Bolt Beranek and Newman Inc.

Report No. 5349

Measurement and Analysis of ARPANET Traffic:
May 1982 and March 1983

Dr. Stephen N. Cohn

May 1983

Prepared for:
Defense Communications Agency

Report No. 5349

MEASUREMENT AND ANALYSIS OF ARPANET TRAFFIC: May 1982 and March 1983

Dr. Stephen N. Cohn

May 1983

Prepared for:

Defense Communications Agency

Errata

28 September 1983

p. 24 Subsequent to publishing this report, I discovered that TENEX and TOPS 20 TCP/IP implementations regularly ping (send Internet Control Message Protocol echo requests to) all gateways they know about. Interaction of this process with the ARPANET End-to-End protocol produces many connection setup and reset control messages, which explains the increase in control messages noted.

p. 27 The pinging of gateways by TENEX and TOPS 20 TCP/IP codes, mentioned above, skews these results as the default pinging interval is 37 seconds. Hence, these hosts would appear in the table as corresponding with the gateways they were pinging even if they did not send or receive data messages. As I had not identified the problem when these measurements were made, I do not know the exact magnitude of the resulting contamination, but when this behavior was identified, there were approximately 30 hosts and 23 gateways involved.

Table of Contents

FIGURES

1 INTRODUCTION

This report documents the basic properties of the traffic transported over the ARPANET during the weeks 10-14 May 1982 and 21-25 March 1983, and the changes in the nature of ARPANET traffic between these two periods. It also attempts to associate changes observed in the ARPANET traffic with their probable causes. During this interval the size and topology of the subnetwork has varied little. The number of hosts active over comparable intervals of time is similar for the two periods. However, two significant changes did occur between the two measurement periods. Conversion from Network Control Protocol (NCP) to Transmission Control Protocol (TCP) and Internet Protocol (IP) as the host end-to-end protocol, mandated by the Defense Communications Agency to occur 1 January 1983, had been largely completed by the March 1983 period. In addition, most of the IMPs on the ARPANET were converted to run the new software version 4305 instead of version 4266, which should not have directly influenced the traffic offered by hosts, but may have affected IMP performance and thereby the subnetwork-overhead component of the total traffic.

Detailed properties of the ARPANET traffic and performance are measured by installing special software packages to run in each IMP and record information about each packet that enters the IMP. Periodically each IMP sends off a message containing the summary of its measurements over the preceding interval.

Typically this interval is 7 minutes. These messages are collected at the Network Operations Center (NOC), compiled and later moved to another computer for processing to produce reports on the traffic transiting individual IMPs and lines and the aggregate traffic and performance of the network as a whole. All of the measurements presented in this report were made using the Cumulative Statistics (CUMSTATS) package which is described in [1].

In the sections that follow we present and discuss statistics on the many parameters CUMSTATS measures. These are compared for the two periods mentioned and subdivided within each period into peak, average, and low traffic levels, where possible. We attempt to explain the major components and controlling processes as well as the significant changes. This report presents only the aggregate statistics for the whole network, although data is measured on an IMP-by-IMP basis. The characteristics presented below include the following:

1) size distribution of messages and packets sent by hosts;
2) size distribution of packets sent out on the inter-node trunks (the concept of a message does not exist at the IMP-IMP level;)
3) utilization of the inter-node trunks;
4) round-trip delay for messages between the source host's IMP and the destination host's IMP;
5) level of traffic between hosts connected to different

IMPs;

6) network mean levels of traffic offered by hosts.

7) the distribution of IMPs according to the number of other IMPs to which messages are sent during the measurement period.

2 OVERHEAD TRAFFIC COMPONENTS ON THE ARPANET

The measured properties of traffic on the ARPANET, shown in the tables which follow, distinguish between the various components of traffic that traverse the subnetwork. In particular, the concept of overhead can lead to much confusion because overhead occurs at many different levels in the protocol structure. The following description aims at clarifying these distinctions. Understanding these distinctions is essential to correctly interpreting the data reported below. The following definitions are adopted, with modifications, from [1].

There are six classes of overhead in the IMP subnet:

a) Packet headers. Each packet contains 128 bits of header. The header contains information needed by the IMPs in order to get the packet from source to destination and to reassemble messages at the destination.

b) Packet framing. Each data packet and RFNM which is transmitted on a line contains 72 bits of framing.

c) Routing update packets. Packet size is variable, but currently averages 174 bits (including framing) on the ARPANET.

d) Line up/down protocol packets, or Hello's and I-Heard-You's. These are 152 bits long (including framing).

e) End-end control packets. Most of these packets are RFNMs, but there are also other sorts of end-end control packets; for example, GET-A-BLOCKs and RESETs. These are all 232 bits long (including framing), except for REQ8s which are 200 bits long.

f) NULL packets. These are link-level acknowledgements sent only when there is no outgoing packet on which to piggyback the acknowledgment. NULLs are 152 bits long (including framing).

As used below, the term "data packet" shall mean any packet which is not a routing message, a link-level control packet, or an end-end control packet. The term "data bit" shall mean any bit which is in a data packet but is not part of the packet header or packet framing. A "data word" is any word consisting solely of data bits [1].

Within this framework, the subnetwork considers host-to-host protocol overhead (transport layer in the ISO reference model), and NCP or TCP/IP headers for example, as data, although discussions of the overall efficiency of the end-to-end communications would identify these headers as overhead. In the tables showing measurements of ARPANET traffic in this report we will adhere to the concept of data stated above. Thus a TCP/IP message consisting of one or two characters would have one 16-bit word of user data, originating from a terminal perhaps, and 20 words of TCP/IP header appended by the TAC or host, yielding 21 words of data from the subnetwork point of view. Higher level protocols such as TELNET and FTP also introduce their own overhead. TELNET's initial connection setup and option negotiation, much of it transparent to the user, adds to the overhead component of a data communication. However, user data for a remote host does not incur any extra TELNET overhead except for the duplication of TELNET's escape character, where it occurs in the data stream. Thus if the example above used an established TELNET connection, the user's data would incur virtually no additional overhead from TELNET and the subnetwork would still see a message with 21 words of data.

3 ARPANET TRAFFIC DURING MAY 1982 AND MARCH 1983

This report presents statistics on the measurements of traffic on the ARPANET during the 5 weekdays 21-25 March 1983 combined with some specialized measurements the following week to clarify ambiguity introduced by limitations in the measurement techniques. Also presented are statistics from similar measurements conducted during the 5 weekdays 10-14 May 1982. The most significant change between these two periods is the adoption of TCP/IP as the principal transport protocol. Most hosts on the ARPANET converted to TCP/IP 1 January 1983 as mandated by the Defense Communications Agency. During the May 1982 measurement period it is clear that most network communications used NCP. However, during the March 1983 measurements, 53 hosts (out of 284 enabled hosts) were still permitted to use NCP, and as we shall discuss below, this represents a significant portion of the traffic on the ARPANET during this period. (Note that 39 of the 284 enabled hosts were completely inactive during the entire measurement period 21-25 March 1983.)

The tables which follow generally present various characteristics of traffic in six columns. Each of the six columns represents a distinct, although overlapping, measurement interval selected as representative of typical or extreme examples of the traffic on the ARPANET.

The first column, labeled Mar 83/peak h, contains data recorded 21 March 1983 between 13:15 and 14:15. This 1-hour period had the highest average inter-node throughput of all hours during the 120 hours recorded during the week. Inter-node throughput essentially indicates the level of user traffic that must be transmitted at least once on an inter-node trunk. It includes overhead associated directly with user messages (e.g. headers and RFNMs,) but excludes routing, line up/down, NULL, and other end-to-end control packets. We believe this 1-hour period well represents the peak traffic level currently on the ARPANET. It is the highest level of traffic ever measured.

The second column shows the corresponding figures for the hour with the highest inter-node throughput during the week of 10-14 May 1982. It occurred between 14:00 and 15:00 12 May 1982.

The third and fourth columns represent the same quantities for the March 1983 and May 1982 periods respectively, averaged over the 5 days, 16 hours per day (08:00-24:00).

Column five contains the March 1983 values averaged over 5 days, 24 hours per day.

The last column represents the 1-hour period with the lowest inter-node traffic during the March 1983 measurements. This occurred between 05:30 and 06:30 21 March 1983. This probably does not represent the lowest traffic level on the ARPANET, which is likely to happen on weekends, but nevertheless provides an indication of the range of fluctuations in traffic on the network during the week. Graphs of 15 minute averages of selected indicators of network traffic, for the March 1983 period, appear in Appendix A.

3.1 Message and Packet Size Statistics

Appreciable growth in the size of the average packet on the network between the two measurement periods indicates the principal impact of the conversion from NCP to TCP/IP for ARPANE communications. Tables 1 and 2 show this shift dramatically.

16 bit words per packet	Mar 83 peak h	May 82 peak h	Mar 83 5 day 16 h/d	May 82 5 day 16 h/d	Mar 83 5 day 24 h/d	Mar 83 slow h
0-1	2043	737	90637	71337	123536	87
2-3	1542	12267	117487	697266	128229	12
4-7	73394	536956	3183798	24727490	3543199	227
8-15	213199	413714	14031170	21179930	19371139	11908
16-31	727843	119177	38583796	5771383	45937323	12695
32-63	266804	167882	12921180	10223208	16180651	4070
total	1284825	1250733	68928068	62670614	85284077	2900

Table 1. SIZE DISTRIBUTION OF PACKETS FROM HOSTS
Each row shows the number of packets, of a given range length, submitted to IMPs by all hosts on the network, each of the measurement periods.

16 bit words per packet	Mar 83 peak h	May 82 peak h	Mar 83 5 day 16 h/d	May 82 5 day 16 h/d	Mar 83 5 day 24 h/d	Mar slo
0-1	0.16	0.06	0.13	0.11	0.14	0
2-3	0.12	0.98	0.17	1.11	0.15	0
4-7	5.71	42.93	4.62	39.46	4.15	0
8-15	16.59	33.08	20.36	33.80	22.71	41
16-31	56.65	9.53	55.98	9.21	53.86	4
32-63	20.77	13.42	18.75	16.31	18.97	1

Table 2. SIZE DISTRIBUTION OF PACKETS FROM HOSTS (PERCENT)
Each row shows the percent of packets, of a given rang length, submitted to IMPs by all hosts on the network each of the measurement periods.

Tables 1 and 2 show the size distribution of packets sent to IMPs from hosts, averaged over the whole network. In May 1982, more than 70 percent of the packets on the network were between 4 and 15 16-bit words in length. This reflects the short, 2.5-word NCP header length. Most of the messages on the ARPANET contain very little user data, usually one or several characters. Hence it is the protocol header that controls the distribution of message lengths. TCP and IP each have minimum header lengths of 10 words so that the TCP/IP message begins with a minimum of 20 words (320 bits) of header. Thus the TCP/IP header is 280 bits longer than the corresponding NCP header.

[(20 - 2.5) * 16 = 280].

This produces a substantial increase in the average size of packets because most messages on the ARPANET fit into a single packet, as will be discussed below. Tables 1 and 2 clearly show the shift in the packet size distribution toward the 16-to-31-word bin, for the March 1983 data. The 20-25 percent of packets in the 4-15 word range, for the March 1983 data, have three possible sources. First, final packets of multi-packet messages can be of any length, and some of them undoubtedly contribute to the count of packets in this size range. The second, and probably dominant source, however, is the population of NCP-enabled hosts that were still active on the network in March 1983. As of 21 March 1983 they numbered 53. The third source is messages sent with only an IP header. Internet Control Message Protocol (ICMP) echo requests and echos between hosts and internet gateways are 14 words long (10 words of IP header and 4 data words).

The NCP-enabled hosts represent a significant portion of the offered traffic. Tables 3 and 4 show the number and distribution of message sizes, for the peak hour during the March 1983 period; note that 4.17 percent of the packets are final packets of multi-packet messages. (We are considering packets not messages. A 5-packet message yields 5 packets, etc., so that this value is less than the percentage of multi-packet messages.) If we assume that the final packets of multi-packet messages are uniformly distributed between 1 and 63 words, then 19 percent of them, or 0.8 percent of all packets, should fall in the size range 4-15 words. The remaining packets in this range, approximately 21 percent of the packets on the network, may be attributed to the NCP-enabled hosts and internet gateways. Independent examination of the traffic from IMPs with NCP-enabled hosts and no gateways shows that NCP traffic accounts for at least 10.6 percent of the messages on the network. We will assume 11 percent NCP traffic in calculations below, assured only that the actual value, during this peak hour, is greater than 10.6 percent.

Tables 3 and 4 detail the distribution of message sizes by the number of packets. All messages on the ARPANET must be between 1 and 8 packets in length. Table 4 shows immediately that during all periods measured single-packet messages dominate

Packets per message	Mar 83 peak h	May 82 peak h	Mar 83 5 day 16 h/d	May 82 5 day 16 h/d	Mar 83 5 day 24 h/d	Mar 83 slow h
1	1105428	1150855	59898940	55407973	73880921	264869
2	24435	11300	1390832	707846	1638105	3617
3	7924	4177	332492	215617	379086	704
4	1629	582	80163	48877	102353	508
5	18323	2702	713175	289267	901258	2586
6	643	440	22969	21756	26244	20
7	186	299	14769	17161	15590	18
8	433	5522	140283	413449	225924	75
total	1159001	1175877	62593623	57121946	77169481	272397

Table 3. SIZE DISTRIBUTION OF MESSAGES FROM HOSTS
Each row shows the number of messages, n packets in length, submitted to IMPs by all hosts on the network, for each of the measurement periods.

Packets per message	Mar 83 peak h	May 82 peak h	Mar 83 5 day 16 h/d	May 82 5 day 16 h/d	Mar 83 5 day 24 h/d	Mar 83 slow h
1	95.38	97.87	95.69	97.00	95.74	97.24
2	2.11	0.96	2.22	1.24	2.12	1.33
3	0.68	0.36	0.53	0.38	0.49	0.26
4	0.14	0.05	0.13	0.09	0.13	0.19
5	1.58	0.23	1.14	0.51	1.17	0.95
6	0.06	0.04	0.04	0.04	0.03	0.01
7	0.02	0.03	0.02	0.03	0.02	0.01
8	0.04	0.47	0.22	0.72	0.29	0.03

Table 4. SIZE DISTRIBUTION OF MESSAGES FROM HOSTS (PERCENT)
Each row shows the percent of messages, n packets in length, submitted to IMPs by all hosts on the network, for each of the measurement periods.

multi-packet messages. However, the peak period statistics show that the percentage of single-packet messages has dropped from 97.9 percent in May 1982 to 95.4 percent in March 1983.

The conversion from NCP to TCP/IP also affects the distribution of single versus multi-packet messages. Analysis of the May 1982 peak hour column, in Table 4, reveals that only 6 percent of the packets on the network are non-final packets, hence full, 63-word packets. Table 2 shows, for the same period, 13.4 percent of the packets falling in the size range 32-63 words. Six percent are non-final packets, leaving 7.4 percent of the packets as final packets, probably distributed toward the shorter end of the range. Since a TCP/IP message is approximately 17 words longer than the same message under NCP, any message with 47 or more words in its final packet would become one packet longer. This corresponds to slightly more than half of the range, easily explaining the increase in the proportion of multi-packet messages.

The distinct peak at five packets, most noticeable in the March 1983 data, results in part from the widespread adoption of a 512 byte TCP window size. The 320 bit TCP/IP header combined with 512 * 8 = 4096 yields 4416 bits, or 4 full, 1008-bit packets and 1 short final packet. Another contribution, which may be important, is the maximum IP datagram size imposed on internet communications by some host implementations. The minimum value a host may impose as the maximum IP datagram size for incoming internet messages is 576 bytes. Adding the 20-byte IP header yields 4768 bits requiring 5 ARPANET packets.

ARPANET TACs impose a 256-byte TCP window size, limiting messages sent to TACs to 3 packets; however, due to the "Silly window syndrome" in TCP implementations, 2 packet messages are more likely. (Silly window syndrome refers to the oscillation of the available window size between near maximum and small values yielding alternately large and small messages.) Possibly this is the source of the peak observed for 2-packet messages in the March 1983 data.

The totals shown in Table 3 for the number of messages imply that the total use of the ARPANET is roughly comparable for the two periods. The May 1982 peak hour produced somewhat higher traffic, in terms of messages, than the corresponding hour in March 1983. The totals also indicate that the peak usage level is approximately four times the minimum usage during the work week.

3.2 Mean Values and Rates of Offered Traffic

Combining the data from the first 4 tables, and normalizing by the time intervals measured, yields the network-wide values shown in Table 5. The first 6 rows represent only that component of traffic offered directly by hosts to IMPs. Comparing corresponding entries for the May 1982 and March 1983 values shows that messages per second, packets per second, packets per message, and the number of active hosts are all similar, as

Network mean values	Mar 83 peak h	May 82 peak h	Mar 83 5 day 16 h/d	May 82 5 day 16 h/d	Mar 83 5 day 24 h/d	Mar 83 slow h
msgs/sec	307.03	311.50	211.68	208.88	177.08	72.16
pkts/sec	340.36	331.33	233.10	229.17	195.70	76.83
kb/sec	141.64	60.22	94.60	51.32	79.53	27.89
pkts/msg	1.11	1.06	1.10	1.10	1.11	1.06
bits/msg	461.34	193.33	446.91	245.71	449.16	386.59
bits/pkt	416.16	181.76	405.84	223.95	406.42	363.10
IMPs rep.	88	92	88	92	92	87
live hosts	210	202	245	>216	245	169
inter-node thrpt kb/s	269.37	173.36	181.93	128.76	149.26	57.71
round trips millions	1.17	1.18	63.29	57.81	78.17	.28
RT delay ms	370.73	285.35	263.01	205.20	253.18	219.15

Table 5. MEAN HOST AND INTER-NODE TRAFFIC
The first six entries all refer to messages delivered to IMPs from hosts, averaged over all of the IMPs from which measurements were collected, shown in the seventh line. The remaining quantities are defined as follows:

live hosts: the number of hosts submitting one or more messages to the network. (In the case of 5-day data for May 82, the value is probably slightly higher than the value, 216, shown.)
Inter-node thrpt kb/s: sum of all data packet and RFNM bits (including ARPANET headers and link level framing bits) between hosts homed to different IMPs, which are elements of completed messages. Each packet and RFNM is counted only once regardless of the number of hops required to deliver it to the destination IMP.
Round trips: total number, in millions, of completed messages that transited the network during the measurement period.
RT delay: timing of R-T delay begins when a packet is fully delivered to the IMP by the host, except that for multi-packet messages it begins when the allocate from the destination IMP is received. It ends when the corresponding RFNM is received and queued for output to the host. The values shown are the averages of round-trip delays for all completed messages.

discussed above. The major differences, attributable to the TCP/IP conversion, are the bits per average message and bits per average packet, which have more than doubled for the peak hour data, resulting in more than doubling the average bits per second offered by hosts to the network. To see that this increase in bits per second is largely due to the TCP/IP conversion, consider taking the May 1982 peak hour message rate, adding 280 bits (the net increase in header length, refer to section 3.1,) to 89 percent of the messages (to account for the 11 percent remaining NCP traffic), and adding this to the May 1982 bit rate. Performing this calculation yields:

$$60{,}221 + (311.5 * 0.89 * 280) = 137{,}847$$

which is 97 percent of the March 1983 value, and accounts for 95 percent of the increase. Similar calculations can be made for the 5-day data, noting, however, that Table 2 suggests that the proportion of NCP traffic is probably higher than was recorded during the peak hour.

The inter-node throughput, shown in Table 5 in kilobits-per-second units, is another measure of the traffic carried by the ARPANET. Inter-node throughput measures each packet sent by a host to another host only once, regardless of how many hops it takes to deliver the packet. It also excludes intra-node traffic, that is, packets destined for hosts connected to the same IMP as the source host. However, inter-node throughput encompasses the subnetwork overhead directly associated with the handling of each message, including 128 bits of ARPANET header and 72 bits of link-level framing per packet, and 232 bits for the RFNM associated with each message. Consequently, inter-node throughput indicates a smaller proportional increase between the two measurement periods. Plots of the inter-node throughput versus time follow in Appendix A.

Average round-trip delay is also shown in Table 5. This value represents the speed with which network communications are completed, and is included for comparison. We discuss delay, in detail, in the next section.

3.3 Round-Trip Delays

The CUMSTAT measurement package measures the round-trip time for each message submitted by hosts on an IMP. The round-trip time reflects the speed with which network communications are completed. Since the elapsed time is measured by the source IMP, the times should be accurate. The round-trip timing begins when a single-packet message, sent from the source host, is fully received by the IMP, or in the case of multi-packet messages, when the IMP receives an allocate message from the destination IMP. Timing ends when the source IMP receives the RFNM for the message from the destination IMP and has queued it to be sent to the source host.

Several factors influence the round-trip time. The primary limiting factors are propagation delay of the inter-node trunks between the source and destination, and transmission and processing time in the several IMPs along the route. As the traffic rises and the trunks along the path begin to saturate, queuing delays can add substantially to the round-trip delay observed between a pair of IMPs. Between different measurement periods the change in average network delay has several possible sources. First, if the traffic level increases then the queuing delays will be higher, on average, yielding higher round-trip delays. Table 5 indicates this clearly. A second possible cause of network delay could be changes in the sources or destinations of the traffic which, without increasing the absolute level, can increase the delay by concentrating the traffic on a smaller number of trunks, or increasing the average number of hops required to deliver each message to its destination. Between May 1982 and March 1983 changes in the traffic and topology, combined with the effects of routing on a more heavily utilized network, have increased the average number of hops per message, during the peak hour, from 5.74 to 6.05. (This is determined by dividing totals for non-control packets from Table 7 by totals in Table 1.) Third, the increase in the size of the average packet increases the transmission delay for a packet on each hop, and as shown in Table 5, the mean size of packets sent by hosts has roughly doubled. (When subnetwork headers are considered, the increase in packet length is approximately 50 percent.) A fourth possible effect is a change in the processing delay associated with the switch to the version 4305 IMP.

Table 5 shows that mean round-trip delay during the peak hour has increased from 285 milliseconds to 371 milliseconds between May 1982 and March 1983. About 34 milliseconds of the increase can be attributed directly to increasing the average message length by 280 bits and assuming 6 hops (RFNMs are still the same length, 232 bits). Most of the remaining increase in the delay probably results from increased queuing delays due to higher utilization of the inter-node trunks.

The average queuing delay per packet, which is not measured, is nevertheless strongly biased toward the queuing delays on the busiest trunks since the delays and packet flow are highest there. On the ARPANET the cross-country trunks generally carry the highest loads. Between May 1982 and March 1983 the load on these cross-country trunks has increased, as indicated by the increase in utilization of the most heavily loaded trunk, shown in Table 9. Plots of the average utilization of the principal cross-country trunks during the day, presented in Appendix A, show utilization averaged over 15-minute periods exceeding 65 percent. This may not seem high, but keep in mind that the delay in a queuing system with random arrivals approaches infinity as the utilization approaches unity. At the 60-percent level, small changes in the utilization impose large changes in the queuing delay. Thus much of the increase in average round-trip delay probably results from the increased utilization of the cross-country trunks.

no. of hops on min-hop path	29 Mar 1983 14:00-15:00 no. of msgs	% of msgs	delay (ms)	25 Mar 1983 05:30-06:00 no. of msgs	% of msgs	delay (ms)
0	135277	12.73	31.24	25629	17.55	34.07*
1	94812	8.92	91.43	11299	7.74	49.74
2	79260	7.46	170.22	5882	4.03	123.52
3	34686	3.26	239.69	5629	3.85	114.70
4	134628	12.67	254.11	10523	7.21	192.69
5	85052	8.00	299.33	12458	8.53	233.27
6	53539	5.04	323.64	9667	6.62	255.09
7	70608	6.64	448.94	14740	10.09	254.84
8	130909	12.32	428.25	14490	9.92	319.67
9	146027	13.74	526.35	13013	8.91	316.15
10	63767	6.00	491.72	18050	12.36	325.15
11	17782	1.67	632.73	2201	1.51	422.84
12	8837	0.83	1096.10	545	0.37	342.72
13	2390	0.22	888.21	318	0.22	353.89
14	3523	0.33	547.09	1579	1.08	376.87
avg RT delay			313.63			216.56
round trips	1061097			146023		
T I-N kb/s		222.54			76.71	
line ut w/ovr		0.226			0.077	
max avg ln ut		0.585			0.349	

Table 6. END-TO-END ROUND-TRIP DELAY BY MIN-HOP PATH LENGTH
Timing of R-T delay begins when a packet is fully delivered to the IMP by the host, except that for multi-packet messages it begins when the allocate from the destination IMP is received. Timing ends when the corresponding RFNM is received and queued for output to the host. The values tabulated are lower bounds.

T I-N kb/s: sum of all data packet and RFNM bits (including ARPANET headers and link level framing bits) between hosts homed to different IMPs, which are elements of completed messages. Each packet and RFNM is counted only once regardless of the number of hops required to deliver it to the destination IMP.

line util w/ ovr: sum over all inter-node trunks of the average bit rate, including all overhead (ARPANET headers, link-level framing, routing, line up/down protocol, and end-to-end control), divided by the combined capacity of all of the trunks.

max avg ut: average utilization over time period of trunk with highest utilization, including all overhead.

*: The measured value is 81.66, but this includes 4604 8-packet messages transiting one particular node, suffering an average delay of 299 ms. Eliminating this node from the average yields the tabulated value for 21,025 messages.

Limitations inherent in the present version of the CUMSTATS package restrict interpretation of delay data for all but the quietest periods during the May 1982 and 21-25 March 1983 measurements, except in terms of network average delay. However, a special series of experiments 29 March 1983 and a carefully selected 30-minute period on 25 March 1983 allow more detailed analysis. Table 6 provides a breakdown of the round-trip delays, for these two periods, in terms of the logical separation of the source and destination IMPs. The leftmost column indicates the number of hops separating the communicating IMPs along the minimum-hop path between them. The packets traversed a route at least this length, but depending on the routing, could have taken longer routes. The two periods represent busy and relatively idle states respectively, however, the peak hour, 21-25 March 1983 is substantially busier than the busy hour during this period. (Note that three IMPs from which data were collected for the March 1983 peak hour did not provide data for this measurement. During the former period they provided 16.3 kb/s of inter-node traffic. Hence, to compare the traffic levels between the 29 March period and March 1983 peak hour, 16.3 kb/s should be subtracted from the 269.4 kb/s yielding 253 kb/s for the March 1983 peak hour, versus 222 kb/s for the measurement 29 March 1983.)

The values quoted in the columns marked delay are actually lower bounds for the observed delay. During the measurement process round-trip times were truncated modulo the measurement time unit. The time units are 25.6 ms and 6.4 ms for the 29 March and 25 March measurements, respectively. Upper bounds are larger than the lower bounds by the value of the time unit. The additional entries are provided for comparison to other measurements. Line utilization is described in the next section.

3.4 Traffic on Inter-Node Trunks

16 bit words per packet	Mar 83 peak h	May 82 peak h	Mar 83 5 day 16 h/d	May 82 5 day 16 h/d	Mar 83 5 day 24 h/d	Mar 83 slow h
0-1	241	347	10308	20222	12823	49
2-3	8686	7090	458294	328116	585861	2697
4-7	568	3067	20601	131827	23370	16
8-15	1560	2661	96716	125312	135261	864
16-31	4011	579	198383	28089	237540	796
32-63	1633	872	70455	48345	87890	329
total	16701	14619	854759	681913	1082748	4754

Table 7. SIZE OF PACKETS SENT ON INTER-NODE TRUNKS (000's)
Each row shows the number of packets (in thousands), of a given range in length, sent on inter-node trunk lines, for each of the measurement periods. Packets requiring several hops to reach their destination are counted once on each hop. Packets falling in the 0-1 and 2-3 word bins are primarily control packets, including RFNMs, and are expected to account for about half of the packets on inter-node trunks when such a high proportion of messages are a single packet in length.

The traffic measured on the inter-node trunks includes all of the inter-node traffic offered by hosts, which has already been described. It also includes the RFNMs and other end-to-end control packets. It excludes line up/down protocol packets and routing update packets, as these are measured separately. Furthermore, in the network-wide statistics, each packet is counted once each time it is transmitted. Hence, ignoring possible retransmissions due to packets lost by errors of various kinds, a packet which is routed on a 4-hop path to its destination is counted 4 times. Tables 7 and 8 show the size distribution of packets sent on trunks. Notice that approximately half of the packets are 0-3 words in length. These are control packets, primarily RFNMs.

16 bit words per packet	Mar 83 peak h	May 82 peak h	Mar 83 5 day 16 h/d	May 82 5 day 16 h/d	Mar 83 5 day 24 h/d	Mar 83 slow h
0-1	1.45	2.38	1.21	2.97	1.18	1.04
2-3	52.01	48.50	53.62	48.12	54.11	56.74
4-7	3.41	20.98	2.41	19.33	2.16	0.34
8-15	9.34	18.21	11.32	18.38	12.49	18.18
16-31	24.02	3.97	23.21	4.12	21.94	16.76
32-63	9.78	5.97	8.24	7.09	8.12	6.93

Table 8. SIZE OF PACKETS SENT ON INTER-NODE TRUNKS (PERCENT)
Each row shows the percent of packets, of a given range in length, sent on inter-node trunk lines, for each of the measurement periods. Packets requiring several hops to reach their destination are counted once on each hop. Packets falling in the 0-1 and 2-3 word bins are primarily control packets, including RFNMs, and are expected to account for about half of the packets on inter-node trunks when such a high proportion of messages are a single packet in length.

Table 8 shows that the proportion of 0-word control packets declined between May 1982 and March 1983. The single control message with 0 data words is a REQ8, which is a request by the source host's IMP for buffer allocation from the destination host's IMP, so that a multi-packet message (of any size: 2 to 8 packets) can be sent. At first it would seem that an increase in multi-packet messages should lead to an increase in REQ8s. However, a destination IMP can also allocate buffers for the next multi-packet message, after receiving a multi-packet message, by sending an allocation piggybacked on a RFNM to the source IMP. This relieves the source IMP from having to request allocation explicitly for each in a series of multi-packet messages. However, after a set timeout period, if the source IMP does not need the allocation for a multi-packet message, perhaps because it has only single-packet messages to send, it must give the allocation back to the destination IMP by sending the corresponding control message. Thus we might expect that as the proportion of multi-packet messages increases, fewer REQ8s will be sent as it is more likely that the source IMP will already have an outstanding allocation for a multi-packet message.

