

1. Consider the following code executed on two processes (initially, flag has the value 0):

```
P1:          P2:
... produce  while (flag == 0);
flag = 1;    ... consume
```

Will this code work as expected on an x86 machine? What about on ARM v8?

2. Two-phase locking is a commonly used technique to implement transactions in databases. What are the two phases? What is the advantage of two-phase locking? Is deadlock still possible with two-phase locking? If so, why? If not, why not?
3. Scalable Reliable Multicast (SRM) uses negative acknowledgements (NACKs) (as opposed to positive acknowledgements from every receiver to the sender indicating that the message has been received) to indicate that a message has not been received by a particular receiver. Why does the protocol broadcast these NACKs? What problem does it solve and what does it create if any?
4. Independent checkpointing has the advantage of not requiring coordination among the processes in the distributed system during failure-free operation. What are the disadvantages? Why would coordinated checkpointing be preferred?