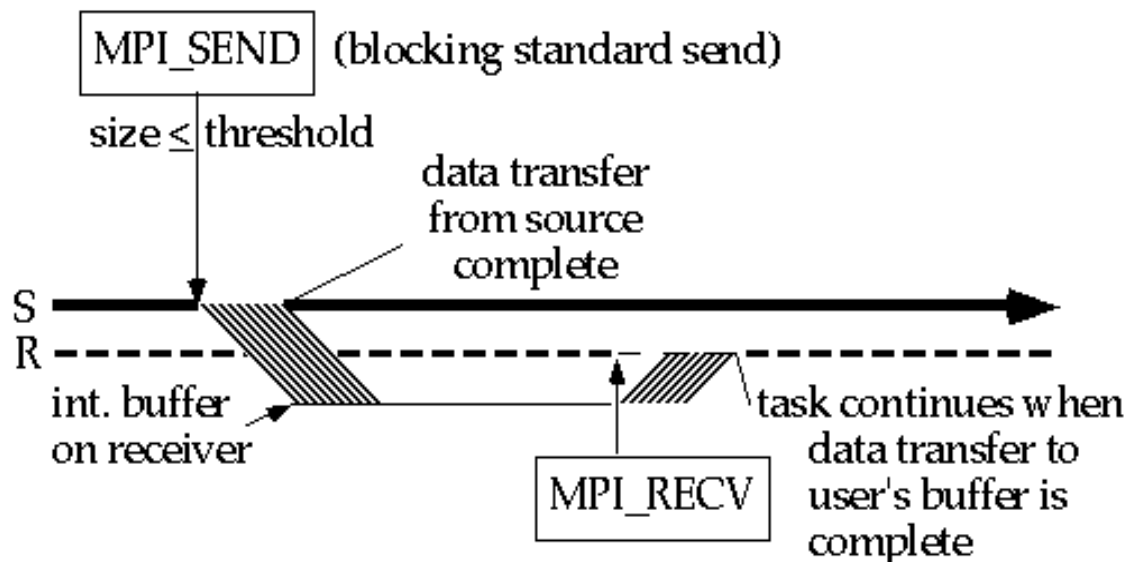


Blocking Send

message size < threshold:

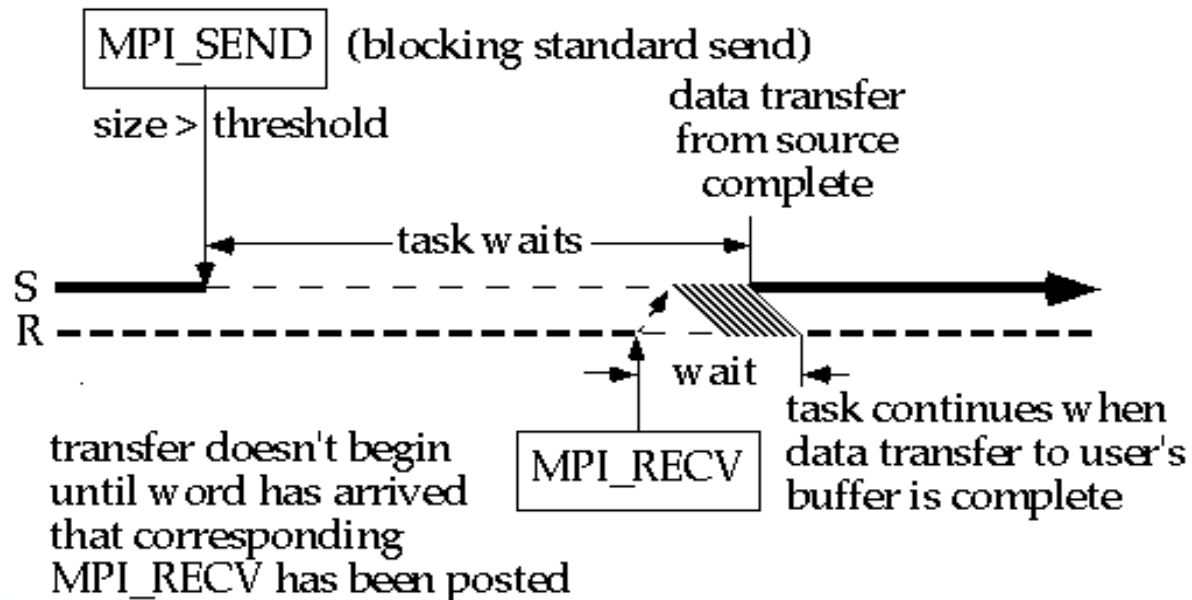
MPI_Send copies the message over the network into system buffer on the receiving node. Send then returns, and the sending task can continue computation. The system buffer is attached when the program is started -- the user does not need to manage it in any way. There is one system buffer per task that will hold multiple messages. The message will be copied from the system buffer to the receiving task when the receive call is executed.



Blocking Send

message size > threshold:

The behavior is essentially the same as for synchronous mode. System overhead is incurred from copying the message data from the sender's message buffer onto the network, and from copying the message data from the network into the receiver's message buffer. The sender must wait for the receive to be executed and for the handshake to arrive before the message can be transferred. The receiver also incurs some synchronization overhead in waiting for the handshake to complete.



Buffered Send

Copies the data from the message buffer to a user-supplied buffer (user-level buffer), and then returns. The sending task can then proceed with calculations that modify the original message buffer. The data will be copied from the user-level buffer over the network once the "ready to receive" notification has arrived.