Table 8 also shows an increase in the proportion of control messages in the 2-3 word category. Calculations based on the peak-hour data in Table 4 reveal that RFNMs should account for 48.4 percent of the packets in May 1982 and 47.4 percent in March 1983. (The proportion has declined because of the increase in multi-packet messages, since only one RFNM is sent per message.) Hence, the packets in the 2-3 word category are essentially all RFNMs for the May 1982 data, however, the March 1983 period has significantly more of these packets than can be explained by RFNMs. At present, the identity and origin of the excess control packets observed during the March 1983 period is unknown. *

Table 9 summarizes the traffic flows on the inter-node trunks. Recall that data bits are defined to include the NCP or TCP/IP headers as well as the user data. Thus, total data bits and data kilobits per second differ by a factor of two between May 1982 and March 1983. During both measurement periods the ARPANET supported approximately 112 physical trunks. For this analysis each trunk is considered as 2 simplex links yielding 224 logical lines. Line utilization considers the ratio of the bits sent on the line to the line's capacity. For the network as a whole, this becomes the total bits sent divided by the total capacity. Hence, line utilization without overhead is the utilization considering only data bits. Line utilization with overhead includes all overhead: ARPANET headers, link-level framing bits, routing packets, line up/down protocol, NULLs (link level acknowledgments), and end-to-end control packets. Also shown is the average utilization of the single line with the highest average utilization, including overhead, and the average utilization of the 10 most heavily loaded lines. Comparison of the utilization of the most heavily utilized lines, or average utilization with overhead, to the round-trip delays from Table 5, reinforces the conclusion that queuing produces the high round-trip delays. The line utilization with overhead clearly indicates that the average utilization of the trunks has risen

* See Errata

Network aggregate values	Mar 83 peak h	May 82 peak h	Mar 83 5 day 16 h/d	May 82 5 day 16 h/d	Mar 83 5 day 24 h/d	Mar 83 slow h
T pkts	16.7E+6	14.6E+6	854.8E+6	681.9E+6	1.1E+9	4.8E+6
T dat bits	3.2E+9	1.6E+9	150.0E+9	82.0E+9	190.0E+9	0.8E+9
pkts/sec	4424.5	3872.8	2890.6	2493.6	2484.6	1259.5
data kb/s	839.9	425.0	515.7	299.8	436.4	199.6
data b/pkt	189.85	109.74	178.43	120.27	175.68	158.5
T I-N kb/s	269.37	173.36	181.93	128.76	149.26	57.71
line util:						
w/o ovr	0.083	0.039	0.051	0.027	0.042	0.019
w/ ovr	0.243	0.178	0.163	0.120	0.138	0.073
max avg ut	0.607	0.534	0.421	0.330	0.355	0.229
avg ut 10	0.516	0.479	0.369	0.282	0.313	0.179

Table 9. INTER-NODE TRUNK UTILIZATION

During both measurement periods the ARPANET had approximately 112 inter-node trunks. For the purpose of measurements these are considered as 224 simplex links.

T pkts: sum over all inter-node trunks of all packets sent, except routing and line up/down packets. (Packets are counted once for each inter-node hop from source to destination.)

T dat bits: sum over all trunks of all data bits sent on each trunk.

pkts/sec: T pkts/elapsed time.

data kb/s: T dat bits/(1000 * elapsed time).

data b/pkt: average data bits per packet. This value is low because many packets sent on trunks are control packets without any data bits.

T I-N kb/s: sum of all data packet and RFNM bits (including ARPANET headers and link level framing bits) between hosts homed to different IMPs, which are elements of completed messages. Each packet and RFNM is counted only once regardless of the number of hops required to deliver it to the destination IMP.

line util w/o ovr: sum over all inter-node trunks of the average data-bit rate on each trunk, divided by the total capacity of all of the trunks.

line util w/ ovr: same as line util w/o ovr, except it includes all overhead: RFNMs, NULLs, ARPANET headers, link-level framing, routing, line up/down protocol, and end-to-end control.

max avg ut: average utilization over time period of trunk with highest utilization, including all overhead.

avg ut 10: average of the utilizations, including overhead, of the 10 most heavily loaded simplex links.

36 percent between May 1982 and March 1983.

3.5 Number of IMPs With Which Each IMP Corresponds

No. of IMPs to which msgs. sent	Mar 83 peak h		May 82 peak h		Mar 83 5 day 16 h/d		May 82 5 day 16 h/d		Mar 83 5 day 24 h/d		Mar 83 slow h	
0-9	38	43	25	54	1	9	4	9	1	8	52	65
10-18	17	14	40	19	9	19	1	19	7	19	14	13
19-27	15	16	11	10	11	10	1	21	11	9	13	7
28-36	12	10	7	9	6	13		12	7	14	4	2
37-45	4	3	5		14	15	3	14	15	14	2	
46-54					9	7	13	4	8	7	1	
55-63	1	1			12	9	13	4	12	8		
64-72					6	4	12	6	6	7		
73-81					4	1	8	2	5	1		
82-90	1	1	3		14		34		14		1	
91-99			1		2	1	3	1	2	1		
total	88		92		88		92		88		87	

Table 10. NUMBER OF IMPS CORRESPONDING WITH N OTHER IMPS
The leftmost column shows the range in n for the number of IMPs to which other IMPs sent messages. Within each row are 6 pairs of numbers showing the number of IMPs that sent messages to n different IMPs during the time period of each measurement. The left figure in each pair is the number of IMPs that sent at least 1 message to n other IMPs. The right figure in each pair indicates the number of IMPs that sent more than 10 messages (100 messages for 5-day measurement periods) to n other IMPs.

In the course of minute-to-minute operations, IMPs are constantly tearing down existing but inactive end-to-end connections between hosts and creating new connections between other host pairs. Table 10 shows the distribution, as a function of n, of the number of IMPs whose hosts sent messages to hosts at n other IMPs. We show two values for each position in the tab The first is the number of IMPs that sent at least 1 message t n different IMPs. The second is number of IMPs that sent more than 10 messages (100 messages for 5-day measurement periods) n other IMPs. We distinguish these because some implementatio of NCP regularly polled other hosts on the network. The counti used to produce the second value effectively prevents this process from skewing the results. The agreement between the values from corresponding measurements implies that the usage patterns of the ARPANET community changed little between the two periods.*

* See Errata

4 CONCLUSIONS

1) The number of active hosts and the number of messages sent on the ARPANET are similar for comparable intervals of time selected from the two measurement periods, 10-14 May 1982 and 21-25 March 1983.

2) The average length of messages sent by hosts, during the peak hour, has increased from 193 bits in May 1982 to 461 bits in March 1983.

3) Peak hour inter-node throughput has increased from 173 to 269 kb/s from May 1982 to March 1983, resulting in a 36 percent increase in average utilization of inter-node trunks (including overhead).

4) Average round-trip delay during the peak hour has increased from 285 ms to 371 ms between May 1982 and March 1983.

5) The conversion of the host end-to-end protocol from NCP to TCP/IP produced most of each of these effects and its impact on the properties of ARPANET traffic and performance discussed here dominates any effects due to changes in usage patterns by the end users.

5 References

[1] Rosen, E. C. "The BBN Arpanet Network Measurement Center," Bolt Beranek and Newman Inc., Report No. 3799 prepared for Defense Communications Agency, October 1978.

1 APPENDIX A

The figures which follow show the topology of the ARPANET, and display graphs of significant indicators of network traffic and performance for the two measurement periods. Each plot shows the variation of the particular variable over a 24-hour period, plotted at 15-minute intervals. Refer to the text for explanations of the plotted quantities.

