SLAC–PUB–5129
October 1989
(A)

IMPROVING RELIABILITY IN THE SLC CONTROL SYSTEM*

N. HEINEN, N. SPENCER, and J. TINSMAN

Stanford Linear Accelerator Center
Stanford University, Stanford, CA 94309

Abstract

During the past year, considerable emphasis has been placed on improving the overall reliability of the SLC control system. The Errorlog Facility has proven a useful tool to diagnose hardware and software problems. By analyzing the various error messages and their correlations, one can usually determine the software component or hardware module causing faults. Daily summaries help to identify problems so that they can be remedied before they become catastrophic; thereby bringing about a considerable increase in performance. We discuss the various tools we use and our operational experience with them.

Presented at the Accelerator Control: International Conference on Accelerators and Large Experimental Physics Control Systems, Vancouver, Canada, October 30–November 3, 1989.

⋆ Work supported by Department of Energy contract DE–AC03–76SF00515.

1. Introduction

The SLC Control system has been developed over the past several years and during that time has been used to commission and run various components and areas of the linear collider [1,2]. As the goal has changed from commissioning various components of the system to achieving a fully functional facility for high energy physics research, more and more emphasis has been placed on reliability and performance. One notable problem was that often component failures were found and fixed serially whereas it might have been possible to replace marginal components in parallel had they been recognized and identified. Another area that was thought to need improvement was operational procedures to guarantee the accuracy of calculations based on Beam Position Monitor (BPM) and magnet calibrations. This paper describes some of the steps taken by the SLC Controls Group to address these problems.

2. Errorlog Facility

The necessity to handle and report error conditions in a complex system such as the SLC was recognized, and support was provided in the original design of the SLC Control System. Since the host computer for the system is a VAX (currently an 8800) running VMS, the message facility supported by VMS was extended to include messages generated by various components of the SLC system [3]. All error or informational messages are logged chronologically to disk and can then be examined by the Errorlog Facility. All software components, including those that run in the microprocessors under the Intel iRMX system, are expected to report problems that they encounter when controlling or monitoring their associated hardware. Informational messages (such as the fact that a knob has been assigned to a group of hardware devices) are also logged so that one can correlate problems with certain procedures or events [3].

This software has proven valuable over the years in reconstructing events and finding bugs in both software and hardware. However, several problems surfaced that did not allow it to be as useful as was hoped. With time, the logs came to be quite voluminous and it became difficult to go through it by hand to pick out the problem areas. The reason for the large volume was threefold: additional informational messages were constantly being added to help in the diagnosis of problems, recurrent problems were not fixed, and erroneous or duplicate error messages were produced.

Also, in many instances, the error messages were not meaningful to those who had to use them to diagnose or fix problems. In addition to being logged, "nonhelpful" error messages were also being broadcast to the operator consoles. This caused operators to miss the important error messages and to mistrust the system in general. Therefore our goals were:

- Automate the procedure of looking through the log for problems.
- Reduce the "noise level." Find and fix the cause of duplicate or erroneous messages.
- Find and fix marginal hardware that causes error messages to be issued.
- Make sure the message terminology is understandable.
- Set the severity of error messages appropriately so that warning messages can be distinguished from those that are truly errors.
- Get appropriate operations and hardware/software personnel to use the summarized log routinely to find and fix problems.
- Once the broadcast message level is reasonable, add different sound effects to alert the operator to the fact that a problem of a certain nature has just occurred.

3. —Errorlog summary process

To achieve the goals outlined in the last section, the first step that was taken was to develop an automated process to read through the formatted Errorlog file and to summarize the information so that it would be possible to tell how many of what kinds of errors had occurred and which microprocessors or VAX processes had reported the errors. Additionally, teams of people were formed to examine the summarized information and to take appropriate action on solving some of the persistent problems that were identified.

In an attempt to allow different groups to concentrate on the problems related to their area of expertise, an additional level of classification (or prefix) was added to error messages. Currently, there are three classes of interest, each with an associated error summary. The first class is the default case that is examined by a team of hardware/software people from the controls engineering group. The second class is associated with klystrons and these errors are examined by the group responsible for their maintenance. The third class is associated with operational or beam-related problems and these are examined by the Operations group. Additionally, there is an error summary associated with devices that exceed tolerance levels and may be responsible for degraded performance. This summary is examined by the area managers for the various accelerator areas.

