

second edition

PRINCIPLES ^{OF} CONCURRENT ^{AND} DISTRIBUTED PROGRAMMING

M. BEN-ARI

 ADDISON-WESLEY



We work with leading authors to develop the strongest educational materials in computing, bringing cutting-edge thinking and best learning practice to a global market.

Under a range of well-known imprints, including Addison-Wesley, we craft high quality print and electronic publications which help readers to understand and apply their content, whether studying or at work.

To find out more about the complete range of our publishing, please visit us on the World Wide Web at: www.pearsoned.co.uk

Algorithm 5.3: Bakery algorithm without atomic assignment	
boolean array[1..n] choosing \leftarrow [false, ..., false] integer array[1..n] number \leftarrow [0, ..., 0]	
loop forever	
p1:	non-critical section
p2:	choosing[i] \leftarrow true
p3:	number[i] \leftarrow 1 + max(number)
p4:	choosing[i] \leftarrow false
p5:	for all <i>other</i> processes j
p6:	await choosing[j] = false
p7:	await (number[j] = 0) or (number[i] \ll number[j])
p8:	critical section
p9:	number[i] \leftarrow 0

Algorithm 5.3 differs from Algorithm 5.2 by the addition of a boolean array choosing. If choosing[i] is true, a process is in the act of choosing a ticket number, and other processes must wait for the choice to be made before comparing numbers.

Each global variable is written to by exactly one process, so there are no cases in which multiple write operations to the same variable can overlap. It is also reasonable to assume that if a set of read operations to a variable does not overlap a write operation to that variable, then they all return the correct value. However, if read operations overlap a write operation, it is possible that inconsistent values may be obtained. Lamport has shown that the bakery algorithm is correct even if such read operations return arbitrary values [36].

5.4 Fast algorithms

In the bakery algorithm, a process wishing to enter its critical section must read the values of all the elements of the array number to compute the maximum ticket number, and it must execute a loop that contains await statements that read the values of all the elements of the two arrays. If there are dozens of processes in the system, this is going to be extremely inefficient. The overhead may be unavoidable: if, in general, processes attempt to enter their critical sections at short intervals, there will be a lot of contention, and a process really will have to query the state of all the other processes. If, however, contention is low, it makes sense to search for an algorithm for the critical section problem that is very efficient, in the sense that only if there is contention does the process incur the significant overhead of querying the other processes.

An algorithm for the critical section problem is *fast*, if in the absence of contention a process can access its critical section by executing pre- and postprotocols consisting of a *fixed* (and small) number of statements, none of which are await statements. The first fast algorithm for mutual exclusion was given by Lamport in [40], and this paper initiated extensive research into algorithms that are efficient under various assumptions about the characteristics of a system, such as the amount of contention and the behavior of the system under errors. Here we will describe and prove Lamport's algorithm restricted to two processes and then note the changes needed for an arbitrary number of processes.

Outline of the fast algorithm

Here is the outline of the algorithm, where we assume that the process identifiers p and q are encoded as nonzero integer constants so that their values can be assigned to variables and compared with the values of these variables:

Algorithm 5.4: Fast algorithm for two processes (outline)	
integer gate1 \leftarrow 0, gate2 \leftarrow 0	
p	q
loop forever non-critical section p1: gate1 \leftarrow p p2: if gate2 \neq 0 goto p1 p3: gate2 \leftarrow p p4: if gate1 \neq p if gate2 \neq p goto p1 critical section p6: gate2 \leftarrow 0	loop forever non-critical section q1: gate1 \leftarrow q q2: if gate2 \neq 0 goto q1 q3: gate2 \leftarrow q q4: if gate1 \neq q if gate2 \neq q goto q1 critical section q6: gate2 \leftarrow 0

(To simplify the proof of correctness, we have given an abbreviated algorithm as was done in the previous chapters.)

