

Network Working Group
Request for Comment: 619
NIC 21990

W. Naylor
H. Opderbeck
UCLA-NMC
March 7, 1974

Mean Round-Trip Times in the ARPANET

In one of our current measurement projects we are interested in the average values of important network parameters. For this purpose we collect data on the network activity over seven consecutive days. This data collection is only interrupted by down-time or maintenance of either the net or our collecting facility (the "late" Sigma-7 or, in future, the 360/91 at CCN).

The insight gained from the analysis of this data has been reported in Network Measurement Group Note 18 (NIC 20793):

L. Kleinrock and W. Naylor
"On Measured Behavior of the ARPA Network"

This paper will be presented at the NCC '74 in Chicago.

In this RFC we want to report the mean round-trip times (or delays) that were observed during these week-long measurements since we think these figures are of general interest to the ARPA community. Let us first define the term "round trip time" as it is used by the statistics gathering program in the IMPs. When a message is sent from a source HOST to a destination HOST, the following events, among others, can be distinguished ($T(i)$ is the time of event i):

- T(1): The message is passed from the user program to the NCP in the source HOST
- T(2): The proper entry is made in the pending packet table (PPT) for single packet messages or the pending leader table (PLT) for multiple packet messages after the first packet is received by the source IMP
- T(3): The first packet of the message is put on the proper output queue in the source IMP (at this time the input of the second packet is initiated)
- T(4): The message is put on the HOST-output queue in the destination IMP (at this time the reassembly of the message is complete)
- T(5): The RFNM is sent from the destination IMP to the source IMP

T(6): The RFNM arrives at the source IMP

T(7): The RFNM is accepted by the source HOST

The time intervals $T(i)-T(i-1)$ are mainly due to the following delays and waiting times:

- T(2)-T(1):
- HOST processing delay
 - HOST-IMP transmission delay for the 32-bit leader
 - Waiting time for a message number to become free (only four messages can simultaneously be transmitted between any pair of source IMP - destination IMP)
 - Waiting time for a buffer to become free (there must be more than three buffers on the "free buffer list")
 - HOST-IMP transmission delay for the first packet
 - Waiting time for an entry in the PPT or PLT to become available (there are eight entries in the PPT and twelve in the PLT table)
- T(3)-T(2):
- Waiting time for a store-and-forward (S/F) buffer to become free (the maximum number of S/F-buffers is 20).
 - Waiting time for a logical ACK-channel to become free (there are 8 logical ACK-channels for each physical channel).
 - For multiple packet messages, waiting time until the ALLOCATE is received (unless an allocation from a previous multiple-packet message still exists; such an allocation is returned in the RFNM and expires after 125 msec)
- T(4)-T(3):
- Queuing delay, transmission delay, and propagation delay in all the IMPs and lines in the path from source IMP to destination IMP
 - Possibly retransmission delay due to transmission errors or lack of buffer space (for multiple packet messages the delays for the individual packets overlap)
- T(5)-T(4):
- Queuing delay in the destination IMP
 - IMP-HOST transmission delay for the first packet
 - For multiple-packet messages, waiting time for reassembly buffers to become free to piggy-back an ALLOCATE on the RFNM (if this waiting time exceeds one second then the RFNM is sent without the ALLOCATE)
- T(6)-T(5):
- Queuing delay, transmission delay, and propagation delay for the RFNM in all the IMPs and lines in the path from destination IMP to source IMP

T(7)-T(6): -Queuing delay for the RFNM in the source IMP
 -IMP-HOST transmission delay for the RFNM

IMP processing delays are not included in this table since they are usually very small. Also, some of the abovementioned waiting times reduce to zero in many cases, e.g. the waiting time for a message number to become available and the waiting time for a buffer to become free.

