

AD-A059 635

UNIVERSITY COLL LONDON (ENGLAND) DEPT OF STATISTICS --ETC F/6 17/2
MEASUREMENT OF PDP-9A USAGE JANUARY - JUNE 1977.(U)

DEC 77 A V STOKES


N00014-77-6-0005

UNCLASSIFIED


TR-42

NL

1 OF 1
AD
A059635



END
DATE
FILMED
12-78
DDC



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

AD A059635

DDC FILE COPY

LEVEL II

(D)

DDC
OCT 4 1978
RESOLVABLE
F

This document has been approved
for public release and sale; its
distribution is unlimited.

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM	
1. REPORT NUMBER TR-42	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER	
4. TITLE (and Subtitle) Measurement of PDP-9A Usage January - June 1977		5. TYPE OF REPORT & PERIOD COVERED Technical, 1977	
7. AUTHOR(s) Dr. A.V. Stokes		6. PERFORMING ORG. REPORT NUMBER	
(10) Adrian V. Stokes		8. CONTRACT OR GRANT NUMBER(s) (5) N-00014-77-G-005 0005	
9. PERFORMING ORGANIZATION NAME AND ADDRESS University College London Gower Street, London WC1, ENGLAND		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS	
11. CONTROLLING OFFICE NAME AND ADDRESS Defence Advanced Research Projects Agency 1400 Wilson Boulevard, Arlington, Virginia 22209, USA		12. REPORT DATE 31 December, 1977	
13. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Office of Naval Research		14. SECURITY CLASS. (of this report) UNCLASSIFIED	
(12) 24 p.		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE	
16. DISTRIBUTION STATEMENT (of this Report) No. 1 This document has been approved for public release; its distribution is unlimited. (9) Technical rept.			
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)			
18. SUPPLEMENTARY NOTES			
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Measurements of Network Usage, Packet Switched Networks, ARPANET			
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This report, which may be considered an addendum to TR-30, examines the pattern of usage of the INDRA PDP-9A over the first six months of 1977.			

LEVEL

ABSTRACT

This Report, which may be considered an addendum to
TR-30, examines the pattern of usage of the INDRA PDP-9A
over the first six months of 1977.

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

78 09 11 051

4059632

INDEX

1. Introduction	3
2. System Overview and Previous Measurements	4
3. Measurements Made and Comparison with Previous Data	6
3.1 Terminal Usage	9
3.2 File Transfer Usage	14
4. Conclusions	18
Acronyms	20
References	21
Acknowledgements	21

APP. 25	on	<input checked="" type="checkbox"/>
POS	25 Nov	<input type="checkbox"/>
UNCLASS		<input type="checkbox"/>
JCS.1 HQ		
BY		
DISTRIBUTION/AVAILABILITY CODES		
Dist.	AVAIL.	and/or SPECIAL
A		

INDEX TO FIGURES

1. Current INDRA Configuration	5
2. Monitoring Performed by Month	7
3. Monitoring Performed by Time of Day	8
4. PDP-9A Port Usage by Time of Day	10
5. RL 360/195 Usage by Time of Day	11
6. Usage of RL 360/195 by Various Research Groups	12
7. Usage of RL 360/195 from Various ARPANET Hosts	13
8. File Transfers Performed by Month	15
9. File Transfers Performed by ARPANET Host	16
10. File Transfers Performed by Time of Day	17

1. Introduction

This report may be considered to be an addendum to INDRA TR-30 (Ref. 1). The previous report evaluated measurements made of the London node of ARPANET over periods up to the end of 1976 (the periods varied according to the precise measurements that were being made). This report extends the data for use of the PDP-9A for six months, that is, up to the end of June 1977 and compares and contrasts the usage of that machine (and, more importantly, the Rutherford Laboratory (RL) 360/195 that it front-ends) with the previous period considered (January to December 1976). In order to make this report self-contained as far as possible, some information contained in the previous report is repeated herein, but very few of the data are given in any detail and, for such detail, reference is made to the previous report.