Figure 1 shows the first page of the report that is produced daily. It allows one to see at a glance whether there were problems the previous day and where in the accelerator they occurred. It shows the most frequent errors binned by type of facility (e.g., magnets) and physical area of the accelerator. The period shown is for 24 hours starting at 6 a.m. on September 25, 1989.

Figure 2 shows the correlation of the most frequent errors and the microprocessors or processes that reported them for the same time period as fig. 1. By looking at this table, one tries to identify the most likely cause of problems. With experience, it has become possible to identify which of several possible components are malfunctioning by the correlation of different kinds of errors. Each component appears to have an identifiable failure signature.

Figure 3 shows the "Hot Spot List": the worst errors and which microprocessor or process reported them. The errors are binned in two-hour time intervals. The time period shown is the same as for figs. 1 and 2. The time history allows one to tell when a problem started or when it was corrected. The problems that appear high on this list are usually the first ones to be attacked, as their correction should provide the biggest increase in performance or reliability. Often the errors are not of a "fatal" nature; i.e., one can continue to operate, but performance is degraded.

In an attempt to track the errors over time and to look for seasonal or other correlations, a history file is kept for each month to keep track of the number of errors by facility (e.g., magnets, klystrons) and area of the accelerator. Figure 4 shows a histogram of CAMAC errors for the month of September 1989. The large peak on September 28 was related to problems encountered when installing a new version of the online database.

4. Additional summaries

Once the Errorlog summary report was generated daily and used to check for hardware/software failures, emphasis was placed on other operational procedures that, if improved, would lead to more reliable operation of the accelerator. These were related to:

- Microprocessor reliability
- Beam Position Monitor calibration history
- Magnet calibration history

A project to upgrade all microprocessors from Intel 8086 boards to Intel 80386 boards has been underway for the past year. This upgrade provides both additional processor speed and additional available memory. All other hardware components of the microprocessor assemblies have been carefully checked and upgraded with a view toward increased reliability. In order to track microprocessor reliability, a daily report was generated listing micros that were rebooted and why it was perceived necessary to reboot them. Figure 5 shows a list of microprocessors that were rebooted (IPLed) during the 24-hour period beginning at 6 a.m. on October 9, 1989. The person rebooting a micro is prompted "Who are you and why are you doing this?" and this information is logged. The reboot history is reviewed daily to determine whether some particular micro has been causing problems. Some reboots have been linked to software errors associated with improper handling of an external event that occurs infrequently, while others have been traced to faulty hardware components.

Another area affecting machine performance is that of calibrations. Feedback and other optimization processes rely on accurate Beam Position Monitor (BPM) measurements which are a function of correct calibrations. Figure 6 shows the last time various groups of BPMs were calibrated. Those groups that have been calibrated within the last 24 hours are marked with an asterisk. This table is printed and checked daily to encourage frequent calibration so that BPM measurements remain accurate.

It is also desirable that magnets be calibrated regularly so that they can be set accurately and problems that may exist over parts of their range can be found and fixed. Once per week an additional histogram is produced in the daily report that indicates how long it has been since magnets in various areas of the accelerator have been calibrated.

5. Operational experience

Since its inception, the Errorlog summary has proven its usefulness as an overall diagnostic for the SLC Control System. The number of errors that appeared in early logs numbered in the tens of thousands. These numbers have been reduced by an order of magnitude. Many software bugs and cases of poor error handling have been found and fixed. Hardware systems have been upgraded and intermittently unreliable components have been found and repaired.

Whereas in the past, software personnel used the Errorlog to find software problems, currently hardware personnel have also become familiar with its use in tracking down faulty hardware components. The Errorlog has been used to spot check on a Friday for a problem that is just beginning and may lead to long down times over the weekend. Many problems that were discovered occurred as infrequently as once every 24 hours, and then only under normal operating conditions. These problems would have been very difficult to find with stand-alone diagnostics in the lab.