If the source and destination HOSTs are attached to the same IMP, this table can be simplified as follows:

T(2)-T(1): as before
T(3)-T(2): for multiple packet messages: waiting time until
 reassembly space becomes available (there are up to 66
 reassembly buffers)
T(4)-T(3): for multiple packet messages: HOST-IMP transmission delay
 for packets 2,3,...
T(5)-T(4): as before
T(6)-T(5): 0
T(7)-T(6): as before

Up to now we have neglected the possibility that a single packet message is rejected at the destination IMP because of lack of reassembly space. If this occurs, the single packet message is treated as a request for buffer space allocation and the time interval T(3)-T(2) increased by the waiting time until the corresponding "ALLOCATE" is received.

The round trip time (RTT) is now defined as the time interval T(6)-T(2). Note that the RTT for multiple packet messages does include the waiting time until the ALLOCATE is received. It does, however, not include the source HOST processing delay (i.e. delays in the NCP), the HOST-IMP transmission delay, and the waiting time until a message number becomes available. Note also, that the RFNM is sent after the first packet of a multiple packet message has been received by the destination HOST.

Let us now turn to the presentation of the average round trip times as they were measured during continuous seven-day periods in August and December '73. In August, an average number of 2935 messages/minute were entering the ARPANET. The overall mean round trip delay for all these messages was 93 milliseconds (msec). The corresponding numbers for December were 2226 messages/minute and 200 msec. An obvious question that immediately arises is: why did the average round trip delay more than double while the rate of incoming messages decreased? The answer to this question can be found in the large round trip delays for the status reports that are sent from each IMP to the NCC. Each IMP sends, on the average, 2.29 status reports per minute to the NCC. Since there

were 45 sites connected to the net in December, a total of 103.05 status reports per minute were sent to the NCC. Thus 4.63 percent of all messages that entered the net were status reports.

The average round trip delay for all these status reports in December was 1.66 sec. This number is five to ten times larger than the average round-trip delay for status reports we observed in August. It is not yet clear what change in the collection of status reports caused this increase. One reason appears to be that the number of these reports was doubled between August and December. Since the large round-trip delays of these status reports distort the overall picture somewhat, we are going to present the December data - wherever appropriate - with and without the effect of these delays. (We should point out here that the traffic/delay picture is distorted by the accumulated statistics messages which were collected to produce this data. We have, however, ignored this effect since these measurement messages represent less than 0.3% of the total traffic.) The overall mean round trip delay without the status reports in December is 132 msec. This value is still more than 35 msec larger than the corresponding value for August. However, before we shall attempt to explain this difference we will first present the measured data.

Table 1 shows the mean round trip delay as a function of the number of hops over the minimum-hop path. This minimum number of hops was calculated from the (static) topology of the net as it existed in August and December of last year. The actual number of hops over which any given message travels may, of course, be larger due to network congestion, line failures or IMP failures. In fact, for August we observed a minimum mean path length of 3.24 while the actual measured mean path length was 3.30; in December we observed 4.02 and 4.40, respectively. (See Network Measurement Group Note #18 for an explanation of the computation of actual mean path length.) As expected we observe a sharp increase of the mean round trip delay as the minimum number of hops is increased. Note, however, that the mean round trip delay is not a strictly increasing function of the minimum number of hops.

Table 2 gives the mean round trip delay for messages from a given site. The December data is presented with and without the large delays incurred by the sending of status reports to the NCC. Table 3 shows the mean round trip delay for messages to a given site. The largest round trip delays, in December, were incurred by messages sent to the NCC-TIP since these messages include all the status reports.

