (SELF IN s(c)) => m.acq >>

APROC Signal(c) = << VAR hs: SET Thread |
% Remove at least one thread from c. In practice, usually just one.
  IF hs <= s(c) /\ hs # {} => s(c) := s(c) - hs [*] SKIP FI >>

APROC Broadcast(c) = << s(c) := {} >>

END Condition

```

As we saw in handout 14, it's not necessary to have a single condition for each set of shared variables. We want enough condition variables so that we don't wake up too many threads whose conditions are not yet satisfied, but not so many that the cost of doing all the `signals` is excessive.

Implementing read/write lock using condition variables

This example shows how to use easy concurrency to make more complex locks and scheduling out of basic mutexes and conditions. We use a single mutex and condition for all the read-write locks here, but we could have separate ones for each read-write lock, or we could partition the locks into groups that share a mutex and condition. The choice depends on the amount of contention for the mutex.

Compare the code with `ForgetfulRWL`; the differences are highlighted with boxes. The `<<...>>` in `ForgetfulRWL` have become `m.acq ... m.rel`; this provides atomicity because shared variables are only touched while the lock is held. The other change is that each guard that could block (in this example, each one that doesn't have `[*]SKIP`) is replaced by a look that tests the guard and does `c.wait` if it doesn't hold. The release operations do the corresponding signal or broadcast operations.

```

MODULE RWLockImpl EXPORT L, New = % implements ForgetfulRWL

TYPE L = Int WITH { rAcq:=ReadAcquire , rRel:=ReadRelease,
                   wAcq:=WriteAcquire, wRel:=WriteRelease }

VAR rw : L -> Int := {}


|   |                    |
|---|--------------------|
| m | := Mutex.New()     |
| c | := Condition.New() |



% ABSTRACTION FUNCTION ForgetfulRWL.rw = rw

APROC New() -> L = << VAR l | ~ rw(l) => rw(l) := {}, RET l >>

PROC ReadAcquire (l) = 

|       |      |
|-------|------|
| m.acq | DO ~ |
|-------|------|

 rw >= 0 => 

|              |
|--------------|
| c.wait(m) OD |
|--------------|

 rw(l) := rw(l)+1; 

|       |
|-------|
| m.rel |
|-------|


PROC WriteAcquire(l) = 

|       |      |
|-------|------|
| m.acq | DO ~ |
|-------|------|

 rw = 0 => 

|              |
|--------------|
| c.wait(m) OD |
|--------------|

 rw(l) := -1; 

|       |
|-------|
| m.rel |
|-------|



PROC ReadRelease (l) =


|       |
|-------|
| m.acq |
|-------|

 rw(l) := rw(l)-1; IF rw(l) = 0 => c.signal [*] SKIP FI; m.rel
PROC WriteRelease(l) =


|       |
|-------|
| m.acq |
|-------|

 rw(l) := 0; 

|             |
|-------------|
| c.broadcast |
|-------------|

; 

|       |
|-------|
| m.rel |
|-------|



END RWLockImpl

```

This is the prototypical example for scheduling resources. There are mutexes (just *m* in this case) to protect the scheduling data structures, conditions (just *c* in this case) on which to delay threads that are waiting for a resource, and logic that figures out when it's all right to allocate a resource (the read or write lock in this case) to a thread.

A FIFO buffer

In this section, we give a spec and code of a simple unbounded buffer that could be used as a communication channel between two threads. This is the prototypical example of a *producer-consumer* relation between threads.

```

MODULE Buffer[T] EXPORT Produce, Consume =

VAR b : SEQ T := {}

APROC Produce(t) = << b := b + {t} >>
APROC Consume() -> T = VAR t | << b.size > 0 => t := b.head; b := b.tail; RET t >>

END Buffer

```

The code is another example of easy concurrency.

```

MODULE BufferImpl[T] EXPORT Produce, Consume =

VAR b : SEQ T := {}
    m := Mutex.New()
    c := Condition.New()

% ABSTRACTION FUNCTION Buffer.b = b

PROC Produce(t) = 

|       |
|-------|
| m.acq |
|-------|

 IF b = {} => c.signal [*] SKIP FI; b := b + {t}; 

|       |
|-------|
| m.rel |
|-------|


PROC Consume() -> T = VAR t |


|       |      |
|-------|------|
| m.acq | DO ~ |
|-------|------|

 b.size > 0 => 

|              |
|--------------|
| c.wait(m) OD |
|--------------|

 t := b.head; b := b.tail; 

|       |
|-------|
| m.rel |
|-------|

 RET t

END BufferImpl

```

Implementing `Mutex` with memory

The usual way to implement `Mutex` is to use an atomic test-and-set operation; we saw this in the `MutexImpl` module above. If such an operation is not available, however, it's possible to implement `Mutex` using only atomic read and write operations on memory. This requires an amount of storage linear in the number of threads, however. We give a fair algorithm due to Petersen⁵ for two threads; if thread `h` is competing for the mutex, we write `h*` for its competitor.

```
MODULE Mutex2Impl EXPORT Acquire, Release =  
  
VAR req          : Thread -> Bool := {* -> false}  
    lastReq      : Int  
  
PROC Acquire() =  
  [a1] req(SELF) := true;  
  [a12] lastReq := self;  
  DO [a2] (req(SELF*) /\ lastReq = SELF) => SKIP OD [a3]  
  
PROC Release() = req(self) := false  
  
END Mutex2Impl
```

This is hard concurrency, and it's tricky to show that it works. To see the idea, consider first a simpler version of `Acquire` that ensures mutual exclusion but can deadlock:

```
PROC Acquire0() =  
  [a1] req(SELF) := true;  
  DO [a2] req(SELF*) => SKIP OD [a3]
```

We get mutual exclusion because once `req(h)` is true, `h*` can't get from `a2` to `a3`. Thus `req(h)` acts as a lock that keeps the predicate `h*.pc = a2` true once it becomes true. Only one of the threads can get to `a3` and acquire the lock.

Of course, `Acquire0` is no good because it can deadlock—if both threads get to `a2` then neither can progress. `Acquire` avoids this problem by making it a little easier for a thread to progress: even if `req(h*)`, `h` can take `(a2, a3)` if `lastReq ≠ h`. Intuitively this maintains mutual exclusion because:

If both threads are at `a2`, only the one \neq `lastReq`, say `h`, can progress to `a3` and acquire the lock. Since `lastReq` won't change, `h*` will remain at `a2` until `h` releases the lock.

Once `h` has acquired the lock with `h*` not at `a2`, `h*` can only reach `a2` by setting `lastReq := h*`, and again `h*` will remain at `a2` until `h` releases the lock.

It ensures progress because the only place to get stuck at in the `DO`, and whichever thread is not in `lastReq` will get past it. It ensures fairness because the first thread to get to `a2` is the one that will get the lock first.

⁵ G. Peterson, A new solution to Lamport's concurrent programming problem using small shared variables. *ACM Trans. Programming Languages and Systems* **5**, 1 (Jan. 1983), pp 56-65.

There is lots more to say about implementing `Mutex` efficiently, especially in the context of shared-memory multiprocessors.⁶

Multi-word clock

Often it's possible to get better performance by avoiding locking. Algorithms which do this are called 'wait-free'; we gave a brief discussion in handout 14. Here we present a wait-free algorithm due to Lamport⁷ for reading and incrementing a clock, even if clock values do not fit into a single memory location that can be read and written atomically.

We begin with the spec. It says that a `Read` returns some value that the clock had between the beginning and the end of the `Read`. As we saw in handout 8, where this spec is called `LateClock`, it takes a prophecy variable to show that this spec is equivalent to the simpler spec that just reads the clock value.

