

REPLY TO "ERRORS IN 'PROCESS SYNCHRONIZATION IN DATABASE SYSTEMS'"

G. Schlageter
FernUniversität
Praktische Informatik
Postfach 940
D 5800 Hagen

As I read the introduction of "Errors in 'Process Synchronization in Database Systems'" by Bernstein, Casanova, Goodmann /BCG/ I thought that everything was wrong in my paper - the authors in fact present a global and extremely rough criticism. Fortunately, after reading the rest of the paper, this first impression proved to be false.

I think it is not of interest to the readers of SIGMOD to go through all the details of /BCG/, so I shall concentrate on essential points only.

I admit that in my theory of serializability the special situation of a "dead update" is considered non-serializable. This situation occurs if, on some object x,

```
process A reads
process B updates
process A updates
process C updates.
```

The interpretation of this schedule in my theory is that the update of B is lost, and thus this situation is considered unacceptable. By the update of C this "local disturbance" is obscured, and one can consider the above schedule as serializable. If one does so, then theorem 3.1 of my paper, and of course theorems 3.2 and 3.6 are too stringent (for this special situation only). From a practical point of view, dead updates make no sense, and it is debatable whether the above situation must be considered acceptable.

This is the only point of criticism which is of interest, and certainly a valuable comment. All the rest of the discussion in /BCG/ is based on misinterpretations. Just three examples:

1. Weak consistency

I never suggested that weak consistency is equivalent to consistency, or that processes get consistent views if a schedule is weakly consistent. The contrary has been demonstrated in my paper, and it is not necessary to find new examples to show this fact again.

I have included the discussion on weak consistency for the only reason that, at the time when I produced my paper, many people considered weak consistency as a sufficient condition for correct synchronization.

2. Distinction of test actions and test-and-read actions

All of the problems discussed in /BCG/ are introduced by the fact that /BCG/ considers a test action (t^-) as a read action. However, t^- is not a read in the sense of the definition of weak consistency, and thus the discussion of example 3 is erroneous.

I want to stress that the interpretation of /BCG/ "if P_i and P_j each read consistent data, then under the conditions of the theorem any database produced by their execution is also consistent" (p.24) cannot be derived from my paper, because I clearly distinguish between consistency and weak consistency. Example 4 demonstrates a fact which is well contained in my paper.

3. Adaptation

Obviously, theorem 4.7 has the meaning: "If the system must guarantee weak consistency for all possible situations, i.e. in general, then ...". All persons involved in this paper so far, had no doubt about this meaning. Of course, one may easily construct special cases where less synchronization is required.

As to example 6, the same arguments hold as in point 2 above (incorrect interpretation of t^-).

Though the criticism of /BCG/ cannot be accepted in the given form, I am - like /BCG/ - of the opinion that other approaches and models have to be discussed. In the meantime serializability theory has made good progress in this direction.

/BCG/ P.A. Bernstein, M.A. Casanova, N. Goodman:
 "Errors in 'Process Synchronization in Database
 Systems'", ACM SIGMOD RECORD 11(1981), 9-29