Similarly, even though the klystron group were already using daily reports with the same information, they have found it easier to see problems using the new summary since problems are sorted by "worst case."

Since the overall error message "noise level" has been reduced, a programmable sound generator has been added to signal the type of problem that has just occurred by the use of different sounds. This allows operations personnel to take corrective action more quickly and improve uptime.

6. Future enhancements

In the future, we hope to use the statistics that we are gathering to look for long-term correlations or trends. Additionally, we will be using our trouble entry and report system (CATER) to keep track of ongoing problems and to keep track of the "signatures" for certain component failures. This system will allow us to search for a keyword and to find out what we did in the past to solve problems where that particular keyword applied.

7. Acknowledgments

We wish to thank M. Breidenbach, J. Humphrey, T. Schalk, N. Phinney, and T. Himel for their ideas, support and encouragement on this project. We also wish to acknowledge the efforts of all of the hardware and software personnel of the SLC Controls Engineering group who worked diligently to find and solve problems. We also thank S. Castillo and H. Kirby for their software contributions.

References

[1] R. E. Melen, "Design and Performance of the Stanford Linear Collider Control System," IEEE Trans. Nuc. Sci., NS–32 (1985) 230.

[2] N. Phinney, "Report on the SLC Control System," IEEE Trans. Nuc. Sci., NS–32 (1985) 2117–2119.

[3] N. Spencer et al., "Error Message Recording and Reporting in the SLC Control System," IEEE Trans. Nuc. Sci., NS–32 (1985) 2120–2122.

Figure captions

1. The most frequent errors binned by type of facility and physical area of the accelerator. The period shown is for 24 hours starting at 6 a.m. on September 25, 1989. This is the first page of a report printed and distributed daily. One can see at a glance whether there were problems the previous day and where in the accelerator they occurred.

2. A correlation of the most frequent errors and the microprocessors or processes that reported them for the same time period as fig. 1. By looking at this table, one tries to identify the most likely cause of problems.

3. The "Hot Spot List": the worst errors and which microprocessor or process reported them. The errors are binned in two-hour time intervals. The time period shown is the same as that for figs. 1 and 2. One can tell from the time history when a problem started or was corrected.

4. A histogram of CAMAC errors for the month of September 1989. The large peak on September 28 was related to problems encountered when installing a new version of the online database.

5. Microprocessors rebooted (IPLed) and comments entered by the person rebooting them when he was prompted "Who are you and why are you doing this?" This list is reviewed daily to determine whether some particular micro has been causing problems.

6. The last time various groups of Beam Position Monitors (BPMs) were calibrated. It is printed and checked daily to encourage calibrating frequently so that BPM measurements remain accurate.

XSTAT 1.92   Most Frequent Area and FacilitY Errors for errors_sep25.out
From 25-SEP-1989 06:00:00   to 26-SEP-1989 06:00:00   PREFIX: ALL

RANK	FACILITY	COUNT	RANK	AREA NAME
1	TIME	1276	1	LINAC
2	BPMO	801	2	OTHER
3	PIOP	654	3	COWS
4	MGNT	186		ARCS
5	MPG	153		FAST FEEDBACK
6	KLYS	114		FINAL FOCUS
7	CAM	93		DAMPING RINGS
8	FDBK	72		
9	FBK	42		
10	OPS	39		
11	SYSTEM	37		
12	SLCNET	35		
13	GPIB	18		
14	CRAT	17		
15	RAMM	13		
16	CUD	12		
17	MSG	11		
18	LGPS	9		
19	MPS	7		
20	HSTB	6		
21	DATABASE	5		
22	MICR	5		
23	MODL	4		
24	CRR	4		
25	FOR	4		
26	POLL	3		
27	COMF	3		
28	STAT	2		
29	STR	1		
30	CONFIG	1		
31	BOOT86	1		

Fig 1

XSTAT 1.92 Most Frequent Errors Table for errors_sep25.out
From 25-SEP-1989 06:00:00 to 26-SEP-1989 06:00:00 PREFIX: $