```

MODULE Clock EXPORT Read =
VAR t          : Int := 0                               % the current time
THREAD Tick() = DO << t := t + 1 >> OD                  % demon thread advances t
PROC Read() -> Int = VAR t1: Int |
    [R1] << t1 := t >>; [R2] << VAR t2 | t1 <= t2 /\ t2 <= t => RET t2 >> [R3]
END Clock

```

The code is based on the idea of doing reads and writes of the same multi-word data in opposite orders. `Tick` writes `hi2`, then `lo`, then `hi1`. `Read` reads `hi1`, then `lo`, then `hi2`; if it sees different values in `hi1` and `hi2`, there must have been a carry during the read, so `t` must have taken on the value `hi2 * base + lo`. The function `T` expresses this idea. The atomicity brackets in the code are the largest ones that are justified by big atomic actions.

```

MODULE ClockImpl EXPORT Read =
VAR base       := 2**32                                % constant
TYPE Word      = Int SUCHTHAT (\ i: Int | i IN 0 .. base-1)
VAR lo         : Word := 0
    hi1        : Word := 0
    hi2        : Word := 0
THREAD Tick() = DO VAR newLo: Word, newHi: Word |
    << newLo := lo + 1 // base; newHi := hi1 + 1 >>;
    IF newLo # 0 => lo := newLo >>
    [*] << hi2 := newHi >>; << lo := newLo >>; << hi1 := newHi >>
    FI

```

⁶ J. Mellor-Crummey and M. Scott, Algorithms for scalable synchronization of shared-memory multiprocessors. *ACM Transactions on Computer Systems* **9**, 1 (Feb. 1991), pp 21-65. A. Karlin et al., Empirical studies of competitive spinning for a shared-memory multiprocessor. *ACM Operating Systems Review* **25**, 5 (Oct. 1991), pp 41-55.

⁷ L. Lamport, Concurrent reading and writing of clocks. *ACM Transactions on Computer Systems* **8**, 4 (Nov. 1990), pp 305-310.

```

PROC Read() -> Int = VAR tLo: Word, tH1: Word, tH2: Word, t1Hist: Int |
  << tH1 := h1; t1Hist := T(lo, hi1, hi2) >>;
  << tLo := lo >>;
  << tH2 := h2; RET T(tLo, tH1, tH2) >>

FUNC T(l: Int, h1: Int, h2: Int) -> Int = h2 * base + (h1 = h2 => 1 [*] 0)

END ClockImpl

```

Given this code for reading a two-word clock atomically starting with atomic reads of the low and high parts, it's obvious how to apply it recursively $n-1$ times to read an n word clock.

Implementing mutexes and condition variables in an operating system

This section presents code for mutexes and condition variables based on the Taos operating system from DEC SRC. Instead of spinning like `SpinLock`, it explicitly queues threads waiting for locks or conditions. The code for mutexes has a fast path that stays out of the kernel in `Acquire` when the mutex is available, and in `Release` when no other thread is waiting for the mutex. There is also a fast path for `Signal`, for the common case that there's nobody waiting on the condition. There's no fast path for `wait`, since that always requires the kernel to run in order to reschedule the processor (unless a `Signal` sneaks in before the kernel gets around to the rescheduling).

Notes on the code for mutexes:

1. `MutexImpl` maintains a queue of waiting threads, blocks a waiting thread using `Deschedule`, and uses `Schedule` to hand a ready thread over to the scheduler to run.
2. `SpinLock` and `ReleaseSpinLock` acquire and release a global lock used in the kernel to protect thread queues.
3. The loop in `Acquire` serves much the same purpose as a loop that waits on a condition variable. If the mutex is already held, the loop calls `KernelQueue` to wait until it becomes available, and then tries again. `Release` calls `KernelRelease` if there's anyone waiting, and `KernelRelease` allows just one thread to run. That thread returns from its call of `KernelQueue`, and it will acquire the mutex unless another thread has called `Acquire` and slipped in since the mutex was released (roughly).
4. There is clumsy code in `KernelQueue` that puts the thread on the queue and then takes it off if the mutex turns out to be available. This is not a mistake; it avoids a race with `Release`, which calls `KernelRelease` to take a thread off the queue only if it sees that the queue is not empty. `KernelQueue` changes `q` and looks at `s`; `Release` uses the opposite order to change `s` and look at `q`.

This opposite-order access pattern often works in hard concurrency, that is, when there's not enough locking to do the job in a straightforward way. We saw another version of it in `Mutex2Impl`, which sets `req(h)` before reading `req(h*)`. In this case `req(h)` acts like a lock to keep $h^*. \$pc = a_2$ from changing from true to false.

The boxes show how `Acquire` and `Release` differ from the versions in `SpinLock`.

```

MODULE MutexImpl EXPORT M, New = % implements ForgetfulMutex

TYPE AH          = Mutex.AH
M                = Int WITH {acq:=Acquire, rel:=Release}
VAR s            : M -> AH          := {}
q                : M -> SEQ Thread := {}

APROC New() -> M = << VAR m | ~ s!m => s(m) := available, q(m) := {}; RET m >>

PROC Acquire(m) = VAR t: AH |
  DO <<t := s(m); s(m) := held>>; IF t#held => RET [*] SKIP FI; KernelQueue() OD

PROC Release(m) = s(m) := available; IF q(m) # {} => KernelRelease(m) [*] SKIP FI

% KernelQueue and KernelRelease run in the kernel.

PROC KernelQueue(m) =
% This is just a delay until there's a chance to acquire the lock.
% Queuing SELF before testing s(m) ensures that Release doesn't miss us.
% The spin lock keeps KernelRelease from getting ahead of us.
  SpinLock();
  q(m) := q(m) + {SELF};
  IF s(m) = available =>
    q(m) := q(m).reml % try again at Acquire
    [*] Deschedule(SELF) % wait, then try again
  FI;
  ReleaseSpinLock()

PROC KernelRelease(m) =
  SpinLock();
  q(m) # {} => Schedule(q(m).head); q(m) := q(m).tail;
  ReleaseSpinLock()
  % The newly scheduled thread competes with others to acquire the mutex.

