Exam

Data replication and consistency maintenance

M2

UNIVERSITÉ DE LORRAINE

★ Exercice 1.

▷ **Question 1.** Consider the two scenarios below where four processes operate on the same data item x and execute a sequence of operations. Operations executed by each process are given along the horizontal time axis. The symbol $W(x)a$ means that the corresponding process executed a write to data item x with the value a. The symbol $R(x)b$ means that the corresponding process read the value of data item x and returned b.

```
P1: W(x)1                    W(x)3
P2:       R(x)1   W(x)2
P3:       R(x)1               R(x)3   R(x)2
P4:       R(x)1               R(x)2   R(x)3
```

```
P1: W(x)1
P2:       R(x)1   W(x)2
P3:                           R(x)2   R(x)1
P4:                           R(x)1   R(x)2
```

(a) Give the definition of sequential consistency.

(b) Are the two above executions sequentially consistent ? Justify your answer.

(c) Give the definition of causal consistency.

(d) Are the two above executions causally consistent ? Justify your answer.

(e) Give the definition of FIFO consistency.

(f) Are the two above executions FIFO consistent ? Justify your answer.

▷ **Question 2.** We consider a library as a special kind of a storage system and the following scenario. Pierre borrowed a book on distributed systems a while ago. Now he returns the book at the library's front desk. Afterwards, he wants to borrow another book about user studies. Therefore, he checks the library's index. He notices that the book he has just returned is not listed. Afterwards, he remembered that one of his colleagues mentioned some time ago about a good book on user studies. Therefore, he goes and asks his colleague about the user studies book he mentioned. He then returns to the library and rechecks the library index. Now the book about distributed systems is listed again. Assuming that there were no other clients in the library except Pierre, which of the consistencies listed below are violated ?

(a) Linearizability

(b) Sequential consistency

(c) Monotonic read consistency

(d) Read-your-writes consistency

▷ **Question 3.** Consider three concurrently executing processes P_1, P_2 and P_3. The statements executed by each process are given along the vertical axis.

```
Process P1      Process P2      Process P3
-----------------------------------------
x=1;            y=1;            z=1;
print(y,z);     print(x,z);     print(x,y);
```

Assume that each variable is initialised to 0. An assignment corresponds to a write operation and a print statement corresponds to a simultaneous read of its two arguments. Various execution sequences are possible, one of them is given below together with the generated display.

```
x=1;
print(y,z);
y=1;
print(x,z);
z=1;
print(x,y);

Prints: 001011
```

Is 000000 a legal output for a sequential execution ? And 001001 ? And 111111 ? Justify your answer. If the display is allowed, give the corresponding execution.

★ Exercice 2.

▷ **Question 1.** What means "ensuring causality" ? Give an example where causality is not respected.

▷ **Question 2.** Consider the collaborative editing scenario in the figure below.

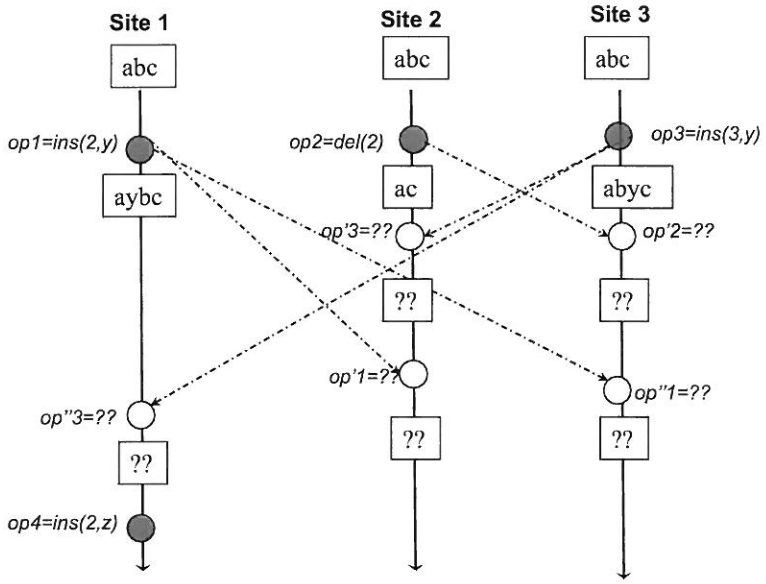

FIGURE 1 – Operational Transformation Scenario

(a) We suppose that initially the state vectors on the three sites are are $[0, 0, 0]$. What are the state vectors associated with each operation (op1, op2, op3, op4, op'1, op"1, op'2, op'3,op"3). Re-draw the figure by adding the state vectors.

(b) Being given the following transformation functions

```
T(Ins(p1,c1),Ins(p2,c2)) :-
   if p1<p2 return Ins(p1, c1)
   else return Ins(p1 + 1, c1)

T(Ins(p1,c1),Del(p2)) :-
   if p1 <= p2 return Ins(p1,c1)
   else return Ins(p1 - 1,c1)

T(Del(p1),Ins(p2,c2)) :-
   if p1 < p2 return Del(p1)
   else return Del(p1 + 1)
```

```
T(Del(p1),Del(p2)) :-
   if p1 < p2 return Del(p1)
   else if p1 > p2 return Del(p1 - 1)
   else return Id()
```

complete the figure with the missing operations and states.

(c) Are these transformation functions respecting condition C2 ? Justify your answer.

▷ **Question 3.** We would like to use So6 algorithm for synchronisation. We suppose that two users execute the scenario illustrated below. The first user performs operation op_1, while second user performs the sequence of operations op_2, op_3. Afterwards the two users initiate the synchronisation procedure as illustrated in the table.

User1	User2
op_1	op_2 op_3
synchronise	
	synchronise
synchronise	

(a) What are the transformations performed at each synchronise step ?

(b) What are the operations executed at each site at each synchronise step ?

★ **Exercice 3.** Let us consider the document content (see figure below right part) and the identifiers associated to each line of the document (see figure below left part) and generated by using Logoot algorithm. Each tuple of the identifier is of the form $< p, s, h >$ where p is an integer in the interval [1,9], s is the identifier of the site that generated the identifier and h is the clock of site s. Suppose site 7 inserts 5 lines between the second and the third lines, i.e. between line "Stars :" and "Johnny Galecki". What are the possible identifiers associated to these 5 lines ?

```
<0,NA,NA>
<1,1,1>
<2,1,2><3,2,3><3,6,3>
<2,3,4><5,4,4><4,5,4>
<MAX,NA,NA>
```

```
Title: Big bang theory
Stars:
Johnny Galecki
```