Table 4, finally, gives for each site the mean round trip delays to those three destination IMP/TIP's to which the most messages were sent during the seven-day measurement period in December. Let us first say few words about the traffic distribution which is dealt with in more

detail in Network Measurement Group Note #18. There are several sites which like to use their IMP as a kind of local multiplexer (UTAH, MIT, HARV, CMU, USCT, CCAT, XROX, HAWT, MIT2). For these sites the most favorite destination site is the source IMP itself. For several other sites the most favorite destination site is just one hop away (BBN, AMES, AMST, NCCT, RUTT). Nobody will be surprised that for many sites ISI (ILL, MTRT, ETAT, SDAT, ARPT, RMLT, LONT) or SRI (UCSB, RADT, NBST) is the most favorite site. There are several other sites (SDC, LL, CASE, DOCT, BELV, ABRD, FNWT, LBL, NSAT, TYMT, MOFF, WPAT) which were rather inactive in terms of generating traffic during the seven-day measurement period in December. Most of their messages were status reports sent to the NCC. (Those IMPs, for which the frequency of messages to the NCC-TIP is less than 2.2 messages per minute, were down for some time during the measurement period).

Let us now attempt to give a few explanations for the overall increase in the mean round trip delay between August and December. These explanations may also help to understand the differences in the mean round trip delays for any given source IMP-destination IMP pair as observed in Table 4.

1. Frequency of routing messages. Routing messages are the major source of queuing delay in a very lightly loaded net. In August, a routing message was sent every 640 msec. Since a routing message is 1160 bits long, 3.625 percent of the bandwidth of a 50 kbs circuit was used for the sending of routing messages. For randomly arriving packets this corresponds to a mean queuing delay of 0.42 msec per hop. Between August and December the frequency of sending routing messages was made dependent on line speed and line utilization. As a result, routing messages are now sent on a 50 kbs circuit with zero load every 128 msec. This corresponds to a line utilization of 18.125 percent and a mean queuing delay of 2.10 msec. The queuing delay due to routing messages in a very lightly loaded net in December was therefore five times as large as it was in August.
2. Traffic matrix. The overall mean round trip delay depends on the traffic matrix. If most of the messages are sent over distances of 0 or 1 hop the overall round trip delay will be small. The heavy traffic between AMES and AMST over a high-speed circuit in August contributed to the small overall mean round trip delay.
3. Network topology. The mean round trip delay depends on the number of hops between source-IMP and destination-IMP and therefore on the network topology. Disregarding line or IMP failures, the mean number of hops for a message in August and December was, respectively, 3.24 and 4.02.

4. Averaging. The network load, given in number of messages per minute, represents an average over a seven-day period. Even though this number may be small, considerable queuing delays could have been incurred during bursts of traffic.
5. Host delays. The round trip delay includes the transmission delay of the first packet from the destination-IMP to the destination-HOST; therefore, the mean round trip delay may be influenced by HOST delays that are independent of the network load.

Table 1 Mean Round Trip Delay as a
Function of the Number of Hops

HOPS	#MESSAGES/MINUTE		#SITE PAIRS		MEAN ROUND TRIP DELAY		
	AUG	DEC	AUG	DEC	AUG	DEC WITH STAT RPTS	DEC W/OUT STAT RPTS
0	646.9	378.3	39	45	27	44	41
1	487.6	288.7	86	100	25	65	50
2	191.0	143.1	118	138	70	119	80
3	380.7	226.9	148	168	95	131	112
4	218.5	274.1	176	196	102	167	119
5	276.3	185.6	204	228	109	217	134
6	183.8	136.3	210	258	175	355	167
7	333.6	212.7	218	256	178	301	240
8	156.7	161.1	160	234	222	365	241
9	59.0	160.3	102	208	270	308	218
10	0.6	29.9	40	124	331	939	410
11	1.0	18.9	20	46	344	998	699
12	—	10.2	—	20	—	992	655
13	—	0.01	—	4	—	809	809