ERR NAME/MICRO NAME	LI00	MP00	FF01	CA11	FDBK	CW02	CF22	FF11	CW05	CW08	CF21	FB31	ERRSUM
TIME-E-FID_NONE1,	1172	.	.	.	.	.	.	.	.	.	.	.	1172
MPG-E-VETOXMIT,	.	152	.	.	.	.	.	.	.	.	.	.	152
BPMO-E-ONEQ,	.	.	80	.	.	.	.	26	.	.	.	.	106
CAM-E-SOFT_TO,	.	.	.	78	.	.	.	.	.	.	.	.	78
KLYS-E-C$CAMAC,	.	.	.	.	58	.	.	.	4	.	.	.	62
BPMO-E-NMISKIP,	.	.	.	51	.	.	.	.	.	.	.	.	51
MGNT-E-OPERR,	.	.	.	.	.	43	7	.	6	.	.	.	56
SYSTEM-F-BADESCAPE,	.	.	.	.	.	.	35	.	.	.	.	.	35
BPMO-E-NORESLTS,	.	.	.	.	31	3	.	.	15	5	6	.	60
MGNT-E-C$NOXRESP,	.	.	.	27	.	.	.	.	.	.	.	.	27
BPMO-E-MPGREQ,	.	.	.	.	9	5	.	.	24	1	18	.	57
TIME-E-FID_MISS,	23	.	.	.	.	.	.	.	.	.	.	.	23
BPMO-E-S$WRONGLEN,	.	.	.	.	.	.	.	.	.	20	.	.	20
FBK-E-MDL_KLYS_PTRB	.	.	.	.	17	.	.	.	.	.	.	.	17
BPMO-E-C$NOGATE,	.	.	.	3	.	.	.	.	.	.	.	15	18
SLCNET-E-NOREAD,	.	.	.	5	14	1	.	.	2	.	.	.	22
PIOP-E-DAC_MONO,	.	.	.	.	.	.	.	.	.	.	.	.	0
RAMM-E-S_TOLFAULT,	.	.	.	.	.	.	.	.	.	.	.	.	0
CRAT-E-OFFVER,	.	.	.	12	.	.	.	.	.	.	.	.	12
MSG-E-NOMEM,	.	.	.	11	.	.	.	.	.	.	.	.	11
BPMO-E-S$YYDBLALLOC	.	.	.	.	.	.	.	.	.	.	.	.	0
MGNT-E-CAMERR,	.	.	.	.	.	10	3	.	.	.	.	.	13
BPMO-E-BSEGTIMOUT,	.	.	9	.	.	.	.	9	.	.	.	.	18
OPS-E-MICRTRNUNRN,	.	.	.	.	.	.	.	.	.	9	.	.	9
MICRO SUM	1195	152	89	187	129	62	45	35	51	35	24	15	2019

