

This is the *PlusCal* specification of the deconstructed bakery algorithm in the paper

Deconstructing the Bakery to Build a Distributed State Machine

There is one simplification that has been made in the *PlusCal* version: the registers *localCh*[*i*][*j*] have been made atomic, a read or write being a single atomic action. This doesn't affect the derivation of the distributed bakery algorithm from the deconstructed algorithm, which also makes the simplifying assumption those registers are atomic because they disappear from the final algorithm.

Here are some of the changes made to the paper's notation to conform to *PlusCal*/TLA+. Tuples are enclosed in $\langle \rangle$, so we write $\langle i, j \rangle$ instead of (i, j) . There's no upside down "?" symbol in TLA+, so that's replaced by the identifier *qm*.

The pseudo-code for main process *i* has two places in which subprocesses $\langle i, j \rangle$ are forked and process *i* resumes execution when they complete. *PlusCal* doesn't have subprocesses. This is represented in *PlusCal* by having a single process $\langle i, j \rangle$ executing concurrently with process *i*, synchronizing appropriately using the variable *pc*.

Here is the basic idea:

This pseudo-code for process *i*:

```

main code ;
process j # i \in S
  s1: subprocess code
end process
p2: more main code
    
```

is expressed in *PlusCal* as follows:

```

In process i
  main code ;
  p2: await \A j # i : pc[<<i,j>>] = "s2"
  more main code

In process <i, j>
  s1: await pc[i] = "p2"
  subprocess code ;
  s2: ...
    
```

Also, processes have identifiers and, for reasons that are not important here, we can't use *i* as the identifier for process *i*, so we use $\langle i \rangle$. So, *pc*[*i*] in the example above should be *pc*[$\langle i \rangle$]. In the pseudo-code, process *i* also launches asynchronous processes $\langle i, j \rangle$ to set *localNum*[*j*][*i*] to 0. In the code, these are another set of processes with ids $\langle i, j, "wr" \rangle$.

We could simplify this algorithm by not waiting for *localNum*[*j*][*i*] to equal 0 in subprocess $\langle i, j \rangle$ and having the asynchronous write of 0 not do anything if process *i* has begun the write to *localCh*[*i*][*j*] that sets its value to *number*[*i*]. However, I think I like the algorithm in the paper the way it is because it makes the pseudo-code more self-contained.

Like the pseudo-code shown in the paper, this version of the algorithm represents the *M* action as an atomic step.

EXTENDS *Data*

```

--algorithm Decon{
  variables number = [p ∈ Procs ↦ 0],
    
```

```

    localNum = [p ∈ Procs ↦ [q ∈ OtherProcs(p) ↦ 0]],
    localCh  = [p ∈ Procs ↦ [q ∈ OtherProcs(p) ↦ 0]];

fair process ( main ∈ ProcIds )
{
  ncs:- while ( TRUE ) {
    skip; noncritical section
    M: await ∀ p ∈ SubProcsOf(self[1]) : pc[p] = "test" ;
    with ( v ∈ {n ∈ Nat \ {0} : ∀ j ∈ OtherProcs(self[1]) :
      localNum[self[1]][j] ≠ qm ⇒ n > localNum[self[1]][j]} ) {
      number[self[1]] := v ;
      localNum := [j ∈ Procs ↦
        [i ∈ OtherProcs(j) ↦
          IF i = self[1] THEN qm
          ELSE localNum[j][i]]] ;
    } ;
    L: await ∀ p ∈ SubProcsOf(self[1]) : pc[p] = "ch" ;
    cs: skip; critical section
    P: number[self[1]] := 0 ;
    localNum := [j ∈ Procs ↦
      [i ∈ OtherProcs(j) ↦
        IF i = self[1] THEN qm
        ELSE localNum[j][i]]] ;
  }
}

fair process ( sub ∈ SubProcs ) {
  ch: while ( TRUE ) {
    await pc[⟨self[1]⟩] = "M" ;
    localCh[self[2]][self[1]] := 1 ;
  test: await pc[⟨self[1]⟩] = "L" ;
    localNum[self[2]][self[1]] := number[self[1]] ;
  Lb: localCh[self[2]][self[1]] := 0 ;
  L2: await localCh[self[1]][self[2]] = 0 ;
  L3:- See below for an explanation of why there is no fairness here.
    await (localNum[self[1]][self[2]] ∉ {0, qm}) ⇒
      (⟨number[self[1]], self[1]⟩ ≪
       ⟨localNum[self[1]][self[2]], self[2]⟩)

    The await condition is written in the form A ⇒ B rather than A ∨ B because
    when TLC is finding new states, when evaluating A ∨ B it evaluates B even when
    A is true, and in this case that would produce an error if localNum[self[1]][self[2]]
    equals qm.

  }
}