Table 2 Mean Round Trip Delays for Messages from a Given Site

SITE	#MESSAGES/MINUTE		MEAN ROUND TRIP DELAY		DECEMBER WITH STATUS REPORTS	DECEMBER WITHOUT STATUS REPORTS
	AUGUST	DECEMBER	AUGUST	DECEMBER		
1 UCLA	50.7	40.3	130	282	165	
2 SRI	377.3	147.9	45	189	174	
3 UCSB	80.2	70.3	120	221	161	
4 UTAH	27.0	46.2	136	247	169	
5 BBN	120.4	128.3	110	133	133	
6 MIT	120.6	96.9	126	160	150	
7 RAND	29.3	34.2	127	323	208	
8 SDC	1.7	2.4	521	2068	131	
9 HARV	50.3	96.0	105	88	72	
10 LL	4.4	6.7	201	602	187	
11 STAN	49.7	39.7	173	300	191	
12 ILL	26.8	53.4	158	216	165	
13 CASE	57.6	2.5	138	1592	335	
14 CMU	61.1	59.5	153	220	170	
15 AMES	242.4	114.1	43	120	81	
16 AMST	304.0	163.0	39	94	67	
17 MTRT	89.5	60.0	126	199	142	
18 RADT	27.7	29.1	145	273	160	
19 NBST	98.4	48.2	118	213	152	
20 ETAT	24.1	20.6	119	280	119	
21 LLL	-	6.8	-	721	169	
22 ISI	372.0	304.4	110	147	142	
23 USCT	298.1	210.3	60	92	70	
24 GWCT	10.5	14.1	144	381	102	
25 DOCT	5.5	7.0	236	791	171	
26 SDAT	14.7	22.9	164	322	177	
27 BELV	1.3	2.4	243	1469	466	
28 ARPT	57.9	64.3	84	150	93	
29 ABRD	1.3	2.4	183	1402	554	
30 BBNT	40.8	10.0	75	372	124	
31 CCAT	177.7	86.7	83	147	115	
32 XROX	56.8	71.7	79	136	78	
33 FNWT	2.3	3.5	347	1466	174	
34 LBL	1.2	2.7	384	1653	621	
35 UCSD	11.9	19.3	237	413	205	
36 HAWT	27.5	5.2	654	569	476	
37 RMLT	10.4	13.0	122	387	97	
40 NCCT	-	59.3	-	110	97	
41 NSAT	0.6	3.4	1022	1870	1056	
42 LONT	-	20.8	-	998	848	
43 TYMT	-	3.7	-	1352	157	

44	MIT2	-	5.6	-	720	100
45	MOFF	-	2.4	-	1982	447
46	RUTT	-	22.4	-	271	153
47	WPAT	-	2.7	-	1399	380

Table 3 Mean Round Trip Delay for Messages to a Given Site

		#MESSAGES/MINUTE		MEAN ROUND TRIP DELAY	
SITE		AUGUST	DECEMBER	AUGUST	DECEMBER
1	UCLA	57.1	43.5	134	209
2	SRI	382.3	149.4	45	158
3	UCSB	61.1	59.1	117	138
4	UTAH	28.1	50.4	128	159
5	BBN	160.8	149.2	185	110
6	MIT	150.4	107.1	116	130
7	RAND	22.6	25.0	95	161
8	SDC	1.7	0.8	149	174
9	HARV	59.3	98.3	101	70
10	LL	4.6	5.2	195	202
11	STAN	65.3	40.6	135	162
12	ILL	29.1	69.8	156	149
13	CASE	52.6	4.0	127	262
14	CMU	74.8	68.9	135	165
15	AMES	210.3	117.2	40	75
16	AMST	316.7	135.0	38	86
17	MTRT	77.7	51.7	130	151
18	RADT	23.4	23.9	142	202
19	NBST	92.2	39.5	125	169
20	ETAT	25.4	22.8	110	111
21	LLL	-	3.7	-	185
22	ISI	361.9	299.2	107	130
23	USCT	298.1	190.6	60	68
24	GWCT	10.5	7.3	144	122
25	DOCT	5.5	4.2	236	187
26	SDAT	13.3	19.7	149	177
27	BELV	0.9	0.9	196	285
28	ARPT	55.4	58.3	78	95
29	ABRD	1.3	0.7	183	271
30	BBNT	40.8	6.4	75	159
31	CCAT	177.7	76.3	83	119
32	XROX	56.8	75.3	79	69
33	FNWT	2.3	1.4	347	165
34	LBL	1.2	0.9	384	305
35	UCSD	11.9	24.0	237	157
36	HAWT	27.5	5.0	654	458
37	RMLT	10.4	11.0	122	97
40	NCCT	-	140.1	-	1263
41	NSAT	0.6	1.6	1022	918
42	LONT	-	17.3	-	855
43	TYMT	-	1.6	-	160
44	MIT2	-	3.9	-	83
45	MOFF	-	0.2	-	219
46	RUTT	-	14.7	-	153
47	WPAT	-	0.5	-	282