ARPANET GEOGRAPHIC MAP, JUNE 1982

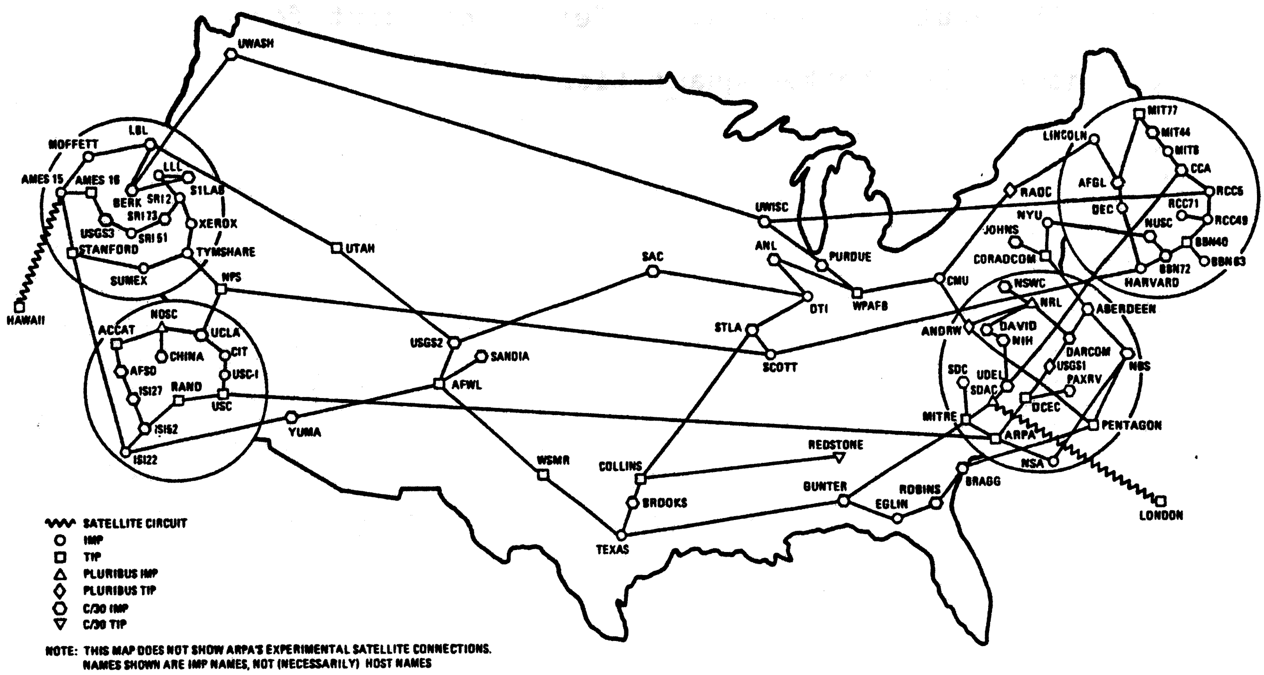

ARPANET topology during the May 1982 measurements.
Figure 1

ARPANET GEOGRAPHIC MAP, MARCH 1983

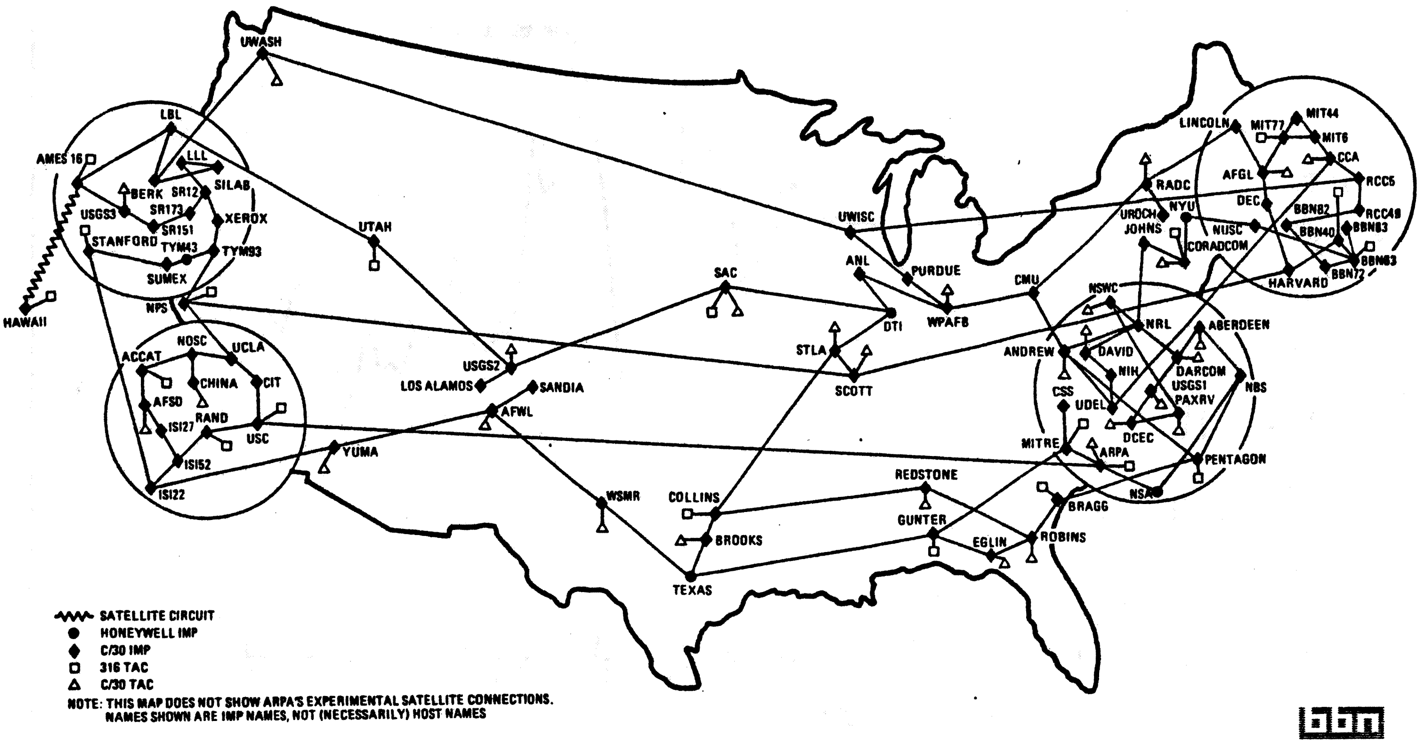

ARPANET topology during the March 1983 measurements.
Figure 2

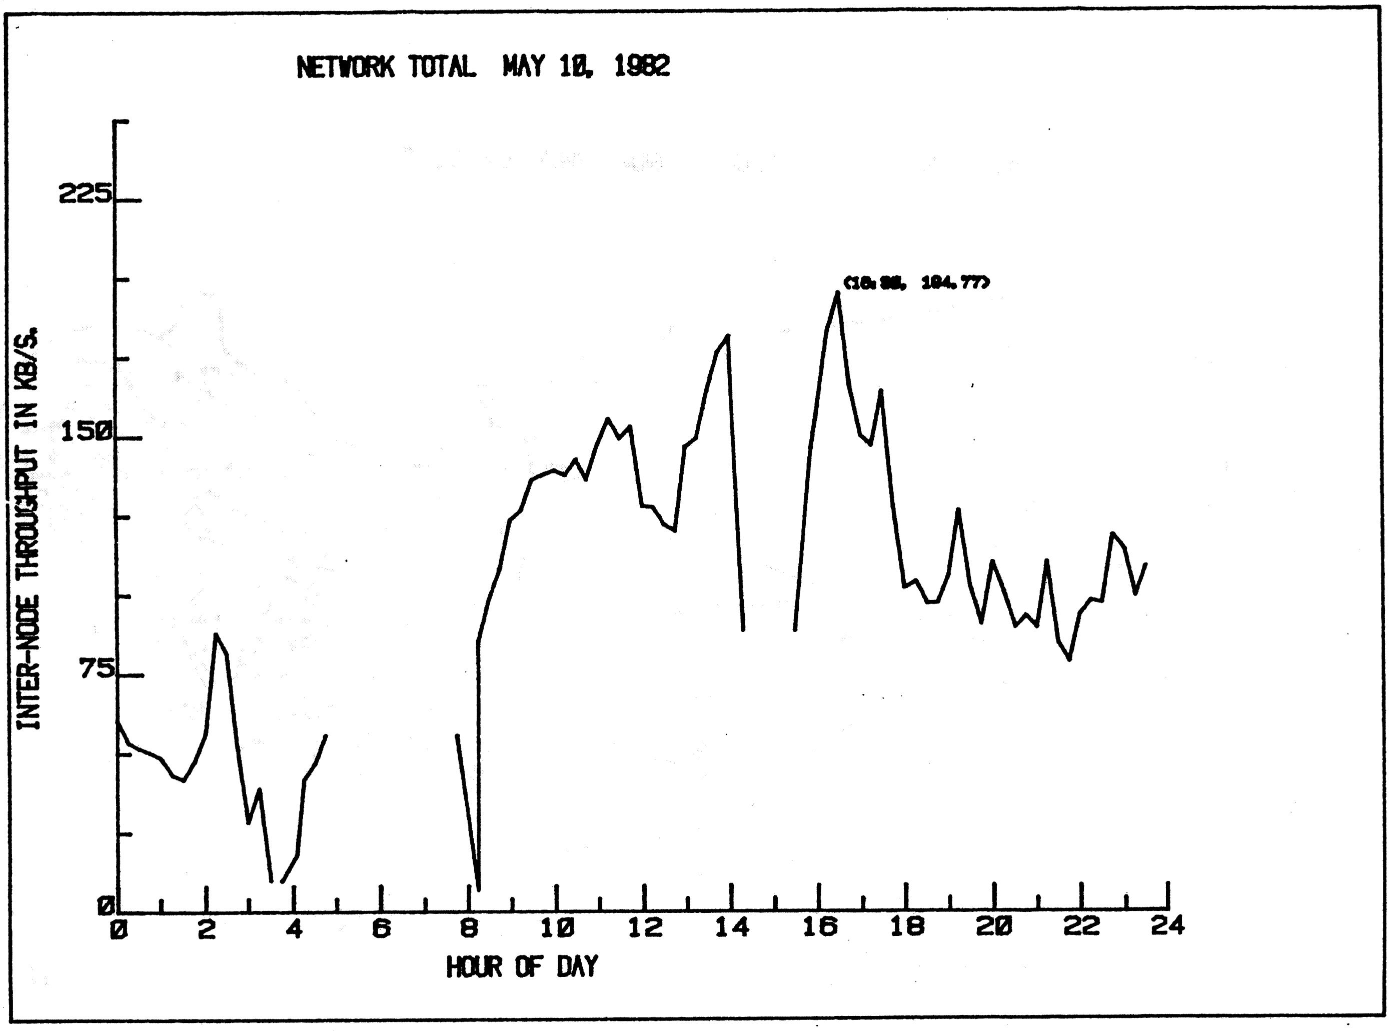

ARPANET Inter-node throughput for 10 May 1982.
Figure 3

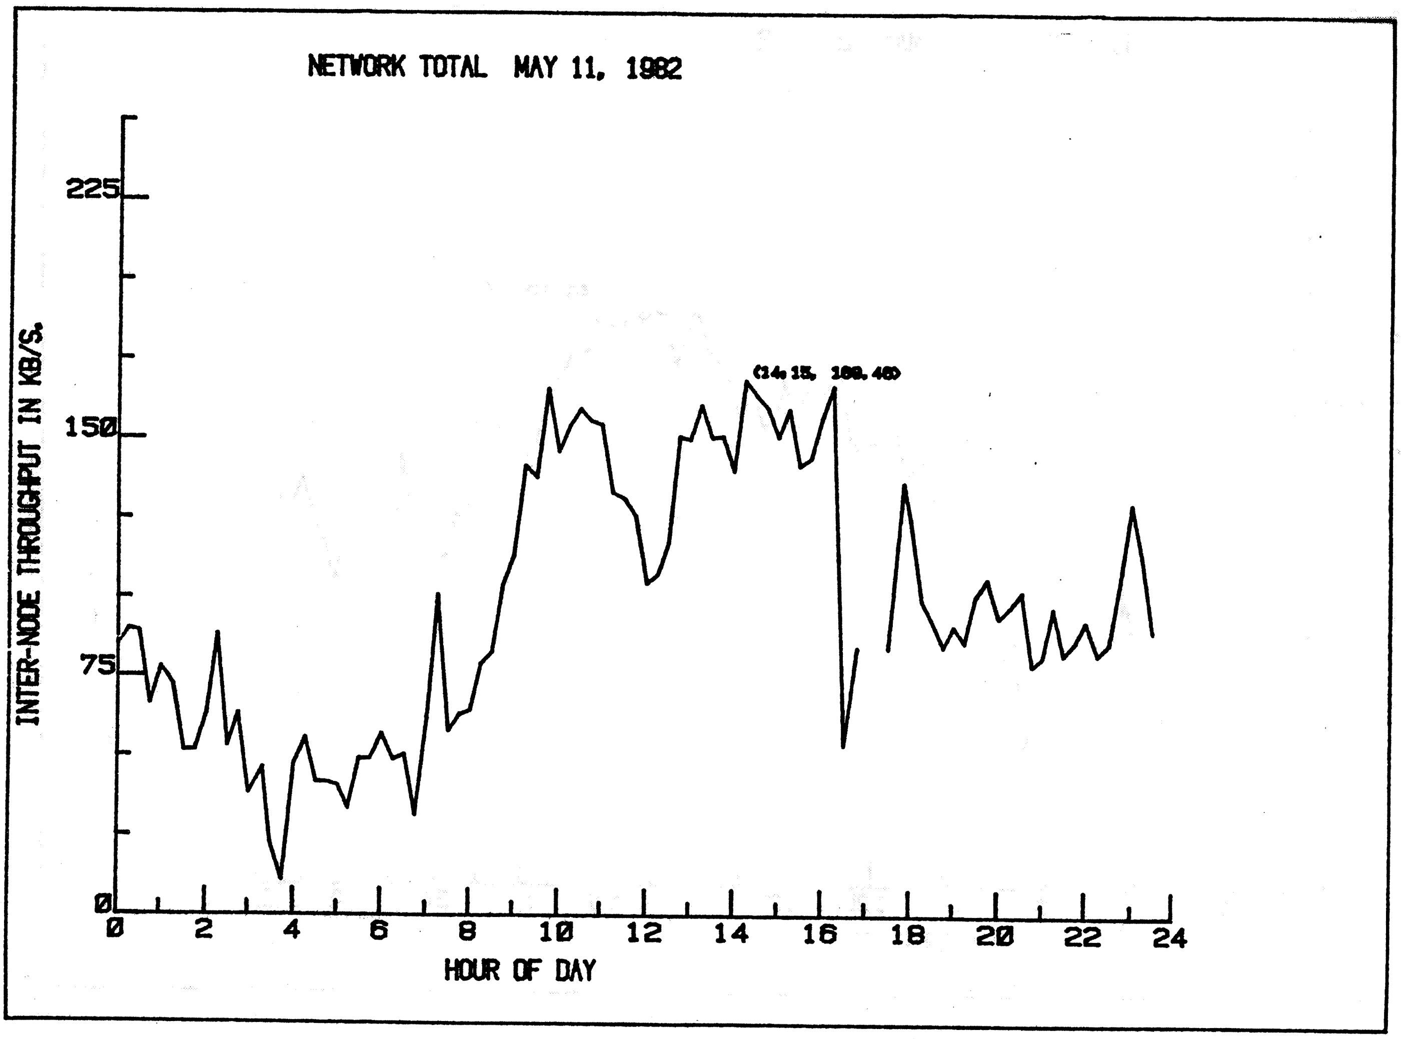

ARPANET Inter-node throughput for 11 May 1982.
Figure 4

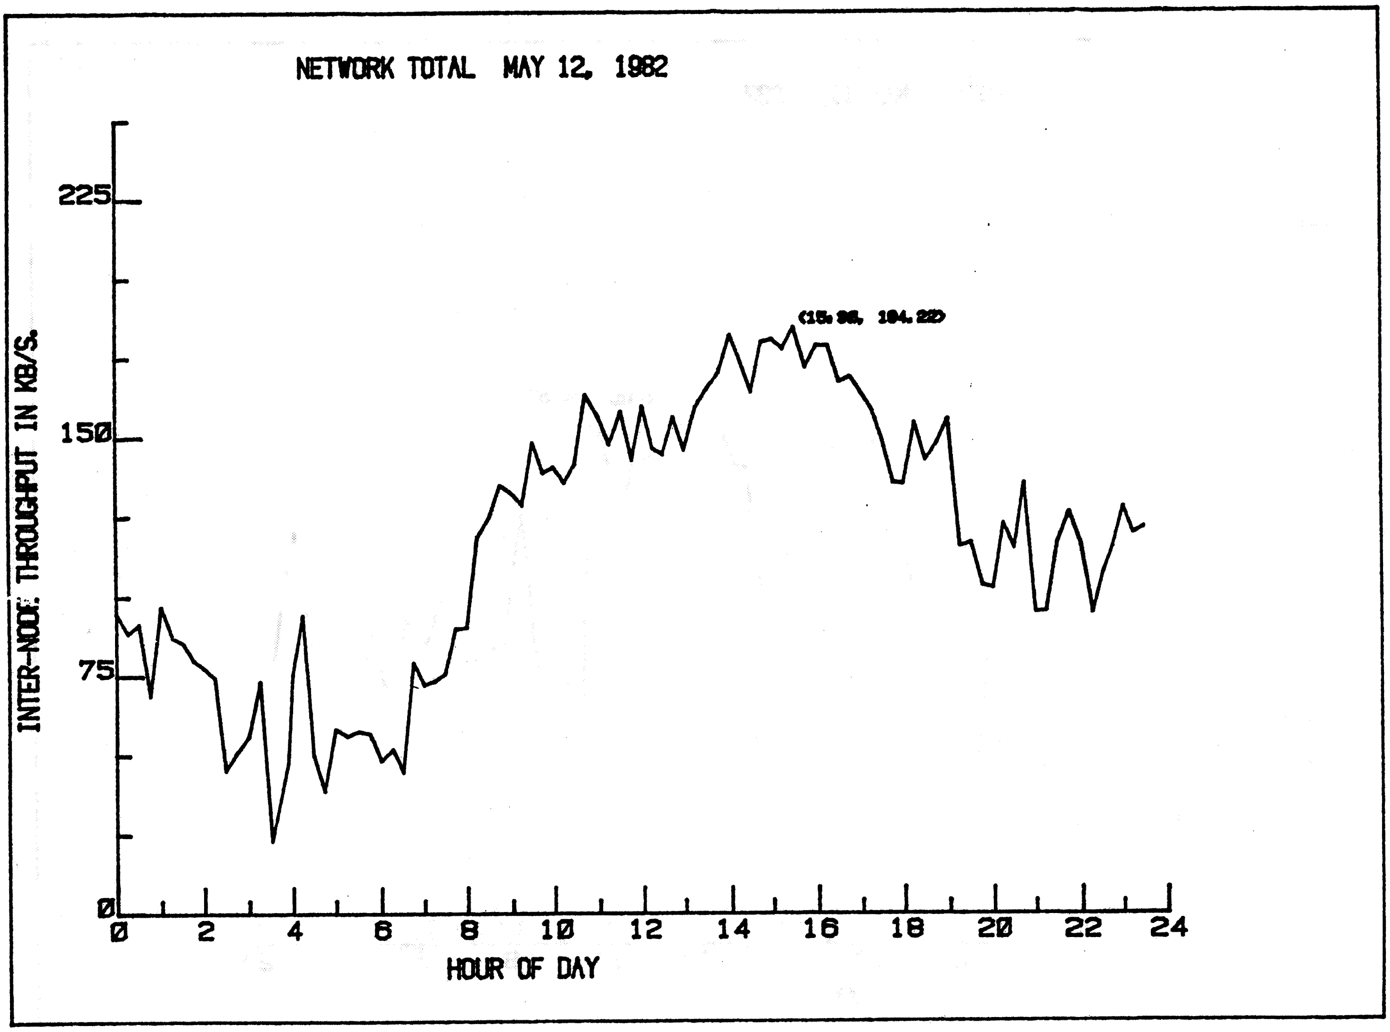

ARPANET Inter-node throughput for 12 May 1982.
Figure 5

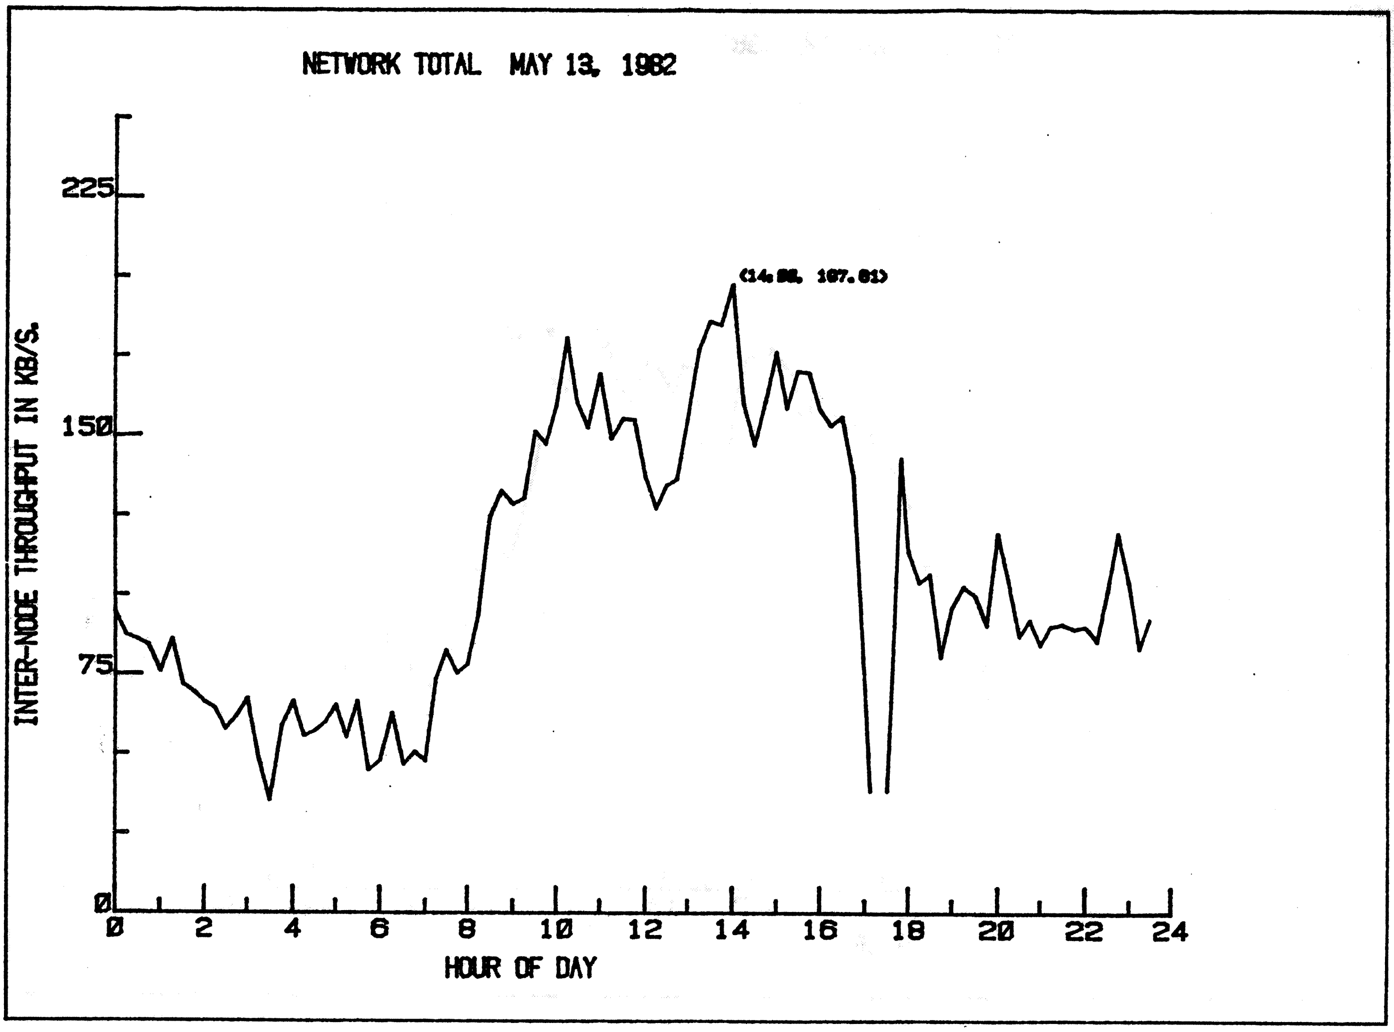

ARPANET Inter-node throughput for 13 May 1982.
Figure 6

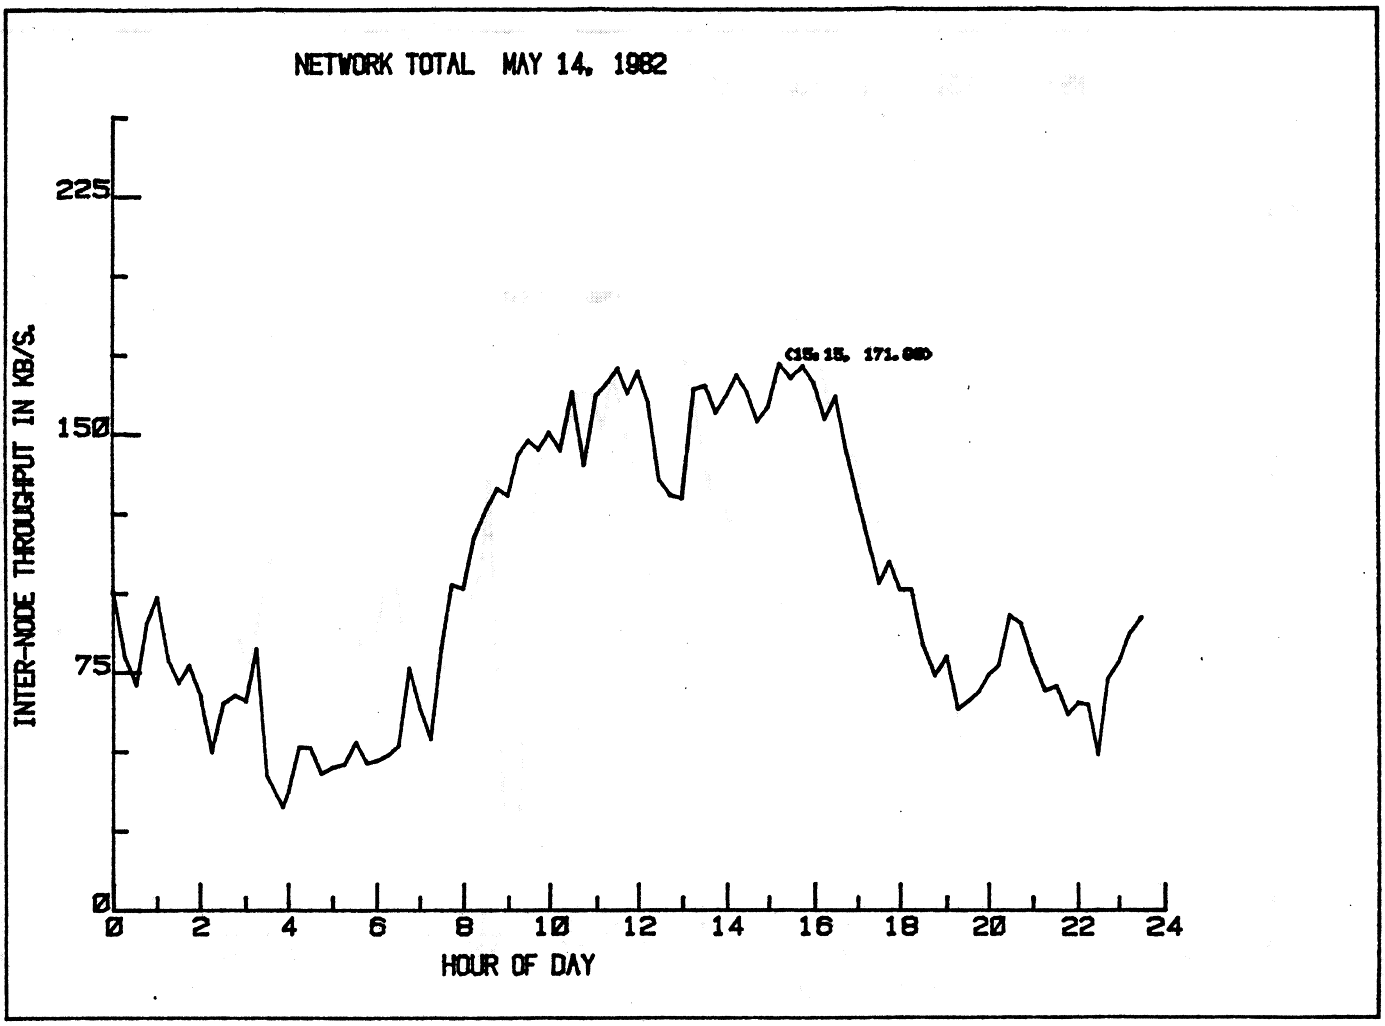

ARPANET Inter-node throughput for 14 May 1982.
Figure 7

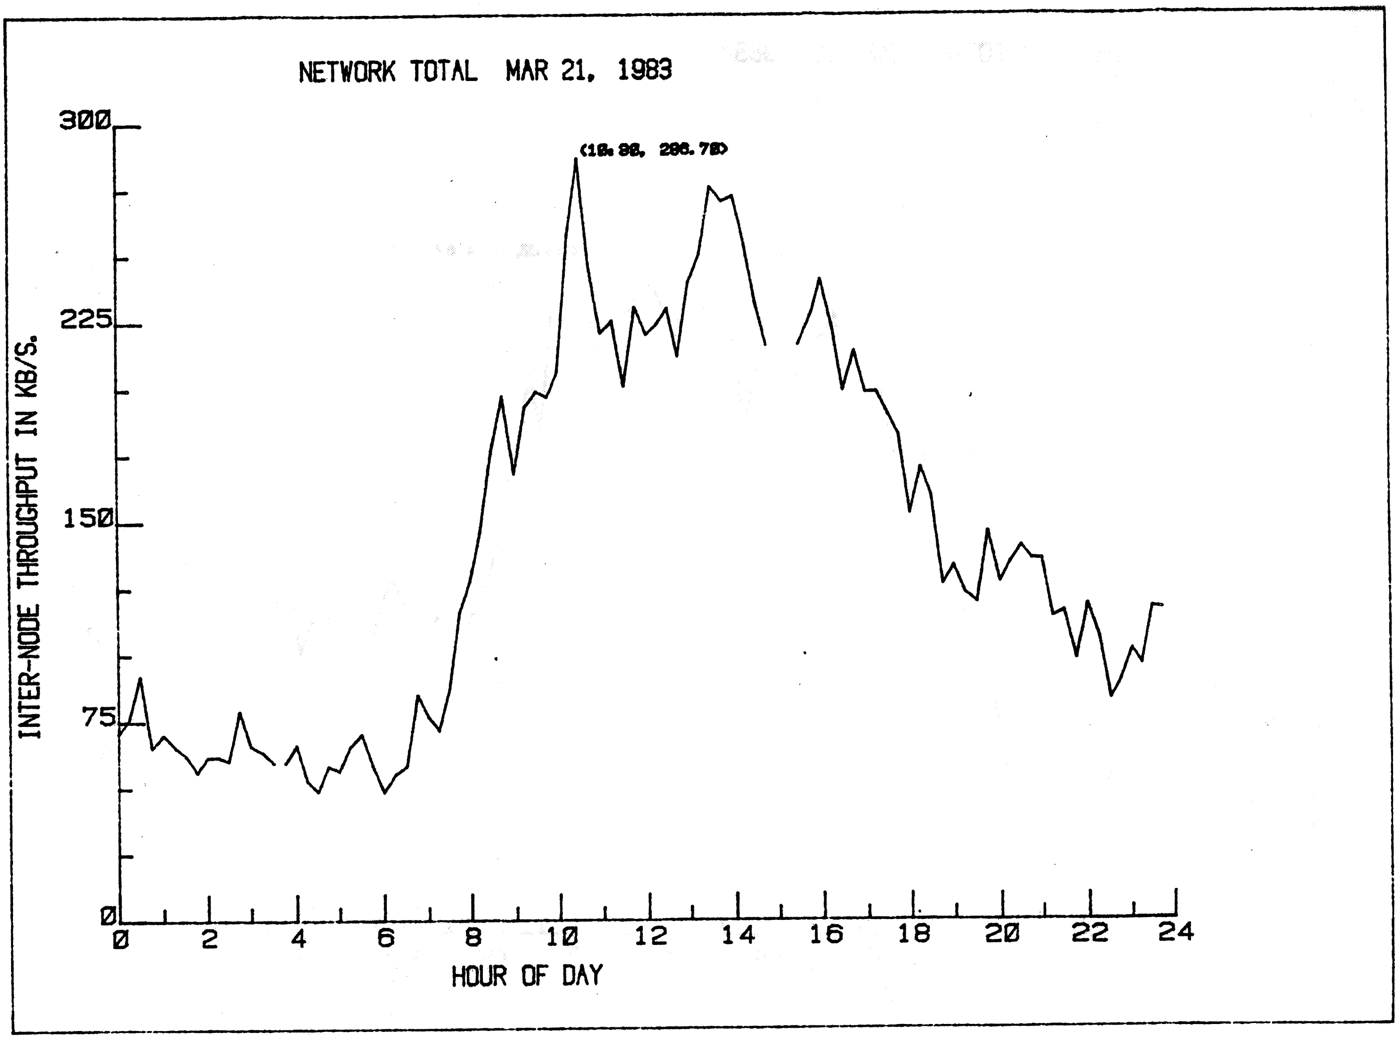

ARPANET Inter-node throughput for 21 March 1983.
Figure 8

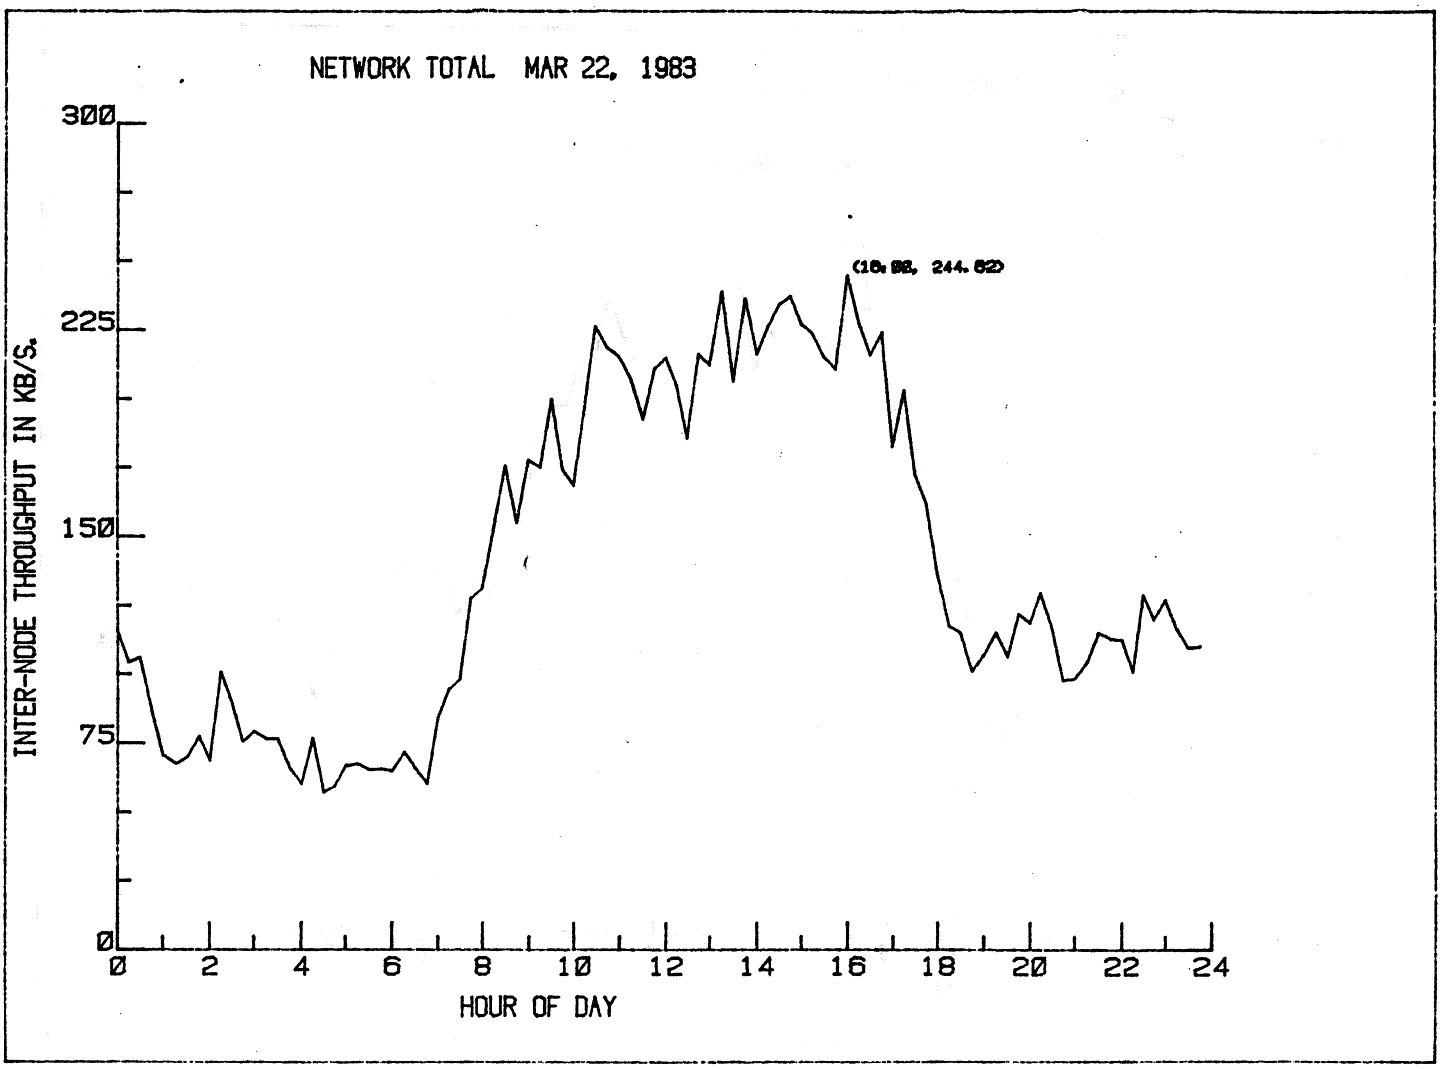

ARPANET Inter-node throughput for 22 March 1983.
Figure 9

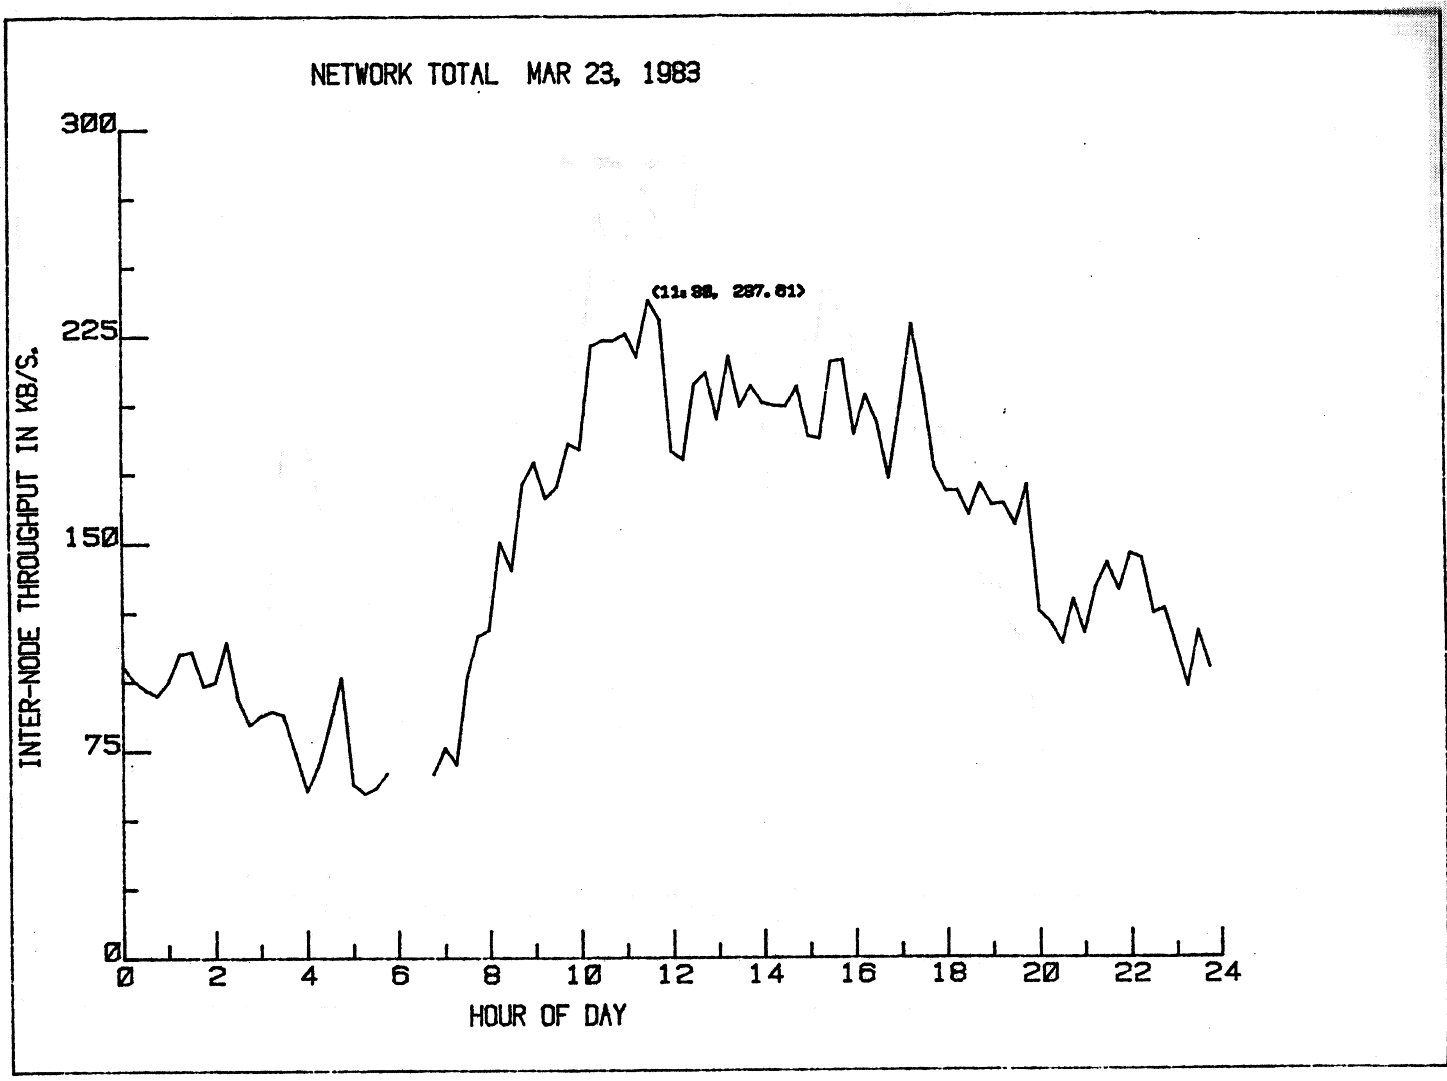

ARPANET Inter-node throughput for 23 March 1983.