```

We allow process $\langle i, j, \text{"wr"} \rangle$ to set $localNum[j][i]$ to 0 only if it has not already been set to qm by process $\langle i \rangle$ in action $M0$. We could also allow it to write 0 after that write of qm but before process $\langle i, j \rangle$ executes statement $test$. Such a write just decreases the possible executions, so eliminating this possibility doesn't forbid any possible executions.

```

fair process (  $wrp \in WrProcs$  ) {
   $wr$ : while ( TRUE ) {
    await  $\wedge localNum[self[2]][self[1]] = qm$ 
            $\wedge pc[\langle self[1] \rangle] \in \{ \text{"ncs"}, \text{"M"} \};$ 
     $localNum[self[2]][self[1]] := 0;$ 
  }
}

```

BEGIN TRANSLATION ($chksum(pcal) = \text{"7827c38d"} \wedge chksum(tla) = \text{"83cb6c12"}$)

VARIABLES $number, localNum, localCh, pc$

$vars \triangleq \langle number, localNum, localCh, pc \rangle$

$ProcSet \triangleq (ProcIds) \cup (SubProcs) \cup (WrProcs)$

$Init \triangleq$ Global variables
 $\wedge number = [p \in Procs \mapsto 0]$
 $\wedge localNum = [p \in Procs \mapsto [q \in OtherProcs(p) \mapsto 0]]$
 $\wedge localCh = [p \in Procs \mapsto [q \in OtherProcs(p) \mapsto 0]]$
 $\wedge pc = [self \in ProcSet \mapsto \text{CASE } self \in ProcIds \rightarrow \text{"ncs"}$
 $\square self \in SubProcs \rightarrow \text{"ch"}$
 $\square self \in WrProcs \rightarrow \text{"wr"}$]

$ncs(self) \triangleq \wedge pc[self] = \text{"ncs"}$
 $\wedge \text{TRUE}$
 $\wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"M"}]$
 $\wedge \text{UNCHANGED } \langle number, localNum, localCh \rangle$

$M(self) \triangleq \wedge pc[self] = \text{"M"}$
 $\wedge \forall p \in SubProcsOf(self[1]) : pc[p] = \text{"test"}$
 $\wedge \exists v \in \{n \in Nat \setminus \{0\} : \forall j \in OtherProcs(self[1]) :$
 $localNum[self[1]][j] \neq qm \Rightarrow n > localNum[self[1]][j]\}$:
 $\wedge number' = [number \text{ EXCEPT } ![self[1]] = v]$
 $\wedge localNum' = [j \in Procs \mapsto$
 $[i \in OtherProcs(j) \mapsto$
IF $i = self[1]$ THEN qm
ELSE $localNum[j][i]$]
 $\wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"L"}]$
 $\wedge \text{UNCHANGED } localCh$