ERRORS SHOWN = 2019 OUT OF 3628 ERRORS IN FILE

Fig. 2

XSTAT 1.92 Error Hot Spot List for errors_sep25.out

From 25-SEP-1989 06:00:00 to 26-SEP-1989 06:00:00 PREFIX: $

RANK	ERROR NAME	MICRO NAME	6-8	8-10	10-12	12-14	14-16	16-18	18-20	20-22	22-24	0-2	2-4	4-6	COUNT
1	TIME-E-FID_NONE1,	LI00	99	125	123	21	.	106	121	102	124	125	105	121	1172
2	MPG-E-VETOXMIT,	MP00	9	21	17	11	5	12	13	9	11	21	6	17	152
3	BPMO-E-ONEQ,	FF01	17	7	27	.	.	1	4	2	2	6	8	6	80
4	CAM-E-SOFT_TO,	CA11	.	3	3	.	1	1	10	25	30	.	5	.	78
5	KLYS-E-C$CAMAC,	FDBK	.	.	1	19	11	12	15	.	.	.	.	.	58
6	BPMO-E-NMISKIP,	CA11	.	.	.	.	.	.	.	7	44	.	.	.	51
7	MGNT-E-OPERR,	CW02	5	1	.	2	.	.	.	.	32	.	.	3	43
8	SYSTEM-F-BADESCAPE,	CF22	.	.	.	35	.	.	.	.	.	.	.	.	35
9	BPMO-E-NORESLTS,	FDBK	.	.	.	4	7	.	7	1	5	1	1	5	31
10	MGNT-E-C$NOXRESP,	CA11	.	.	2	.	.	1	8	4	11	.	1	.	27
11	BPMO-E-ONEQ,	FF11	.	.	12	.	.	1	.	7	1	.	2	3	26
12	BPMO-E-MPGREQ,	CW05	1	2	2	.	.	.	.	1	1	7	4	6	24
13	TIME-E-FID_MISS,	LI00	.	.	.	.	23	.	.	.	.	.	.	.	23
14	BPMO-E-S$WRONGLEN,	CW08	.	.	20	.	.	.	.	.	.	.	.	.	20
15	BPMO-E-MPGREQ,	CF21	.	.	.	8	10	.	.	.	.	.	.	.	18
16	FBK-E-MDL_KLYS_PTRB	FDBK	.	.	1	4	3	6	3	.	.	.	.	.	17
17	BPMO-E-C$NOGATE,	FB31	.	.	.	11	3	1	.	.	.	.	.	.	15
18	BPMO-E-NORESLTS,	CW05	1	4	2	.	.	.	1	.	.	2	1	4	15
19	SLCNET-E-NOREAD,	FDBK	.	.	.	.	.	.	8	.	6	.	.	.	14
20	RAMM-E-S_TOLFAULT,	RAMM	1	.	.	1	2	1	3	2	.	.	2	1	13
21	PIOP-E-DAC_MONO,	LI18	2	11	.	.	.	.	.	.	.	.	.	.	13
22	CRAT-E-OFFVER,	CA11	.	.	.	.	.	.	2	4	6	.	.	.	12
23	MSG-E-NOMEM,	CA11	.	.	.	11	.	.	.	.	.	.	.	.	11
24	MGNT-E-CAMERR,	CW02	2	.	.	.	.	.	.	.	.	.	.	8	10
25	BPMO-E-MPGREQ,	CW04	9	.	1	.	.	.	.	.	.	.	.	.	10
26	BPMO-E-S$YYDBLALLOC	LI16	4	.	6	.	.	.	.	.	.	.	.	.	10
27	BPMO-E-S$YYDBLALLOC	LI08	4	.	6	.	.	.	.	.	.	.	.	.	10
28	BPMO-E-MPGREQ,	FDBK	.	.	.	3	.	2	.	1	2	1	.	.	9
29	BPMO-E-BSEGTIMOUT,	FF11	.	.	.	4	5	.	.	.	.	.	.	.	9
30	BPMO-E-BSEGTIMOUT,	FB69	.	.	.	4	5	.	.	.	.	.	.	.	9
31	OPS-E-MICRTRNUNRN,	CW06	.	.	7	.	.	.	.	.	.	.	.	2	9
32	BPMO-E-CALFAIL,	CF21	.	.	.	4	5	.	.	.	.	.	.	.	9
33	BPMO-E-BSEGTIMOUT,	FF01	.	.	.	4	5	.	.	.	.	.	.	.	9
34	OPS-E-MICRTRNUNRN,	CW08	.	.	.	.	.	9	.	.	.	.	.	.	9
35	BPMO-F-PUBLIC_ERR,	CF21	.	.	.	4	5	.	.	.	.	.	.	.	9
36	TIME-E-ERRINT,	LI16	.	.	1	.	7	.	1	.	.	.	.	.	9
37	BPMO-E-C$MBCDSTAT,	CA11	.	.	.	.	.	.	7	1	.	.	.	.	8