Table 4 Mean Round Trip Delay to the Three Most Favorite Sites

FROM SITE	TO SITE	#MESSAGES/MINUTE		MEAN ROUND TRIP DELAY	
		AUGUST	DECEMBER	AUGUST	DECEMBER
1 UCLA	1 RAND	10.8	9.4	57	92
	26 SDAT	5.6	5.9	157	191
	22 ISI	3.1	3.1	99	146
2 SRI	12 RADT	16.6	19.5	142	163
	17 MTRT	21.9	18.7	140	161
	2 SRI	266.1	17.5	14	69
3 UCSB	2 SRI	8.1	17.8	72	68
	22 ISI	18.1	17.0	75	86
	14 CMU	16.6	11.8	140	152
4 UTAH	4 UTAH	3.5	13.5	136	27
	22 ISI	3.7	4.8	131	165
	5 BBN	4.2	4.1	168	204
5 BBN	40 NCCT	—	81.4	—	105
	5 BBN	12.5	19.7	102	37
	9 HARV	0.5	9.2	22	37
6 MIT	6 MIT	40.6	24.0	81	85
	23 USCT	9.8	13.9	150	173
	9 HARV	1.7	12.0	63	88
7 RAND	1 UCLA	12.5	10.4	54	96
	16 AMST	0.8	2.6	99	190
	40 NCCT	—	2.5	—	1941
8 SDC	40 NCCT	—	2.2	—	2217
	1 UCLA	0.2	0.2	110	136
	8 SDC	0.01	0.01	93	13
9 HARV	9 HARV	7.6	50.5	49	21
	2 MIT	1.6	11.9	62	85
	5 BBN	1.6	9.5	56	37
10 LL	40 NCCT	—	2.2	—	1420
	10 LL	1.5	1.8	238	135
	24 GWCT	0.04	0.6	146	80
11 STAN	14 CMU	3.0	7.0	215	207
	4 UTAH	0.2	5.5	117	117
	6 MIT	6.5	5.0	186	225

12	ILL	22	ISI	13.3	20.3	146	142
		15	AMES	0.8	14.6	109	135
		35	UCSD	6.7	6.5	192	269
13	CASE	40	NCCT	—	2.2	—	1744
		1	UCLA	0.2	0.2	296	400
		2	SRI	7.1	0.01	163	316
14	CMU	14	CMU	13.8	23.4	129	94
		3	UCSB	13.8	9.2	153	166
		11	STAN	3.2	5.1	193	209
15	AMES	16	AMST	205.0	65.8	15	34
		12	ILL	1.2	19.6	115	120
		31	CCAT	3.2	4.6	174	230
16	AMST	15	AMES	176.8	74.3	13	28
		22	ISI	63.6	33.2	50	69
		32	XROX	13.3	17.4	41	60
17	MTRT	22	ISI	26.3	27.5	115	118
		2	SRI	23.8	20.3	137	155
		5	BBN	3.5	4.2	179	133
18	RADT	2	SRI	17.7	21.7	139	156
		1	UCLA	0.4	2.3	265	181
		40	NCCT	—	2.3	—	1618
19	NBST	2	SRI	14.1	12.1	132	163
		22	ISI	29.6	11.8	100	117
		5	BBN	21.6	9.6	71	97
20	ETAT	22	ISI	11.9	11.3	106	107
		24	GWCT	5.0	5.9	99	107
		40	NCCT	—	2.2	—	1602
21	LLL	5	BBN	—	2.9	—	183
		40	NCCT	—	2.2	—	1847
		4	UTAH	—	0.5	—	71
22	ISI	28	ARPT	26.0	38.3	106	104
		23	USCT	69.0	32.7	80	92
		16	AMST	62.0	28.5	53	87
23	USCT	23	USCT	160.9	119.2	19	23
		22	ISI	69.2	34.1	78	91
		6	MIT	12.9	19.6	135	150