$L(self) \triangleq \wedge pc[self] = \text{"L"}$

$$\begin{aligned}
& \wedge \forall p \in \text{SubProcsOf}(\text{self}[1]) : pc[p] = \text{"ch"} \\
& \wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"cs"}] \\
& \wedge \text{UNCHANGED } \langle number, localNum, localCh \rangle \\
cs(\text{self}) \triangleq & \wedge pc[self] = \text{"cs"} \\
& \wedge \text{TRUE} \\
& \wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"P"}] \\
& \wedge \text{UNCHANGED } \langle number, localNum, localCh \rangle \\
P(\text{self}) \triangleq & \wedge pc[self] = \text{"P"} \\
& \wedge number' = [number \text{ EXCEPT } ![self][1] = 0] \\
& \wedge localNum' = [j \in \text{Procs} \mapsto \\
& \quad [i \in \text{OtherProcs}(j) \mapsto \\
& \quad \quad \text{IF } i = \text{self}[1] \text{ THEN } qm \\
& \quad \quad \quad \text{ELSE } localNum[j][i]]] \\
& \wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"ncs"}] \\
& \wedge \text{UNCHANGED } localCh \\
main(\text{self}) \triangleq & ncs(\text{self}) \vee M(\text{self}) \vee L(\text{self}) \vee cs(\text{self}) \vee P(\text{self}) \\
ch(\text{self}) \triangleq & \wedge pc[self] = \text{"ch"} \\
& \wedge pc[\langle self[1] \rangle] = \text{"M"} \\
& \wedge localCh' = [localCh \text{ EXCEPT } ![self][2][self[1]] = 1] \\
& \wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"test"}] \\
& \wedge \text{UNCHANGED } \langle number, localNum \rangle \\
test(\text{self}) \triangleq & \wedge pc[self] = \text{"test"} \\
& \wedge pc[\langle self[1] \rangle] = \text{"L"} \\
& \wedge localNum' = [localNum \text{ EXCEPT } ![self][2][self[1]] = number[self[1]]] \\
& \wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"Lb"}] \\
& \wedge \text{UNCHANGED } \langle number, localCh \rangle \\
Lb(\text{self}) \triangleq & \wedge pc[self] = \text{"Lb"} \\
& \wedge localCh' = [localCh \text{ EXCEPT } ![self][2][self[1]] = 0] \\
& \wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"L2"}] \\
& \wedge \text{UNCHANGED } \langle number, localNum \rangle \\
L2(\text{self}) \triangleq & \wedge pc[self] = \text{"L2"} \\
& \wedge localCh[self[1]][self[2]] = 0 \\
& \wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"L3"}] \\
& \wedge \text{UNCHANGED } \langle number, localNum, localCh \rangle \\
L3(\text{self}) \triangleq & \wedge pc[self] = \text{"L3"} \\
& \wedge (localNum[self[1]][self[2]] \notin \{0, qm\}) \Rightarrow \\
& \quad (\langle number[self[1]], self[1] \rangle \ll \\
& \quad \quad \langle localNum[self[1]][self[2]], self[2] \rangle) \\
& \wedge pc' = [pc \text{ EXCEPT } ![self] = \text{"ch"}]
\end{aligned}$$

$$\begin{aligned}
& \wedge \text{UNCHANGED } \langle \text{number}, \text{localNum}, \text{localCh} \rangle \\
\text{sub}(\text{self}) & \triangleq \text{ch}(\text{self}) \vee \text{test}(\text{self}) \vee \text{Lb}(\text{self}) \vee \text{L2}(\text{self}) \vee \text{L3}(\text{self}) \\
\text{wr}(\text{self}) & \triangleq \wedge \text{pc}[\text{self}] = \text{"wr"} \\
& \wedge \wedge \text{localNum}[\text{self}[2]][\text{self}[1]] = \text{qm} \\
& \wedge \text{pc}[\langle \text{self}[1] \rangle] \in \{ \text{"ncs"}, \text{"M"} \} \\
& \wedge \text{localNum}' = [\text{localNum} \text{ EXCEPT } ![\text{self}[2]][\text{self}[1]] = 0] \\
& \wedge \text{pc}' = [\text{pc} \text{ EXCEPT } ![\text{self}] = \text{"wr"}] \\
& \wedge \text{UNCHANGED } \langle \text{number}, \text{localCh} \rangle \\
\text{wrp}(\text{self}) & \triangleq \text{wr}(\text{self}) \\
\text{Next} & \triangleq (\exists \text{self} \in \text{ProcIds} : \text{main}(\text{self})) \\
& \vee (\exists \text{self} \in \text{SubProcs} : \text{sub}(\text{self})) \\
& \vee (\exists \text{self} \in \text{WrProcs} : \text{wrp}(\text{self})) \\
\text{Spec} & \triangleq \wedge \text{Init} \wedge \square[\text{Next}]_{\text{vars}} \\
& \wedge \forall \text{self} \in \text{ProcIds} : \text{WF}_{\text{vars}}((\text{pc}[\text{self}] \neq \text{"ncs"}) \wedge \text{main}(\text{self})) \\
& \wedge \forall \text{self} \in \text{SubProcs} : \text{WF}_{\text{vars}}((\text{pc}[\text{self}] \neq \text{"L3"}) \wedge \text{sub}(\text{self})) \\
& \wedge \forall \text{self} \in \text{WrProcs} : \text{WF}_{\text{vars}}(\text{wrp}(\text{self}))
\end{aligned}$$