END MutexImpl

```

Notes on the code for conditions:

1. In the code for Condition, the 'event count' `ecSig` deals with the standard 'wakeup-waiting' race condition: the `Signal` arrives after the `m.rel` but before the thread is queued. Note the use of the global spin lock as part of this. It looks as though `Signal` always schedules exactly one thread if the queue is not empty, but other threads that are in `wait` but have not yet acquired the spin lock may keep running; in terms of the spec they are awakened by `Signal` as well.
2. `Signal` and `Broadcast` test for any waiting threads without holding any locks, in order to avoid calling the kernel in this common case. The other event count `ecWait` ensures that this test doesn't miss a thread that is in `KernelWait` but hasn't yet blocked.

```

MODULE ConditionImpl EXPORT C, New = % implements Condition

TYPE C          = Int WITH {wait:=Wait, signal:=Signal, broadcast:=Broadcast}
M              = Mutex.M

VAR ecSig      : C -> Int
ecWait       : C -> Int
q            : C -> SEQ Thread := {}

APROC New() -> C =
  << VAR c | ~ ecSig!c => ecSig(c) := 0; ecWait(c) := 0; q(c) := {}; RET c >>

PROC Wait(c, m) = VAR i := ecSig(c) | m.rel; KernelWait(c, i); m.acq

PROC Signal(c) = VAR i := ecWait(c) |
  ecSig(c) := ecSig(c) + 1; IF q(c) # 0 \ / i # ecWait(c) => KernelSig(c)

PROC Broadcast(c) = VAR i := ecWait(c) |
  ecSig(c) := ecSig(c) + 1; IF q(c) # 0 \ / i # ecWait(c) => KernelCast(c)

PROC KernelWait(c, i: Int) = % internal kernel procedure
  SpinLock();
  ecWait(c) := ecWait(c) + 1;
  % if ecSig changed, there must have been a Signal, so return, else queue
  IF i = ecSig(c) => q(c) := q(c) + {SELF}; Deschedule(SELF) [*] SKIP FI;
  ReleaseSpinLock()

PROC KernelSig(c) = % internal kernel procedure
  SpinLock();
  IF q(c) # {} => Schedule(q(c).head); q(c) := q(c).tail [*] SKIP FI;
  ReleaseSpinLock()

PROC KernelCast(c) =
  SpinLock();
  DO q(c) # {} => Schedule(q(c).head); q(c):= q(c).tail OD;
  ReleaseSpinLock()

END ConditionImpl

```

The implementations of mutexes and conditions are quite similar; in fact, both are cases of a general semaphore.

Proving concurrent modules correct

This section explains how to prove the correctness of concurrent program modules. It reviews the simulation method that we have already studied, which works just as well for concurrent as for sequential modules. Then several examples illustrate how the method works in practice. Things are more complicated in the concurrent case because there are many more atomic transitions, and because the program counters of the threads are part of the state.

Before using this method in its full generality, you should first apply the theorem on big atomic actions as much as possible, in order to reduce the number of transitions that your proofs need to consider. If you are programming with easy concurrency, that is, if your code uses a standard locking discipline, this will get rid of nearly all the work. If you are doing hard concurrency, there will still be lots of transitions, and in doing the proof you will probably find bugs in your program.

The formal method

We use the same simulation technique that we used for sequential modules, as described in handouts 6 and 8. In particular, we use the most general version of this method, presented near the end of handout 8. This version does not require the transitions of the code to correspond one-for-one with the transitions of the spec. Only the external behavior (invocations and responses) must be the same—there can be any number of internal steps. The method proves that every trace (external behavior sequence) produced by the code can also be produced by the spec.

Of course, the utility of this method depends on an assumption that the external behavior of a module is all that is of interest to callers of the module. In other words, we are assuming here, as everywhere in this course, that the only interaction between the module and the rest of the program is through calls to the external routines provided by the module.

We need to show that each transition of the code simulates a sequence of transitions of the spec. An external transition must simulate a sequence that contains exactly one instance of the same external transition and no other external transitions; it can also contain any number of internal transitions. An internal transition must simulate a sequence that contains only internal transitions.

Here, once again, are the definitions:

Suppose T and S are modules with same external interface. An abstraction function F is a function from $states(T)$ to $states(S)$ such that:

Start: If u is any initial state of T then $F(u)$ is an initial state of S .

Step: If u and $F(u)$ are reachable states of T and S respectively, and (u, π, u') is a step of T , then there is an execution fragment of S from $F(u)$ to $F(u')$, having the same trace.

Thus, if π is an invocation or response, the fragment consists of a single π step, with any number of internal steps before and/or after. If π is internal, the fragment consists of any number (possibly 0) of internal steps.

As we saw in handout 8, we may have to add history variables to T in order to find an abstraction function to S (and perhaps prophecy variables too). The values of history variables are calculated in terms of the actual variables, but they are not allowed to affect the real steps.

An alternative to adding history variables is to define an abstraction relation instead of an abstraction function. An abstraction relation R is a relation between $states(T)$ and $states(S)$ such that:

Start: If u is any initial state of T then there exists an initial state s of S such that $(u, s) \in R$.

Step: If u and s are reachable states of T and S respectively, $(u, s) \in R$, and (u, π, u') is a step of T , then there is an execution fragment of S from s to some s' having the same trace, and such that $(u', s') \in R$.

Theorem: If there exists an abstraction function or relation from T to S then T implements S ; that is, every trace of T is a trace of S .

Proof: By induction.

The strategy

The formal method suggests the following strategy for doing hard concurrency proofs.

1. Start with a spec, which has an abstract state.
2. Choose a concrete state for the code.
3. Choose an abstraction function, perhaps with history variables, or an abstraction relation.
4. Write code, identifying the critical actions that change the abstract state.
5. While (checking the simulation fails) do

Add an invariant, checking that all actions of the code preserve it, or

Change the abstraction function (step 3), the code (step 4), the invariant (step 5), or more than one, or

Change the spec (step 1).

This approach always works. The first four steps require creativity; step 5 is quite mechanical except when you find an error. It is somewhat laborious, but experience shows that if you are doing hard concurrency and you omit any of these steps, your program won't work. Be warned.

Owicki-Gries proofs

Owicki and Gries invented a special case of this general method that is sometimes useful.⁸ Their idea is to do an ordinary sequential proof of correctness for each thread h , annotating each atomic command in the usual style with an assertion that is true at that point if h is the only thread running. This proof shows that the code of h establishes each assertion. Then you show that each of these assertions remains true after any command that any other thread can execute while h is at that point. This condition is called 'non-interference'; meaning not that other threads don't interfere with *access* to shared variables, but rather that they don't interfere with the *proof*.

The Owicki-Gries method amounts to defining an invariant of the form

$$h.\$pc = \alpha \implies A_\alpha \wedge h.\$pc = \beta \implies A_\beta \wedge \dots$$

and showing that it's an invariant in two steps: first, that every step of h maintains it, and then that every step of any other thread maintains it. The hope is that this decomposition will pay because the most complicated parts of the invariant have to do with private variables of h that aren't affected by other threads.

⁸ S. Owicki and D. Gries, An axiomatic proof technique for parallel programs. *Acta Informatica* **6**, 1976, pp 319-340.

Prospectus

The remainder of this handout contains example proofs of correctness for several of the examples above: the `RWLockImpl` code for a read/write lock, the `BufferImpl` code for a FIFO buffer, the `SpinLock` code for a mutex (given in two versions), the `Mutex2Impl` code for a mutex used by two threads, and the `ClockImpl` code for a multi-word clock.

The amount of detail in these proofs is uneven. The proof of the FIFO buffer code and the second proof of the `Spinlock` code are the most detailed. The others give the abstraction functions and key invariants, but do not discuss each simulation step.

Read/write locks

We indicate how to prove directly that the module `RWLockImpl` implements `ForgetfulRWL`. This could be done by big atomic actions, since the code uses easy concurrency, but we discuss how to do it directly. The two modules are based on the same data, the variable `rw`. The difference is that `RWLockImpl` uses a condition variable to prevent `Acquire` threads from busy-waiting when they don't see the condition they require. A mutex is used to restrict accesses to `rw`, so that a series of accesses to `rw` can be done atomically.

An abstraction function maps `RWLockImpl` to `Forgetful RWLock`. The interesting part of the state of `ForgetfulRWL` is the `rw` variable. We define that by the identity mapping from `RWLockImpl`.

The mapping for steps is mostly determined by the `rw` identity mapping: the steps that assign to `rw` in `RWLockImpl` are the ones that correspond to the procedure bodies in `ForgetfulRWL`. Then the checking of the state and step correspondences is pretty routine.

There is one subtlety. It would be bad if a series of `rw` steps done atomically in `ForgetfulRWL` were interleaved in `RWLockImpl`. Of course, we know they aren't, because they are always done by a thread holding the mutex. But how does this fact show up in the proof?

The answer is that we need some invariants for `RWLockImpl`. The first, a “dominant thread invariant”, says that only a thread whose name is in `m` (a ‘dominant thread’) can be in certain portions of its code (those guarded by the mutex). The dominant thread invariant is in turn used to prove other invariants called “data protection invariants”.

For example, one data protection invariant says that if a thread (in `RWLockImpl`) is in middle of the assignment statement `rw := rw + 1`, then in fact `rw ≥ 0` (that is, the test is still true). We need this data protection invariant to show that the corresponding abstract step (the body of `ReadAcquire` in `ForgetfulRWLock`) is enabled.

`BufferImpl` implements `Buffer`

The FIFO buffer is another example of easy concurrency, so again we don't need to do a transition-by-transition proof for it. Instead, it suffices to show that a thread holds the lock `m` whenever it touches the shared variable `b`. Then we can treat the whole critical section during which the lock is held as a big atomic action, and the proof is easy. We will work out the

important details of a low-level proof, however, in order to get some practice in a situation that is slightly more complicated but still straightforward, and in order to convince you of the power of the theorem about big atomic actions.

First, we give the abstraction function; then we use it to show that the code simulates the spec. We use a slightly simplified version of `Produce` that always signals, and we introduce a local variable `temp` to make explicit the atomicity of assignment to the shared variable `b`.

Abstraction function

The abstraction function on the state must explain how to interpret a state of the code as a state of the spec. Remember that to prove a concurrent program correct, we need to consider the entire state of a module, including the program counters and local variables of threads. For sequential programs, we can avoid this by treating each external operation as a single atomic action.

To describe the abstraction function, we thus need to explain how to construct a state of the spec from a state of the code. So what is a state of the `Buffer` module above? It consists of:

- A sequence of items `b` (the buffer itself);
- for each thread that is active in the module, a program counter; and
- for each thread that is active in the module, values for local variables.

A state of the code is similar, except that it includes the state of the `Mutex` and `Condition` modules.

To define the mapping, we need to enumerate the possible program counters. For the spec, they are:

P_1 — before the body of `Produce`
 P_2 — after the body of `Produce`
 C_1 — before the body of `Consume`
 C_2 — after the body of `Consume`

or as annotations to the code:

```
PROC Produce(t) = [ $P_1$ ] << b := b + {t} >> [ $P_2$ ]
```

```
PROC Consume() -> T =  
  [ $C_1$ ] << b.size > 0 => VAR t := b.head | b := b.tail; RET t >> [ $C_2$ ]
```

For the code, they are:

- For a thread in `Produce`:

p_1 — before `m.acq`
in `m.acq`—either before or after the action
 p_2 — before `temp := b + {t}`
 p_3 — before `b := temp`
 p_4 — before `c.signal`

in `c.signal`—either before or after the action
 p_5 — before `m.rel`
 in `m.rel`—either before or after the action
 p_6 — after `m.rel`

- For a thread in `Consume`:

c_1 — before `m.acq`
 in `m.acq`—either before or after action
 c_2 — before the test `b.size = 0`
 c_3 — before `c.wait`
 in `c.wait`—at beginning, in middle, or at end
 c_4 — before `t := b.head`
 c_5 — before `temp := b.tail`
 c_6 — before `b := temp`
 c_7 — before `m.rel`
 in `m.rel`—either before or after action
 c_8 — before `RET t`
 c_9 — after `RET t`

or as annotations to the code:

```

PROC Produce(t) = VAR temp |
  [ $p_1$ ] m.acq;
  [ $p_2$ ] temp = b + {t};
  [ $p_3$ ] b := temp;
  [ $p_4$ ] c.signal;
  [ $p_5$ ] m.rel [ $p_6$ ]

PROC Consume() -> T = VAR t, temp |
  [ $c_1$ ] m.acq;
  DO [ $c_2$ ] b.size = 0 => [ $c_3$ ] c.wait OD;
  [ $c_4$ ] t := b.head;
  [ $c_5$ ] temp := b.tail; [ $c_6$ ] b := temp;
  [ $c_7$ ] m.rel;
  [ $c_8$ ] RET t [ $c_9$ ]

```

Notice that we have broken the assignment statements into their constituent atomic actions, introducing a temporary variable `temp` to hold the result of evaluating the right hand side. Also, the PC's in the `Mutex` and `Condition` operations are taken from the specs of those modules (*not* the code; we prove their correctness separately). Here for reference is the relevant code.