24	GWCT	20	ETAT	6.6	10.8	93	91
		40	NCCT	—	2.1	—	1978
		10	LL	0.03	0.5	359	115
25	DOCT	40	NCCT	—	2.3	—	2091
		22	ISI	1.0	1.6	220	118
		15	AMES	1.9	1.2	167	198
26	SDAT	22	ISI	2.9	8.7	154	138
		1	UCLA	5.9	6.0	169	209
		2	SRI	1.0	4.4	182	184
27	BELV	40	NCCT	—	2.2	—	1553
		1	UCLA	0.1	0.2	405	517
		22	ISI	—	0.01	—	325
28	ARPT	22	ISI	27.4	41.6	106	101
		28	ARPT	19.2	13.7	20	35
		2	SRI	3.3	3.3	139	157
29	ABRD	40	NCCT	—	2.2	—	1461
		1	UCLA	0.2	0.2	439	562
		9	HARV	—	0.01	—	112
30	BBNT	5	BBN	24.2	5.1	36	64
		40	NCCT	—	2.1	—	1327
		22	ISI	4.2	1.1	170	217
31	CCAT	31	CCAT	81.9	28.2	15	31
		22	ISI	31.3	23.3	156	171
		5	BBN	7.8	7.3	45	42
32	XROX	32	XROX	20.2	36.4	19	15
		16	AMST	10.5	13.3	69	93
		14	CMU	2.5	3.0	193	251
33	FNWT	40	NCCT	—	2.2	—	2210
		9	HARV	0.01	0.3	208	194
		7	RAND	0.3	0.3	96	171
34	LBL	40	NCCT	—	2.4	—	1814
		41	NSAT	—	0.2	—	1674
		1	UCLA	0.1	0.2	295	478
35	UCSD	12	ILL	6.0	7.5	220	260
		16	AMST	1.7	4.9	120	172
		40	NCCT	—	2.0	—	2183

36 HAWT	36 HAWT	0.04	1.6	17	26
	22 ISI	5.1	1.0	600	623
	15 AMES	2.5	0.8	551	590
37 RMLT	22 ISI	7.5	9.0	68	67
	40 NCCT	-	2.2	-	1918
	28 ARPT	-	1.0	-	63
40 NCCT	5 BBN	-	41.2	-	33
	40 NCCT	-	6.6	-	433
	22 ISI	-	3.2	-	151
41 NSAT	40 NCCT	-	2.2	-	2308
	2 SRI	0.01	0.4	1046	1002
	3 UCSB	0.01	0.2	1169	1018
42 LONT	22 ISI	-	6.1	-	837
	2 SRI	-	3.7	-	884
	4 UTAH	-	2.2	-	921
43 TYMT	40 NCCT	-	2.6	-	1859
	2 SRI	-	0.5	-	79
	3 UCSB	-	0.2	-	74
44 MIT2	44 MIT2	-	2.8	-	18
	40 NCCT	-	2.3	-	1664
	1 UCLA	-	0.2	-	589
46 MOFF	40 NCCT	-	2.2	-	2091
	1 UCLA	-	0.2	-	447
46 RUTT	9 HARV	-	4.3	-	38
	5 BBN	-	3.5	-	93
	22 ISI	-	2.9	-	172
47 WPAT	40 NCCT	-	2.2	-	1643
	3 UCSB	-	0.2	-	301
	1 UCLA	-	0.2	-	671

[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. 12/99]

