Lost Message Detection

As an aside to RFC 533, note that if sending Hosts do uniquely
identify messages on a given link using the extra four bits and
receiving Hosts do look at these bits, a lost message detection
system such as those suggested in RFCs 512 and 516 drops right out of
using of the unique message-id. These extra four bits can be treated
as Hathaway's SCB of RFC 512 providing a 16 element sequence number
on a per connection basis. A 16 element sequence is sufficient as
the IMPs never allow more than four outstanding messages at one time
between a given pair of Hosts. As Hathaway also suggests, the 0
element in the sequence can be used to indicate to the receiving Host
that sequence numbers are not being used.

To summarize, there appear to be three modes of using the message-id
number under Host/Host protocol:

1. The sender can always set the extra four bits to 0 and only
   transmit one message over a given link at a time -- this is slow
   but it allows orderly retransmission of messages without any help
   from the receiver.

2. The receiver can give no help to the sender. In this case it
   doesn't matter whether the sender uses the extra four bits to
   uniquely identify the messages or not -- the sender has no method
   of orderly retransmission, although the sender can accurately
   identify which message was lost if the sender has uniquely
   identified the messages.

3. The sender can have multiple messages outstanding (i.e., RFNMs not
   received) on a given link and the receiver can help the sender.
   In this case, if the sender uses the extra four bits to uniquely
   identify the messages in a way which can be synchronized with the
   receiver (e.g., sequential id numbers), the receiver can reliably
   detect lost messages.

Although it probably will seem insufficient to some, if the sender
and receiver use synchronized unique message-id numbers, very
reliable retransmission schemes are readily available. For instance,
the sender can retransmit the appropriate messages in response to
incomplete transmissions and the receiver can use the unique
message-ids to sort the retransmitted messages into the proper order

with the other received messages. Alternatively, the receiver can
discard all messages received out of order and the sender can back up
and retransmit a message for which an incomplete transmission was
received and all subsequent messages.

        [This RFC was put into machine readable form for entry]
      [into the online RFC archives by Alex McKenzie with 10/99]