Figure 10

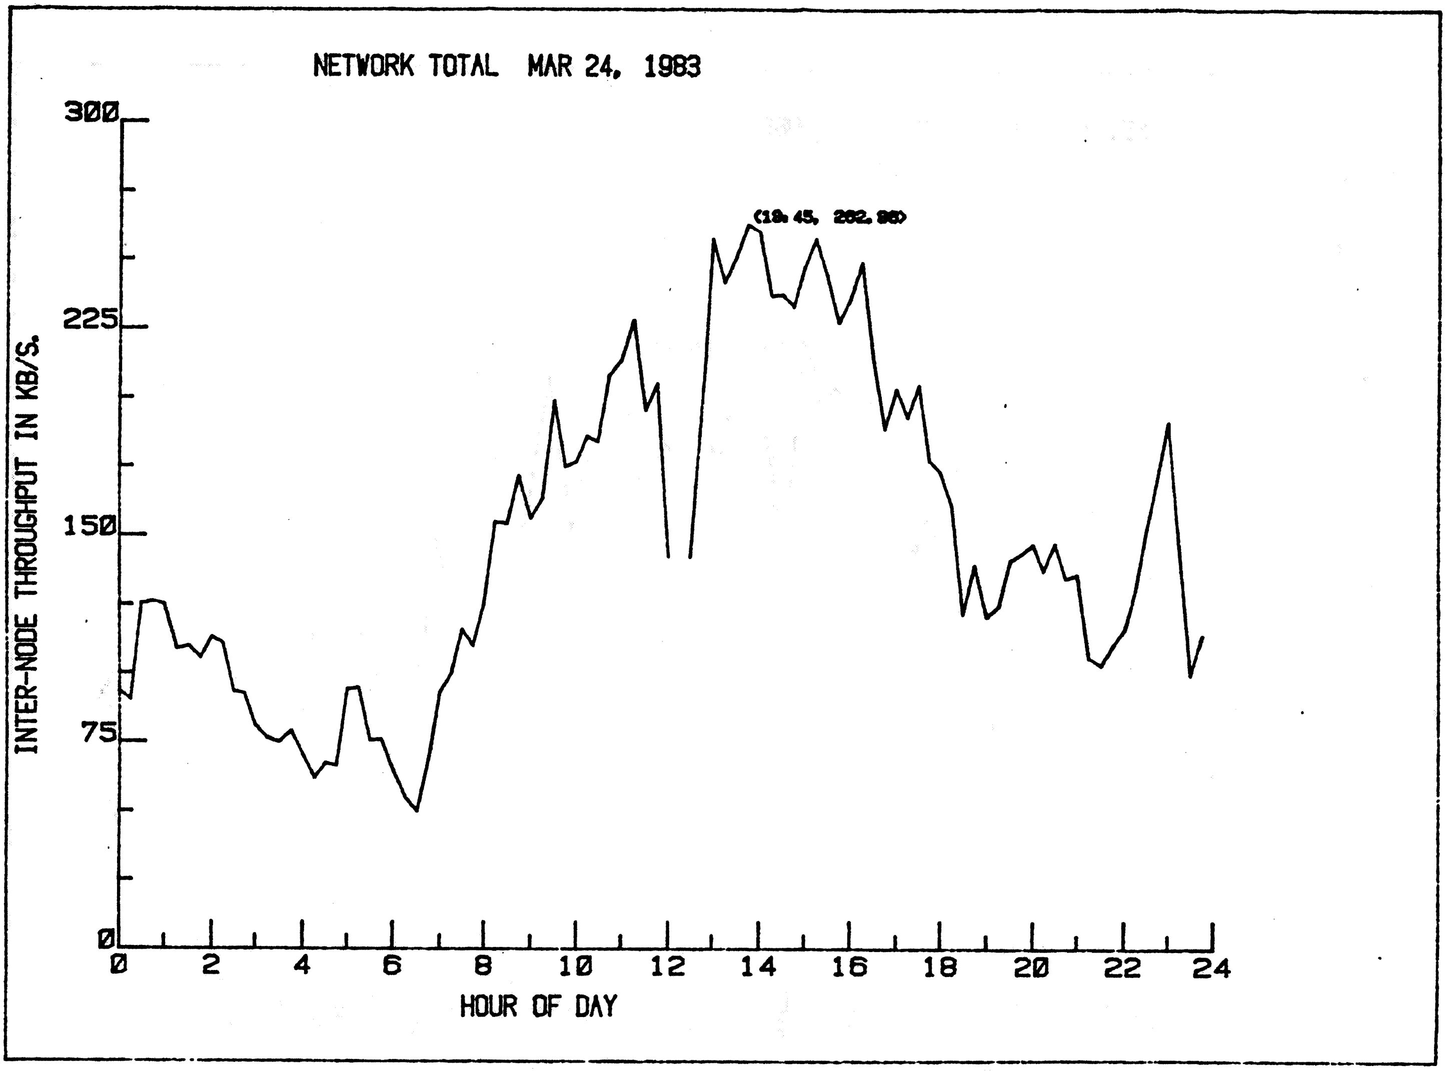

ARPANET Inter-node throughput for 24 March 1983.
Figure 11

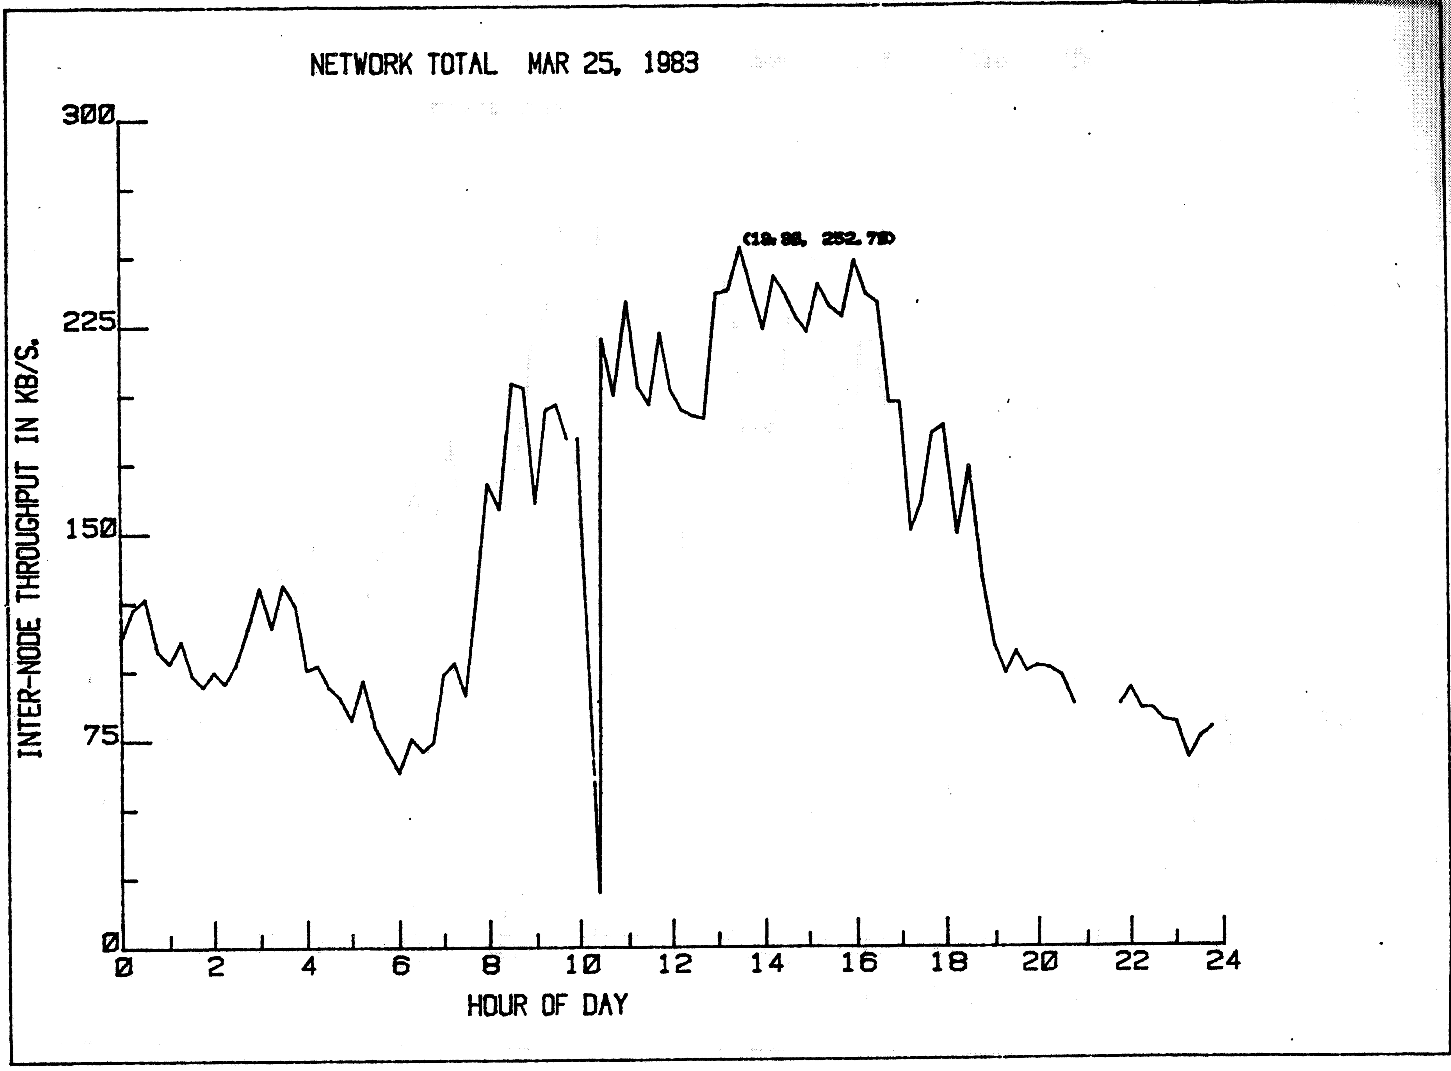

ARPANET Inter-node throughput for 25 March 1983.
Figure 12

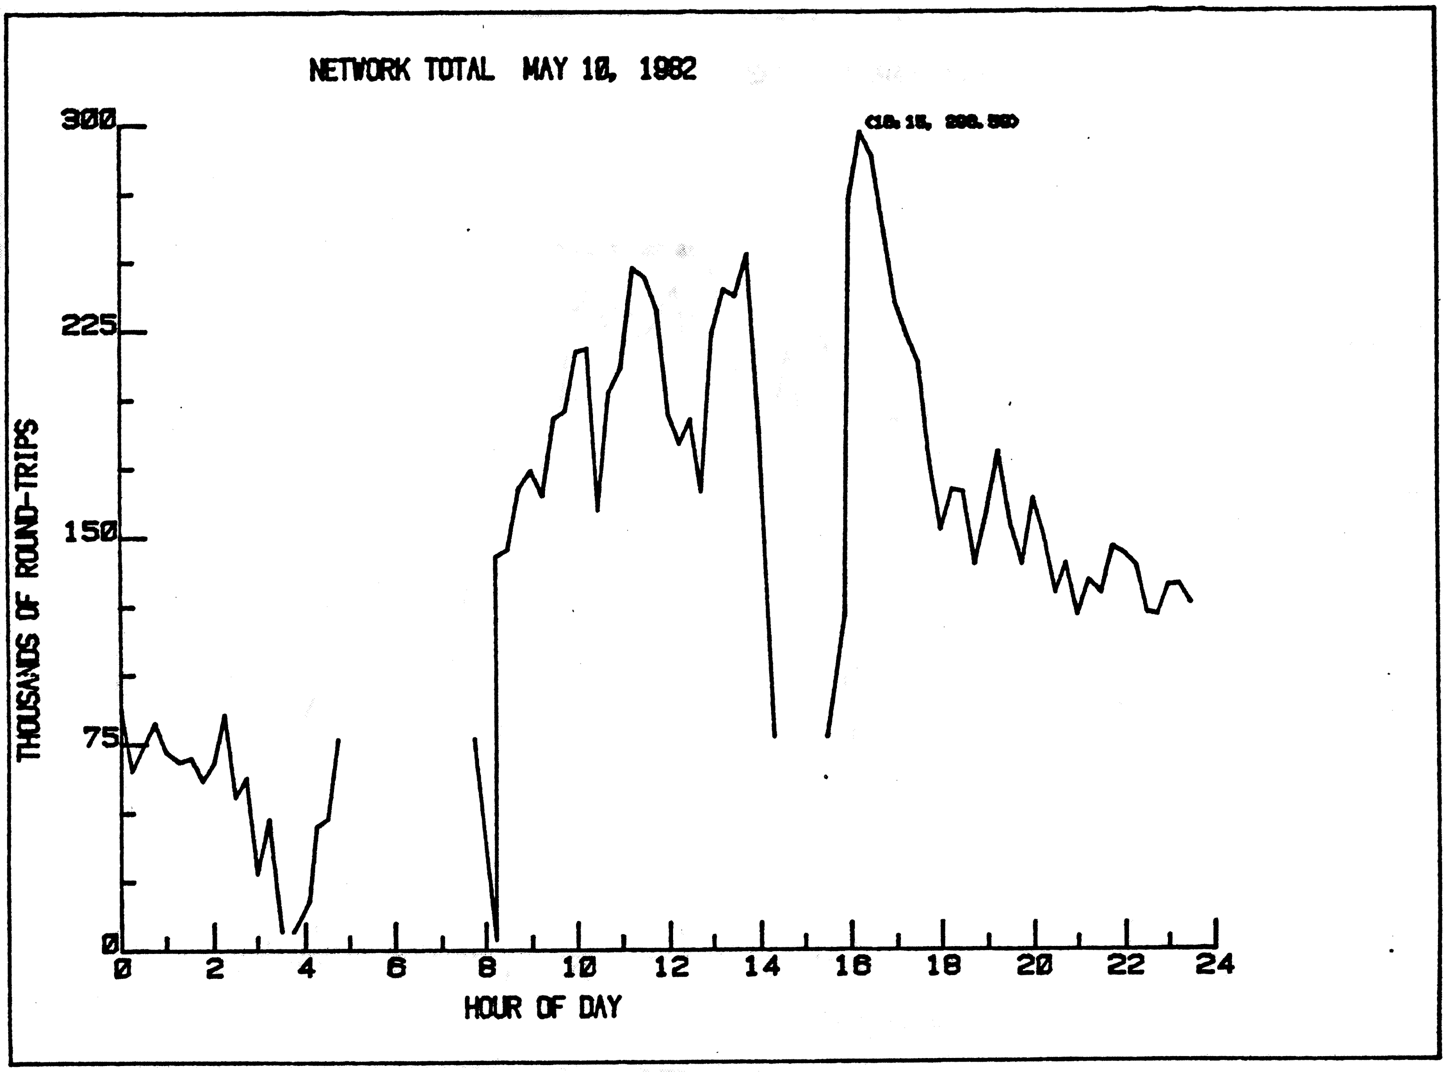

Messages per 15-minute interval on 10 May 1982.
Figure 13

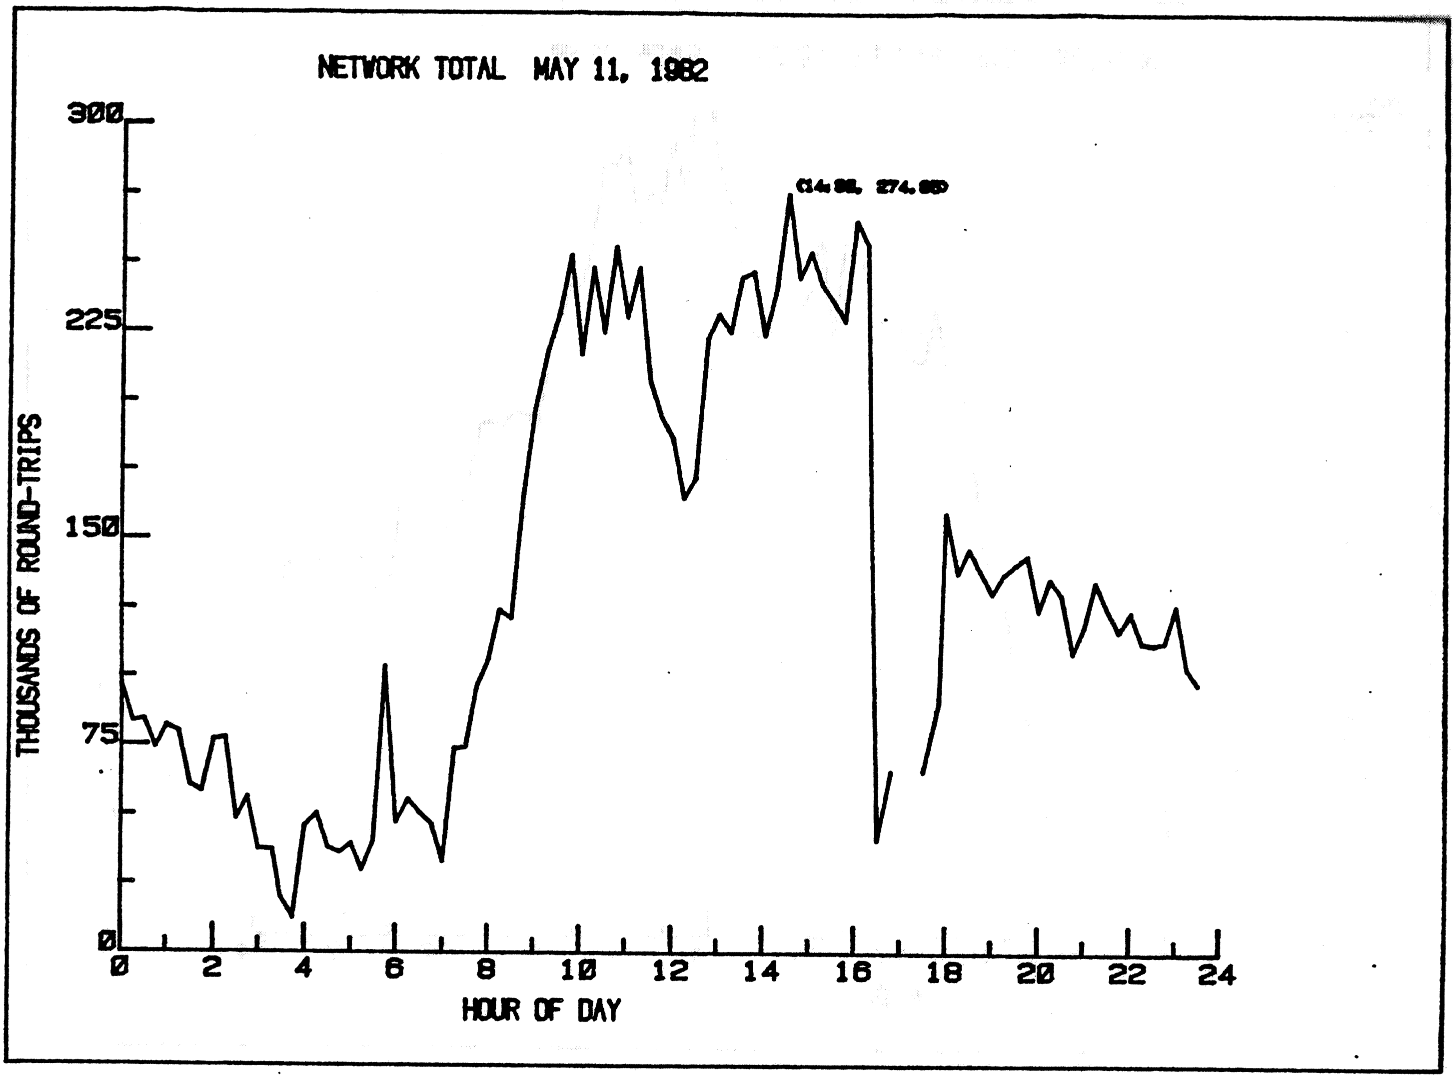

Messages per 15-minute interval on 11 May 1982.
Figure 14

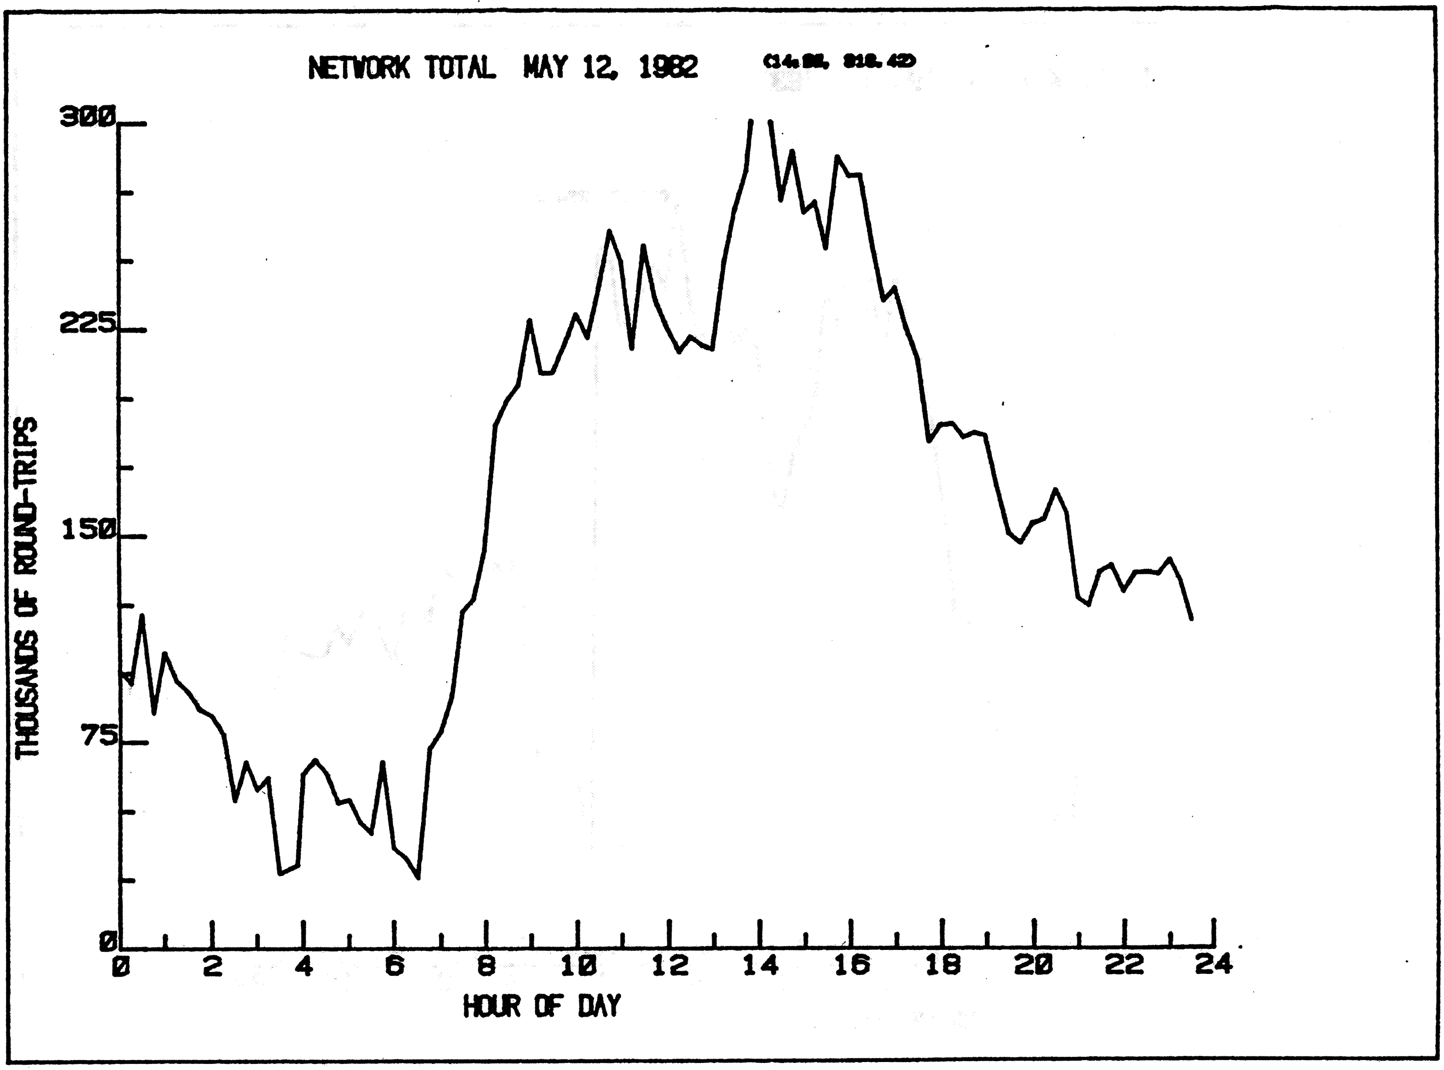

Messages per 15-minute interval on 12 May 1982.
Figure 15

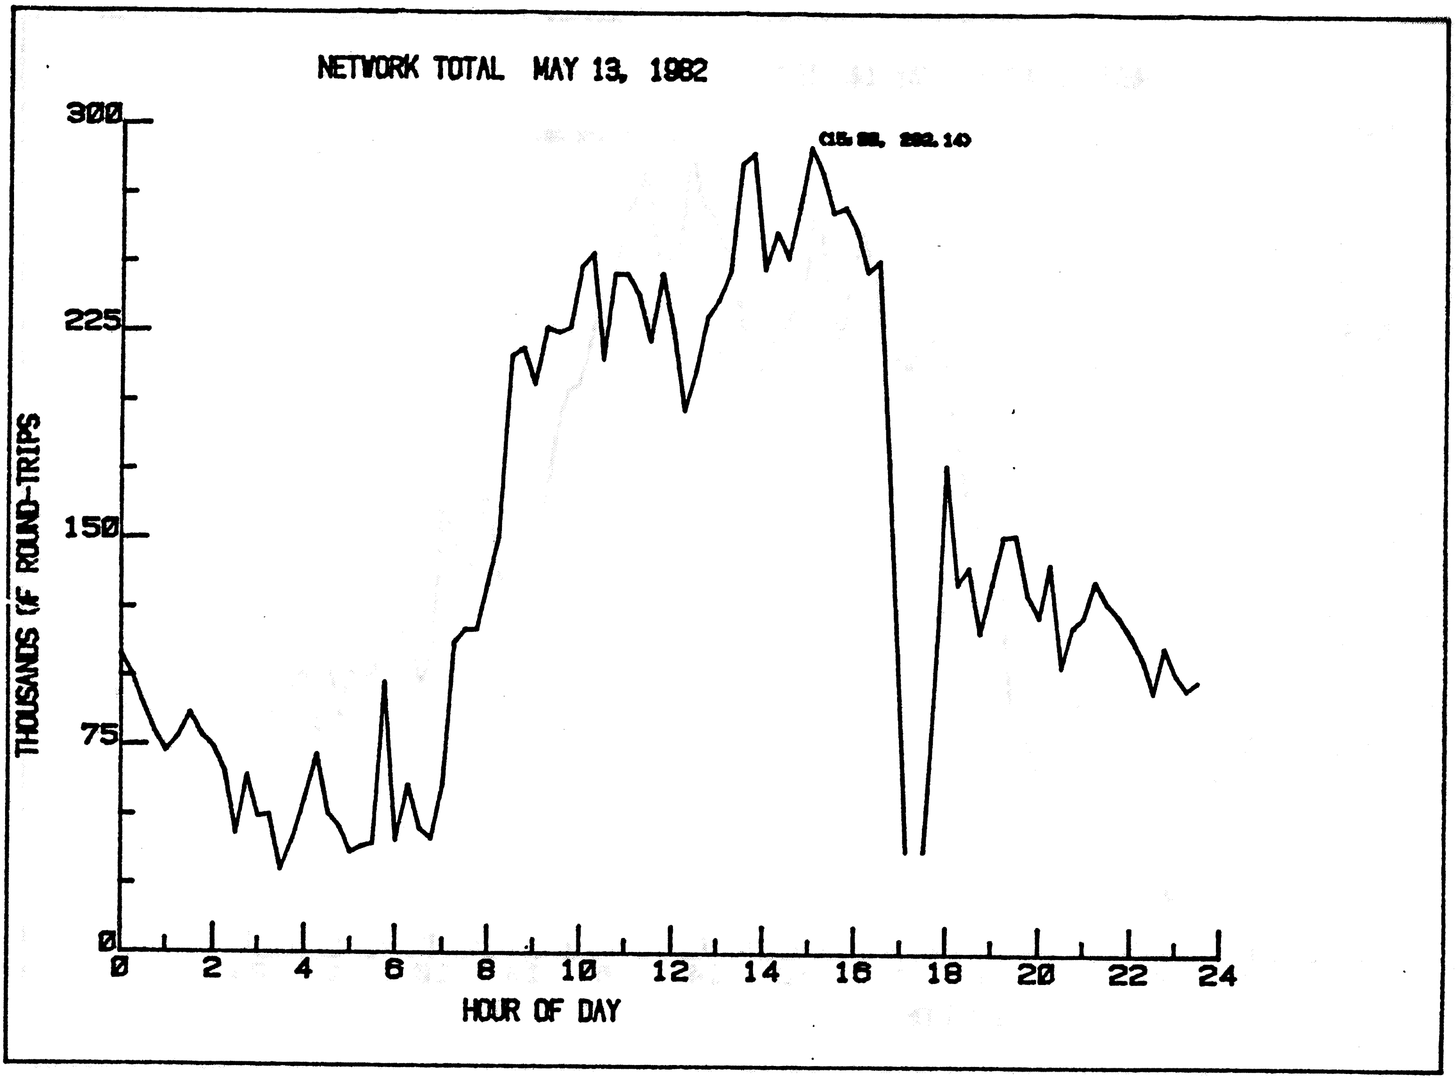

Messages per 15-minute interval on 13 May 1982.
Figure 16

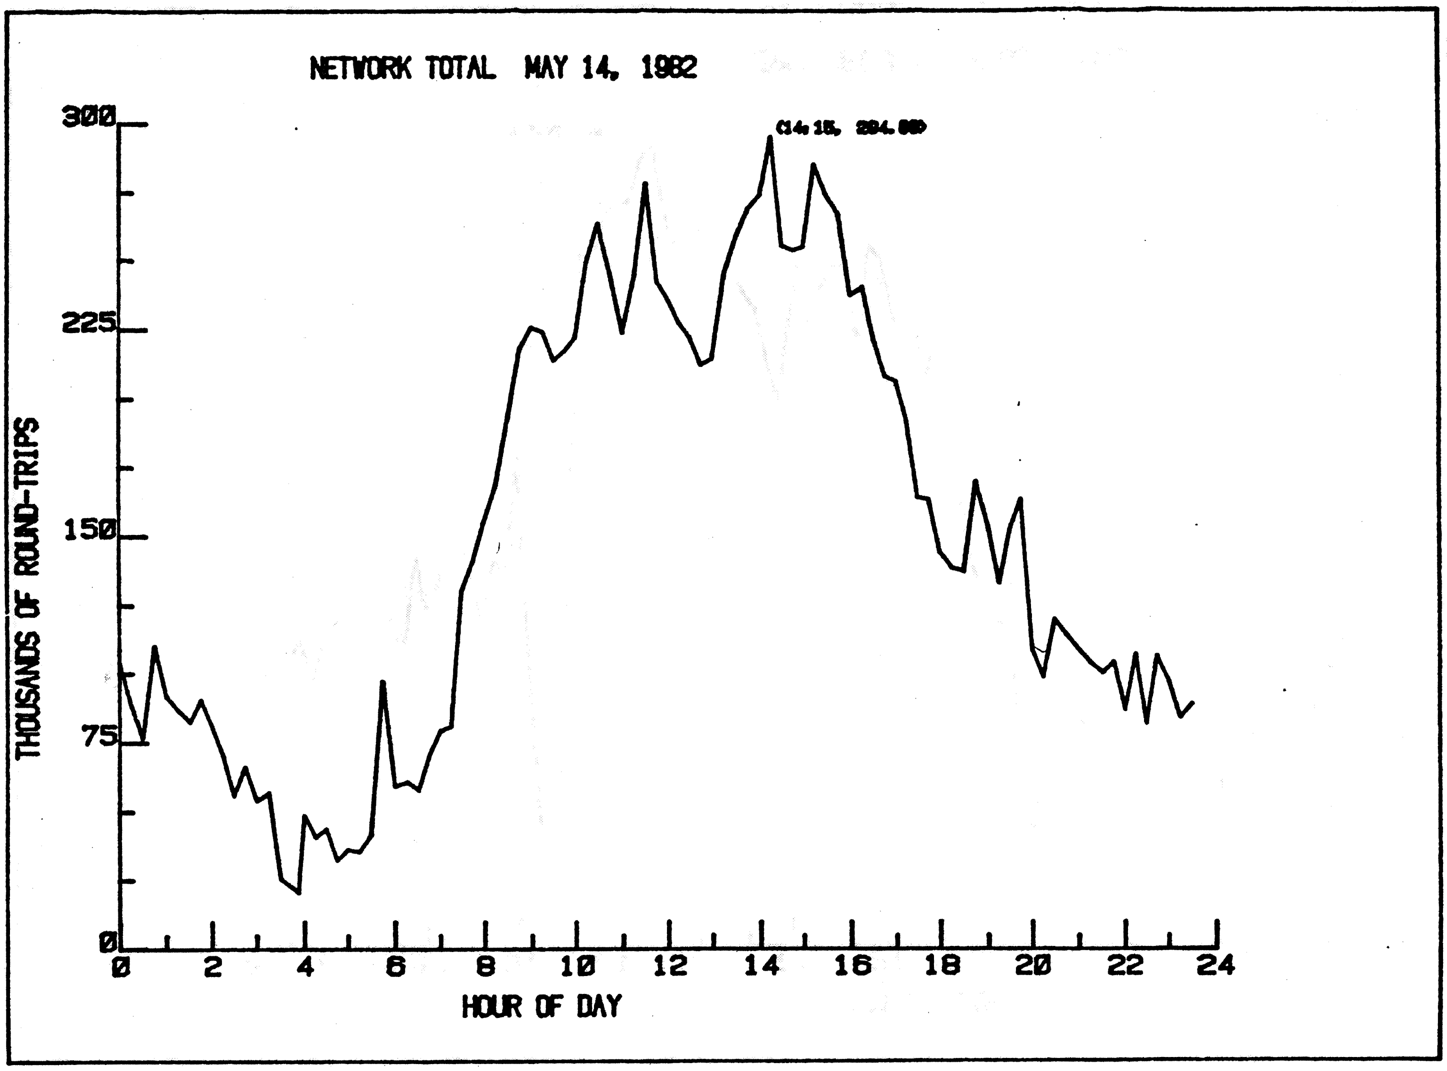

Messages per 15-minute interval on 14 May 1982.
Figure 17

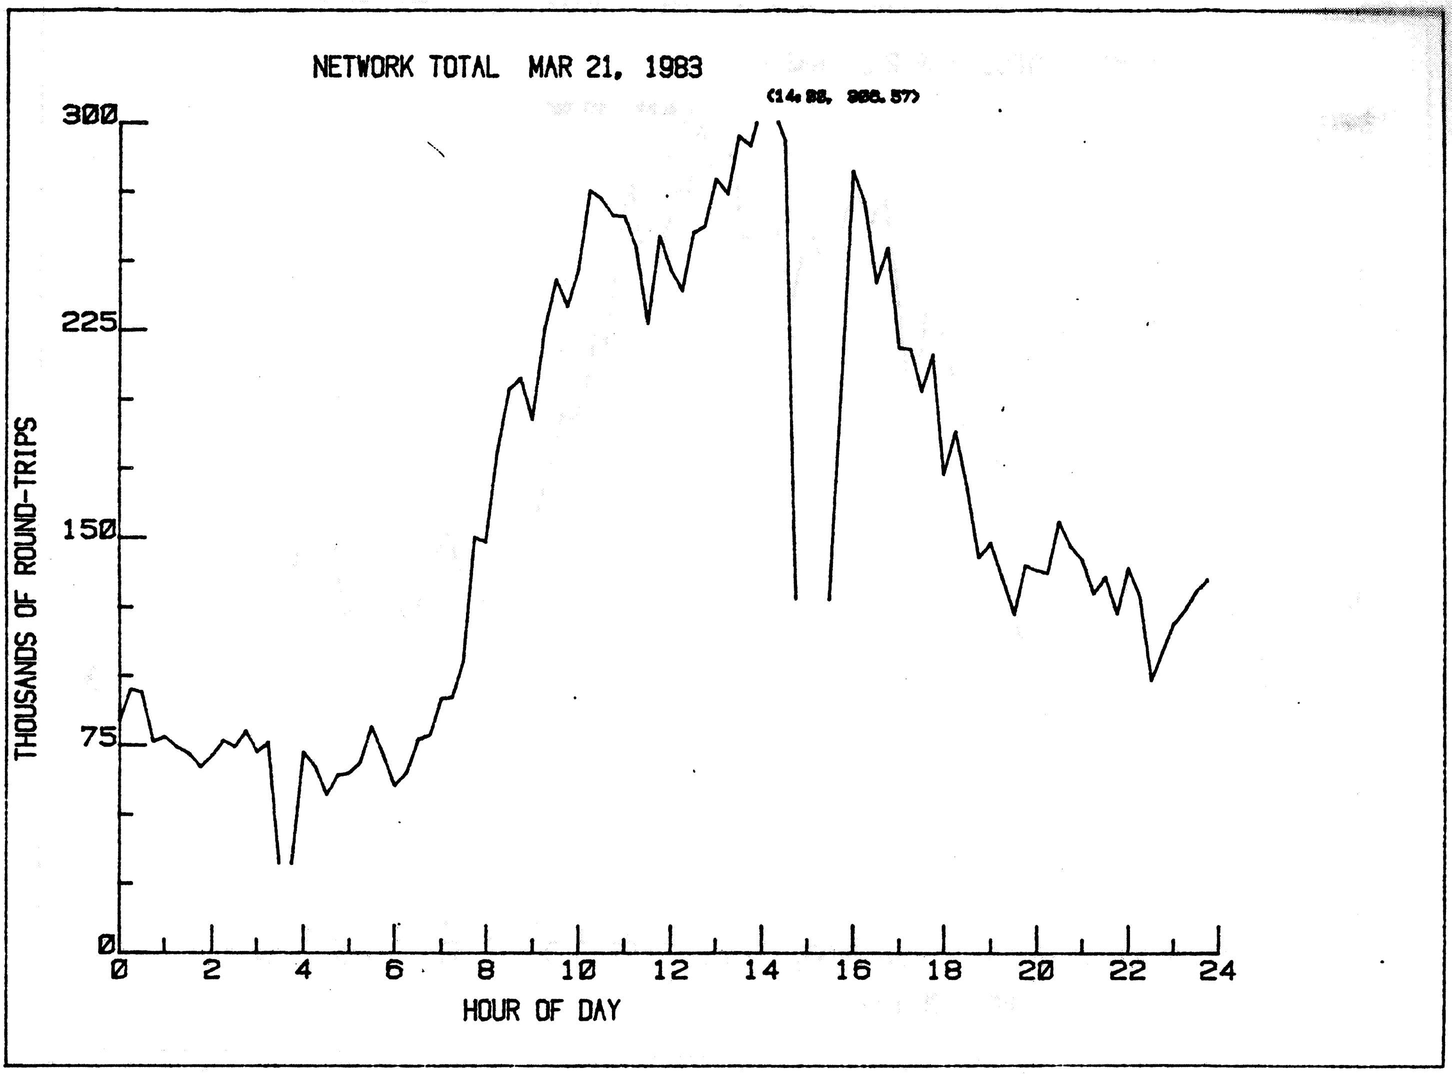

Messages per 15-minute interval on 21 March 1983.
Figure 18

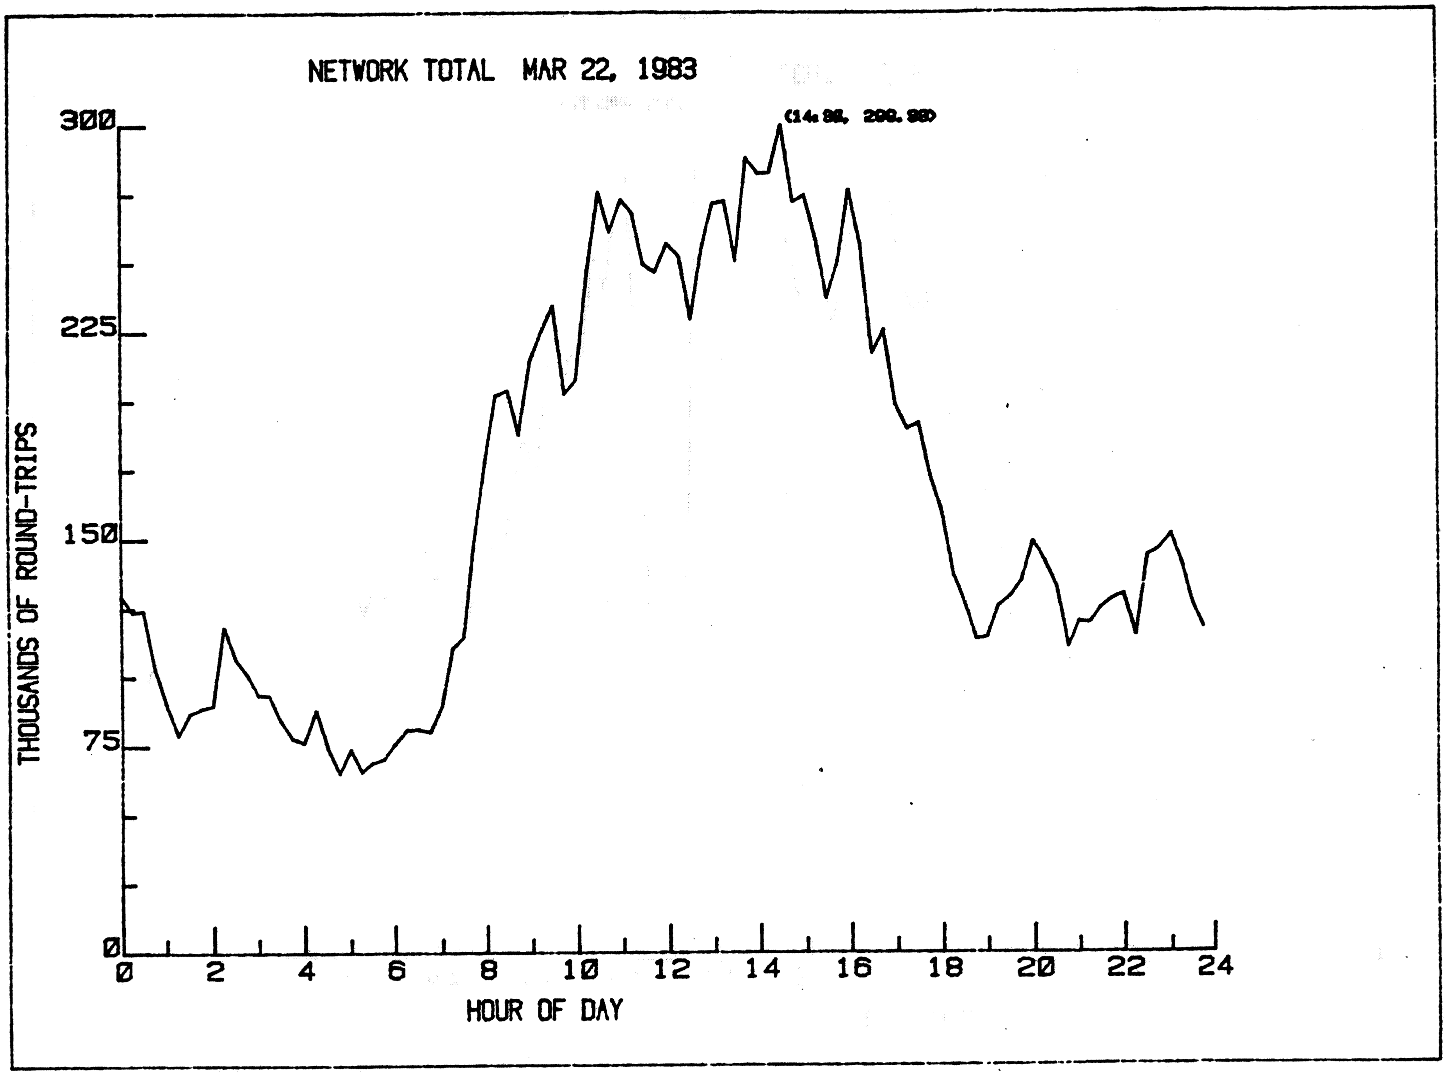

Messages per 15-minute interval on 22 March 1983.
Figure 19

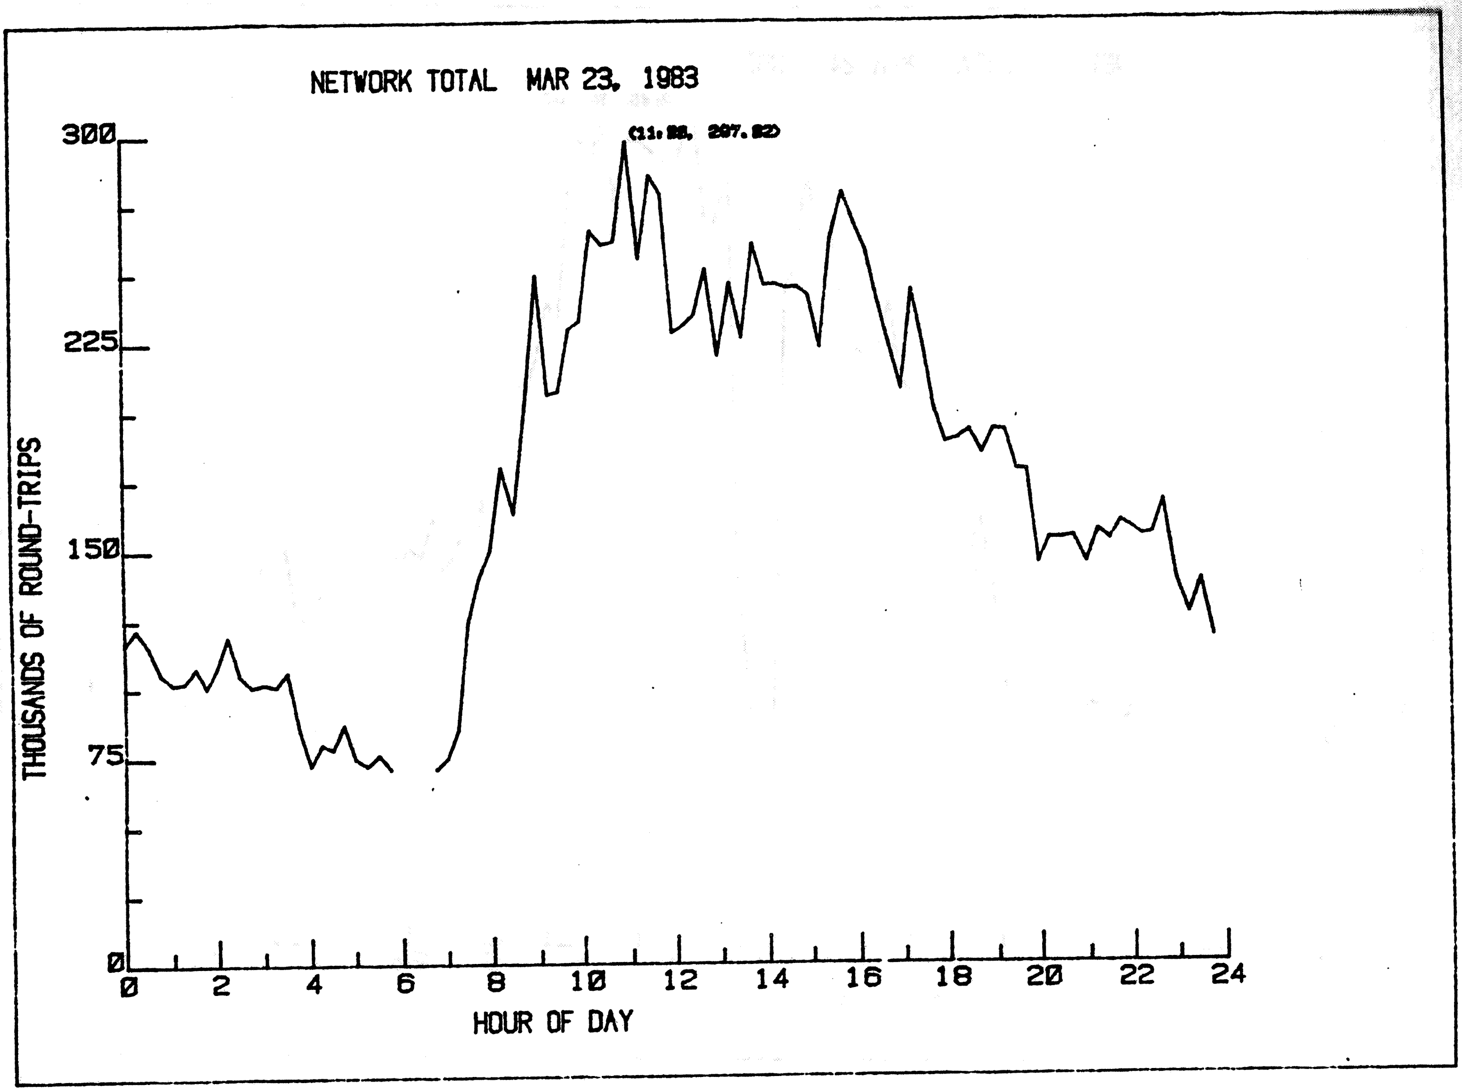

Messages per 15-minute interval on 23 March 1983.
Figure 20

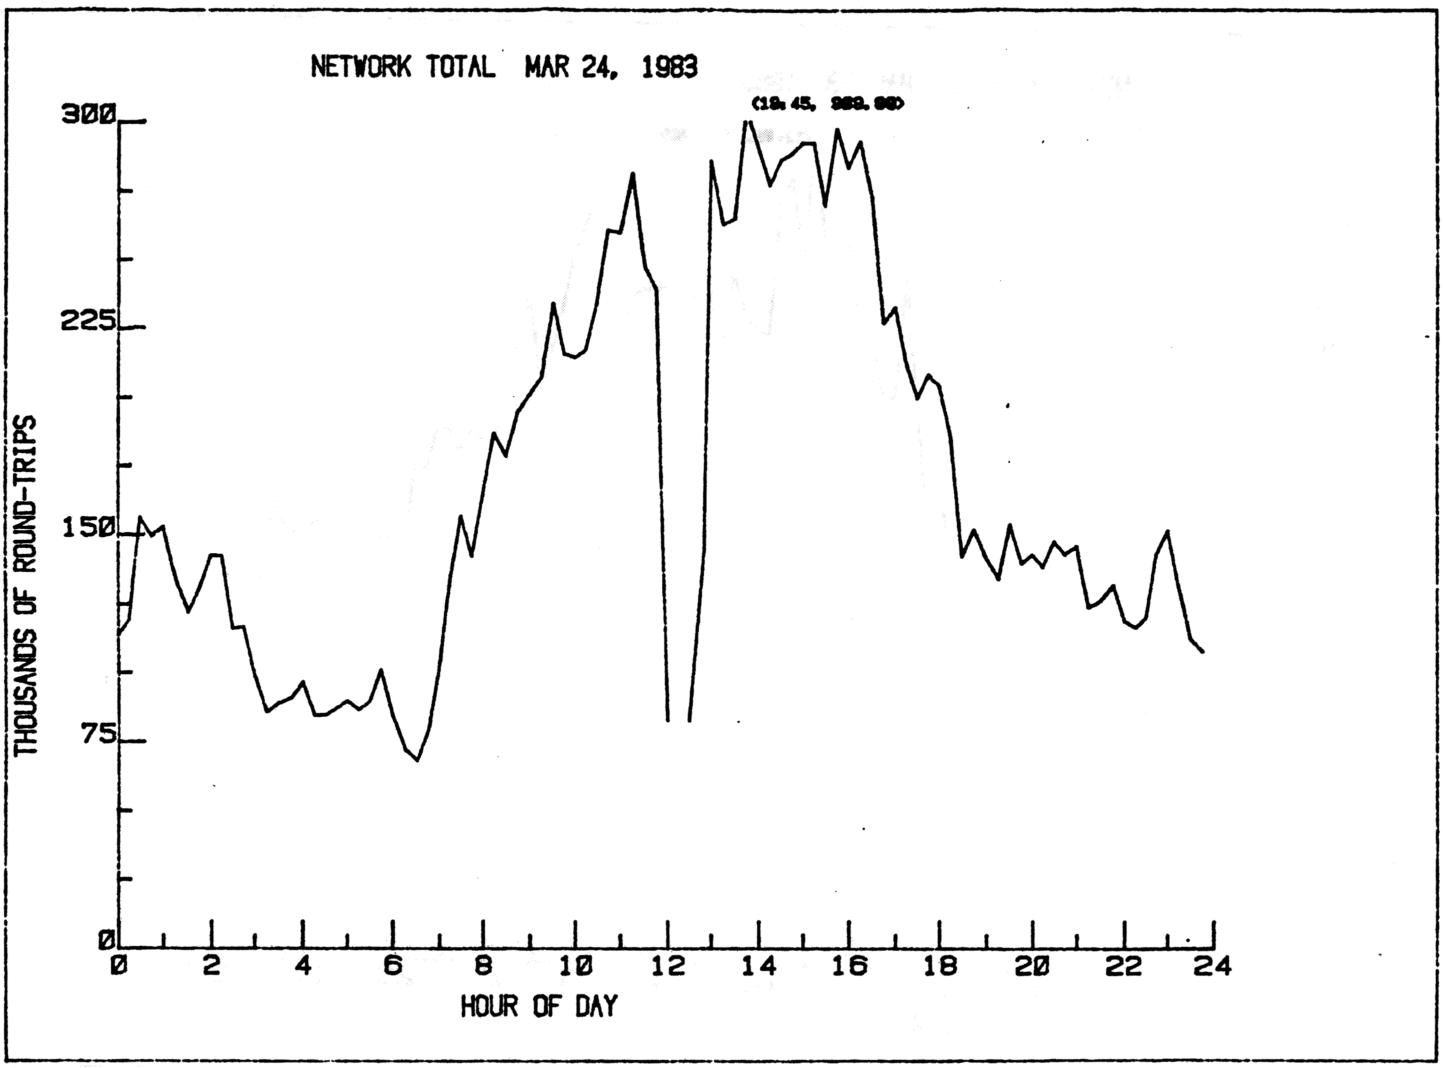