```

APROC Acquire(m) = << s(m) = nil => s(m) := SELF; RET >>
APROC Release(m) = << s(m) = SELF => s(m) := nil ; RET [*] HAVOC >>

APROC Signal(c) = << VAR hs: SET Thread |
  IF hs <= s(c) /\ hs # {} => s(c) := s(c) - hs [*] SKIP FI >>

```

Now we can define the mapping on program counters:

- If a thread h is not in `Produce` or `Consume` in the code, then it is not in either procedure in the spec.
- If a thread h is in `Produce` in the code, then:
 - If $h.\$pc$ is in $\{p_1, p_2, p_3\}$ or is in `m.acq`, then in the spec $h.\$pc = P_1$.
 - If $h.\$pc$ is in $\{p_4, p_5, p_6\}$ or is in `m.rel` or `c.signal` then in the spec $h.\$pc = P_2$.
- If a thread h is in `Consume` in the code, then:
 - If $h.\$pc \in \{c_1, \dots, c_6\}$ or is in `m.acq` or `c.wait` then in the spec $h.\$pc = C_1$.
 - If $h.\$pc$ is in $\{c_7, c_8, c_9\}$ or is in `m.rel` then in the spec $h.\$pc = C_2$.

The general strategy here is to pick, for each atomic transition in the spec, some atomic transition in the code to simulate it. Here, we have chosen the modification of `b` in the code to simulate the corresponding operation in the spec. Thus, program counters before that point in the code map to program counters before the body in the spec, and similarly for program counters after that point in the code.

This choice of the abstraction function for program counters determines how each transition of the code simulates transitions of the spec as follows:

- If π is an external transition, π simulates the singleton sequence containing just π .
- If π takes a thread from a PC of p_3 to a PC of p_4 , π simulates the singleton sequence containing just the body of `Produce`.
- If π takes a thread from a PC of c_6 to a PC of c_7 , π simulates the singleton sequence containing just the body of `Consume`.
- All other transitions π simulate the empty sequence.

This example illustrates a typical situation: we usually find that a transition in the code simulates a sequence of either zero or one transitions in the spec. Transitions that have no effect on the abstract state simulate the empty sequence, while transitions that change the abstract state simulate a single transition in the spec. The proof technique used here works fine if a transition simulates a sequence with more than one transition in it, but this doesn't show up in most examples.

In addition to defining the abstract program counters for threads that are active in the module, we also need to define the values of their local variables. For this example, the only local variables are `temp` and the item `t`. For threads active in either `Produce` or `Consume`, the abstraction function on `temp` and `t` is the identity; that is, it defines the values of `temp` and `t` in a state of the spec to be the value of the identically named variable in the corresponding operation of the code.

Finally, we need to describe how to construct the state of the buffer `b` from the state of the code. Given the choices above, this is simple: the abstraction function is the identity on `b`.

Proof sketch

To prove the code correct, we need to prove some invariants on the state. Here are some obvious ones; the others we need will become clear as we work through the rest of the proof.

First, define a thread h to be *dominant* if $h.\$pc$ is in `Produce` and $h.\$pc$ is in $\{p_2, p_3, p_4, p_5\}$ or is at the end of `m.acq`, in `c.signal`, or at the beginning of `m.rel`, or if $h.\$pc$ is in `Consume` and $h.\$pc$ is in $\{c_2, c_3, c_4, c_5, c_6, c_7\}$ or is at the end of `m.acq`, at the beginning or end of `c.wait` (but not in the middle), or at the beginning of `m.rel`.

Now, we claim that the following property is invariant: a thread h is dominant if and only if $\text{Mutex.s}(m) = h$. This simply says that h holds the mutex if and only if its PC is at an appropriate point. This is the basic mutual exclusion property. Amazingly enough, given this property we can easily show that operations are mutually exclusive: for all threads h, h' such that $h \neq h'$, if h is dominant then h' is not dominant. In other words, at most one thread can be in the middle of one of the operations in the code at any time.

Now let's consider what needs to be shown to prove the code correct. First, we need to show that the claimed invariants actually are invariants. We do this using the standard inductive proof technique: Show that each initial state of the code satisfies the invariants, and then show that each atomic action in the code preserves the invariants. This is left as an exercise.

Next, we need to show that the abstraction function defines a simulation of the spec by the code. Again, this is an inductive proof. The first step is to show that an initial state of the code is mapped by the abstraction function to an initial state of the spec. This should be straightforward, and is left as an exercise. The second step is to show that the effects of each transition are preserved by the abstraction function. Let's consider a couple of examples.

- Consider a transition π from r to r' in which an invocation of an operation occurs for thread h . Then in state r , h was not active in the module, and in r' , its PC is at the beginning of the operation. This transition simulates the identical transition in the spec, which has the effect of moving the PC of this thread to the beginning of the operation. So $AF(r)$ is taken to $AF(r')$ by the transition.
- Consider a transition in which a thread h moves from $h.\$pc = p_3$ to $h.\$pc = p_4$, setting b to the value stored in `temp`. The corresponding abstract transition sets b to $b + \{t\}$. To show that this transition does the right thing, we need an additional invariant:

If $h.\$pc = p_3$, then $\text{temp} = b + \{t\}$.

To prove this, we use the fact that if $h.\$pc = p_3$, then no other thread is dominant, so no other transition can change b . We also have to show that any transition that puts $h.\$pc$ at this point establishes the consequent of the implication — but there is only one transition that does this (the one that assigns to `temp`), and it clearly establishes the desired property.

The transition in `Consume` that assigns to b relies on a similar invariant. The rest of the transitions involve straightforward case analyses. For the external transitions, it is clear that they correspond directly. For the other internal transitions, we must show that they have no abstract effect, i.e., if they take r to r' , then $AF(r) = AF(r')$. This is left as an exercise.

SpinLock implements Mutex, first version

The proof is done in two layers. First, we show that `ForgetfulMutex` implements `Mutex`. Second, we show that `SpinLock` implements `ForgetfulMutex`. For convenience, we repeat the definitions of the two modules, omitting the complication of multiple mutexes and the `s` variable.

```
MODULE Mutex EXPORT Acquire, Release =
VAR m          : (Thread + Null) := nil
APROC Acquire() = << m = nil => m := SELF; RET >>
APROC Release() = << m = SELF => m := nil ; RET [*] HAVOC >>
END Mutex

