

Why is deadlock detection much more expensive in a distributed environment than in a centralized environment? Each non-leaf controller maintains a wait-for graph that contains relevant information from the graphs of the controllers in the subtree below it. In particular, let  $SA$ ,  $Sp$ , and  $Sc$  be controllers such that  $Sc$  is the lowest common ancestor of  $SA$  and  $Sp$  ( $Sc$  must be unique, since we are dealing with a tree).

a. Would you use a deadlock-detection scheme or a deadlock-prevention scheme?