Messages per 15-minute interval on 24 March 1983.
Figure 21

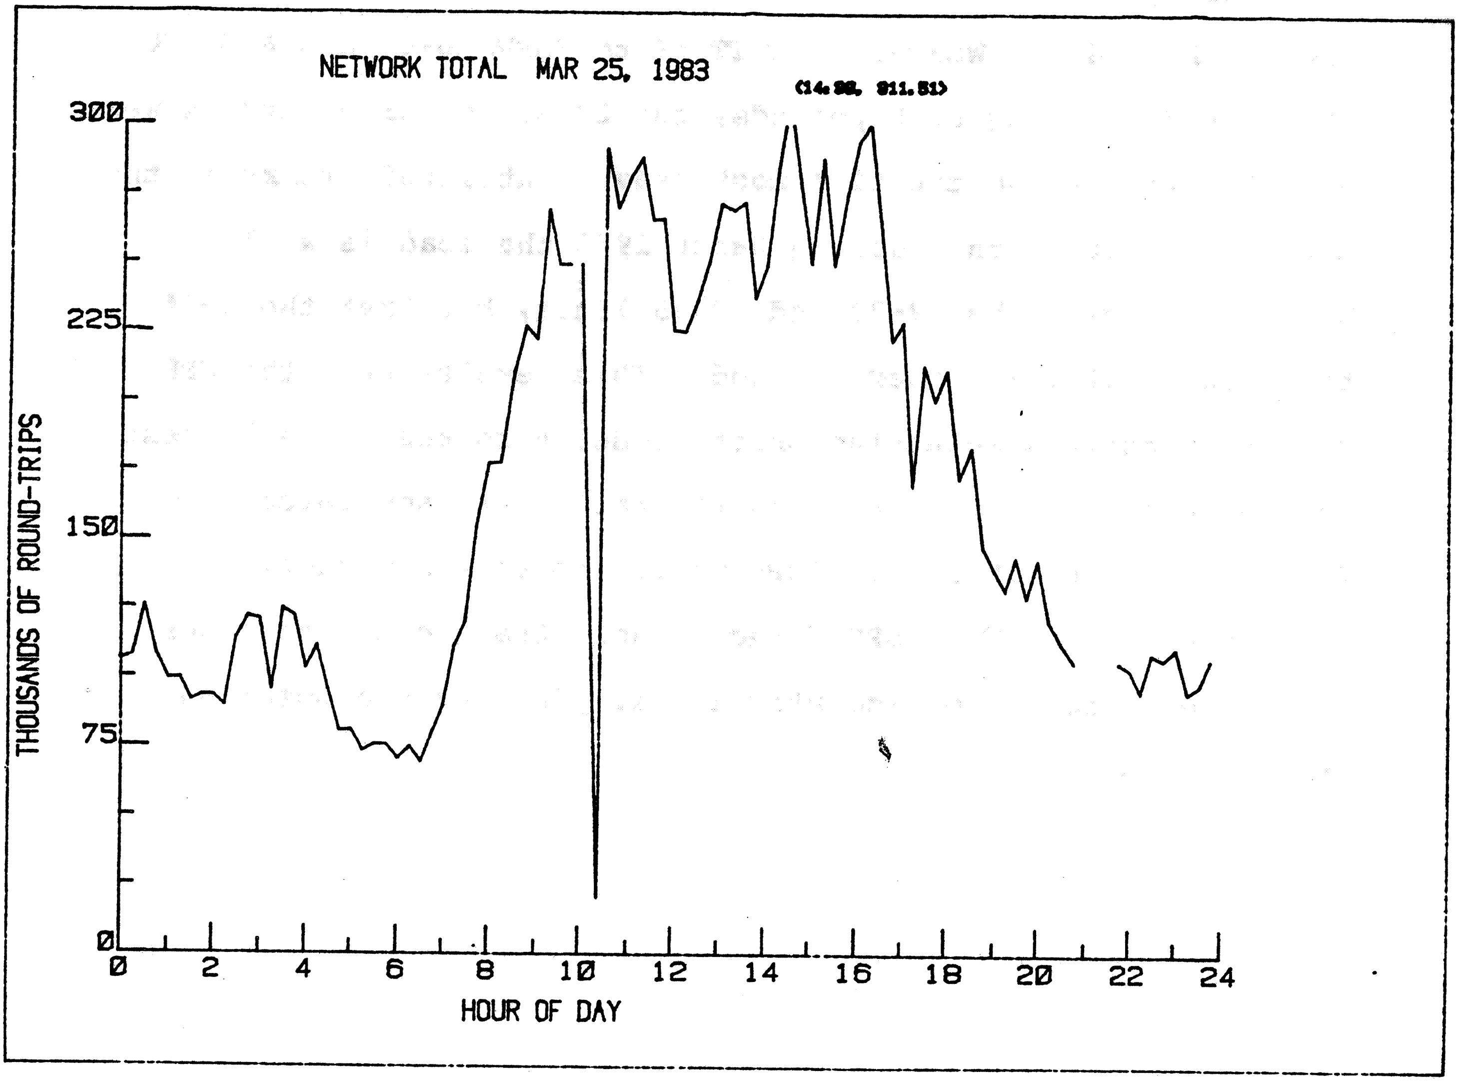

Messages per 15-minute interval on 25 March 1983.
Figure 22

The next series of figures shows the average utilization of selected inter-node trunks. Line utilization is plotted for the principal cross-country trunks: SCOTT-59 to NPS-33, ARPA-28 to USC-23, AFWL-48 to WSMR-74, SCOTT-59 to HARVARD-9, and BBN/RCC-5 to WISC-94. (During both periods, the 23-28 and 33-59 trunks were almost always among the five most heavily utilized trunks in the ARPANET.) Notice that during March 1983 the load is well balanced between the 59-33 and 28-23 lines, but that the 5-94 trunk is relatively underutilized. This results from the SPF routing algorithm selecting shortest-delay routes. The largest sources and sinks of cross-country traffic are separated from IMPs 5 and 94 by chains of IMPs which impose a substantial aggregate delay. Thus SPF directs their traffic to the other two, more direct, cross-country links, yielding the observed distribution.