MODULE ForgetfulMutex EXPORT Acquire, Release =
TYPE M          = ENUM[available, held]
VAR m          := available
PROC Acquire() = << m = available => m := held >>
PROC Release() = << m := available >>
END ForgetfulMutex
```

Proof that ForgetfulMutex implements Mutex

These two modules have the same atomicity. The difference is that `ForgetfulMutex` forgets which thread owns the mutex, and so it can't check that the "right" thread releases it. We use an abstraction relation AR . It needs to be multi-valued in order to put back the information that is forgotten in the code. Instead of using a relation, we could use a function and history variables to keep track of the owner and havoc. The single-level proof given later on that `Spinlock` implements `Mutex` uses history variables.

The main interesting relationship that AR must express is:

$s.m$ is non-`nil` if and only if $u.m = \text{held}$.

In addition, AR must include less interesting relationships. For example, it has to relate the `$pc` values for the various threads. In each module, each thread is either not there at all, before the body, or after the body. Thus, AR also includes the condition:

The `$pc` value for each thread is the same in both modules.

Finally, there is the technicality of the special `Havoc` state that occurs in `Mutex`. We handle this by allowing AR to relate all states of `ForgetfulMutex` to the `Havoc` state.

Having defined AR , we just show that the two conditions of the abstraction relation definition are satisfied.

The start condition is obvious. In the unique start states of both modules, no thread is in the module. Also, if u is the state of `ForgetfulMutex` and s is the state of `Mutex`, then we have $u.m = \text{available}$ and $s.m = \text{nil}$. It follows that $(u, s) \in AR$, as needed.

Now we turn to the step condition. Let u and s be reachable states of `ForgetfulMutex` and `Mutex`, respectively, and suppose that (u, π, u') is a step of `ForgetfulMutex` and that $(u, s) \in AR$. If $s.\$havoc$, then it is easy to show the existence of a corresponding execution fragment of `Mutex`, because any transition is possible. So we suppose that $s.\$havoc = \text{false}$. Invocation and response steps are straightforward; the interesting cases are the internal steps.

So suppose that π is an internal action of `ForgetfulMutex`. We argue that the given step corresponds to a single step of `Mutex`, with “the same” action. There are two cases:

1. π is the body of an `Acquire`, by some thread h . Since `Acquire` is enabled in `ForgetfulMutex`, we have $u.m = \text{available}$, and $h.\$pc$ is right before the `Acquire` body in u . Since $(u, s) \in AR$, we have $s.m = \text{nil}$, and also $h.\$pc$ is just before the `Acquire` body in s . Therefore, the `Acquire` body for thread h is also enabled in `Mutex`. Let s' be the resulting state of `Mutex`.

By the code, $u'.m = \text{held}$ and $s'.m = h$, which correspond correctly according to AR . Also, since the same thread h gets the mutex in both steps, the PC's are changed in the same way in both steps. So $(u', s') \in AR$.

2. π is the body of a `Release`, by some thread h . If $u.m = \text{available}$ then `ForgetfulMutex` does something sensible, as indicated by its code. But since $(u, s) \in AR$, $s.m = \text{nil}$ and `Mutex` does `HAVOC`. Since $\$havoc$ in `Mutex` is defined to correspond to everything in `ForgetfulMutex`, we have $(u', s') \in AR$ in this case.

On the other hand, if $u.m = \text{held}$ then `ForgetfulMutex` sets $u'.m := \text{available}$. Since $(u, s) \in AR$, we have $s.m \neq \text{nil}$. Now there are two cases: If $s.m = h$, then corresponding changes occur in both modules, which allows us to conclude $(u', s') \in AR$. Otherwise, `Mutex` goes to `HAVOC`. But as before, this is OK because `HAVOC` corresponds to everything in `ForgetfulMutex`.

The conclusion is that every trace of `ForgetfulMutex` is also a trace of `Mutex`. Note that this proof does not imply anything about liveness, though in fact the two modules have the same liveness properties.

Proof that SpinLock implements ForgetfulMutex

We repeat the definition of `SpinLock`, again simplifying it to a single mutex.

```
MODULE SpinLock EXPORT Acquire, Release =
TYPE M          = ENUM[available, held]
VAR m           := available

PROC Acquire(m) = VAR t: AH |
    DO << t := m; m := held >>; IF t # held => RET [*] SKIP FI OD
PROC Release() = << m := available >>

END SpinLock
```

These two modules use the same basic data. The difference is their atomicity. Because they use the same data, an abstraction function AF will work. Indeed, the point of introducing `ForgetfulMutex` was to take care of the need for history variables or abstraction relations there.

The key to defining AF is to identify the exact moment in an execution of `SpinLock` when we want to say the abstract `Acquire` body occurs. There are two logical choices: The moment when a thread converts $u.m$ from `available` to `held`, or the later moment when the thread discovers that it has done this. Either will work, but to be definite we consider the earlier of these two possibilities.

Then AF is defined as follows. If u is any state of `SpinLock` then $AF(u)$ is the unique state s of `ForgetfulMutex` such that:

- $s.m = u.m$, and
- The PC values of all threads “correspond”.

We must define the sense in which the PC values correspond. The correspondence is straightforward for threads that aren’t there, or are engaged in a `Release` operation. For a thread h engaged in an `Acquire` operation, we say that

- $h.\$pc$ in `ForgetfulMutex`, $s.h.\$pc$, is just before the body of `Acquire` if and only if $h.\$pc$ is in `SpinLock` either (a) at the `DO`, and before the test-and-set, or (b) after the test-and-set with h ’s local variable t equal to `held`.
- $h.\$pc$ in `ForgetfulMutex`, $s.h.\$pc$, is just after the body of `Acquire` if and only if $u.h.\$pc$ is either (a) after the test-and-set with h ’s local variable t equal to `available` or (b) after the $t \# \text{held test}$.

The proof that this is an abstraction function is interesting. The start condition is easy. For the step condition, the invocation and response cases are easy, so consider the internal steps. The `Release` body corresponds exactly in both modules, so the interesting steps to consider are those that are part of the `Acquire`.

`Acquire` in `SpinLock` has two kinds of internal steps: the atomic test-and-set and the test for $t \# \text{held}$. We consider these two cases separately:

The atomic test-and-set, $(u, \text{test-and-set}, u')$. Say this is done by thread h . In this case, the value of m *might* change. It depends on whether the step of `SpinLock` changes m from `available` to `held`. If it does, then we map the step to the `Acquire` body. If not, then we map it to the empty sequence of steps. We consider the two cases separately:

1. The step changes m . Then in `SpinLock`, $h.\$pc$ moves after the test-and-set with h ’s local variable $t = \text{available}$. We claim first that the `Acquire` body in `ForgetfulMutex` is enabled in state $AF(u)$. This is true because it requires only that $s.m = \text{available}$, and this follows from the abstraction function since $u.m = \text{available}$. Then we claim that the new states in the two modules are related by AF . To see this, note that $m = \text{held}$ in both cases. And the new PC’s correspond: in `ForgetfulMutex`, $h.\$pc$ moves to right after the `Acquire` body, which corresponds to the position of $h.\$pc$ in `SpinLock`, by the definition of the abstraction function.

2. The step does not change m . Then $h.\$pc$ in `SpinLock` moves to the test, with $t = \text{held}$. Thus, there is no change in the abstract value of $h.\$pc$.

The test for $t \neq \text{held}$, (u, test, u') . Say this is done by thread h . We always map this to the empty sequence of steps in `ForgetfulMutex`. We must argue that this step does not change anything in the abstract state, i.e., that $AF(u') = AF(u)$. There are two cases:

1. If $t = \text{held}$, then the step of `SpinLock` moves $h.\$pc$ to after the `DO`. But this does not change the abstract value of $h.\$pc$, according to the abstraction function, because both before and after the step, the abstract $h.\$pc$ value is before the body of `Acquire`.
2. On the other hand, if $t = \text{available}$, then the step of `SpinLock` moves $h.\$pc$ to after the `=>`. Again, this does not change the abstract value of $h.\$pc$ because both before and after the step, the abstract $h.\$pc$ value is after the body of `Acquire`.

SpinLock implements Mutex, second version

Now we show again that `SpinLock` implements `Mutex`, this time with a direct proof that combines the work done in both levels of the proof in the previous section. For contrast, we use history variables instead of an abstraction relation.

Abstraction function

We need to be precise about what constitutes a state of the code and what constitutes a state of the spec. A state of the spec consists of:

- A value for m (either a thread or `nil`); and
- for each thread that is active in the module, a program counter.

There are no local variables for threads in the spec.

A state of the code is similar; it consists of:

- A value for m (either `available` or `held`);
- for each thread that is active in the module, a program counter; and
- for each thread that is active in `Acquire`, a value for the local variable t .

Now we have a problem: there is no way to define an abstraction function from a code state to a spec state. The problem here is that the code does not record which thread holds the mutex, yet the spec keeps track of this information. To solve this problem, we have to introduce a history variable. We do this as follows:

- We augment the state of the code with an additional variable, ms , which has type $(\text{Thread} + \text{Null})$; the initial state of ms is `nil`.
- We define the effect of each atomic action in the code on the history variable; written in `Spec`, this results in the following modified code:

