Speedup of Host-IMP Interface

I. Introduction

   In order to make the full performance capabilities of the subnet
   available for interprocess communication, the host's IMP interface
   and the IMP's host interface should operate at the highest speed
   obtainable.

   First, this high throughput will minimize the latency observed when
   RFNM's, control messages, and NVT (network virtual terminal)
   characters are queued behind full sized messages. A full-sized
   message currently ties up a 100 kb interface for almost 100 Msec.
   delaying short messages behind it by 100 Msec. Speeding up the host
   interface to 300 kilobaud will shrink this latency to 30 Msec.

   Secondly, this high-speed operation minimizes the time that the IMP
   buffer and the host core buffer are locked down during message
   transfer. (One being emptied, one being filled). Being able to
   dispose of buffers far faster means that many fewer of them will
   suffice to carry the communications traffic; each buffer can be
   reused far more often.

   Third, high-speed operation makes it possible to improve error
   control: currently, a destination IMP returns an RFNM after
   transmitting the first packet of a multipacket message to the
   destination host. If an error occurs during the transmission of the
   (up to seven) other packets into the destination host, the source
   host will not be informed of the error: it has already been given a
   positive message acknowledgement in the RFNM. The alternative,
   holding off the RFNM until all packets have been transmitted into the
   destination host, would add another 80 Msec. to the round trip
   message - RFNM time with the current 100 kilobaud interface. A
   higher speed interface will reduce this delayed - RFNM cost to a more
   acceptable value, making it practical to eliminate this source of
   undetected message transmission errors.

Fourth, a high speed interface will permit greater host
communications bandwidth. (Currently limited to 100 kilobaud). This
increase in bandwidth will be essential for communications between
hosts at a "network-structured" site, where different hosts on the
same IMP are specialized to perform different parts of a computation.

Clearly, any new or retrofitted host interfaces should be very high
speed, and existing host interfaces should be adjusted to operate at
their maximum speed, which is in excess of 300 kilobaud.

II. Experimental Results

In support of the above predictions, the BBN TENEX staff performed an
experiment in cooperation with the BBN IMP group to determine how
fast the System A (BBN-TENEX) and System B (BBNB) distant interfaces
would operate.

Results are as follows:

The Host-to-IMP connection is synchronized by a two-way handshake
which has an available burst bandwidth of 1 bit/(2225 nsec + 3
nsec/ft.*<cable length>ft) For our cable length, this results in a
bandwidth of 310 kilobaud.

The IMP-to-Host connection is synchronized by a four-way handshake
which has an available burst bandwidth of 1 bit/(1350 nsec + 6
nsec/ft.*<cable length>ft.) which results in a bandwidth of 290
kilobaud for our installation.

Both System A and System B are now operating at this higher interface
speed.

Since the propogation delay time through a distant host driver-
receiver pair amounts to 250 nsec, it is expected that local host
interfaces (<30ft) can be operated at speeds substantially faster
than our 300 kilobaud.

In addition to the above measurements of hardware speed, new results
were obtained in measurements of file transfer performance, i.e. the
CPU time and real time used per megabit of information transmitted
over the network.

This experiment involved the movement of one-megabit data files to
and from an FTP User process in System B communicating with the FTP
Server Process in System A. The results are summarized in the
followiing table:

Operation	Byte Size	Type	Bandwidth	User CPU seconds/ megabit
Get	8	ASCII	47Kbaud	7.9
Send	8	ASCII	50Kbaud	7.9
Get	32	LocalByte	43Kbaud	1.80
Send	32	LocalByte	38Kbaud	1.70
Get	36	Image	79Kbaud	1.85
Send	36	Image	85Kbaud	.95

The 36-bit bandwidth of around 80Kbaud is a great improvement from
the (typically 25Kbaud measured before the speedup of the interface
hardware. The CPU time use has also decreased somewhat from that
reported in RFC #557 by Barry Wessler: this demonstrates continued
improvement of system efficiency between the TENEX version 1.31 and
TENEX version 1.32.

In conclusion, the BBN-TENEX staff recommends that all host-IMP
interfaces in the network be speeded up to the fastest operation
obtainable.

        [This RFC was put into machine readable form for entry]
         [into the online RFC archives by Alan Whinery, 1/02]