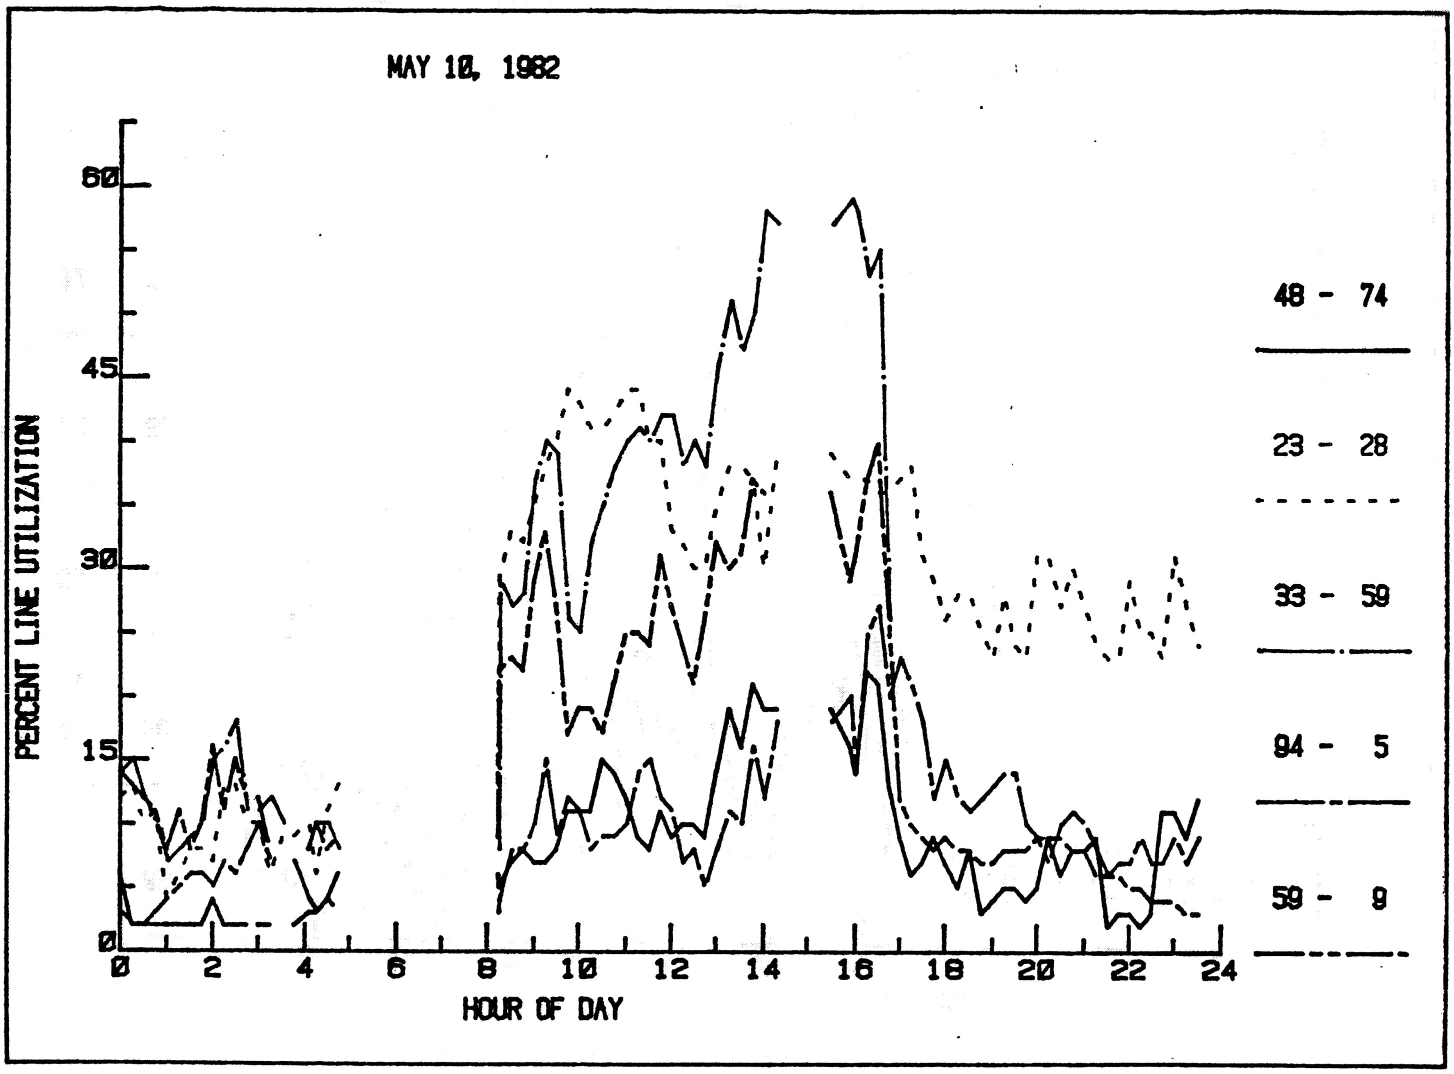

Cross-country trunk utilization on 10 May 1982.
Figure 23

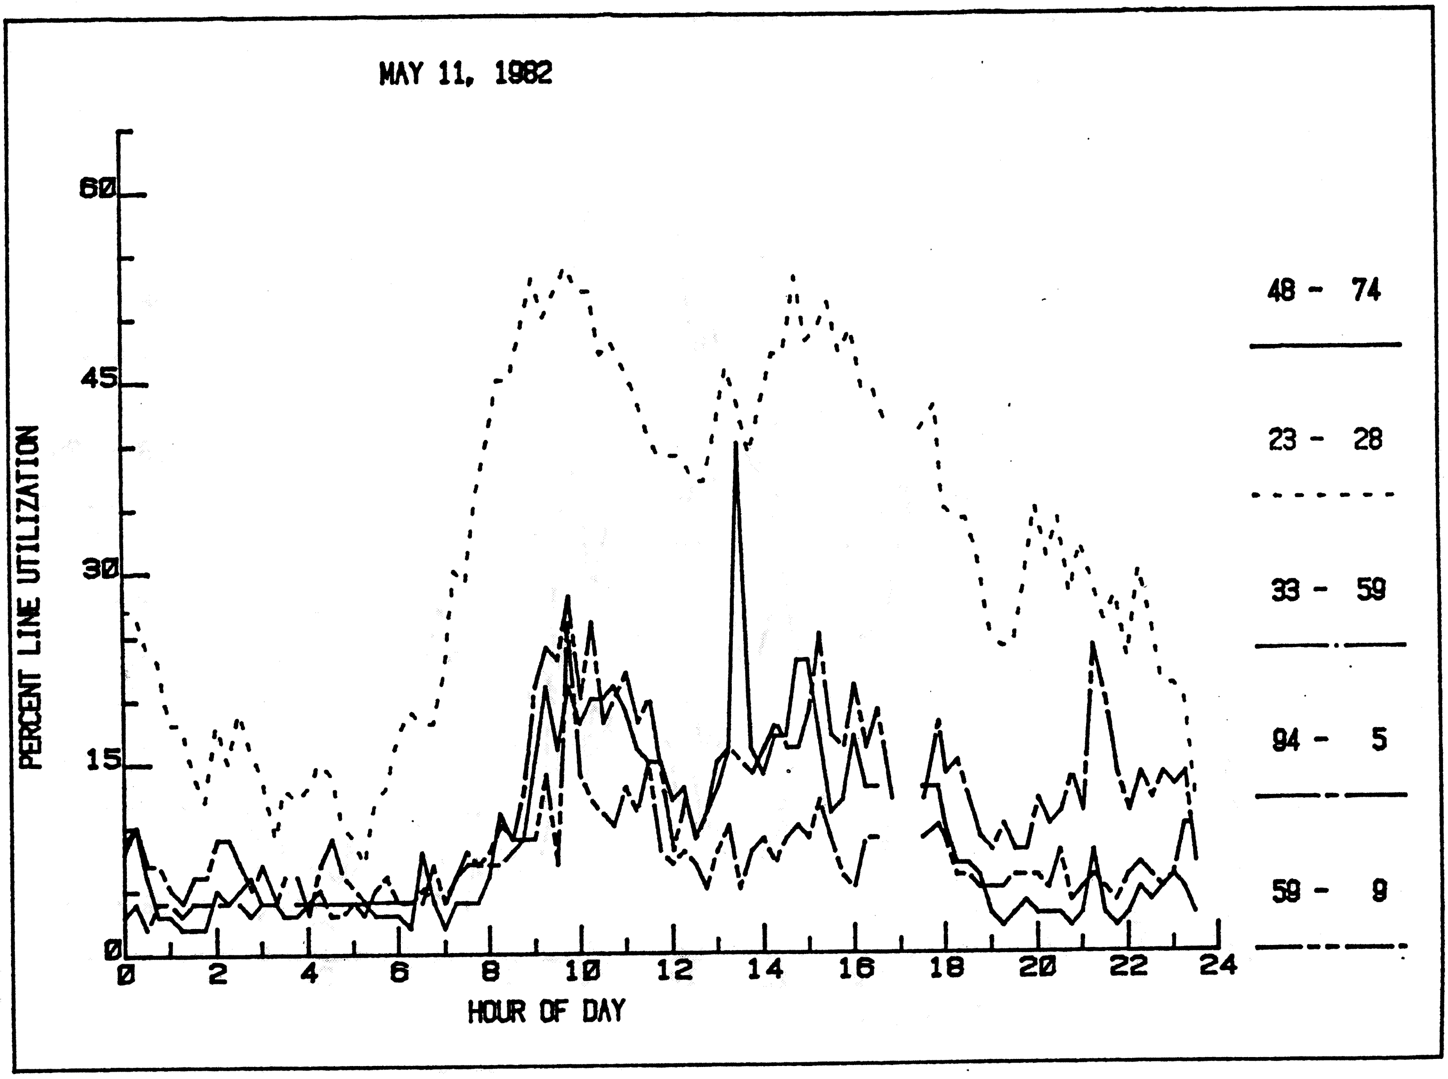

Cross-country trunk utilization on 11 May 1982.
Figure 24

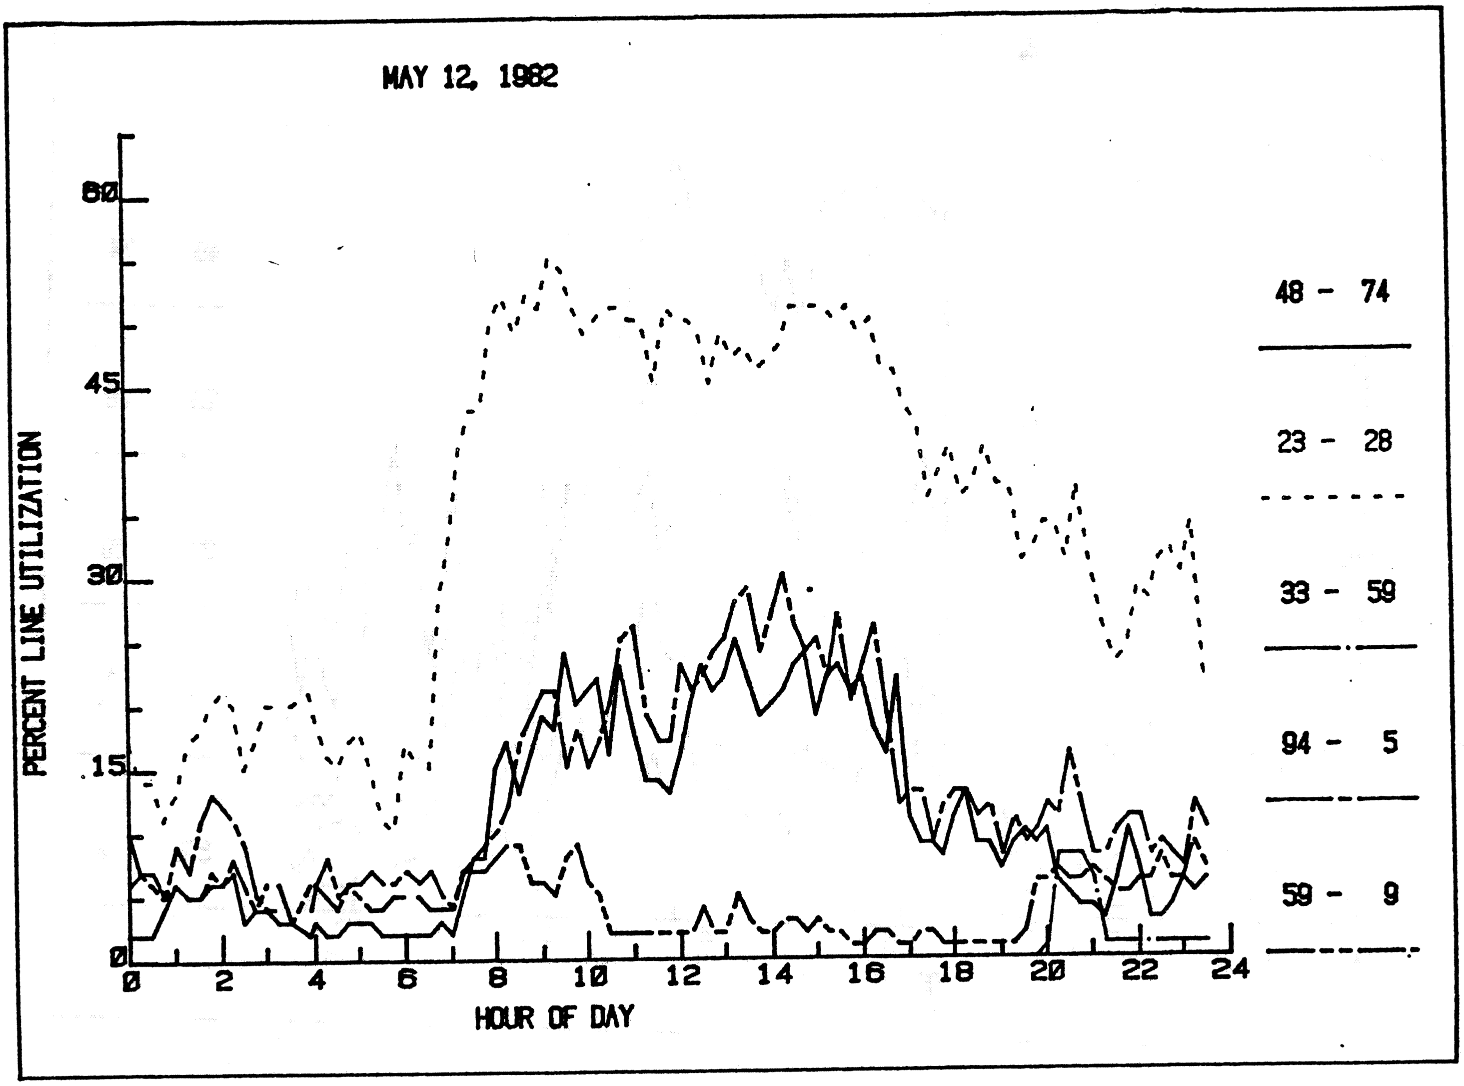

Cross-country trunk utilization on 12 May 1982.
Figure 25

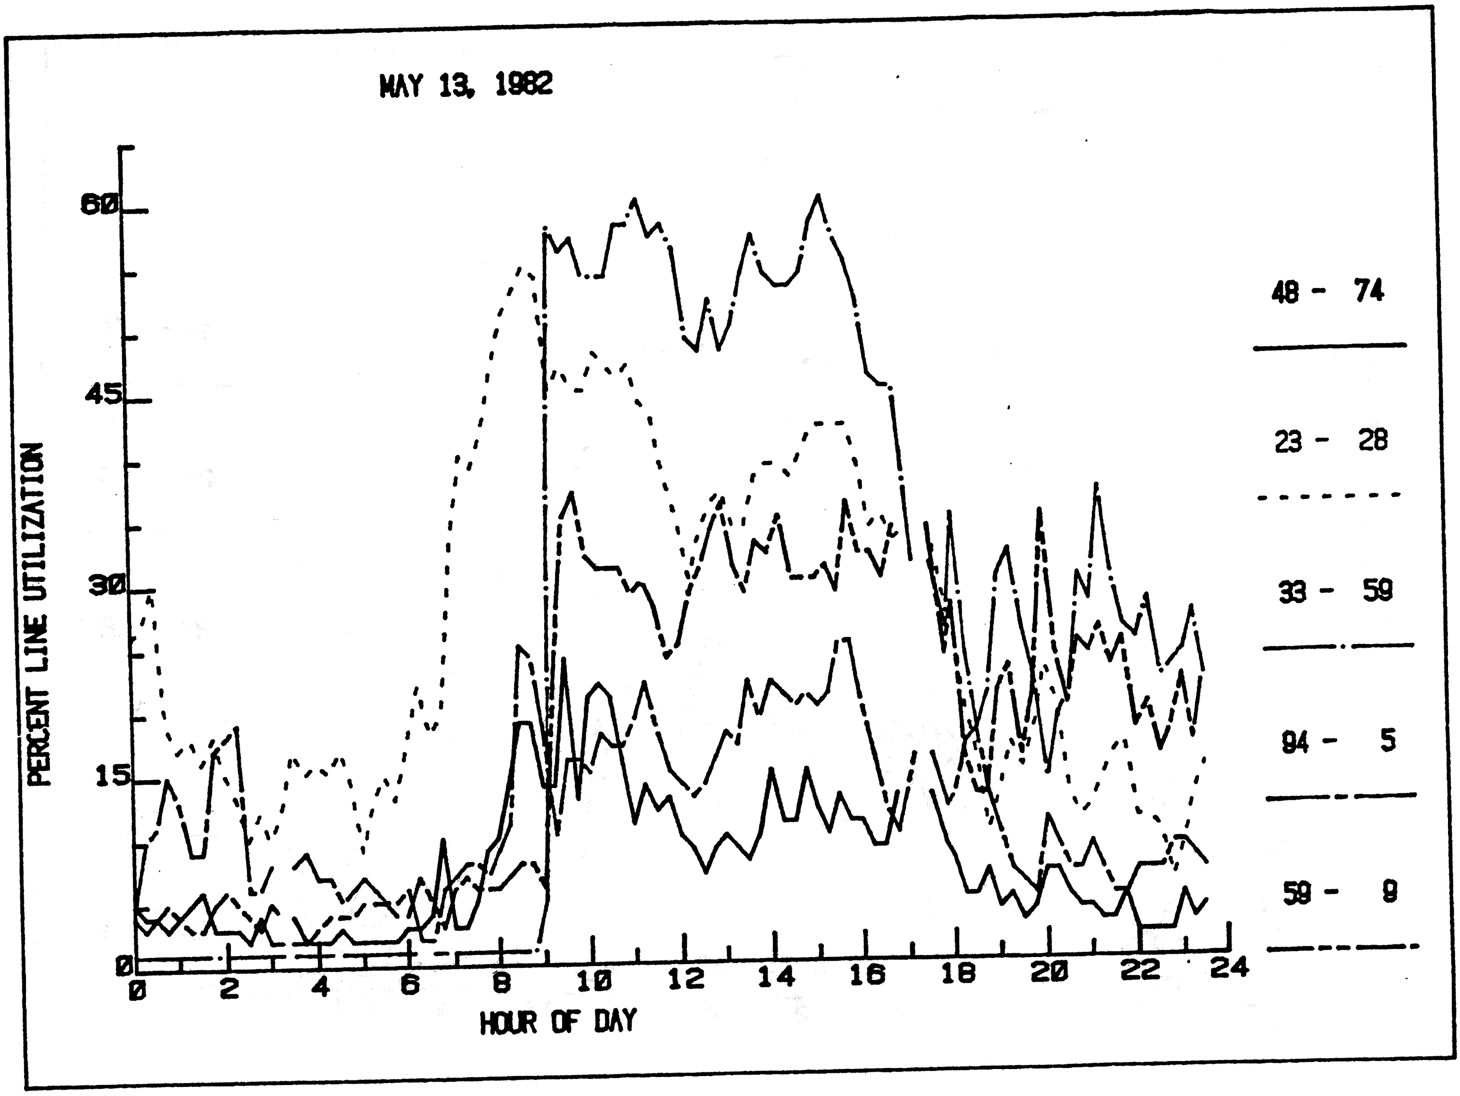

Cross-country trunk utilization on 13 May 1982.
Figure 26

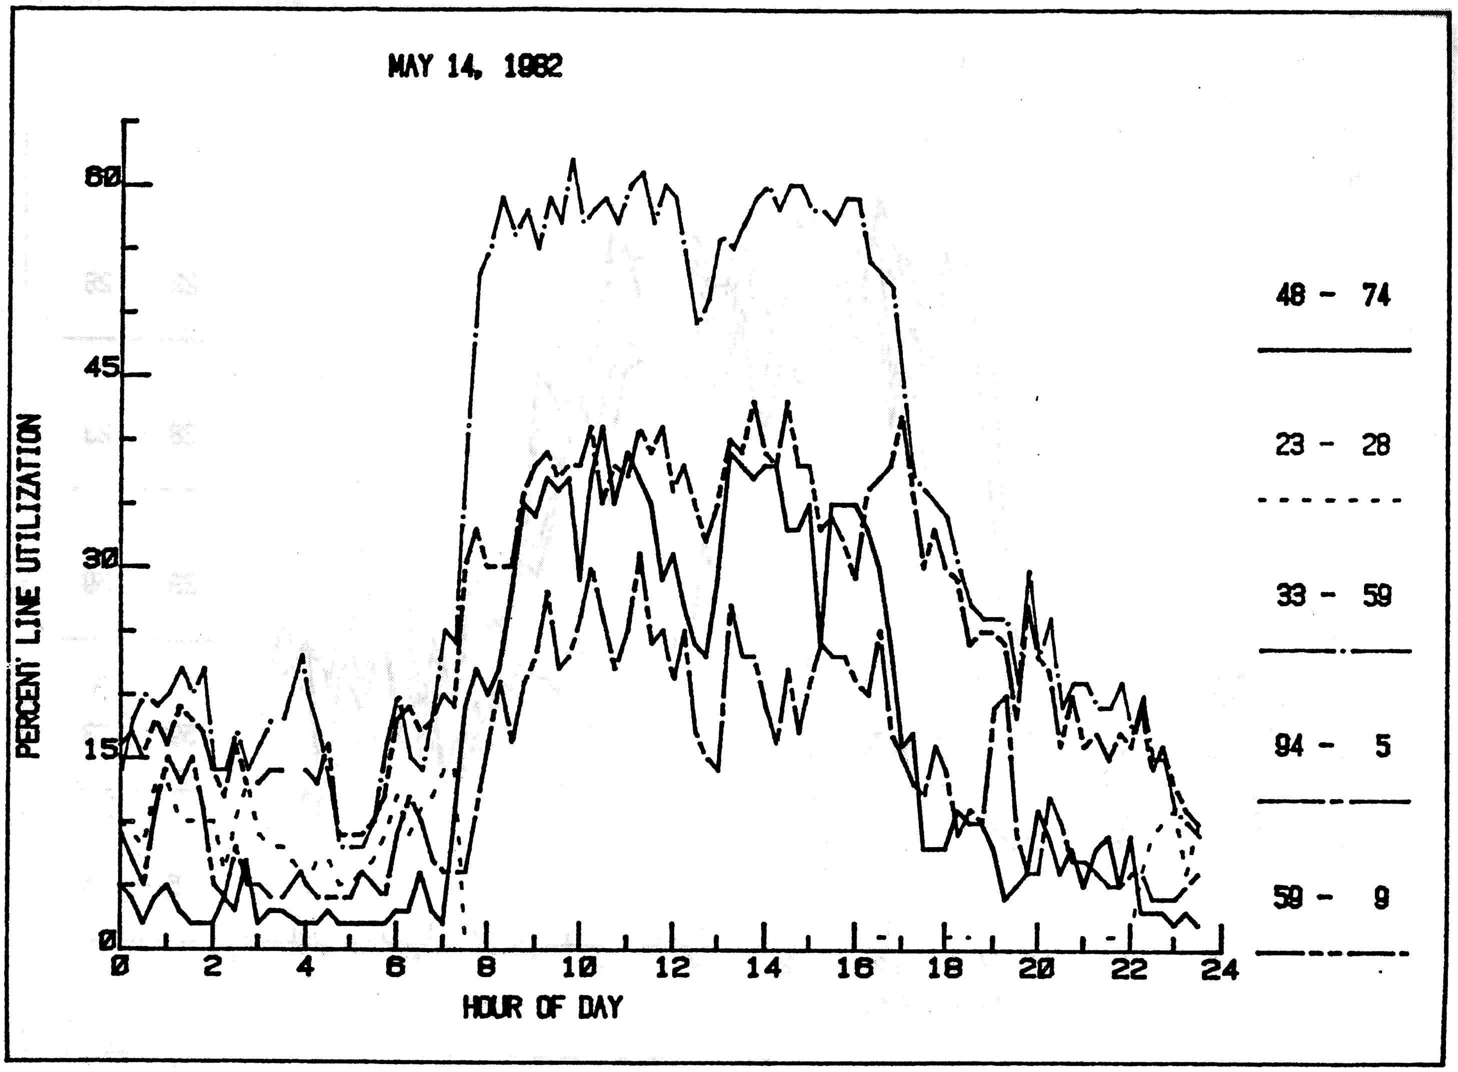

Cross-country trunk utilization on 14 May 1982.
Figure 27

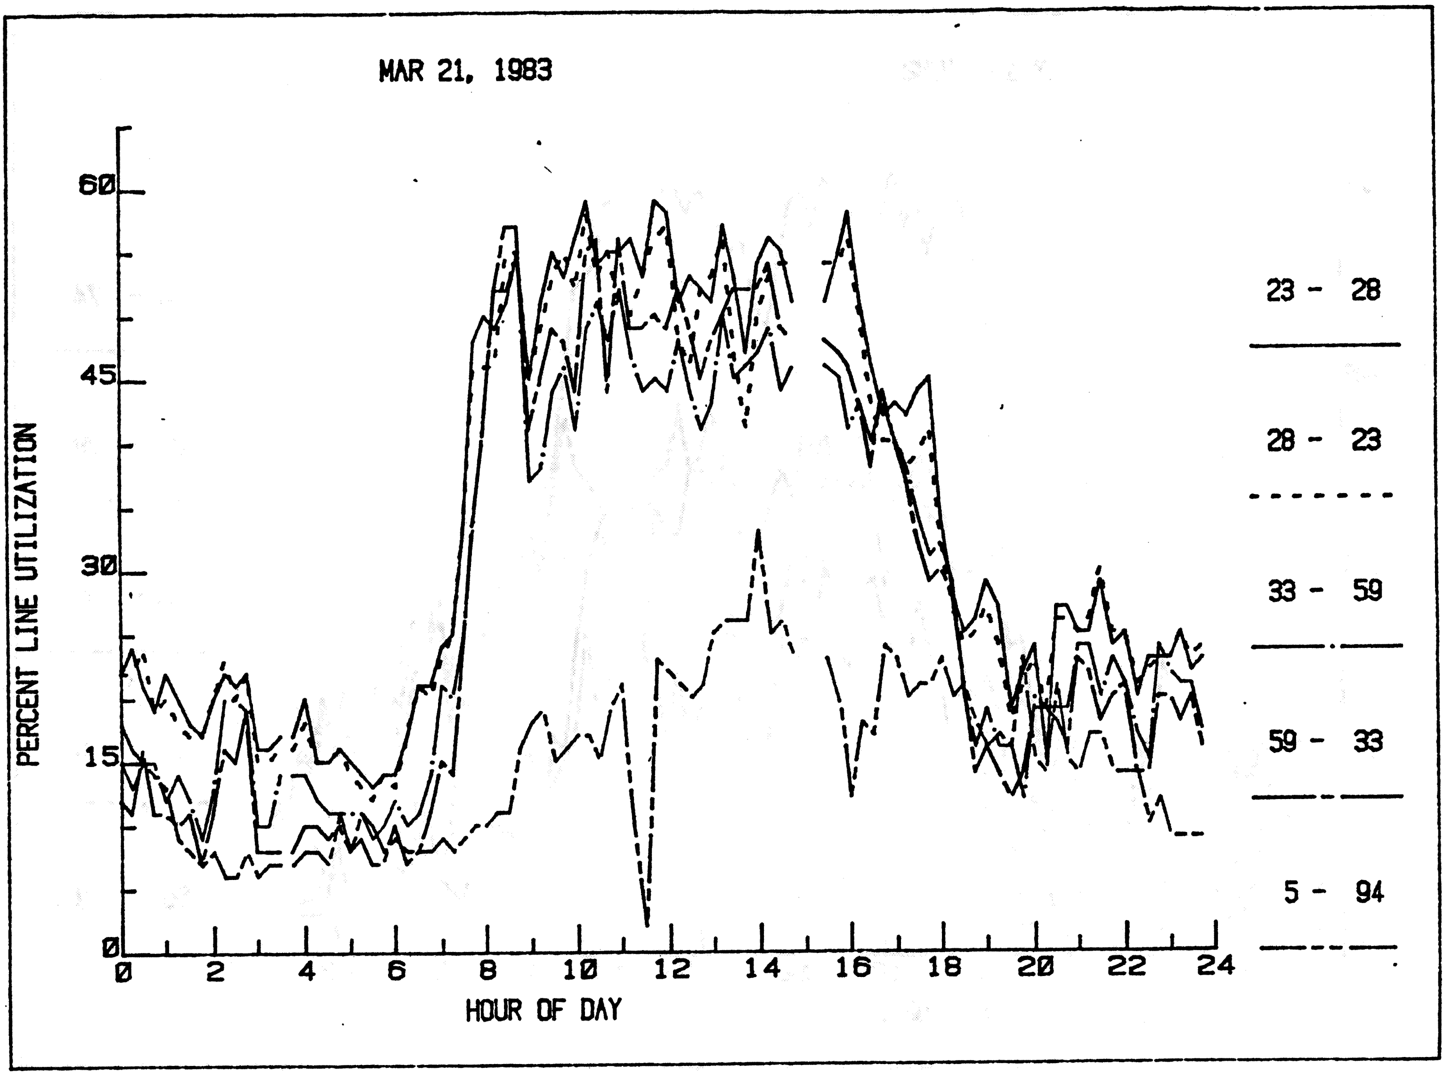

Cross-country trunk utilization on 21 March 1983.
Figure 28

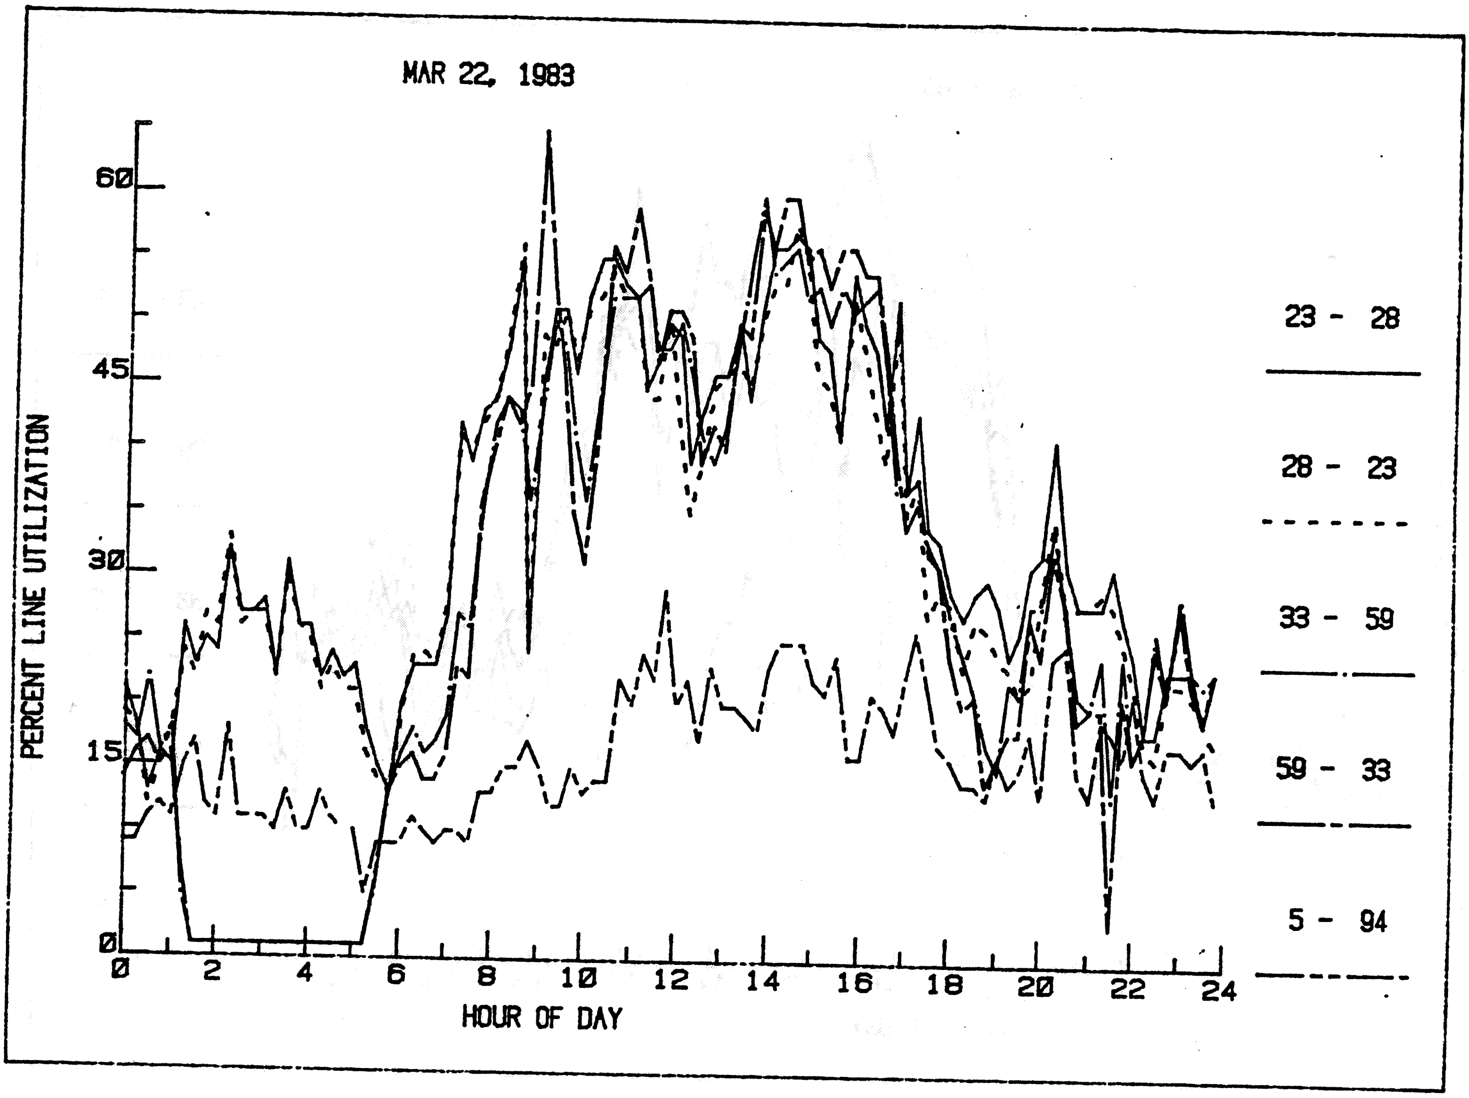

Cross-country trunk utilization on 22 March 1983.
