

Programming Applications with Databases

Exercise Set 4

1. Explain shortly the meaning and reasoning for ACID in transactions.
[1p]
2. Explain shortly *savepoints* concept in DBMS. Using the example presented at <https://learn.microsoft.com/en-us/sql/t-sql/language-elements/save-transaction-transact-sql>, create a new one based on *AdventureWorksLT* database. Propose and implement a use case that includes the context in which a stored procedure is executed.
[1p]
3. Create three examples showing the anomalies of dirty read, non-repeatable read, and phantom read, respectively.
[2p]
4. Execute samples 1 and 2 from the example *02-poziomy-izolacji-blokady.sql* from the attached materials. Present and explain the locks as discussed during the lecture.
[2p]
5. Execute the example *03-zakleszczenie.sql* from the attached materials several times on different isolation levels. Check when the deadlock occurs and when it does not, and explain why it happens.
[2p]
6. Explain the locking hints concept. Prepare a small example where the current transaction isolation level is serializable and `SELECT ... WITH (NOLOCK)` is executed. Show the difference in applied locks when `NOLOCK` is introduced and when it is not.
[1p]
7. Based on a small (theoretical) example, explain the difference between pessimistic and optimistic concurrency control mechanisms.
[1p]

Paweł Rajba