```

PROC Acquire () = VAR t: AH |
  DO <<t := m; m := held>>; IF t # held => << ms := SELF >>; RET [*] SKIP FI OD;

PROC Release () = << m := available; ms := nil >>

```

You can easily check that these additions to the code satisfy the constraints required for adding history variables.

Now we can proceed to define the abstraction function. First, we enumerate the program counters. For the spec, they are:

A_1 — before the body of `Acquire`
 A_2 — after the body of `Acquire`
 R_1 — before the body of `Release`
 R_2 — after the body of `Release`

For the code, they are:

- For a thread in `Acquire`:

a_1 — before the `VAR t`
 a_2 — after the `VAR t` and before the `DO` loop
 a_3 — before the test-and-set in the body of the `DO` loop
 a_4 — after the test-and-set in the body of the `DO` loop
 a_5 — before the assignment to `ms`
 a_6 — after the assignment to `ms`

- For a thread in `Release`:

r_1 — before the body
 r_2 — after the body

The transitions in `Acquire` may be a little confusing: there's a transition from a_4 to a_3 , as well as transitions from a_4 to a_5 .

Here are the routines in `Mutex` annotated with the PC values:

```

APROC Acquire() = [A1] << m = nil => m := SELF >> [A2]

```

```

APROC Release() = [R1] << m # SELF => HAVOC [*] m := nil >> [R2]

```

Here are the routines in `SpinLock` annotated with the PC values:

```

PROC Acquire () = [a1] VAR t := AH |
  [a2] DO [a3] << t := m; m := held >>;
  [a4] IF t # held => [a5] << ms := SELF >>; [a6] RET [*] SKIP FI OD;

```

```

PROC Release () = [r1] << m := available; ms := nil >> [r2]

```

Now we can define the mappings on program counters:

- If a thread is not in `Acquire` or `Release` in the code, then it is not in either in the spec.

- $\{a_1, a_2, a_3, a_4, a_5\}$ maps to A_1 , a_6 maps to A_2
- r_1 maps to R_1 , r_2 maps to R_2

The part of the abstraction function dealing with the global variables of the module simply defines m in the spec to have the value of ms in the code. As in handout 8, we just throw away all but the spec part of the state.

Since there are no local variables in the spec, the mapping on program counters and the mapping on the global variables are enough to define how to construct a state of the spec from a state of the code.

Once again, the abstraction function on program counters determines how transitions in the code simulate sequences of transitions in the spec:

- If π is an external transition, π simulates the singleton sequence containing just π .
- If π takes a thread from a_5 to a_6 , π simulates the singleton sequence containing just the transition from A_1 to A_2 .
- If π takes a thread from r_1 to r_2 , π simulates the singleton sequence containing just the transition from R_1 to R_2 .
- All other transitions simulate the empty sequence.

Proof sketch

As in the previous example, we will need some invariants to do the proof. Rather than trying to write them down first, we will see what we need as we do the proof.

First, we show that initial states of the code map to initial states of the spec. This is easy; all the thread states correspond, and the initial state of ms in the code is `nil`.

Next, we show that transitions in the code and the spec correspond. All transitions from outside the module to just before a routine's body are straightforward, as are transitions from the end a routine's body to outside the module (i.e., when a routine returns). The transition in the body of `Release` is also straightforward. The hard cases are in the body of `Acquire`.

Consider all the transitions in `Acquire` before the one from a_5 to a_6 . These all simulate the null transition, so they should leave the abstract state unchanged. And they do, because none of them changes ms .

The transition from a_5 to a_6 simulates the transition from A_1 to A_2 . There are two cases: when $ms = \text{nil}$, and when $ms \neq \text{nil}$.

1. In the first case, the transition from A_1 to A_2 is enabled and, when taken, changes the state so that $m = \text{SELF}$. This is just what the transition from a_5 to a_6 does.
2. Now consider the case when $ms \neq \text{nil}$. We claim this case is impossible—that if a thread is at a_5 , then $ms = \text{nil}$. But even though this invariant is what we want, it's too weak to prove itself inductively; for that, we need the following, stronger invariant:

If `m = available` then `ms = nil`, and

If a thread is at `a5`, or at `a4` with `t = available`, then `ms = nil`, `m = held`, there are no other threads at `a5`, and for all other threads at `a4`, `t = held`.

Given this invariant, we are done: we have shown the appropriate correspondence for all the transitions in the code. So we must prove the invariant. We do this by induction. It's vacuously true in the initial state, since no thread could be at `a4` or `a5` in the initial state. Now, for each transition, we assume that the invariant is true before the transition and prove that it still holds afterwards.

The external transitions preserve the invariant, since they change nothing relevant to it.

The transition in `Release` preserves the first part of the invariant because afterwards both `m = available` and `ms = nil`. To prove that the transition in `Release` preserves the second part of the invariant, we assume the precondition for `Release` in the spec is true (otherwise, anything is allowed). If the precondition is true, `ms ≠ nil`, and since we're assuming the invariant is true before the transition, this implies that no thread is at `a4` with `t = available` or at `a5`. After the transition to `r2` it's still the case that no thread is at `a4` with `t = available` or at `a5`, so the invariant is still true.

Now we consider the transitions in `Acquire`. The transitions from `a1` to `a2` and from `a2` to `a3` obviously preserve the invariant. The transition from `a4` to `a5` puts a thread at `a5`, but `t = available` in this case so the invariant is true after the transition by induction. The transition from `a4` to `a3` also clearly preserves the invariant.

The transition from `a3` to `a4` is the first interesting one. It certainly preserves the first part of the invariant, since it doesn't change `ms` and only changes `m` to `held`. Now we assume the second part of the invariant true before the transition. There are two cases:

1. Before the transition, there is a thread at `a5`, or at `a4` with `t = available`. Then we have `m = held` by induction, so after the transition both `t = held` and `m = held`. This preserves the invariant.
2. Before the transition, there are no threads at `a5` or at `a4` with `t = available`. Then after the transition, there is still no thread at `a5`, but there is a new thread at `a4`. (Any others must have `t = held`.) Now, if this thread has `t = held`, the second part of the invariant is true vacuously; but if `t = available`, then we have:

`ms = nil` (since when the thread was at `a3` `m` must have been `available`, hence the first part of the invariant applies);

`m = held` (as a direct result of the transition);

there are no threads at `a5` (by assumption); and

there are no other threads at `a4` with `t = available` (by assumption). So the invariant is still true after the transition.

Finally, assume a thread h is at a_5 , about to transition to a_6 . If the invariant is true here, then h is the only thread at a_5 , and all threads at a_4 have $t = \text{held}$. So after it makes the transition, the invariant is vacuously true, because there is no other thread at a_5 and the threads at a_4 haven't changed their state.

We have proved the invariant. The invariant implies that if a thread is at a_5 , $ms = \text{nil}$, which is what we wanted to show.

Mutex2Impl implements Mutex

First we show that the simple, deadlocking version `Acquire0` maintains mutual exclusion. Recall that we write h^* for the thread that is the partner of thread h . Here is the code for `Acquire0`.

