

CSE 486/586





































Sequential Consistency

- Your storage appears to process all requests in a single interleaved ordering (single client), where...
 - ...each and every process's program order is preserved (single copy),
 - ...and each process's program order is only *logically* preserved, i.e., it doesn't need to preserve its physical-time ordering.
- It works as if all clients are reading out of a single copy.
 - This meets the expectation from an (isolated) client, working with a single copy.
 - Linearizability meets the expectation of all clients even if they all know what others are doing.
 - Both sequential consistency and linearizability provide an illusion of a single copy.

CSE 486/586

Sequential Consistency vs. Linearizability · Both should behave as if there were only a single copy, and a single client. It's just that SC doesn't preserve the physical-time order, but just the program order of each client. • Difference You (NY) x.write(5) Friend (CA) x.write(2) read(x)? - Linearizability: Once a write is returned, the system is obligated to make the result visible to all clients based on physical time. I.e., the system has to return 5 in the example. - Sequential consistency: Even if a write is returned, the system is not obligated to make the result visible to other clients immediately. I.e., the system can still return 2 in the

Implementing Sequential Consistency

· In what implementation would the following happen?

a.read()->A a.read()->B

23

- P1: a.write(A)
- P2: a.write(B)
- P3: a.read()->B a.read()->A
- P4:
- Possibility
 - P3 and P4 use different copies.
 - In P3's copy, P2's write arrives first and gets applied.
 - In P4's copy, P1's write arrives first and gets applied.
 - Writes are applied in different orders across copies.
 - This doesn't provide sequential consistency.

CSE 486/586

Implementing Sequential Consistency

- When implementing a consistency model, we need to think about how to handle writes and how to handle reads
- · Handling writes
 - Single-client, per-process single-copy: Write synchronization happens (or writes are applied) in the same order everywhere across different copies, while preserving each process's logical write order.
 - The synchronization does not have to be complete at the time of return from a write operation. (I.e., actual writes on different copies can be done at different times.)

· Handling reads

 Single-client, per-process single-copy: A read from a process should be done on a copy that already has applied the process's latest write. And all reads should be processed by the program order.

С







