

4) The pseudo-code (incomplete) presented below implements a producer-consumer pattern by using a shared array *BUF* of size *N*. Knowing that synchronization is based in semaphores (*S1* and *S2*), complete the missing parts in the pseudo-code.

1. operation B.produce(v) is	1. operation B.consume() is
2. S1. _____;	2. _____;
3. BUF[in].put(v);	3. r = BUF[out]._____;
4. _____;	4. _____;
5. S2. _____;	5. _____;
6. return ();	6. return (r);
8. end operation ;	8. end operation ;

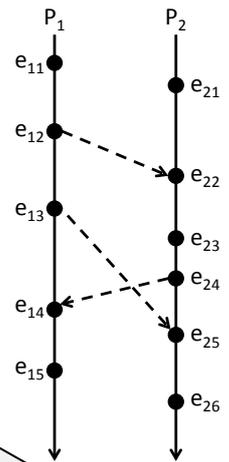
5) Mark each of the following statements as *true* (T or V) or false (F). (*Note: wrongly marking a statement will reduce your grade in this question*).

- Obstruction free programs always exhibit progress of the program as a whole.
- Lock-free programs may suffer from starvation.
- Obstruction freedom provides more guaranties than lock freedom.
- Lock freedom provides more guaranties than wait freedom.
- Obstruction-free programs may deadlock.
- Wait-free programs may deadlock.
- Under low contention, lock-free programs behave as wait-free programs.
- Wait-free programs may suffer from starvation.
- Lock-free programs may deadlock.
- Obstruction-free programs may suffer from starvation.
- Under high contention, lock-free programs behave as obstruction-free programs.
- All waif-free programs are also lock-free programs.
- Obstruction-free programs always generate linearizable executions.
- Wait-free programs always exhibit progress of the program as a whole.
- All linearizable programs are sequential consistent.
- Linearizability is a local property.
- Sequential consistency is a global property.
- The composition of linearized executions is linearizable.
- The composition of linearized executions is linearizable.
- All sequential consistent executions are also linearizable.

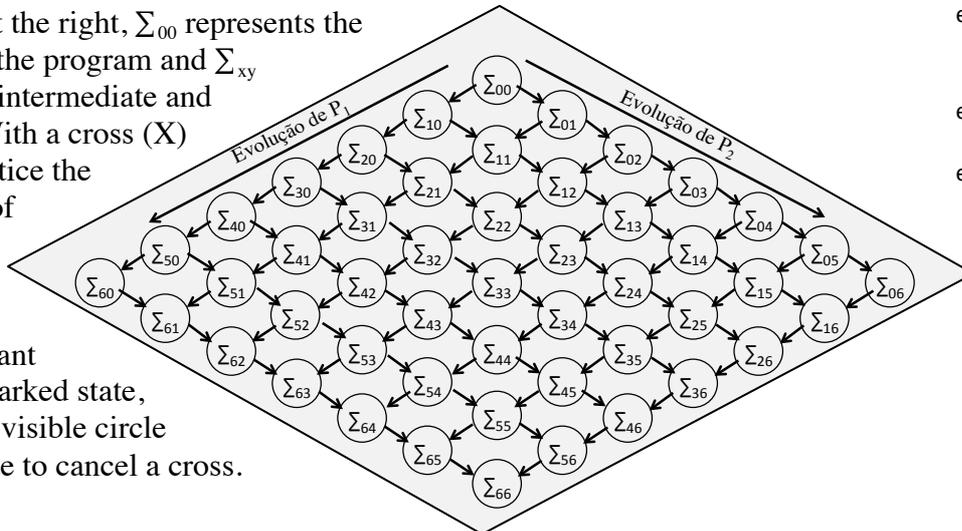
6) Consider a program with multiple threads. In which of the following situations **at least one thread may be progressing** in its intended computation? Mark each case with Progress (P) or No Progress (N).

- deadlock livelock starvation busy-waiting
- all threads simultaneously execute a *lock(X)* over the same lock variable *X*
- all threads simultaneously execute a *compareAndSwap(X,v)* over the same variable *X*

7) Consider that the processes P1 and P2 belong to the same program. Processes P1 and P2 interact by exchanging messages as shown in the figure at the right side of this question. The full line arrows denote time and the dashed line arrows denote communication events.



a) In the lattice at the right, Σ_{00} represents the initial state of the program and Σ_{xy} represents the intermediate and final states. With a cross (X) mark in the lattice the invalid states of this program.



If you make a mistake and want to *unmark* a marked state, draw a clearly visible circle around the state to cancel a cross.

b) Knowing that ' $i(e_{nm})$ ' and ' $r(e_{nm})$ ' correspond to the invocation and reply of na event ' e_{nm} ' that occurred during the execution of the program, show:

i) A sequential history that is **complete and consistent**. If there is none, write "NONE" or "NÃO HÁ".

ii) Another sequential history that is **complete and consistent** and different form i), in which at least three events are in different positions than in i). If there is none, write "NONE" or "NÃO HÁ".

iii) Explain when two sequential histories are equivalent. If they cannot be equivalent, write "NONE" or "NÃO HÁ".

iv) A concurrent history that is **equivalent** to the history i) above. If there is none, write "NONE" or "NÃO HÁ".

v) A concurrent history that is **not equivalent** to the history i) above. If there is none, write "NONE" or "NÃO HÁ".

8) Consider the code shown at the right that presents the definition of the class *Node* and part of the class *LockFreeStack*, in particular its private fields and the code relevant to the implementation of the method *push()*. The method *backoff.backoff()* does a random pause.

The class *AtomicReference* has two relevant methods, namely *get()* that returns the reference and *compareAndSet(old, new)* with the usual semantics.

Fill the fields below with the code that should be present in lines 8, 9 and 10.

8:

9:

10:

```

1  public class LockFreeStack<T> {
2      AtomicReference<Node> top = new AtomicReference<Node>(null);
3      static final int MIN_DELAY = ...;
4      static final int MAX_DELAY = ...;
5      Backoff backoff = new Backoff(MIN_DELAY, MAX_DELAY);
6
7      protected boolean tryPush(Node node){
8          Node oldTop = ;
9          node.next = ;
10         return();
11     }
12     public void push(T value) {
13         Node node = new Node(value);
14         while (true) {
15             if (tryPush(node)) {
16                 return;
17             } else {
18                 backoff.backoff();
19             }
20         }
21     }

```

```

1  public class Node {
2      public T value;
3      public Node next;
4      public Node(T value)
5          value = value;
6          next = null;
7      }
8  }

```

9) Considere um sistema de memória transacional, em que duas transações estão a executar concorrentemente. Considere ainda as seguintes transações:

- T₁: R(50, a); R(100, b); W(120, 8); R(110, c)
- T₂: W(50, 3); R(130, d); W(110,8); R(120, e)
- T₃: R(50, f); R(100, g); R(120, h)

Knowing that *T_i(N) / T_j(M)* mean that transactions *T_i* and *T_j* execute in parallel starting respectively in time 'i' and 'j', fill the table below.

<i>T_i // T_j</i>	Read Set T _i	Read Set T _j	Write Set T _i	Write Set T _j	T _i commit or abort? Why?	T _j commit or abort? Why?
T ₁ (5) // T ₂ (7)						
T ₁ (8) // T ₃ (4)						
T ₂ (2) // T ₃ (8)						

END!!!