```
PROC Acquire0() =
  [a1] req(SELF) := true;
  DO [a2] req(SELF*) => SKIP OD [a3]
```

Intuitively, we get mutual exclusion because once `req(h)` is true, h^* can't get from a_2 to a_3 . It's convenient to define

```
FUNC Holds0(h: Thread) = RET req(h) /\ h.$pc # a2
```

Abstractly, h has the mutex if `Holds0(h)`, and the transition from a_2 to a_3 simulates the body of `Mutex.Acquire`. Precisely, the abstraction function is

```
Mutex.ms = (Holds0.set = {} => nil [*] Holds0.set.choose)
```

Recall that if P is a predicate, $P.set$ is the set of arguments for which it is true.

To make precise the idea that `req(h)` stops h^* from getting to a_3 , the invariant we need is

```
Holds0.set.size <= 1 /\ (h.$pc = a2 ==> req(h))
```

The first conjunct is the mutual exclusion. It holds because, given the first conjunct, only (a_2, a_3) can increase the size of `Holds0.set`, and h can take that step only if `req(h*) = false`, so `Holds0.set` goes from `{}` to `{h}`. The second conjunct holds because it can never be `true ==> false`, since only the step $(a_1, \text{req}(h) := \text{true}, a_2)$ can make the antecedent true, this step also makes the consequent true, and no step away from a_2 makes the consequent false.

This argument applies to `Acquire0` as written, but you might think that it's unrealistic to fetch the shared variable `req(SELF*)` and test it in a single atomic action; certainly this will take more than one machine instruction. We can appeal to big atomic actions, since the whole sequence from a_2 to a_3 has only one action that touches a shared variable (the fetch of `req(SELF*)`) and therefore is atomic.

This is the right thing to do in practice, but it's instructive to see how to do it by hand. We break the last line down into two atomic actions:

```
VAR t | DO [a2] << [t := req(SELF*)] >>; [a21] << t => SKIP >> OD [a3]
```

We examine several ways to show the correctness of this; they all have the same idea, but the details differ. The most obvious one is to add the conjunct $h.\$pc \neq a_{21}$ to `Hold0`, and extend the mutual exclusion conjunct of the invariant so that it covers a thread that has reached a_{21} with $t = \text{false}$:

```
(Hold0.set + {h | h.$pc = a21 /\ h.t = false}).size <= 1
```

Or we could get the same effect by saying that a thread acquires the lock by reaching a_{21} with $t = \text{false}$, so that it's the transition (a_2, a_{21}) with $t = \text{false}$ that simulates the body of `Mutex.Acquire`, rather than the transition to a_3 as before. This means changing the definition of `Hold0` to

```
FUNC Hold0(h: Thread) =
  RET req(h) /\ h.$pc # a2 /\ (h.$pc = a21 ==> h.t = false)
```

Yet another approach is to make explicit in the invariant what h knows about the global state. One purpose of an invariant is to remember things about the global state that a thread has discovered in the past; the fact that it's an invariant means that those things stay true, even though other threads are taking steps. In this case, $t = \text{false}$ in h means that either $\text{req}(h^*) = \text{false}$ or h^* is at a_2 or a_{21} , in other words, $\text{Hold}(h^*) = \text{false}$. We can put this into the invariant with the conjunct

```
h.$pc = a21 /\ h.t = false ==> Hold(h*) = false
```

and this is enough to ensure that the transition (a_{21}, a_3) maintains the invariant.

We return from this digression on proof methodology to study the non-deadlocking `Acquire`:

```
PROC Acquire() =
  [a11] req(SELF) := true;
  [a12] lastReq := self;
  DO [a2] (req(SELF*) /\ lastReq = SELF) => SKIP OD [a3]
```

We discussed liveness informally earlier, and we don't attempt to prove it. To prove mutual exclusion, we need to extend `Hold0` in the obvious way:

```
FUNC Hold(h: Thread) = req(h) /\ h.$pc # a12 /\ h.$pc # a2
```

and add $\neg h.\$pc = a_{12}$ to the antecedent of the invariant. We also need to add something to the invariant to express the fact that h won't find $\text{lastReq} = h^*$ as long as h^* holds the lock. This is

```
Hold0.set.size <= 1
/\ (h.$pc = a2 /\ h.$pc = a12 ==> req(h))
/\ (Hold(h*) /\ h.$pc = a2 ==> lastReq = h)
```

The last conjunct holds because (a_{12}, a_2) makes it true, and the only way to make it false is for h^* to do `lastReq := SELF`, which it can only do from a_{12} , so that $\text{Hold}(h^*)$ is false. With this invariant it's obvious that (a_2, a_3) maintains the invariant.

clockImpl implements clock

The invariant that makes this work is based on the idea that `Read` might complete before the next `Tick`, and therefore the value `Read` would return by reading the rest of the shared variables must be between `t1Hist` and `clock.t`. We can write this most clearly by annotating the labels in `Read` with assertions that are true when the PC is there.

```
% ABSTRACTION FUNCTION clock.t = T(lo, hi1, hi2), clock.Read.t1 = Read.t1Hist
% The PC correspondence is R1 <-> r1, R2 <-> r2, r3, R3 <-> r4

PROC Read() -> Int = VAR tLo: Word, tH1: Word, tH2: Word, t1Hist: Int |
  [r1] << tH1 := hi1; t1Hist := T(lo, hi1, hi2) >>;
  [r2] % I2: T(lo, tH1, hi2) IN t1Hist .. T(lo, hi1, hi2)
      << tLo := lo; >>
  [r3] % I3: T(tLo, tH1, hi2) IN t1Hist .. T(lo, hi1, hi2)
      << tH2 := hi2; RET T(tLo, tH1, tH2) >>
  [r4] % I4: $a IN t1Hist .. T(lo, hi1, hi2)
```

The whole invariant is thus

$$h.\$pc = r_2 \implies I_2 \wedge h.\$pc = r_3 \implies I_3 \wedge h.\$pc = r_4 \implies I_4$$

The steps of `Read` clearly maintain this invariant, since they don't change the value before `IN`. The steps of `Tick` maintain it by case analysis.