Figure 29

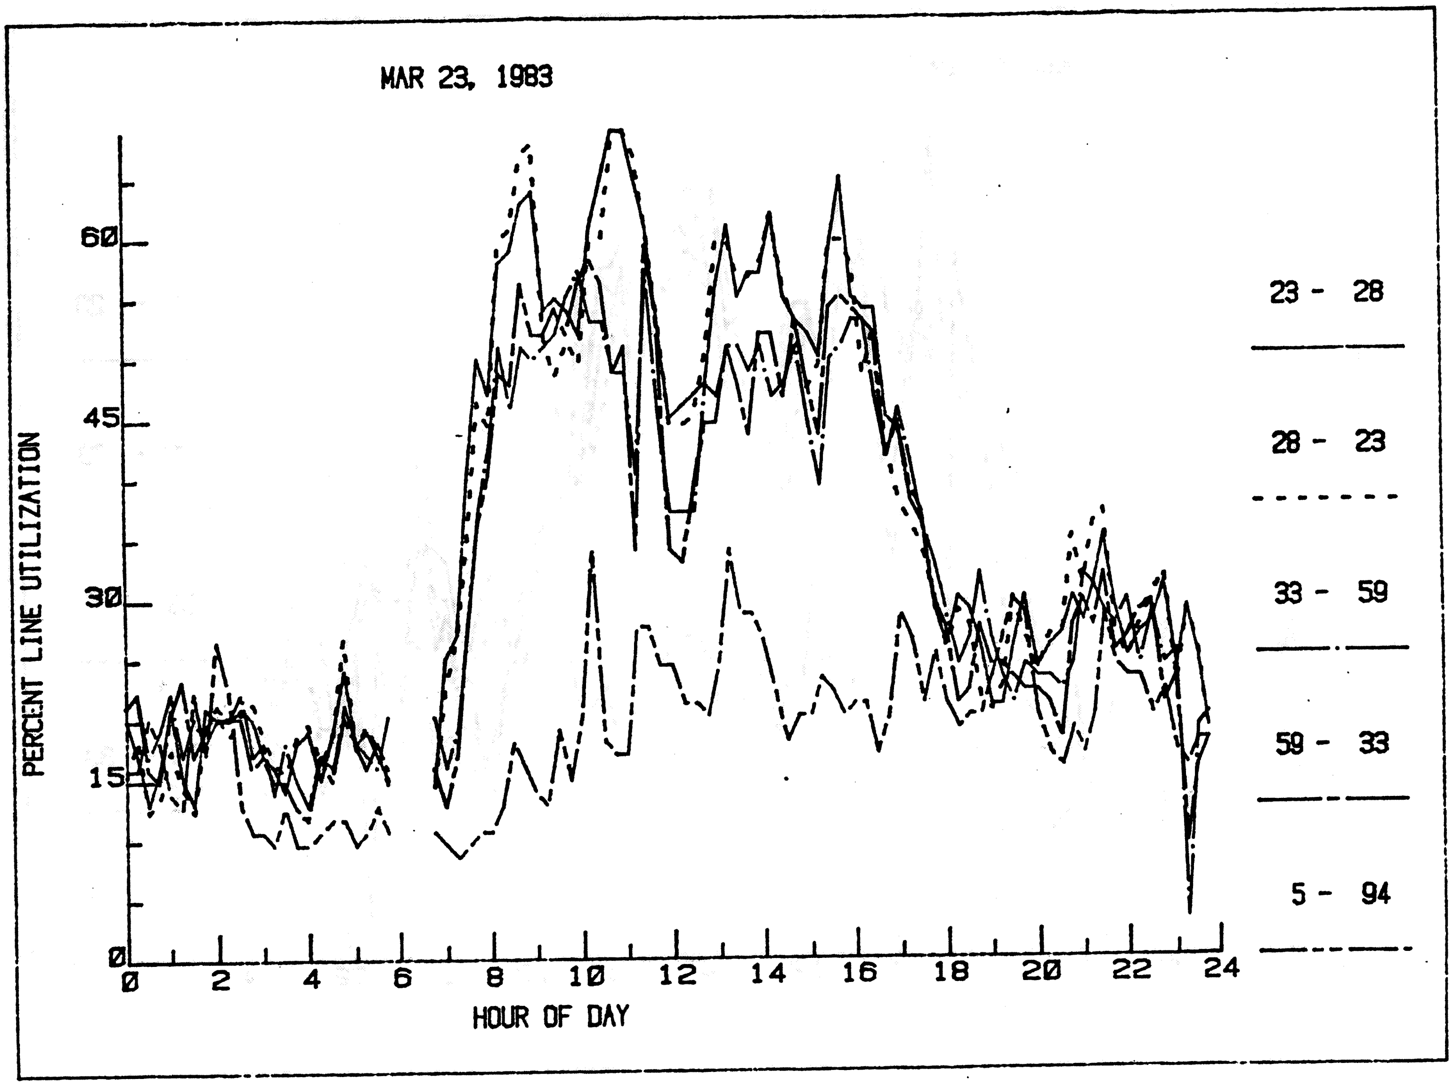

Cross-country trunk utilization on 23 March 1983.
Figure 30

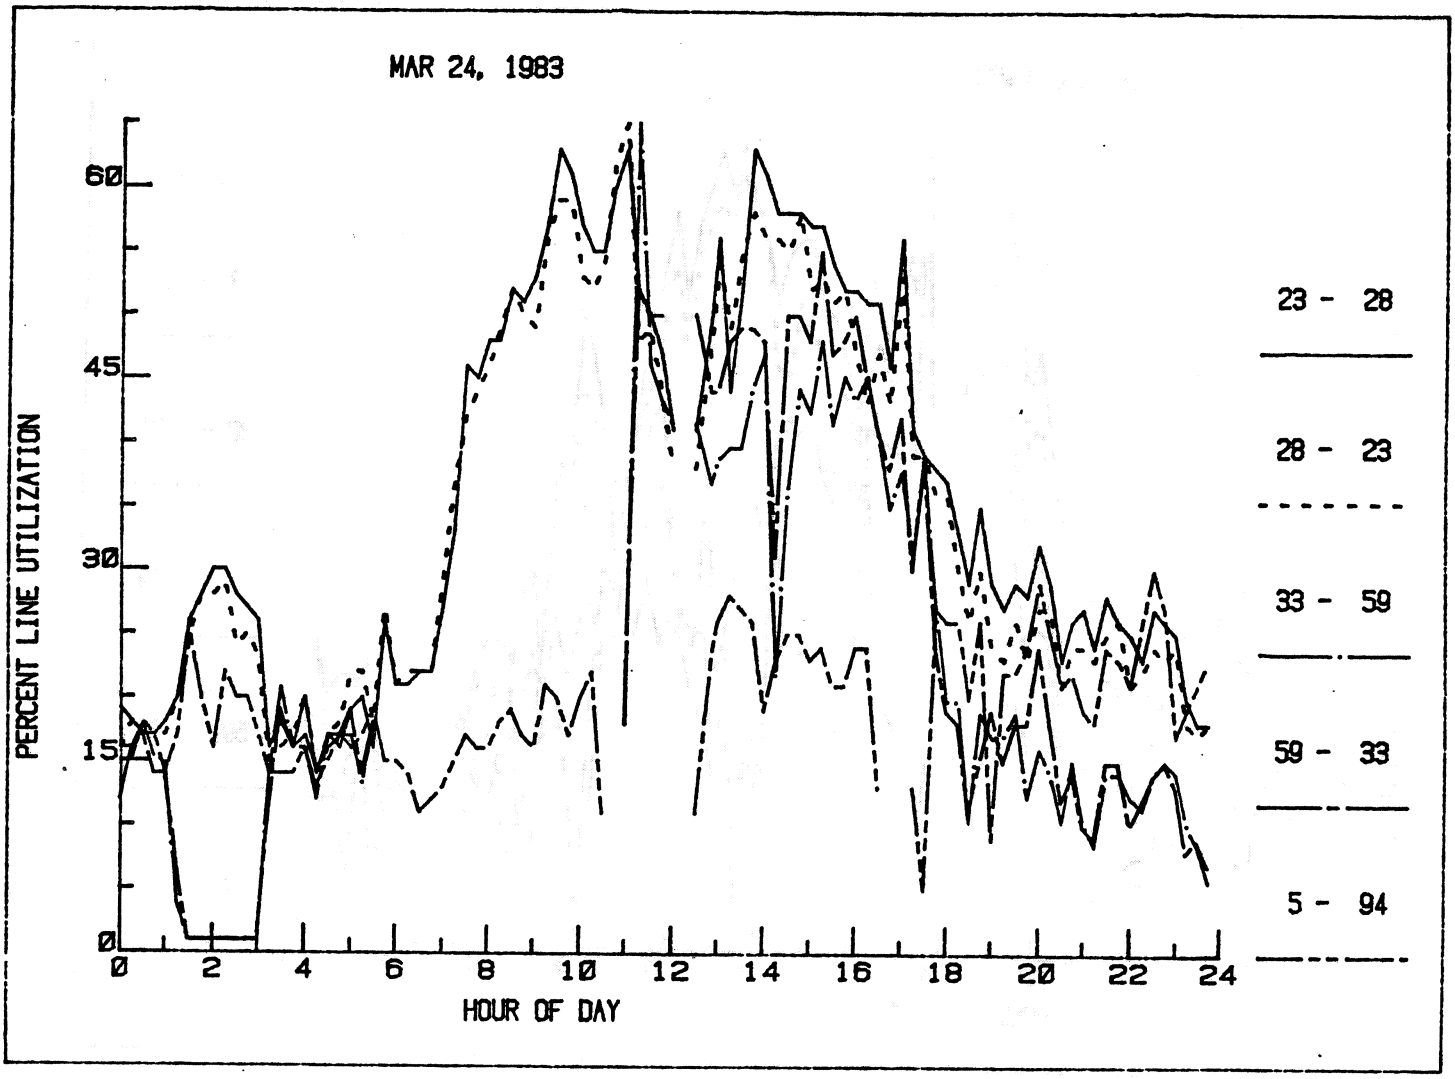

Cross-country trunk utilization on 24 March 1983.
Figure 31

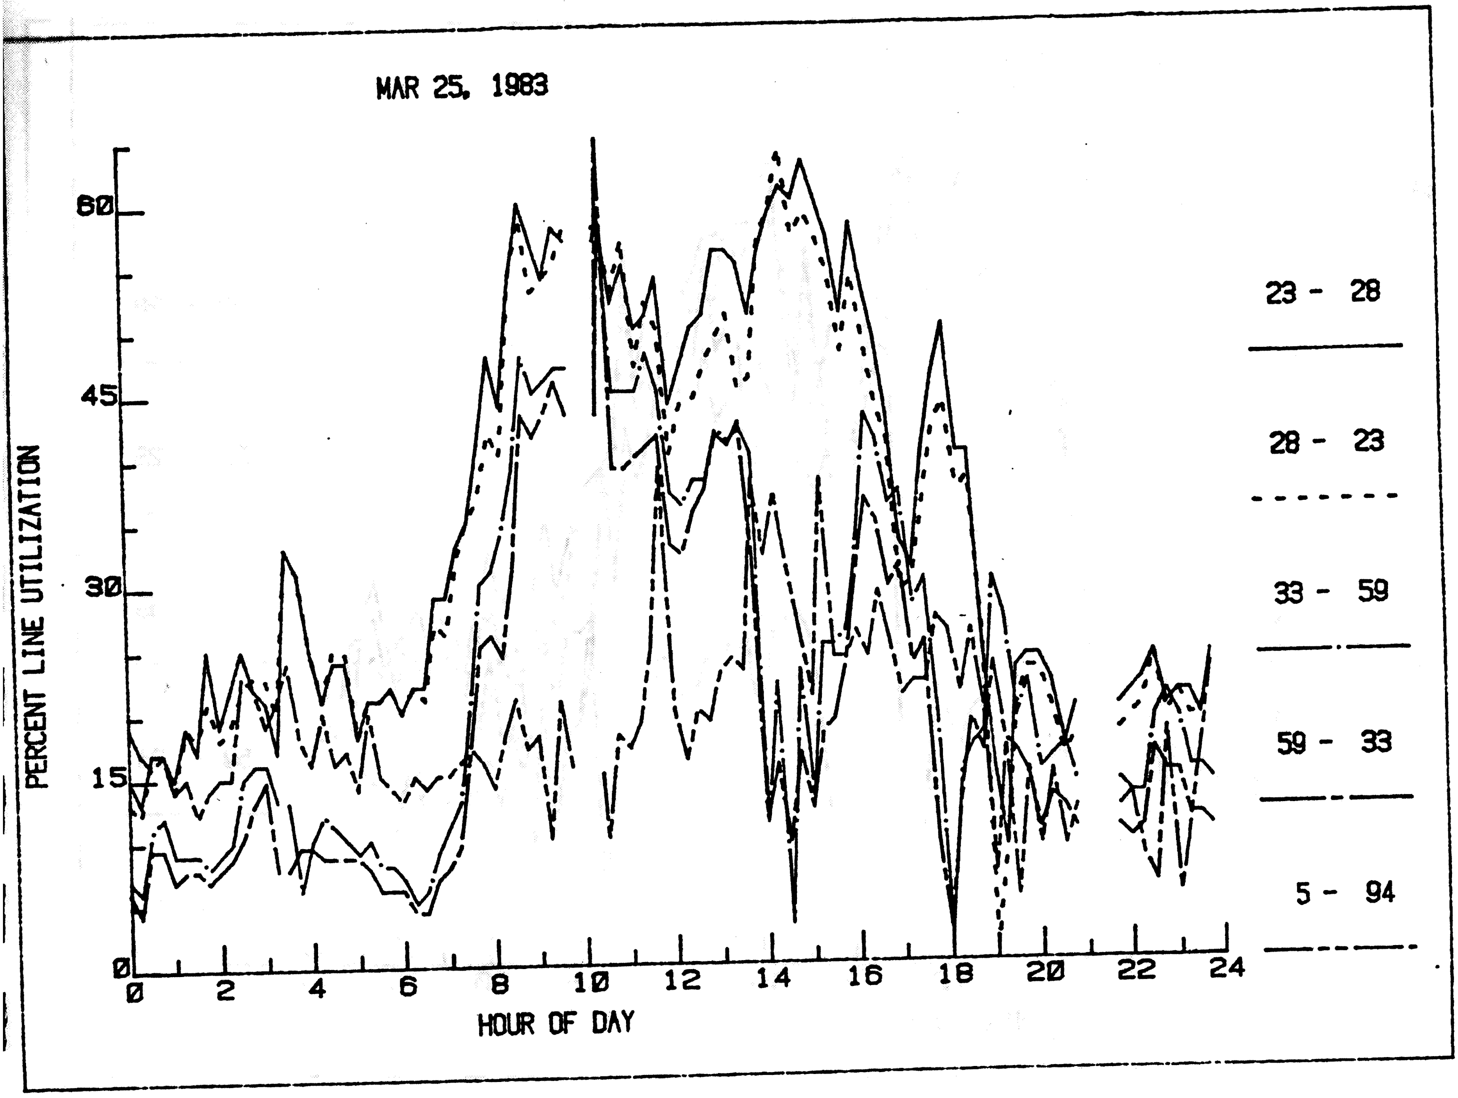

Cross-country trunk utilization on 25 March 1983.
Figure 32

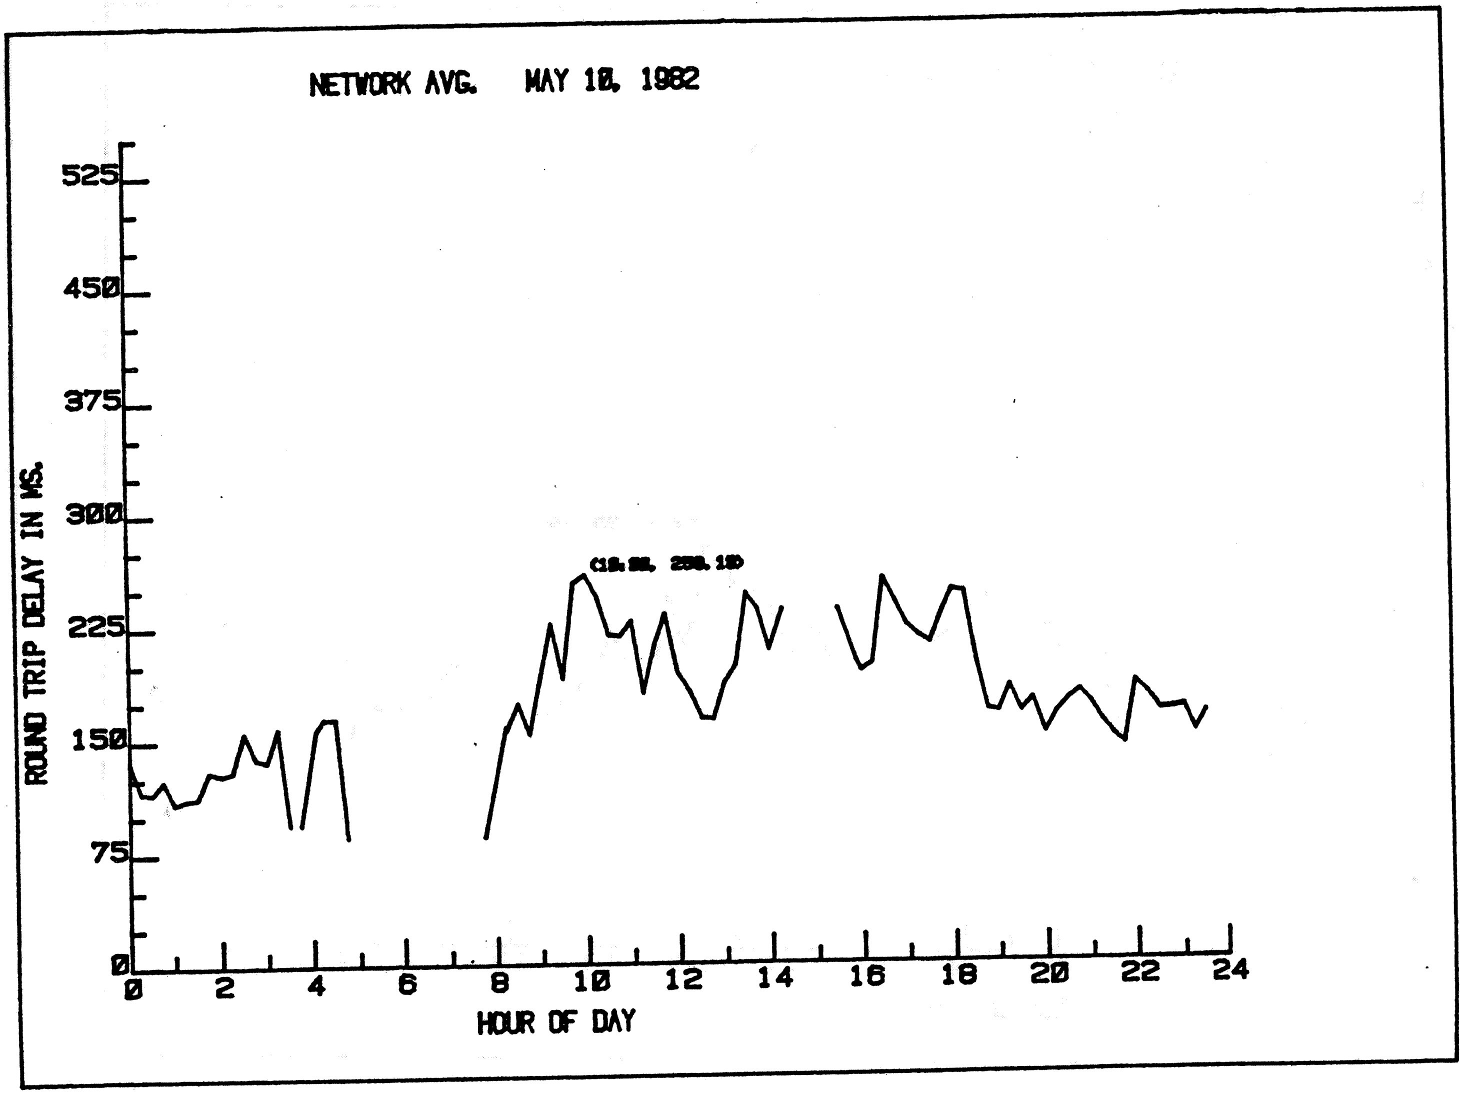

Average ARPANET round-trip delay on 10 May 1982.
Figure 33

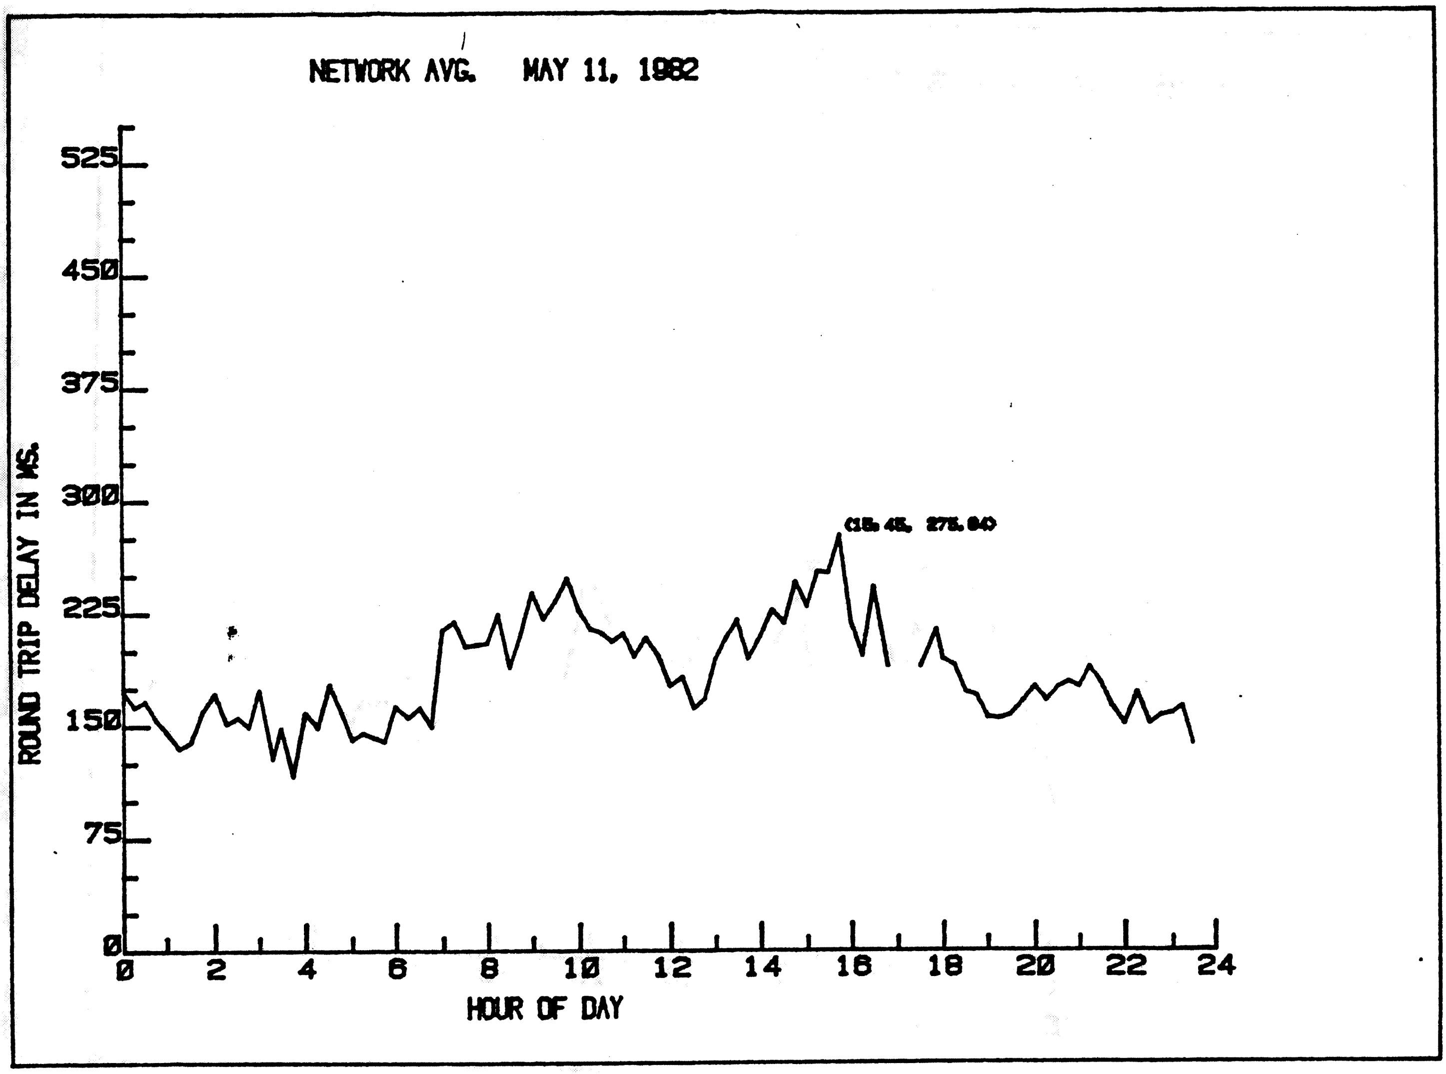

Average ARPANET round-trip delay on 11 May 1982.
Figure 34

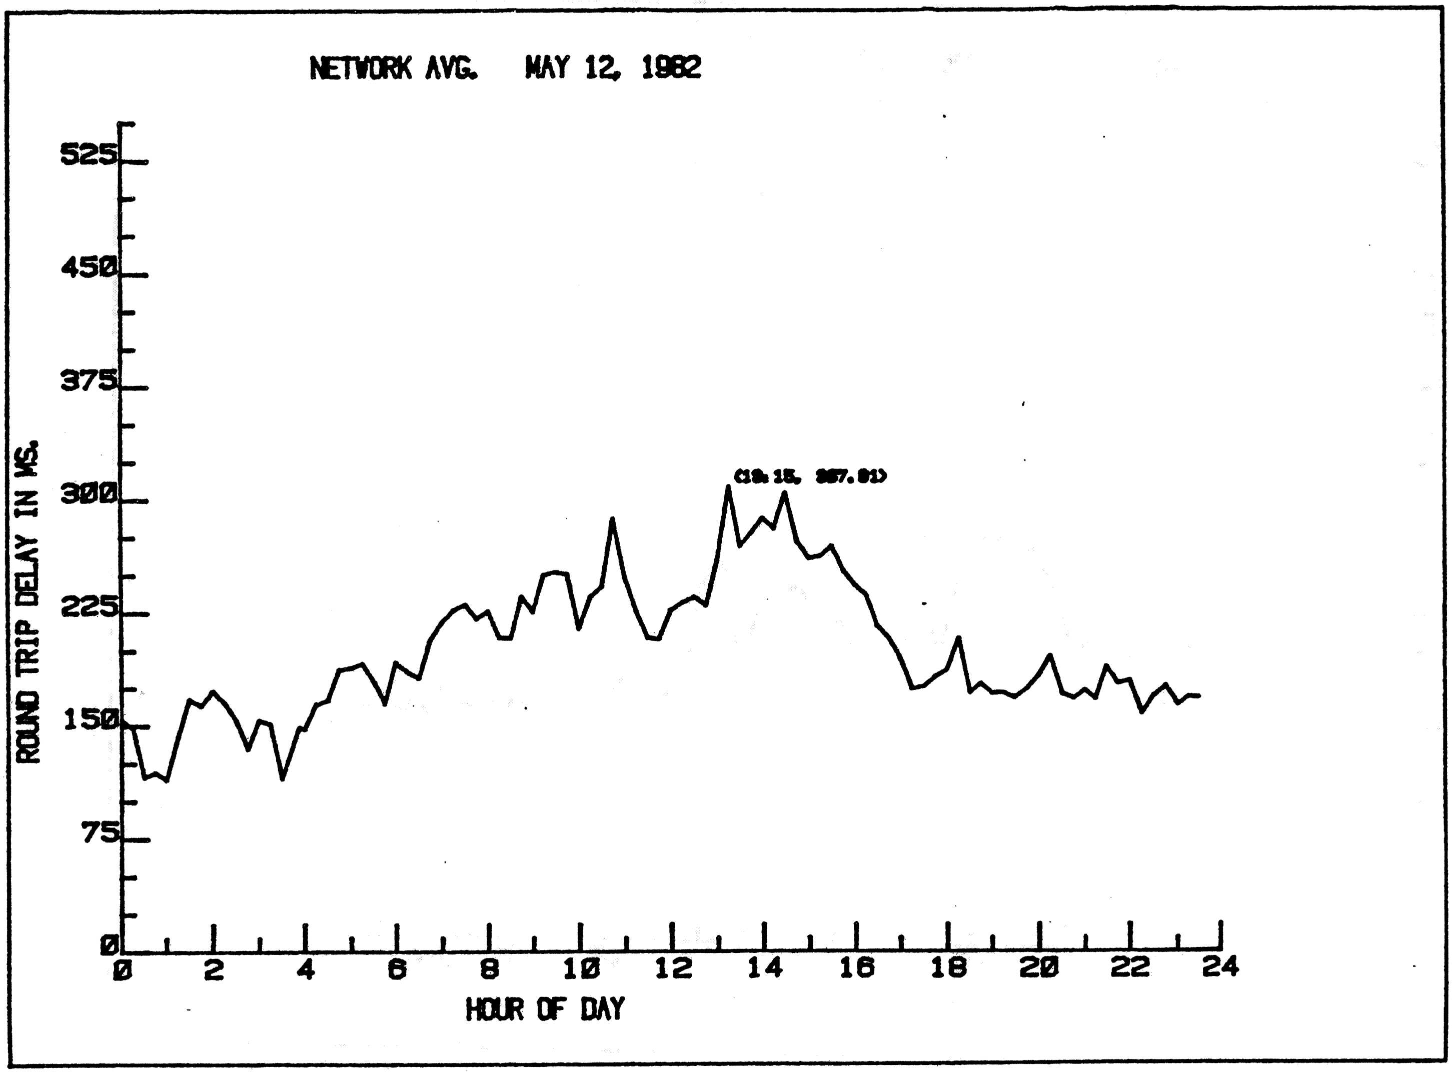

Average ARPANET round-trip delay on 12 May 1982.
Figure 35

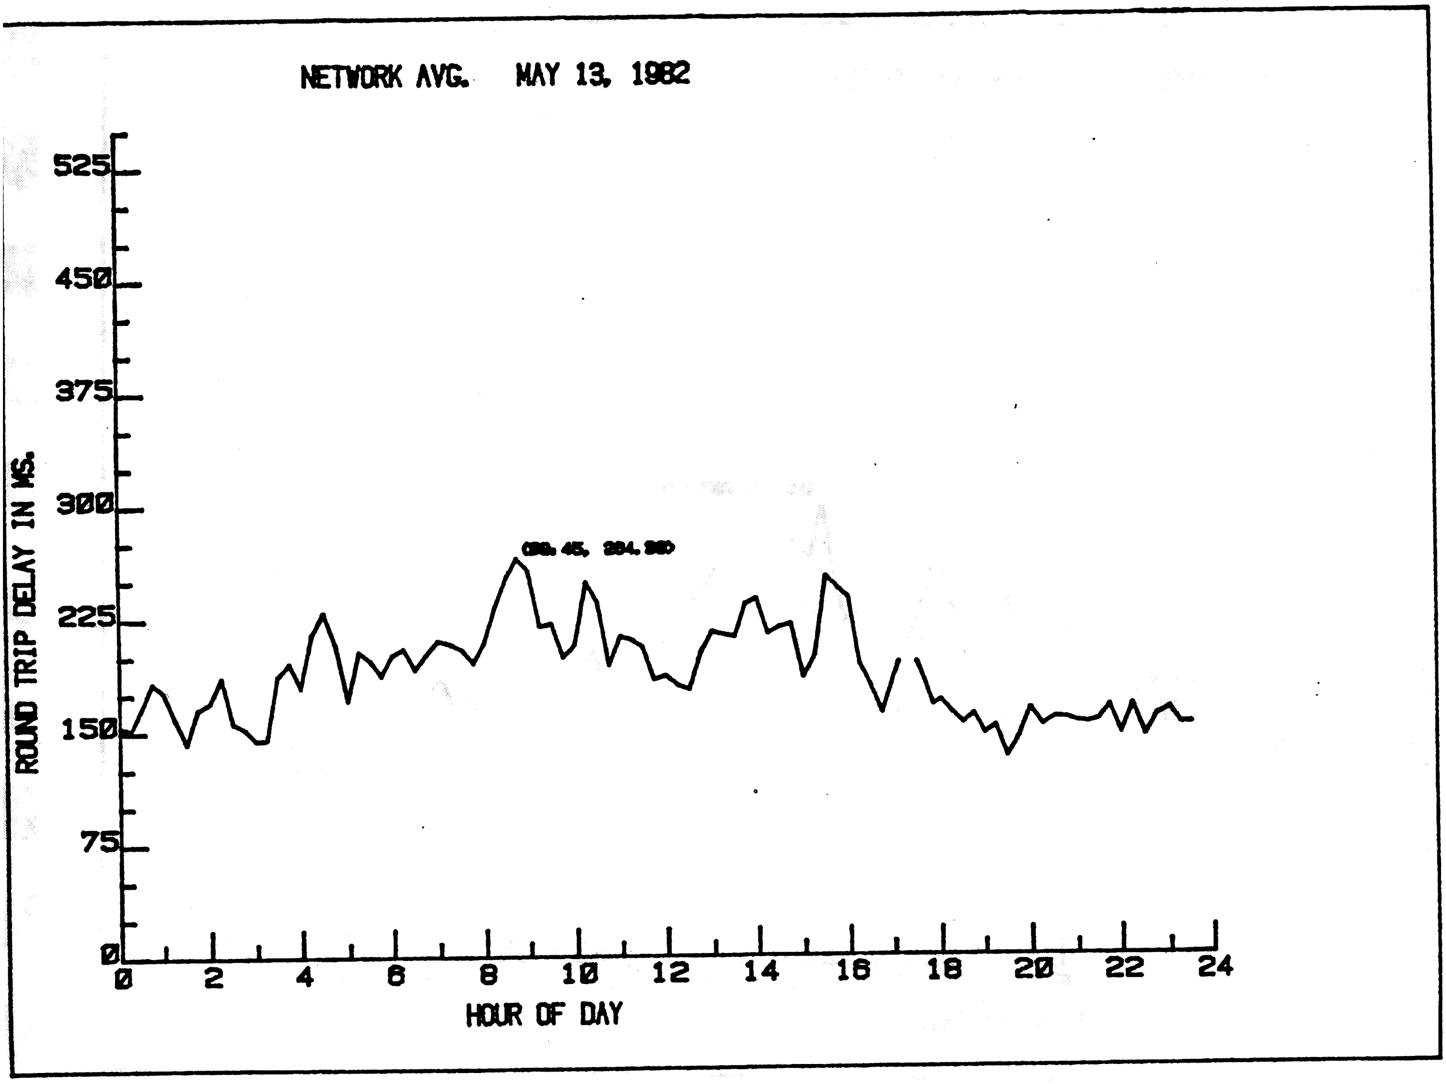

