

NIC NCP EXPERIMENT

For the past couple of weeks, the NIC NCP has been keeping statistics on total incoming messages, incoming host-host control opcodes, and size of outgoing messages. The results have been rather enlightening and, I think, should be carefully considered by future implementors of NCPs for servers. The statistics will be presented in a rather qualitative fashion, since they were reset each time the system came up, but they represent a total of about 100 hours of uptime, most of it during the working day.

The total numbers of incoming and outgoing messages were almost identical. There were about 5% more outgoing. There were slightly over half as many incoming control opcodes processed as incoming messages; on the assumption that no incoming control message had more than one opcode, slightly over half the incoming messages were control messages.

The Opcode statistics were somewhat variable. In all cases the ALL opcode accounted for the great majority, from a low of about 50% on weekends to a high of 98% on a busy weekday. Almost all of the remainder were NOPs. No other opcode ever accounted for more than 5%.

The output message statistics were taken as $\log_2(\text{message size})$: this included 1 word of buffer header, 1 word of IMP header, and 1 word of host header. As might be expected, 95% of all outgoing messages had 1 to 4 PDP-10 words (36-bit) of data. However, if one multiplies the count for each bucket by the average message size for the bucket, the result is that only 75% of all outgoing data was in the smallest message size: the remaining data was spread out fairly evenly between the other buckets.

I would draw the following conclusions from these statistics. First, half the messages on the network appear to be ALLs. This suggests that NCPs should give some thought to processing control messages efficiently. Second, 95% of the messages are very short. This suggests that elaborate buffering and queuing schemes are not likely to be valuable, since the hypothetical gain in efficient use of the IMP is probably swamped by the overhead within the host. Third, a sufficiently large fraction of all data is in large messages (presumably file transfers) that it is also necessary to deal with this situation efficiently, e.g. a NCP which always sent 1-character messages would not be satisfactory.

The ARPANET has been in vigorous operation for a year or two, and many NCPs have been written during this time (including a rewrite of the TENEX NCP, which probably handles more traffic than all other NCPs combined); to my knowledge, no one has bothered to gather these statistics before. The total time invested in putting these measurements into the NIC system was about half an hour (10 instructions). I find it regrettable that even those of us presumably engaged in "computer science" have not found it necessary to confirm our hypotheses about network operation by experiment an to improve our theories on the basis of evidence.

[This RFC was put into machine readable form for entry]
[into the online RFC archives by Alex McKenzie with]
[support from GTE, formerly BBN Corp. 10/99]