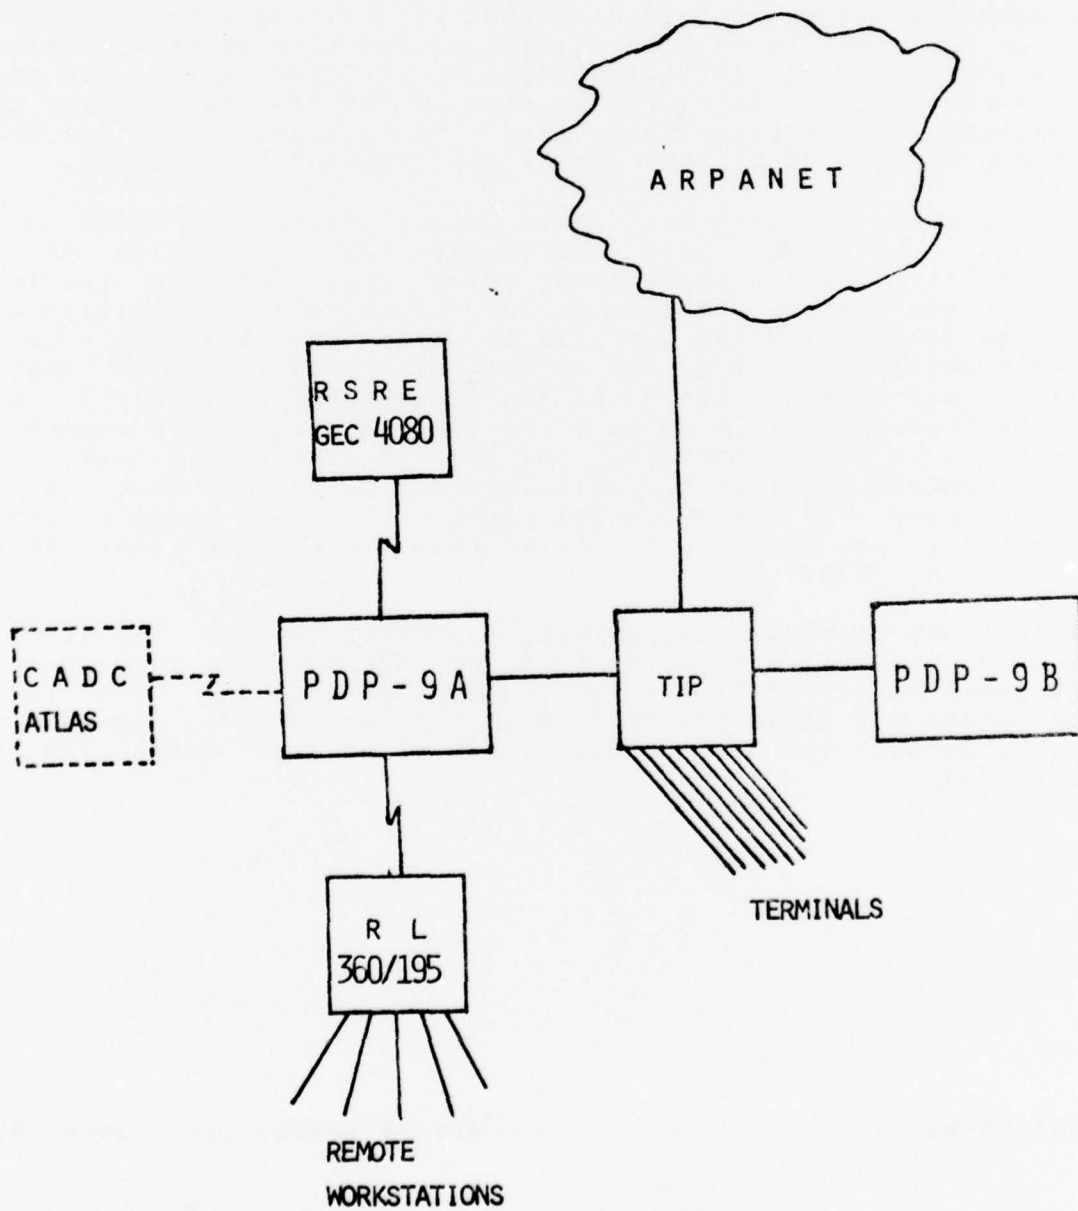
2. System Overview and Previous Measurements