Average ARPANET round-trip delay on 13 May 1982.
Figure 36

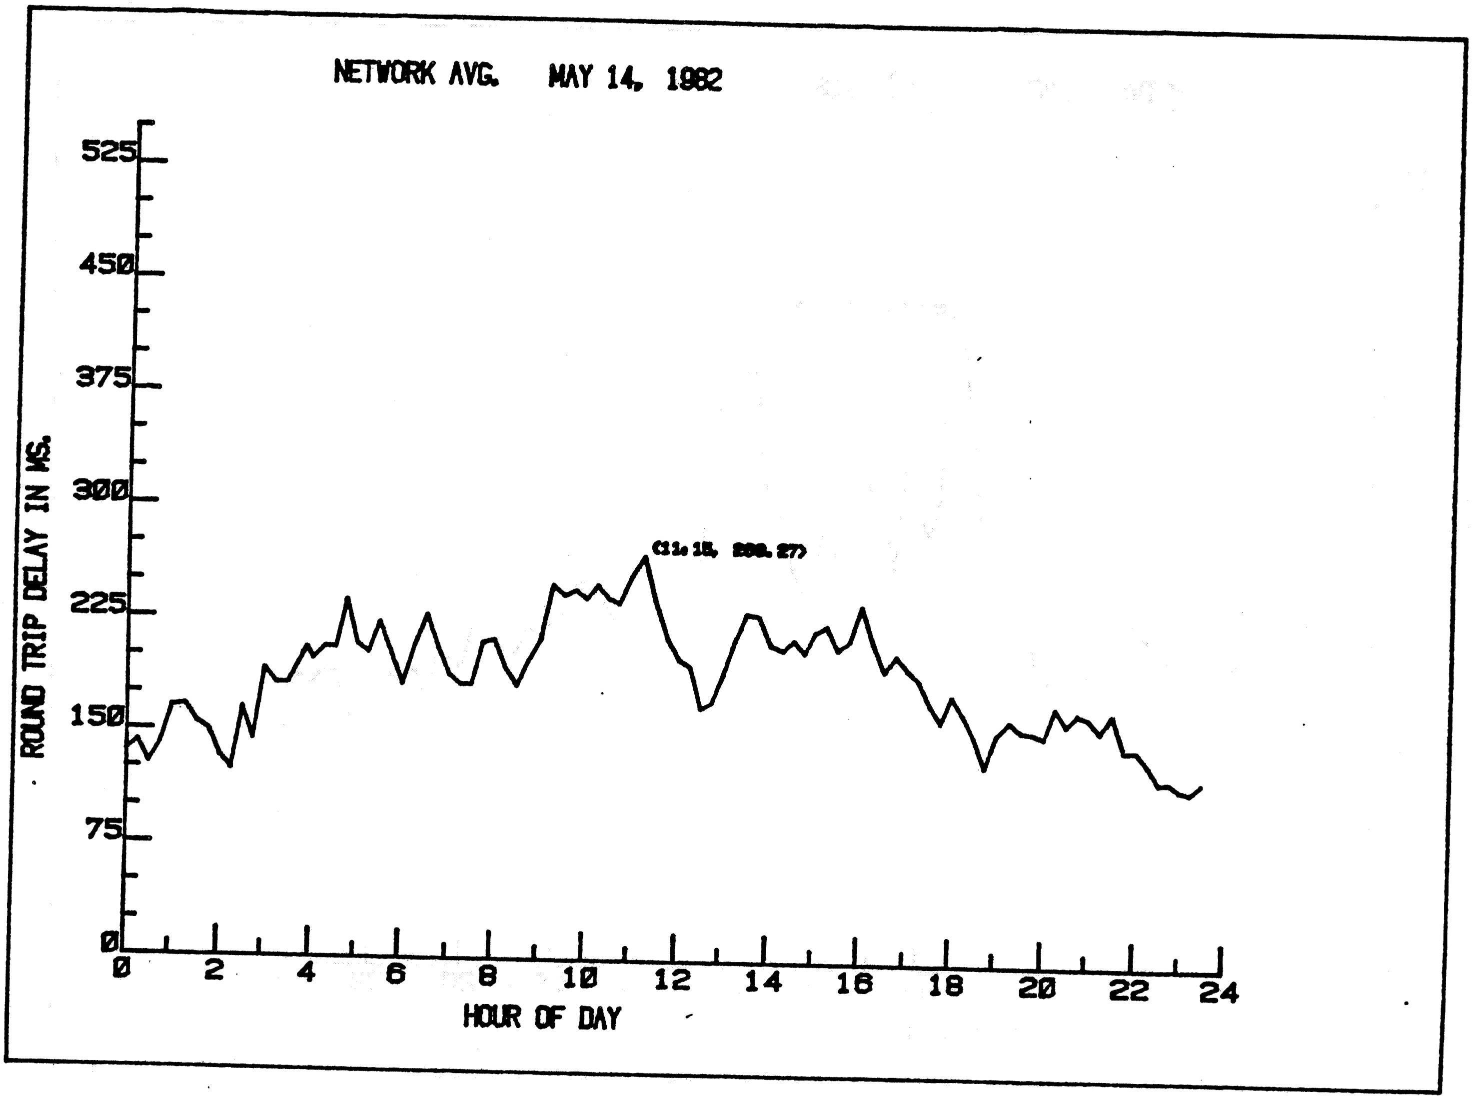

Average ARPANET round-trip delay on 14 May 1982.
Figure 37

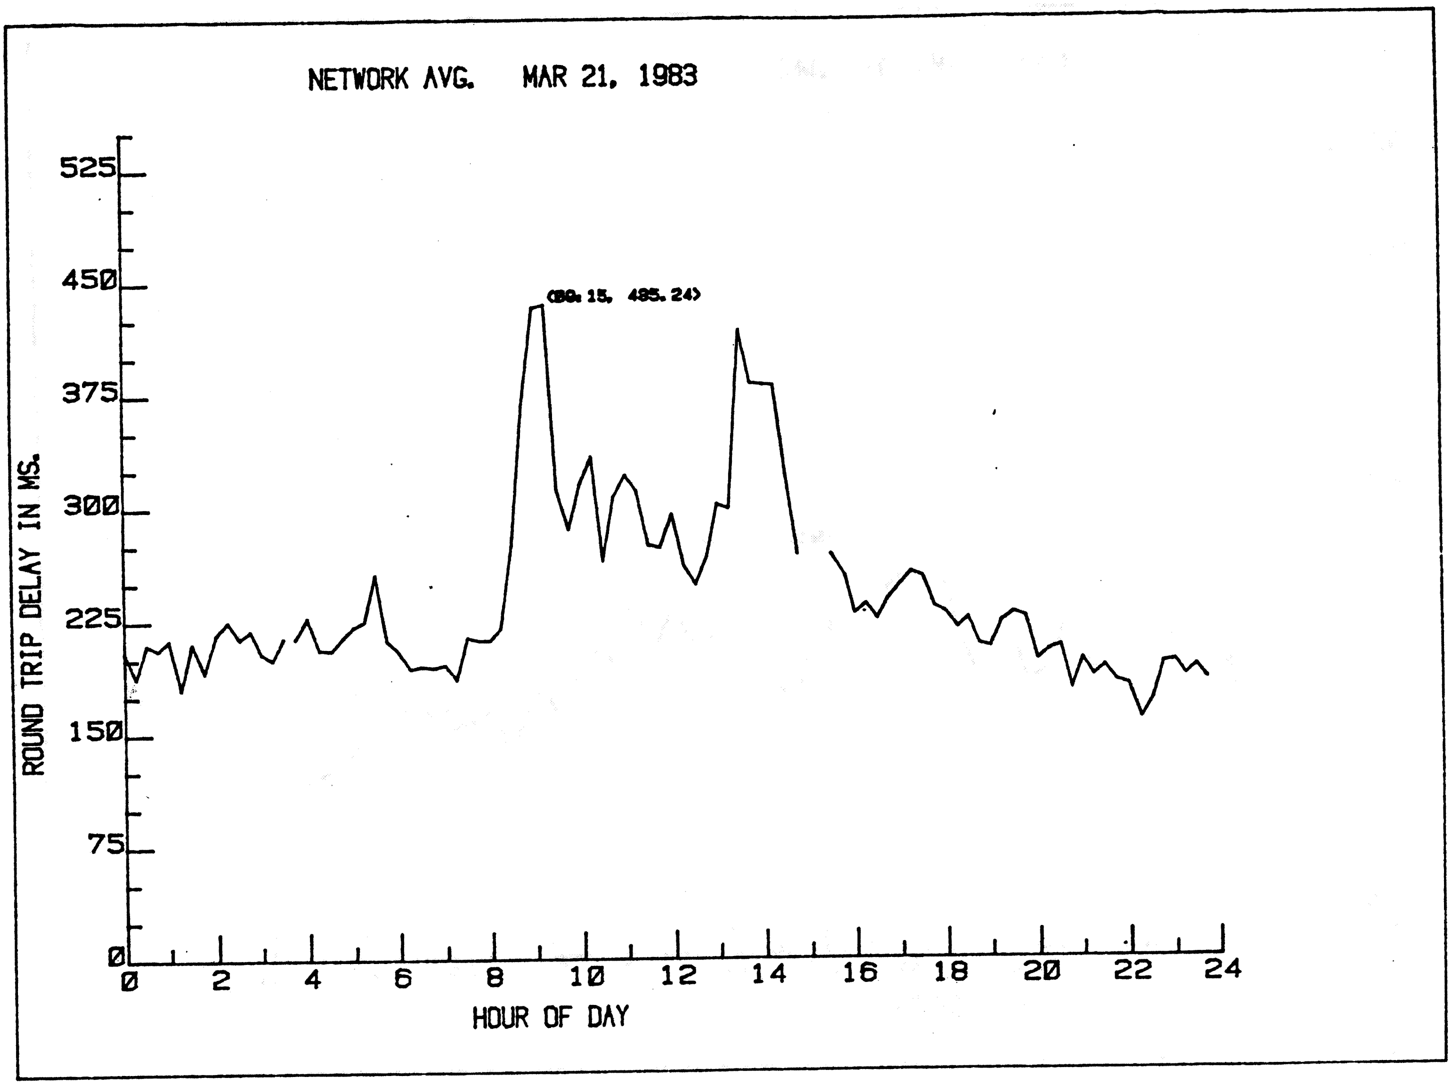

Average ARPANET round-trip delay on 21 March 1983.
Figure 38

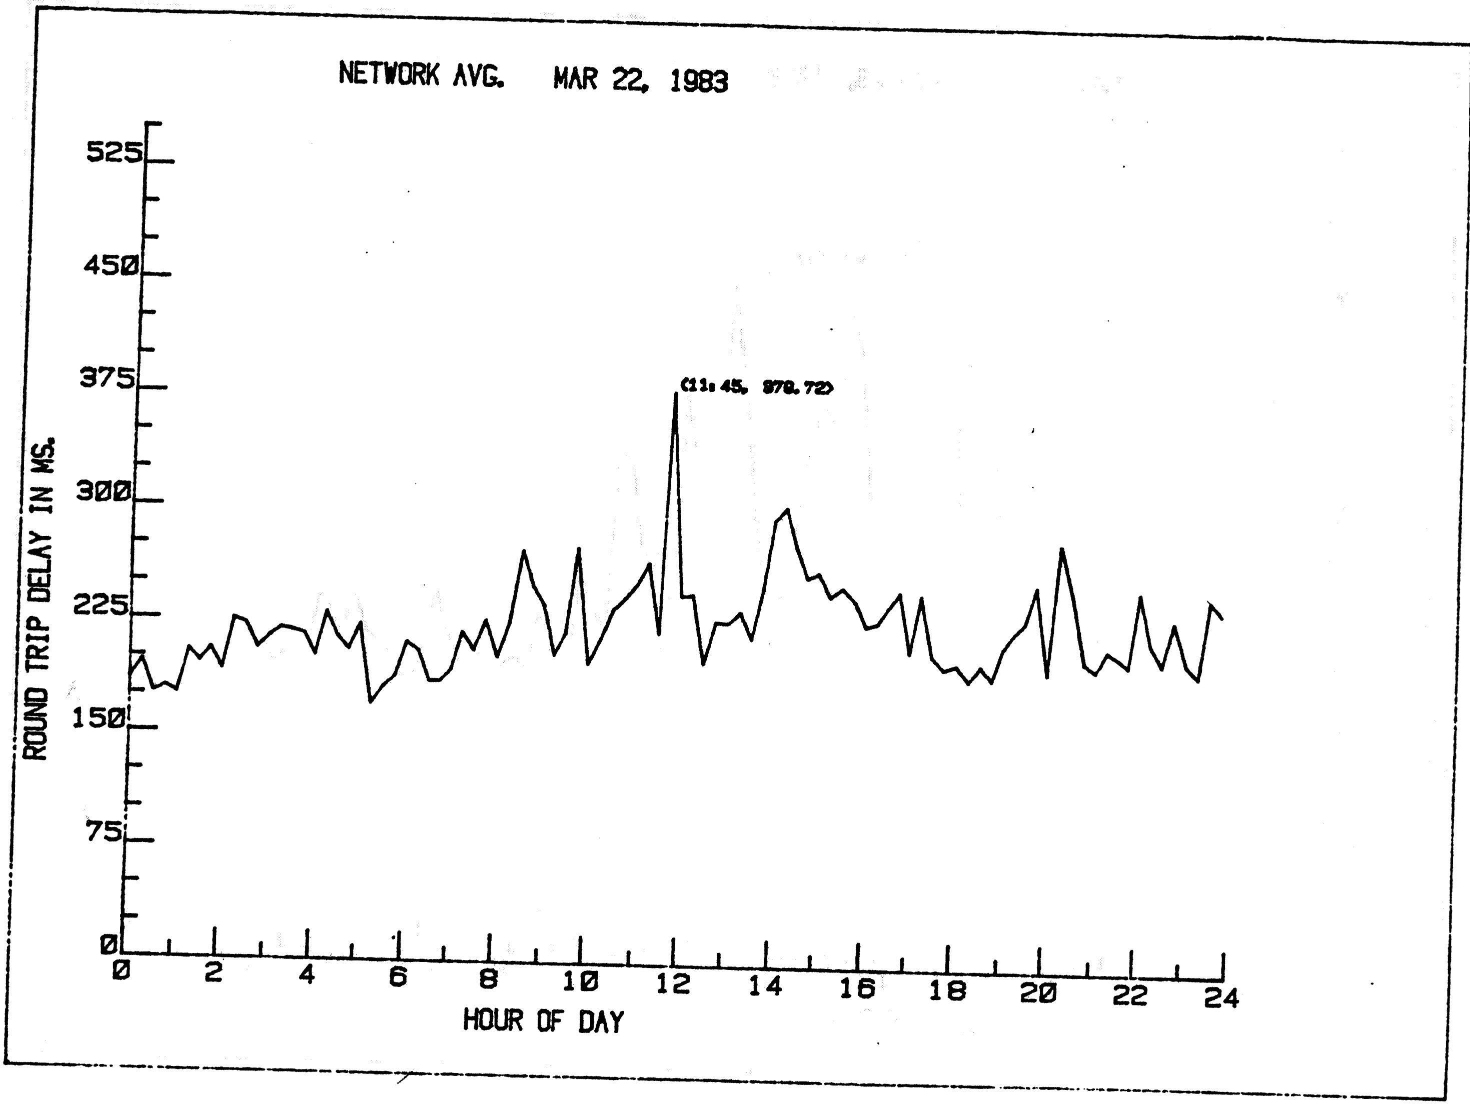

Average ARPANET round-trip delay on 22 March 1983.
Figure 39

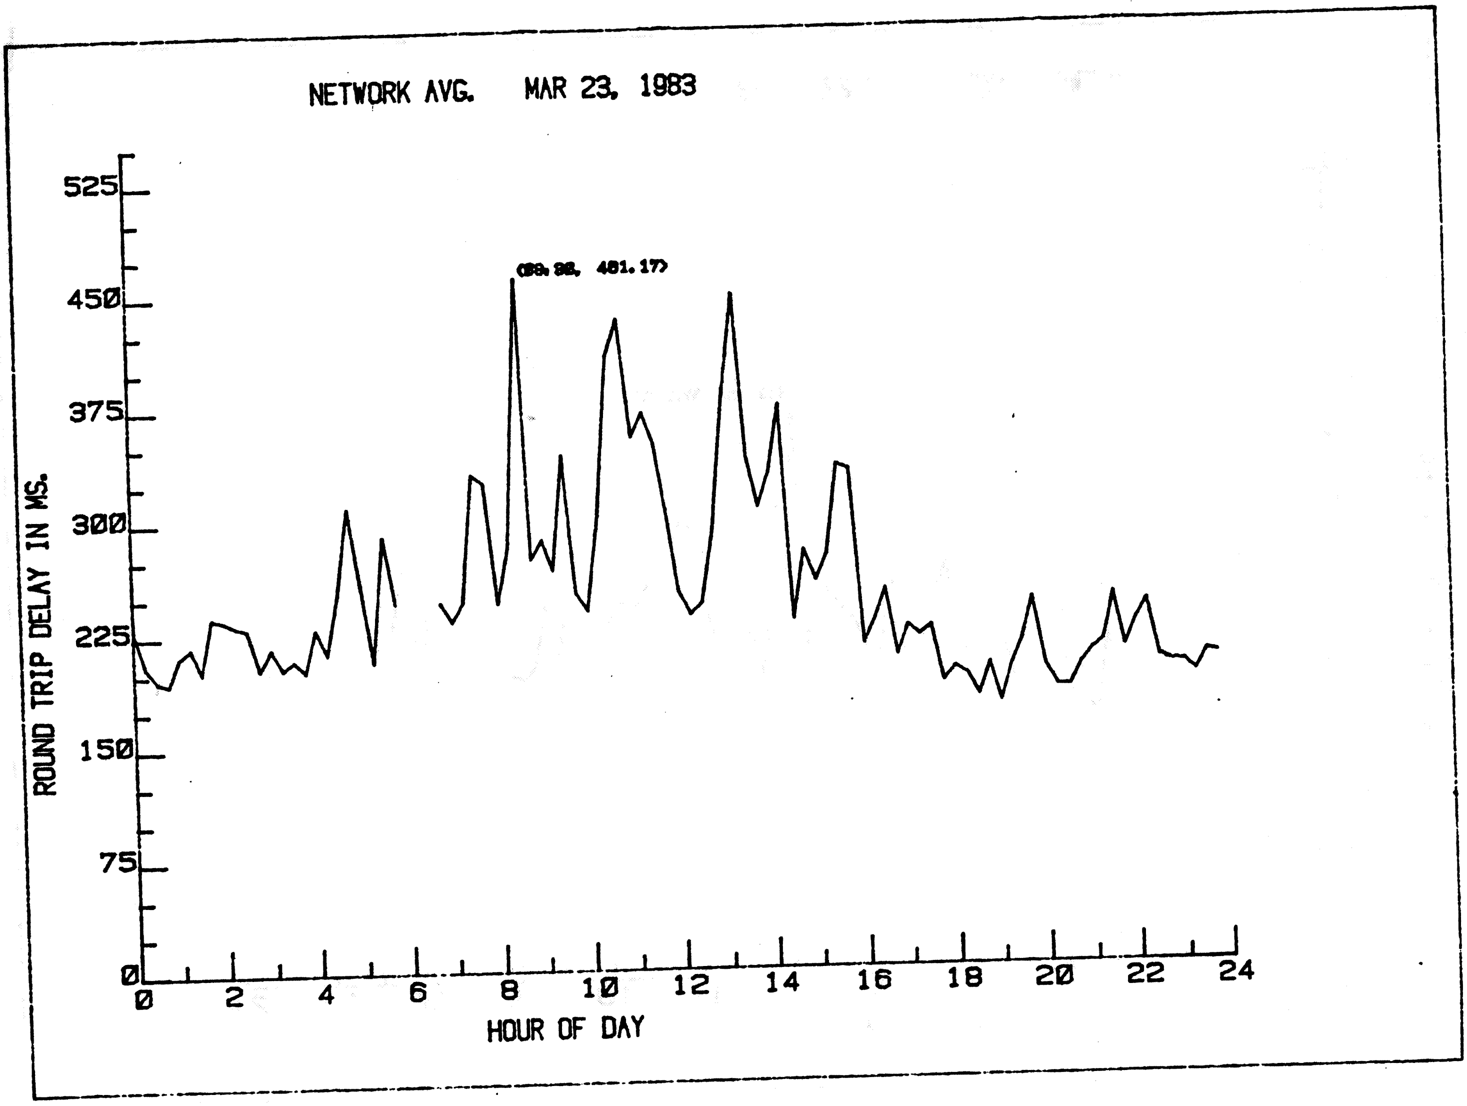

Average ARPANET round-trip delay on 23 March 1983.
Figure 40

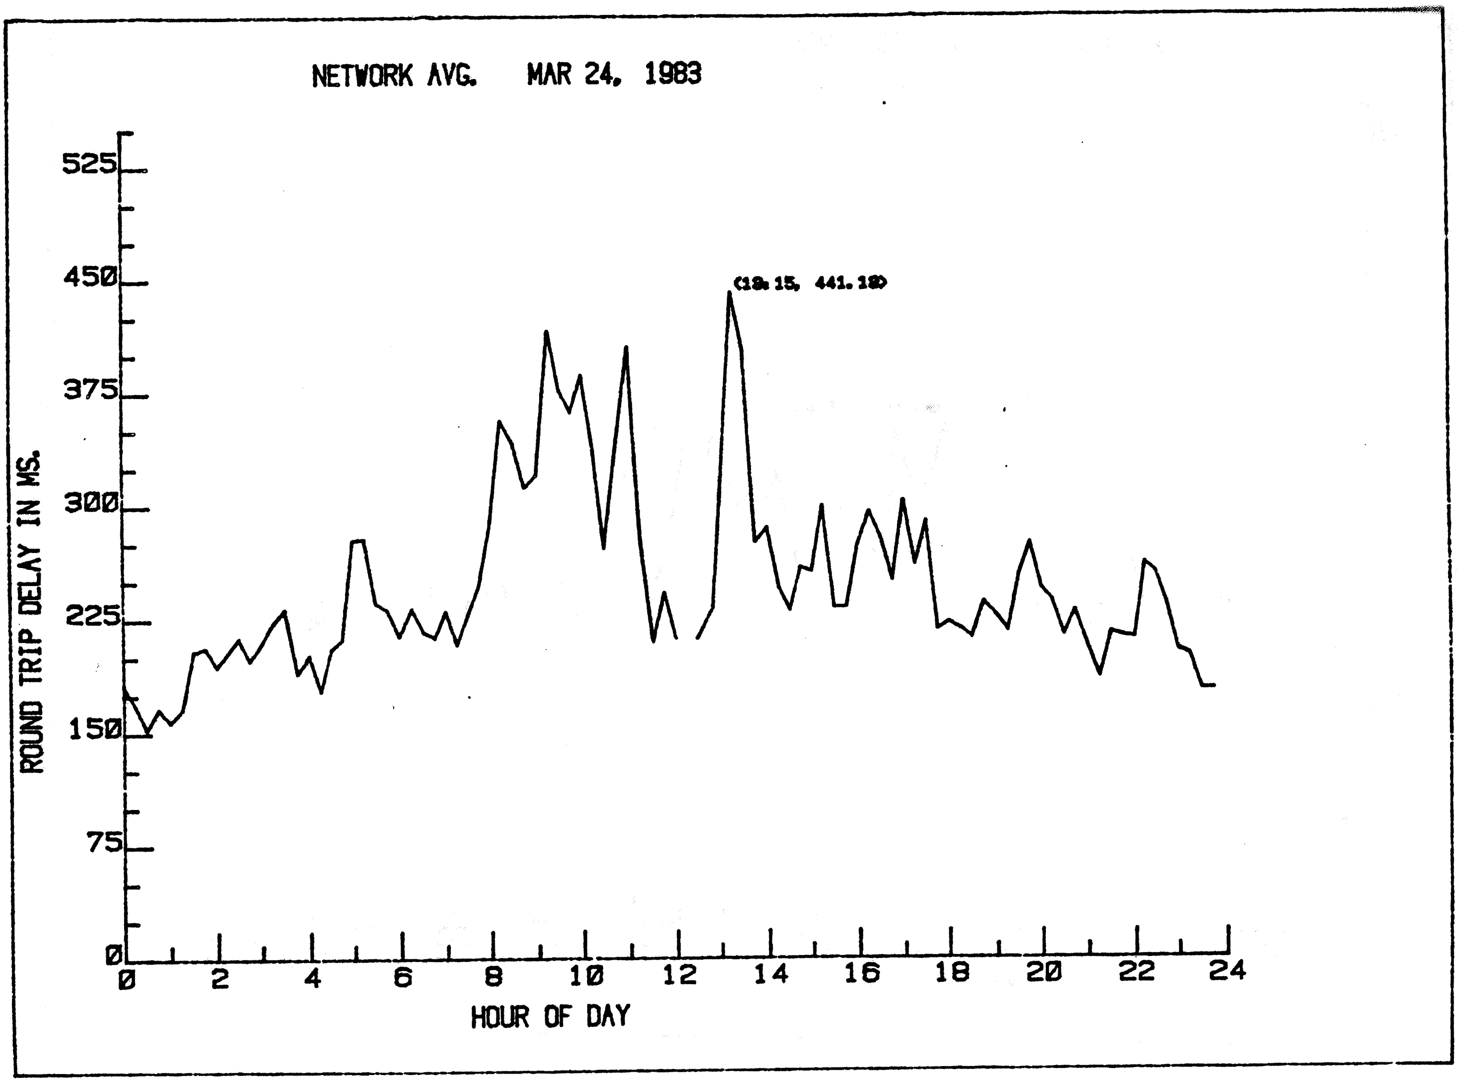

Average ARPANET round-trip delay on 24 March 1983.
Figure 41

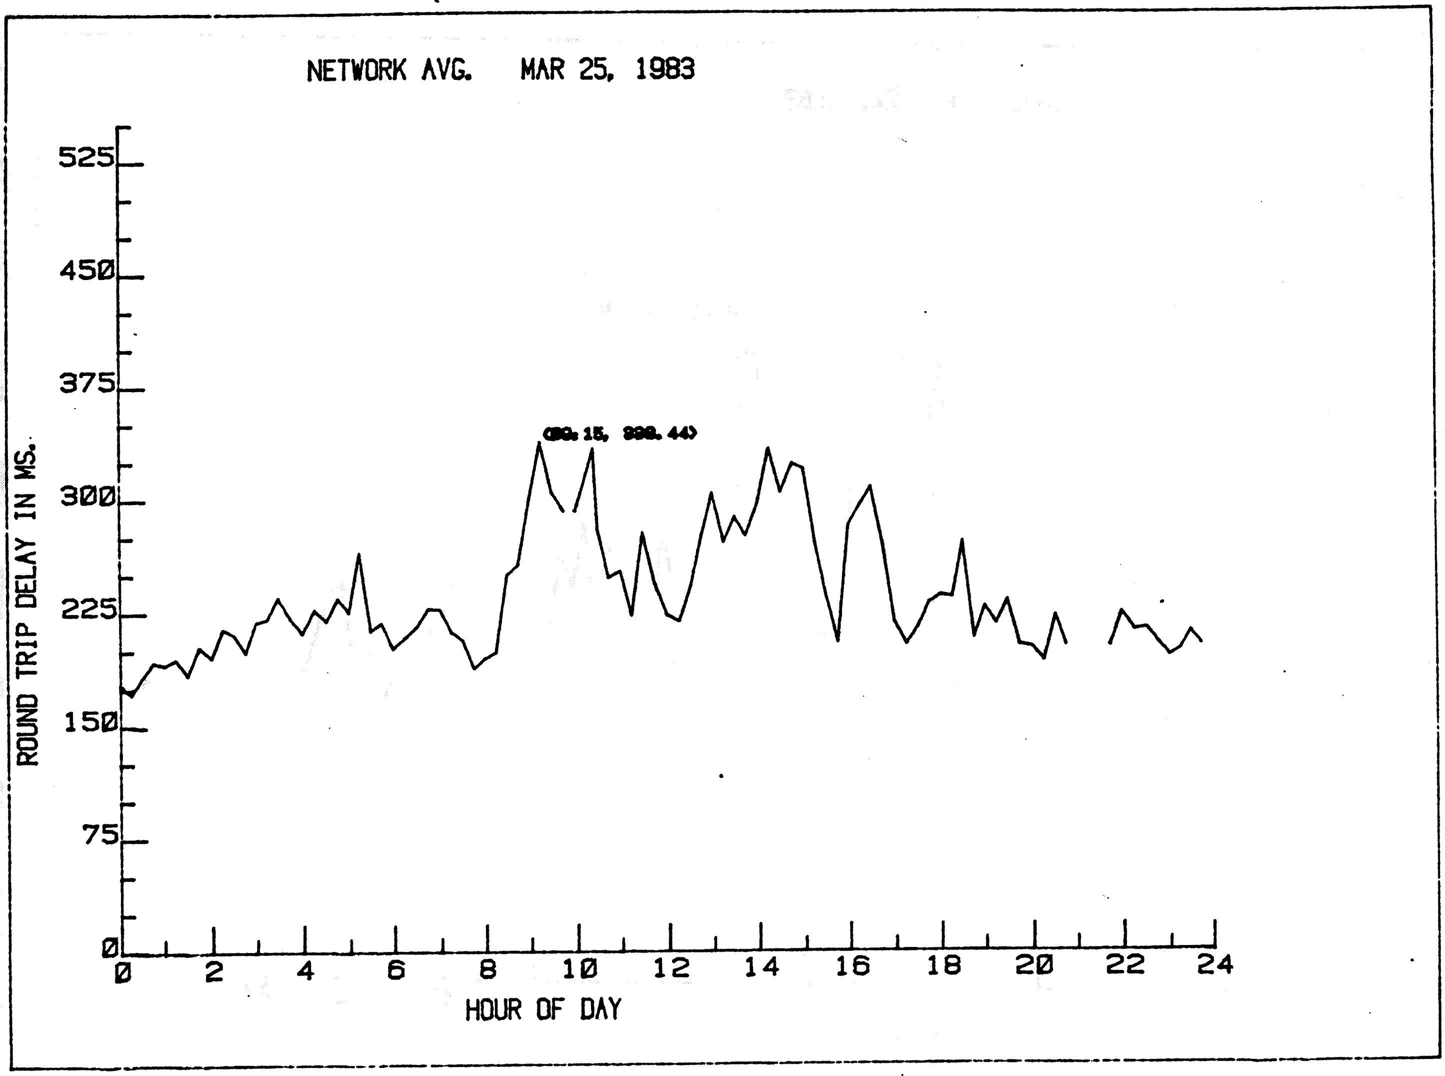

Average ARPANET round-trip delay on 25 March 1983.
Figure 42