

9 Concurrent and Distributed Systems (RNW)

This pseudocode, executing in process P_j , employs buffering to impose ordering:

```
receive(M from Pi) {           // Message M received from process Pi
  S = getSeq(M);              // Extract sequence number S
  if (S == nextSeq(Sji)) {    // If S is the next sequence number:
    deliver(M);               // Deliver M to current process (Pj)
    Sji = flush(HBQ, Sji);    // Deliver backlog from HBQ; update Sji
  } else holdback(HBQ, M);    // Else: Hold back M for future delivery
}
```

- (a) Explain what ordering model(s) this pseudocode implements. [2 marks]
- (b) Write pseudocode (with comments) for the following functions, to be used on the sender (P_i) or receiver (P_j), which accept M (a message), and S (a sequence number): [8 marks]

Receiver	<code>receive_reliably(M)</code>	Reliably receive M from P_i .
Sender	<code>send_reliably(M)</code>	Reliably send M to P_j .
Sender	<code>process_ack(S)</code>	Handle a received ACK for S from P_j .
Sender	<code>timeout(S, M)</code>	Process a timeout for S and M .

As needed, employ the following additional utility functions:

<code>drop(M)</code>	Drop received M without delivering.
<code>setSeq(M, S)</code>	Set sequence number S on message M .
<code>transmit_msg(M)</code>	Transmit message M to P_j .
<code>transmit_ack(S)</code>	Transmit an ACK with sequence number S to P_i .
<code>sched_timeout(S, M)</code>	Schedule <code>timeout(S, M)</code> to run in 5 ms.
<code>cancel_timeout(S)</code>	If scheduled, cancel timeout for S .

- (c) Define the *happens-before* relationship. [2 marks]
- (d) The pseudocode above imposes ordering on pair-wise communications. Assuming reordering but no message loss, write pseudocode (with comments) for the following functions supporting *causal ordering* for group communications: [8 marks]

Receiver	<code>receive_causally(M)</code>	Causally receive from the group.
Sender	<code>send_causally(M)</code>	Causally send to the group.

As needed, employ the following additional utility functions:

<code>getVec(M)</code>	Retrieves the version vector from a message.
<code>setVec(M, V)</code>	Set vector V on message M .
<code>testVec(LV, RV)</code>	Returns whether vector RV only differs from LV in that it has exactly one entry one greater than the corresponding entry in LV .
<code>updateVec(V)</code>	Returns V with the local vector entry incremented.
<code>transmit_group(M)</code>	Transmits message M to the entire group.