The concept of the algorithm is displayed graphically in Figures 5.1–5.3. A process is represented by a stick figure. The non-critical section is at the left of each diagram; a process must pass through two gates to reach the critical section on the right side of the diagram. Figure 5.1 shows what happens in the absence of contention. The process enters the first gate, writing its ID number on the gate [p1, (a)]. (The notation correlates the line number p1 in the algorithm with diagram (a) in the figure.) It then looks at the second gate [p2, (b)]; in the absence of contention, this gate will not have an ID written upon it, so the process writes its ID on that gate also [p3, (c)]. It then looks back over its shoulder at the first gate to check




Figure 5.1: Fast algorithm—no contention

if its ID is still written there [p4, (d)]. If so, it enters the critical section [p6, (e)]. Upon leaving the critical section and returning to the non-critical section, it erases its ID from the second gate [p6, (f)].

In the absence of contention, the algorithm is very efficient, because a process can access its critical section at the cost of three statements that assign a constant to a global variable (p1, p3, p6), and two if statements that check that the value of a global variable is not equal to a constant (p2, p4).

The case of contention at the second gate is shown in Figure 5.2. The process enters the first gate [p1, (a)], but when it checks the second gate [p2, (b)], it sees that the other process q has already written its ID, preparing to enter the critical section. Since q has gotten through both gates sooner, process p returns to before the first gate [p2, (c)] to try again [p1, (d)].

In the case of contention at the first gate, the algorithm is a bit more complex (Figure 5.3). Initially, the algorithm proceeds as before [p1, (a)], [p2, (b)], [p3, (c)], until looking back over its shoulder, process p perceives that process q has entered the first gate and written its ID [p4, (d)]. There are now two possibilities:




Figure 5.2: Fast algorithm—contention at gate 2




Figure 5.3: Fast algorithm—contention at gate 1

the ID of process p is still written on the second gate; if so, it continues into the critical section [p5, (e)]. Or, process q has passed process p and written its ID on the second gate; if so, the new process defers and returns to the first gate [p5, (f)].

The outline of the algorithm is not correct; in the exercises you are asked to find a scenario in which mutual exclusion is not satisfied. We will now partially prove the outline of the algorithm and then modify the algorithm to ensure correctness.

Partial proof of the algorithm

To save turning pages, here is the algorithm again:

Algorithm 5.5: Fast algorithm for two processes (outline)	
integer gate1 \leftarrow 0, gate2 \leftarrow 0	
p	q
loop forever non-critical section p1: gate1 \leftarrow p p2: if gate2 \neq 0 goto p1 p3: gate2 \leftarrow p p4: if gate1 \neq p p5: if gate2 \neq p goto p1 critical section p6: gate2 \leftarrow 0	loop forever non-critical section q1: gate1 \leftarrow q q2: if gate2 \neq 0 goto q1 q3: gate2 \leftarrow q q4: if gate1 \neq q q5: if gate2 \neq q goto q1 critical section q6: gate2 \leftarrow 0

Lemma 5.4 The following formulas are invariant:

$$p5 \wedge gate2 = p \rightarrow \neg(q3 \vee q4 \vee q6) \quad (5.1)$$

$$q5 \wedge gate2 = q \rightarrow \neg(p3 \vee p4 \vee p6) \quad (5.2)$$

We *assume* the truth of this lemma and use it to prove the following lemma:

Lemma 5.5 The following formulas are invariant:

$$p4 \wedge gate1 = p \rightarrow gate2 \neq 0 \quad (5.3)$$

$$p6 \rightarrow gate2 \neq 0 \wedge \neg q6 \wedge (q3 \vee q4 \rightarrow gate1 \neq q) \quad (5.4)$$

$$q4 \wedge gate1 = q \rightarrow gate2 \neq 0 \quad (5.5)$$

$$q6 \rightarrow gate2 \neq 0 \wedge \neg p6 \wedge (p3 \vee p4 \rightarrow gate1 \neq p) \quad (5.6)$$

Mutual exclusion follows immediately from invariants (5.4) and (5.6).