In this Report, we consider measurements made of the usage of the INDRA PDP-9 (designated PDP-9A) which front-ends the Rutherford Laboratory 360/195. The configuration in which we are interested is shown in Fig. 1.

Usage of the system may take many forms. First, there is single usage of the PDP-9A itself. Very little can be done on this machine except reading "Help" files and so this usage follows a similar pattern to that of the RL 360, except where users have been able to connect to the PDP-9 but unable to log in to the 360 for any reason. The majority of usage of the PDP-9 is merely to access the dual-processor 195 at the Rutherford Laboratory. There are (effectively) three logical ports from the PDP-9A to the 360 and we measure usage of the 195 through these ports, including port usage data. A second form of usage is that of File Transfers both to and from the RL machine and we show some measurements on these. Finally, we have usage of ARPANET from the 360 (which uses a separate logical port on which, in fact, a number of separate interactions may be multiplexed).

In all cases, we describe usage of the various systems in a similar way to that followed in Ref. 1 in order to facilitate comparison. Similarly, as far as possible, all graphs are drawn to the same scale as in that document.

FIGURE 1: CURRENT INDRA CONFIGURATION



3. Measurements Made and Comparison with Previous Data

The measurements made are exactly the same as those made in the previous year since, at this level, there were no changes to the program in the PDP-9. However, one small change to the system on the RL machine considerably helped. This modification produced a date/time stamp just before the system was taken down; as this happened fairly often, this provided an excellent cross-check on the date/time stamp of the PDP-9. This was in error on a sufficient number of occasions to make it worthwhile reiterating the proposal made in Ref. 1 that the PDP-9 should check the time given to the system via an independent source.

The measurements taken were distorted significantly by the move of the RL machine giving a very reduced time on that machine, particularly during February (31%). In addition, due to the need for extra system development time on PDP-9A, this machine was unavailable for an appreciable time. The preferred time for this work was 1300 to 1500, although it was extended on occasion.

Figure 2 shows the amount of monitoring performed by month (and includes, on the same histogram an indication of RL availability). While appreciably lower than for the previous year for the reasons outlined above, the overall availability was 73%. As in the previous year, it is clear that this figure is an underestimate due to the loss of data either by the PDP-9 running out of paper tape (on which the log is punched) or physical loss of the tapes and it is quite clear that this method of recording statistics is unsatisfactory. The MTBF of the PDP-9 was just over fourteen hours (after allowing for an incident which caused it to crash and auto-restart over 200 times before being corrected), an appreciable improvement over the nine hours recorded for 1976.

The analysis of monitoring performed, broken down by time of day (Fig. 3) is very similar to that for 1976 with the very obvious difference in the early afternoon (the minimum for 1976 being just under 60% while for 1977 it is only about 30%). It is also noticeable that the maximum figure has risen from about 72% to nearly 80%.

FIGURE 2: MONITORING PERFORMED BY MONTH





FIGURE 3: MONITORING PERFORMED BY TIME OF DAY



3.1 Terminal Usage

It is obvious that the non-availability of FDP-9A will affect port usage in a similar manner and this is borne out by Fig. 4 which shows this usage. Making qualitative allowance for the depression of usage over the 1300 to 1700 time slot in accordance with Figure 3, the histograms for 1976 and 1977 bear a remarkable similarity and it is clear that the pattern of usage has stabilized with the exception of weekend usage which has decreased significantly. As would be expected, very similar comments may be made about RI port usage, shown in Fig. 5.