END TRANSLATION

In statement *L3*, the await condition is satisfied if process $\langle i, j \rangle$ reads $\text{localNum}[\text{self}[1]][\text{self}[2]]$ equal to qm . This is because that's a possible execution, since the process could "interpret" the qm as 0. For checking safety (namely, mutual exclusion), we want to allow that because it's a possibility that must be taken into account. However, for checking liveness, we don't want to require that the statement must be executed when $\text{localNum}[\text{self}[1]][\text{self}[2]]$ equals qm , since that value could also be interpreted as $\text{localNum}[\text{self}[1]][\text{self}[2]]$ equal to 1, which could prevent the wait condition from being true. So we omit that fairness condition from the formula *Spec* produced by translating the algorithm, and we add weak fairness of the action when $\text{localNum}[\text{self}[1]][\text{self}[2]]$ does not equal qm . This produces the TLA+ specification *FSpec* defined here.

$$\begin{aligned}
\text{FSpec} & \triangleq \wedge \text{Spec} \\
& \wedge \forall q \in \text{SubProcs} : \text{WF}_{\text{vars}}(\text{L3}(q) \wedge (\text{localNum}[q[1]][q[2]] \neq \text{qm})) \\
\text{TypeOK} & \triangleq \wedge \text{number} \in [\text{Procs} \rightarrow \text{Nat}] \\
& \wedge \wedge \text{DOMAIN } \text{localNum} = \text{Procs} \\
& \wedge \forall i \in \text{Procs} : \text{localNum}[i] \in [\text{OtherProcs}(i) \rightarrow \text{Nat} \cup \{ \text{qm} \}] \\
& \wedge \wedge \text{DOMAIN } \text{localCh} = \text{Procs} \\
& \wedge \forall i \in \text{Procs} : \text{localCh}[i] \in [\text{OtherProcs}(i) \rightarrow \{0, 1\}] \\
\text{MutualExclusion} & \triangleq \forall p, q \in \text{ProcIds} : (p \neq q) \Rightarrow (\{ \text{pc}[p], \text{pc}[q] \} \neq \{ \text{"cs"} \}) \\
\text{StarvationFree} & \triangleq \forall p \in \text{ProcIds} : (\text{pc}[p] = \text{"M"}) \rightsquigarrow (\text{pc}[p] = \text{"cs"})
\end{aligned}$$

Checking the invariant in the appendix of the paper.

$$\text{inBakery}(i, j) \triangleq \vee \text{pc}[\langle i, j \rangle] \in \{ \text{"Lb"}, \text{"L2"}, \text{"L3"} \}$$

$$\begin{aligned} & \vee \wedge pc[\langle i, j \rangle] = \text{"ch"} \\ & \wedge pc[\langle i \rangle] \in \{\text{"L"}, \text{"cs"}\} \end{aligned}$$

$$inCS(i) \triangleq pc[\langle i \rangle] = \text{"cs"}$$

In TLA+, we can't write both $inDoorway(i, j, w)$ and $inDoorway(i, j)$, so we change the first to $inDoorwayVal$. Its definition differs from the definition of $inDoorway(i, j, w)$ in the paper to avoid having to add a history variable to remember the value of $localNum[self[1]][j]$ read in statement $M0$. It's a nicer definition, but it would have required more explanation than the definition in the paper.

The definition of $inDoorway(i, j)$ is equivalent to the one in the paper. It is obviously implied by $\exists w \in Nat : inDoorwayVal(i, j, w)$, and type correctness implies the opposite implication.

$$\begin{aligned} inDoorwayVal(i, j, w) \triangleq & \wedge pc[\langle i \rangle] = \text{"L"} \\ & \wedge pc[\langle i, j \rangle] = \text{"test"} \\ & \wedge number[i] > w \end{aligned}$$

$$\begin{aligned} inDoorway(i, j) \triangleq & \wedge pc[\langle i \rangle] = \text{"L"} \\ & \wedge pc[\langle i, j \rangle] = \text{"test"} \end{aligned}$$