ERRORS SHOWN = 2077 OUT OF 3628 ERRORS IN FILE

Fig. 3

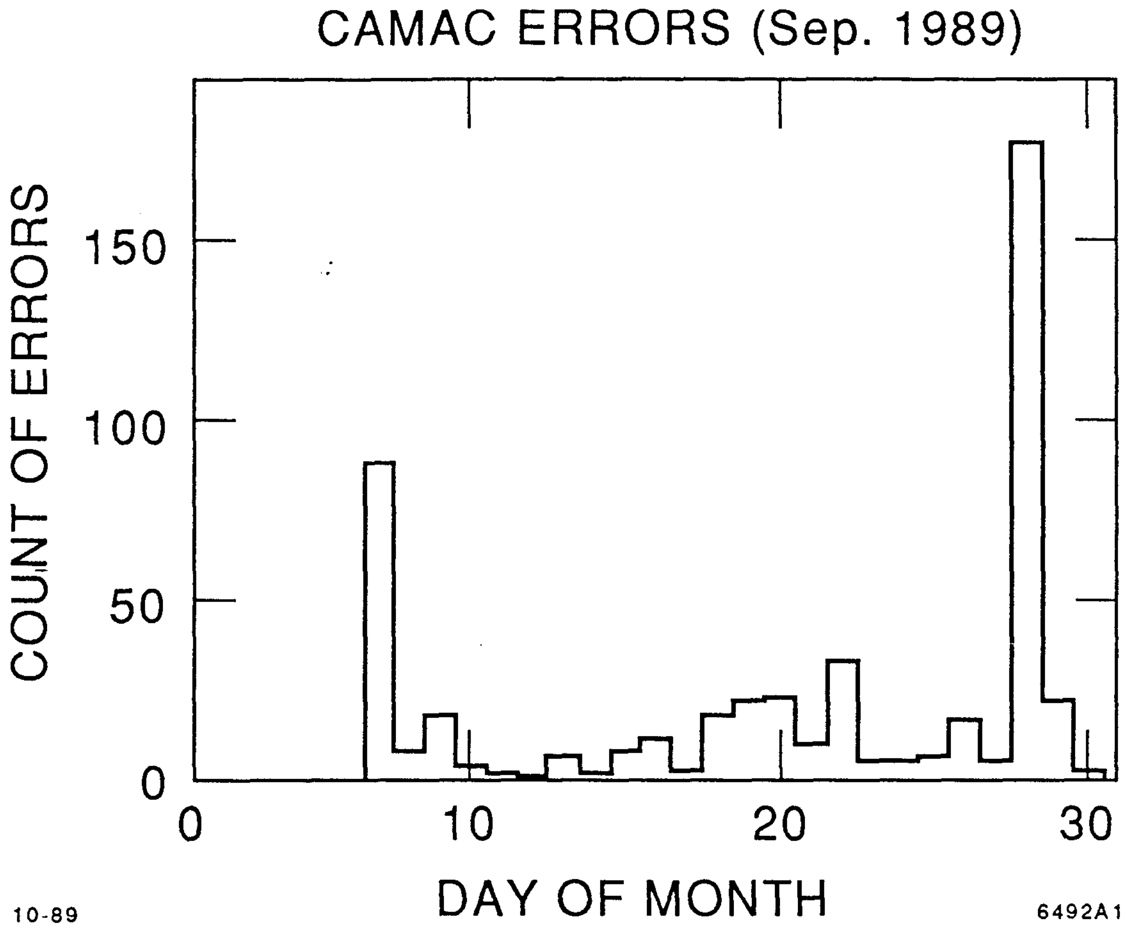
CAMAC ERRORS (Sep. 1989)
COUNT OF ERRORS
150
100
50
0
0
10
20
30
DAY OF MONTH
10-89
6492A1