All the five FDP-9 ports were in use for only 77 minutes on a total of 63 occasions, slightly less than in 1976 but with a similar average time per occasion. However, the number of times when the ports to the 360 were in use (302 times, 952 minutes), while averaging a similar time per occasion as for 1976, is a noticeably higher figure (cf. 422 times, 1152 minutes for the whole year). A user logged in to the FDP-9 but was unable to log in to the 360 51 times - a much lower figure than previously. These data imply that users realized that, if all ports were in use, there was little point in continuing to try (since, on average, the condition persisted for about three minutes) but came back later.

Figure 6 shows the usage of the RI machine broken down by group. Comparison with the 1976 figures shows that the INERA group usage is slightly up, as is the seismic usage. The high energy physics usage is down to a third of that previously recorded as is the use of "quest" accounts. The latter is explicable by the fact that no general quest accounts were provided in 1977 and the "quest" accounts were provided for specific low-usage accounts. "Other" is a large number of accounts using very little time each; a total of 96 accounts were used in the period, exactly the same as for the whole of 1976.

The usage is generally reflected by Fig. 7 showing the source of the terminal interactions. Two points of note are that the usage from the London-TIP is virtually constant, occupying nearly three-quarters of the total time, and the (new) usage from SRI making it the second largest user (but with only 7% of the total).

FIGURE 4: PDP-9A PORT USAGE BY TIME OF DAY




FIGURE 5: RL 360/195 USAGE BY TIME OF DAY




FIGURE 6: USAGE OF RL 360/195 BY VARIOUS RESEARCH GROUPS





FIGURE 7: USAGE OF RL 360/195 FROM VARIOUS ARPANET HOSTS



3.2 File Transfer Usage

So far we have only considered terminal access. This section deals with file transfers and, in what follows, the figures quoted are normalized as the number of events per 10,000 minutes monitored. Making allowance for the down-time of the 360, the histogram (Fig. 8) is very similar to that for 1976 with the inexplicable difference of an anomalously low figure for June.

The host usage (Fig. 9) shows some noticeable differences. In particular, the previous largest user, Stanford Research Institute, has decreased to an insignificant time. This may well be explained by the fact that it was used quite heavily in 1976 for transferring files such as the PDP-9 "Help" files so that they could be edited easily using the Tenex editor, but this usage effectively ceased during 1976. However, the increase in usage of BBN (8.7% to 28.5%) may be related in that, to some extent, usage of EBN may have replaced SRI. The only other noticeable change is the reduction of the MIT-AI usage from 6.9% to virtually zero. However, in general, the production use of the file transfer (particularly from ILLIAC-IV and Lawrence Berkeley Laboratories) seems to have become more consistent.

Although not shown by the statistics gathered, it is clear from examining the records of file transfers that the reliability of the service had noticeable improved; this was due partly to the increased familiarity of the users with the system and partly to the file transfer software.

Figure 10 shows the distribution of FTFs by time of day. Again (cf. interactive use), we note a similar pattern to that seen in the previous year with the exception of the early afternoon due to the increased PDP-9 system development time.

FIGURE 8: FILE TRANSFERS PERFORMED BY MONTH




FIGURE 9: FILE TRANSFERS PERFORMED BY ARPANET HOST





FIGURE 10: FILE TRANSFERS PERFORMED BY TIME OF DAY



4. Conclusions

Perhaps the most significant datum is the considerable increase in the reliability of PDP-9A, making the MTEF fourteen hours, although it is still possible (and it happened once) that it may crash and auto-restart a very large number of times over night. This is usually caused by hardware errors (in the case mentioned, it was an interrupt from the disk) and there is probably little that can be done.

