**COMPUTER SCIENCE TRIPOS Part IB – 2013 – Paper 5**

## 8 Concurrent and Distributed Systems (RNW)

(*a*) The ACID properties are often used to define transactional semantics.

(*i*) Define "atomicity" as used in the ACID context. [1 mark]

(*ii*) Define "durability" as used in the ACID context. [1 mark]

(*b*) Write-ahead logging is a commonly used scheme to accomplish transactional semantics when storing a database on a block storage device, such as a hard disk.

(*i*) Under what circumstances, during write-ahead log recovery, can a transaction in the UNDO list be moved to the REDO list? [2 marks]

(*ii*) Synchronously flushing commit records to disk is expensive. How can we safely reduce synchronous I/O operations on a high-throughput system without sacrificing ACID properties? [2 marks]

(*iii*) Describe two performance changes that might arise from using your solution to part (*b*)(*ii*). [2 marks]

(*c*) (*i*) Transaction records in a write-ahead logging scheme contain five fields: $\langle \mathit{TransactionID}, \mathit{ObjID}, \mathit{Operation}, \mathit{OldValue}, \mathit{NewValue} \rangle$, but storing the complete old and new values can consume significant amounts of space. One strategy that might be employed, for reversible operations being applied to some data such as XOR by a constant, is to store only the constant arguments, rather than the full before and after data. What problems might occur as a result of this design choice? [4 marks]

(*ii*) Write-ahead logging systems must know the actual on-disk sector size for the write-ahead log to behave correctly. An errant disk vendor decides to rebrand its 512-byte sector disks as 2K-sector disks, and adjusts the value reported back to the database system. How might this affect database integrity? [4 marks]

(*iii*) Explain how a database vendor who is aware of the problem described in part (*c*)(*ii*) mitigate this problem in software, and what limitations might there be to this approach. [4 marks]