Fig. 4

```
--------------------------------------------------------------------------------------------------------------------

/DESTINATION=ALL, /FACILITY=(OPS), /SEVERITY=(I), /ERRCODE=ALL, /FORMAT=ALL, /MATCH=AND, /STRING=(IPL), /BINSIZE=FULL,
/SLICES=SUMS, /PRIMARY=NONE, /MICRO=NONE, /UNIT=NONE, /SECONDARY=NONE, /ITEM=ALL, /OPER=NONE

--------------------------------------------------------------------------------------------------------------------
9-OCT-1989 07:43:06  %OPS-I-MSGTXT,       V007  SCP_07 - IPL: CA01 - HO NO
9-OCT-1989 07:44:15  %OPS-I-MSGTXT,       V007  SCP_07 - IPL: CA02 - PSH NO
9-OCT-1989 07:45:28  %OPS-I-MSGTXT,       V007  SCP_07 - IPL: CA11 - PSH NO
9-OCT-1989 07:45:56  %OPS-I-MSGTXT,       V004  SCP_04 - IPL: LI04 - DER SUB BOOSTER FAULT
9-OCT-1989 07:46:29  %OPS-I-MSGTXT,       V007  SCP_07 - IPL: CA12 - PSH NO
9-OCT-1989 08:27:08  %OPS-I-MSGTXT,       V007  SCP_07 - IPL: CA11 - NAN BACK TO NEW JOB
9-OCT-1989 11:11:34  %OPS-I-MSGTXT,       V003  SCP_03 - IPL: LI17 - GAS
9-OCT-1989 11:12:55  %OPS-I-MSGTXT,       V003  SCP_03 - IPL: LI17 - GAS
9-OCT-1989 11:43:19  %OPS-I-MSGTXT,       V003  SCP_03 - IPL: LI17 - AAN
9-OCT-1989 13:58:55  %OPS-I-MSGTXT,       V005  SCP_05 - IPL: LI04 - HVS
9-OCT-1989 15:58:50  %OPS-I-MSGTXT,       V007  SCP_07 - IPL: LI06 -
9-OCT-1989 16:07:58  %OPS-I-MSGTXT,       V004  SCP_04 - IPL: LI06 - KGG FOR NORB
9-OCT-1989 16:09:08  %OPS-I-MSGTXT,       V004  SCP_04 - IPL: LI06 -
9-OCT-1989 16:45:07  %OPS-I-MSGTXT,       V007  SCP_07 - IPL: FB73 - JXK, TEST PMPLETE
9-OCT-1989 16:49:23  %OPS-I-MSGTXT,       V002  SCP_02 - IPL: LI04 - las deleted
9-OCT-1989 16:50:16  %OPS-I-MSGTXT,       V002  SCP_02 - IPL: LI04 -
9-OCT-1989 17:08:55  %OPS-I-MSGTXT,       V002  SCP_02 - IPL: LI04 -
9-OCT-1989 17:37:34  %OPS-I-MSGTXT,       V004  SCP_04 - IPL: DR13 - KGG SOME ADC ERRORS ON YCORS
9-OCT-1989 17:42:59  %OPS-I-MSGTXT,       V004  SCP_04 - IPL: DR11 - KGG SOME ADC
9-OCT-1989 17:44:20  %OPS-I-MSGTXT,       V004  SCP_04 - IPL: DR11 -
9-OCT-1989 17:53:51  %OPS-I-MSGTXT,       V004  SCP_04 - IPL: DR11 -
9-OCT-1989 18:01:36  %OPS-I-MSGTXT,       V003  SCP_03 - IPL: DR13 - LAS
9-OCT-1989 21:27:47  %OPS-I-MSGTXT,       V005  SCP_05 - IPL: FB31 - JLT ENERGY FEEDBACK CAMAC PBLMS
9-OCT-1989 21:37:29  %OPS-I-MSGTXT,       V004  SCP_04 - IPL: LI06 - KGG - RED BPMS
9-OCT-1989 23:47:35  %OPS-I-MSGTXT,       V002  SCP_02 - IPL: FF01 - jlt
```

Fig. 5

Summary of SLC BPMD Calibrations 10-OCT-1989 06:35

Unit #	BPMD Label	Time Calibrated
* 3	SLC Extrct Elec	10-OCT-1989 00:00
* 4	SLC Extrct Posi	9-OCT-1989 23:38
10	SLC IP Posi	8-OCT-1989 07:35
* 11	BSY Fast Fdbk	9-OCT-1989 21:25
* 12	SLC IP 2 Beam	10-OCT-1989 00:31
15	Coll Fast Fdbk	8-OCT-1989 07:37
30	SLC L1Inj e-RateLm	22-SEP-1989 13:34
40	SLC Scav e-	8-OCT-1989 07:34
41	SLC LInjecte- 1	8-OCT-1989 07:32
42	SLC HInjecte- 2	8-OCT-1989 07:33
43	SLC e+ Return	8-OCT-1989 07:34
44	SLC Inject Posi	8-OCT-1989 07:34

NOTE: * Indicates that BPMD has been calibrated within the last 24 hrs.

Fig. 6