Interim Revision to Appendix F of BBN Report 1822

Over the past few months we have become aware that there has been
some confusion as to how to operate a Host connected to an IMP as a
Very Distant Host (or VDH). Therefore, next time BBN Report 1822
("Specifications for the Interconnection of a Host and an IMP") is
revised, we will include additional information on how the IMP side
of a VDH connection works and how the Host side may operate most
efficiently. As an interim measure, we are distributing this RFC
which takes the form of a (logical) update to Appendix F of BBN
Report 1822.

On page F-6 on Appendix F, delete the second footnote.

On page F-7, find the phrase "... and the odd/even bit is complemented."
on line 17 of the page. Delete the rest of the page and insert the
following text:

In a standard Host to IMP interface, messages are delivered in a
specific order and received in the same order. A Very Distant Host
interface operates similarly in that messages are passed, for
example, from the IMP to its RTP in order; the Host's RTP then
delivers them to its receiving process in the same order. It is
important to note, however, that between these two software
interfaces there is nothing said about ordering. In particular, if
the special interface detects an error in a packet, for example,
the receiving RTP will discard the packet. The next packet may
arrive on another logical channel before the sending RTP
retransmits the discarded and unacknowledged packet, and the
receiver should be prepared to accept this packet out of order.
The protocol described above explicitly permits such out-of-order
behavior between the RTPs, requiring only that the transmit portion
of the RTP fill its channels in sequence (one to channel zero, one
to channel one, one to channel zero, etc.), and that the receive
portion of the RTP empty its channels in sequence. In addition, to
insure correct sequencing, the first channel filled or emptied
after initialization must be channel zero. Null packets use
neither a channel nor a channel number when sent and are not
acknowledged when received.

When packets must be retransmitted until acknowledged, processing
and transmission delay may cause acknowledgement to be delayed for
more than one transmission time. Unnecessary retransmission may
interfere with new transmissions, as well as placing an added

burden on both receiver and transmitter. Therefore, we recommend a
program delay before deciding to retransmit an unacknowledged
packet. This amount of delay should be adjustable, but we
recommend a trial value of 100 msec. Additional efficiency may be
gained if the RTP can notice that the next packet has been
acknowledged while the previous one has not: in this case, it is
clear that the first packet was not correctly received and it may
be retransmitted immediately without waiting for the programmed
delay to expire. This option has not, however, been implemented in
the IMP at this time.