$$Outside(i, j) \triangleq \neg(inDoorway(i, j) \vee inBakery(i, j))$$

$$\begin{aligned} passed(i, j, LL) \triangleq & \text{IF } LL = \text{"L2"} \text{ THEN } \vee pc[\langle i, j \rangle] = \text{"L3"} \\ & \vee \wedge pc[\langle i, j \rangle] = \text{"ch"} \\ & \wedge pc[\langle i \rangle] \in \{\text{"L"}, \text{"cs"}\} \\ & \text{ELSE } \wedge pc[\langle i, j \rangle] = \text{"ch"} \\ & \wedge pc[\langle i \rangle] \in \{\text{"L"}, \text{"cs"}\} \end{aligned}$$

$$\begin{aligned} Before(i, j) \triangleq & \wedge inBakery(i, j) \\ & \wedge \vee Outside(j, i) \\ & \vee inDoorwayVal(j, i, number[i]) \\ & \vee \wedge inBakery(j, i) \\ & \wedge \langle number[i], i \rangle \ll \langle number[j], j \rangle \\ & \wedge \neg passed(j, i, \text{"L3"}) \end{aligned}$$

$$\begin{aligned} Inv(i, j) \triangleq & \wedge inBakery(i, j) \Rightarrow Before(i, j) \vee Before(j, i) \\ & \vee inDoorway(j, i) \\ & \wedge passed(i, j, \text{"L2"}) \Rightarrow Before(i, j) \vee Before(j, i) \\ & \wedge passed(i, j, \text{"L3"}) \Rightarrow Before(i, j) \end{aligned}$$

$$I \triangleq \forall i \in Procs : \forall j \in OtherProcs(i) : Inv(i, j)$$

The following is for testing. Since the spec allows the values of $number[n]$ to get arbitrarily large, there are infinitely many states. The obvious solution to that is to use models with a state constraint that $number[n]$ is at most some value $TestMaxNum$. However, TLC would still not be able to execute the spec because the with statement in action M allows an infinite number of possible values for $number[n]$. To solve that problem, we have the model redefine Nat to a finite set of numbers. The obvious set is $0 \dots TestMaxNum$. However, trying that reveals a subtle problem. Running the model produces a bogus counterexample to the $StarvationFree$ property.

This is surprising, since constraints on the state space generally fail to find real counterexamples to a liveness property because the counterexamples require large (possibly infinite) traces that are ruled out by the state constraint. The remaining traces may not satisfy the liveness property, but they are ruled out because they fail to satisfy the algorithm's fairness requirements. In this case, a behavior that didn't satisfy the liveness property *StarvationFree* but shouldn't have satisfied the fairness requirements of the algorithm did satisfy the fairness requirement because of the substitution of a finite set of numbers for *Nat*.

Here's what happened: In the behavior, two nodes kept alternately entering the critical section in a way that kept increasing their values of *num* until one of those values reached *TestMaxNum*. That one entered its critical section while the other was in its noncritical section, re-entered its noncritical section, and then the two processes kept repeating this dance forever. Meanwhile, a third process's subprocess was trying to execute action *M*. Every time it tried to execute that action, it saw that another process's number equaled *TestMaxNum*. In a normal execution, it would just set its value of *num* larger than *TestMaxNum* and eventually enter its critical section. However, it couldn't do that because the substitution of $0 \dots TestMaxNum$ for *Nat* meant that it couldn't set *num* to such a value, so the enter step was disabled. The fairness requirement on the enter action is weak fairness, which requires an action eventually to be taken only if it's continually enabled. Requiring strong fairness of the action would have solved this problem, because the enabled action kept being enabled and strong fairness would rule out a behavior in which that process's enter step never occurred. However, it's important that the algorithm satisfy starvation freedom without assuming strong fairness of any of its steps.

The solution to this problem is to substitute $0 \dots (TestMax + 1)$ for *Nat*. The state constraint will allow the enter step to be taken, but will allow no further steps from that state. The process still never enters its critical section, but now the behavior that keeps it from doing so will violate the weak fairness requirements on that process's steps.

$$TestMaxNum \triangleq 4$$
$$TestNat \triangleq 0 \dots (TestMaxNum + 1)$$

```
\* Modification History
\* Last modified Tue Nov 16 18:37:41 CET 2021 by merz
\* Last modified Thu Jul 01 12:24:37 CEST 2021 by merz
\* Last modified Wed Apr 28 18:06:24 PDT 2021 by lamport
\* Created Sat Apr 24 09:45:26 PDT 2021 by lamport6
```