The pattern of usage seems to have stabilized noticeably, but it is clear that the amount of usage of the RI machine through the UCL PDP-9 is not particularly high from outside the INDRA group (about 25% but this figure is, in fact, an overestimate as usage is made from a terminal in the PDP-9 room which is not monitored and, from looking at the RI statistics, usage of this terminal is of the same order of magnitude as the other three ports combined; taking this figure, the outside usage would be about 14% of the total).

The file transfer protocol seems to be used as much as last year, but, on examining the figures in detail, it seems to be more reliable and it certainly has been used on a fairly regular basis by a number of groups.

The final conclusions are on the actual recording of data. For many reasons, the method used is unsatisfactory. First, the fact of the operator specifying the time and date frequently leads to these being given incorrectly and, although it is not single, it would seem worthwhile to check that these are correct. This could be done by sending a suitable message to the RI 360 (H+M TIME) and analyzing the reply. Alternatively, there are a number of machines on ARPANET (PDP-10s) which have a particular socket reserved for the date and time. A connection to such a socket prints out a single line and then is closed; however, this has the disadvantage of conversion from the relevant time standard.

A second problem which was not envisaged when the logging program was written was the fact that the Host addresses of a number of machines have changed, either with machines leaving ARPANET and being replaced by new sites or by a simple change of number (for example, SRI changed from 66 to 51). Since the datum recorded by the PDP-9 is the Host number (since it is this which identifies the interaction), this introduces considerable complications in the analysis programs. The suggested modification is for the PDP-9 to have a look-up table of Host names/addresses which is kept up-to-date and for the machine to record the Host name. Luckily, in the data recorded, the problems could be overcome relatively easily since we were considering a short timescale.

The final objection to the method is the use of paper tape. While this is convenient in that there are no problems of corrupt files or of closing a file when the machine crashes, it has many disadvantages, noted above. Therefore, it seems a worthwhile improvement to utilise the fact that the PDP-9 is directly connected to the machine on which the analysis takes place. There are two areas on the system drum (known as SDR and IDR) which are not used for any other purpose. It would seem reasonable to use one of these (e.g. the IDR which, on PDP-9A, has 242 (octal) free blocks) for measurement purposes. When the machine is started, it could write the appropriate job header to the file, then write the measurement data. When the area becomes, say, half full, the file could be sent to the 360 and a new file opened. When the job has run correctly, the old file could be deleted. This would take some development work; it would seem a worthwhile investment on any machine on which serious long term measurement were envisaged, but will not be undertaken on the UCI PDP-9.

ACECNYMS

ARPA	Advanced Research Projects Agency
BL	British Library
BIL	British Library, Lending Division
BLRD	British Library, Research and Development Dept.
bps	Bits per second
CADC	Computer Aided Design Centre (Cambridge)
HEP	High Energy Physics
IMP	Interface Message Processor
INDRA	InterNetwork Display and Remote Access
ISI	Information Sciences Institute
MEDLINE	MEDLARS On-Line
NAM	Network Access Machine
NBS	National Bureau of Standards
NLM	National Library of Medicine
NMM	Network Measurement Machine
PSTN	Public Switched Telephone Network
RI	Rutherford Laboratory
RSRE	Royal Signals and Radar Establishment
SIMP	Satellite Interface Message Processor
STEIN	Short-Term Experimental Information Network
TIP	Terminal Interface Message Processor
UCL	University College London
UCLA	University of California at Los Angeles

References

1. Stokes, A.V., "The INDRA Network Measurement Project", INDRA TR-30, Department of Statistics and Computer Science, University College London, June 1977.

Acknowledgements

We wish to acknowledge support by the British Library under grant SI/G/093. Also, such research would not be feasible without support of the many bodies for the AREA project at University College, in particular, the Science Research Council (B/RG/5981 and 67022), the US Advanced Research Projects Agency (NCCC14-74-C-0240 and NCCC14-77-E-C005) and the UK Ministry of Defence (AT/2047/064).