2010 Paper 7 Question 7

Comparative Architectures

(*a*) Why is a shared second-level (L2) cache typically divided into multiple banks (banked) in a chip multiprocessor? [3 marks]

(*b*) In what situation might a shared second-level cache offer a performance advantage over a memory hierarchy for a chip multiprocessor with private L2 caches? [4 marks]

(*c*) A cache controller in a chip multiprocessor snoops the bus and observes a transaction that refers to a block that its cache contains. The block is held in State M (Modified). The bus transaction has been generated by a processor wishing to read the block. Assuming a MSI (write-back invalidate) cache coherence protocol, what actions will be taken by the cache controller? [6 marks]

(*d*) How does adopting an inclusion policy simplify the implementation of a cache coherence mechanism in a chip multiprocessor with private L1 and L2 caches? [4 marks]

(*e*) How might multiple buses be exploited to enable a greater number of processors to be supported by a snoopy cache coherence protocol? [3 marks]