

UNIVERSITY OF MINNESOTA COMPUTER CENTER
Deadstart Systems Newsletter

12 August 1980

Vol. 6, No. 15

Send all comments, criticisms and contributions to the editor: T. W. Lanzatella,
University Computer Center, 2520 Broadway Drive, Lauderdale, MN 55113.
The University of Minnesota is an equal opportunity educator and employer.

TABLE OF CONTENTS

NOTICE OF CHANGES TO THE SYSTEM 142

PROPOSED CHANGES TO THE SYSTEM 144

 TIMESHARING COBOL - S. Reisman 144

 UNPURGE - D. W. Mears 145

 WRITEUP AND CALLPRG AS SEPARATE PROGRAMS - M. Riviere 146

SYSTEM MAINTENANCE 149

 LAST WEEK'S SYSTEMS GROUP MEETING - T. W. Lanzatella 149

 CALLPRG AND LIBRARY TAPE NEWS - M. Riviere 150

 16 WORD PFC SYSTEM - K. C. Matthews 150

 CYBER DEADSTART DUMP ANALYSIS - K. C. Matthews 152

 CYBER 720 DEADSTART DUMP ANALYSIS - R. A. Williams 153

 TELEX AND TELEX PDP-11 CRASH ANALYSIS - D. W. Mears 154

NOTICE OF CHANGES TO THE SYSTEM

Tom Lanzatella added a change to QAC which allows input files to be attached as local files. This change is in preparation for a utility called CQLOAD which gathers certain jobs out of the input queue and dumps them to a tape which is carried to a Cray site and then loaded into the Cray input queue.

Kevin Matthews added some space-saving code to lTA, a perennially tight program. Kevin also made changes to PFP and lWF in preparation for the 16 word PFC. Additionally, several errors were fixed in DSP.

- 1) The routed bit is no longer accidentally set when a routed file is requeued.
- 2) The TID will not be set to BC when another TID had been previously set.

Also, QDV will now process a file with a bad TID correctly if the file has no EFNT entry.

John Larsen changed the N=* option on MST so that, if specified, the routine runs until a track limit is reached and stops rather than starting over.

Paul Thompson changed ISF to check for a system origin call before running. Paul also added a UN-option and SUN directive to Dumppf. These options work analogously to the UI option and SUI directive and are used almost exclusively for end-of-month permanent file archiving.

Don Mears installed the following changes:

- 1) Lace card processing in the event of a compare error occurring on the last lace card punched was corrected.
- 2) Don added several very innovative changes which correct system behavior when the system is at PP-saturation.
 - a) CPUMTR was changed to not stop calling ISF so that jobs will continue to be scheduled.
 - b) ISJ was changed so that before it loads IAF or IRI on top of itself, another copy is called up to complete the scheduling cycle.
 - c) MAGNET was changed so that it does not stop calling LMT to read labels and check unit status.
 - d) LMT was changed so that it does not stop reading or writing L-format tapes.
 - e) DIS was changed to not respond with PP BUSY to all commands.

The overall effect of these changes is to attempt to keep the system moving despite PP-saturation. The desired effect is to smooth out response-time and so far measurements bear this out.

- 3) Some PSR code from the 518 release was installed which considerably cleans up DIS.
- 4) The "schedule-around-hung-control-point" feature was changed so that the time-out for storage move was reduced from 8 seconds to 2 seconds. This helps the system schedule around control points using L-format tapes.
- 5) The TIELINE facility was installed. This utility will allow Cybers to communicate with IBM mainframes.

Tim Hoffmann installed the following changes:

- 1) A spurious NO JOBS FOUND message was removed from the new MMF-ENQUIRE facility.
- 2) DIVERT processing in EXPORT was enhanced to allow diverting to non-BC sites.
- 3) Several small errors in RELOAD were repaired.

Jeff Drummond added the following changes:

- 1) Program PFPACK was changed in preparation for the 16 word PFC.
- 2) Several new symbols were added to deck CALLCPU so that it assembles without errors.
- 3) Some PSR code from the 518 release was installed which repairs some serious problems with multi-file positioning.

- 4) A new version of the common deck COMCMCF was added which is used by BLOCK/UNBLOCK. The new version will substantially speed up BLOCK/UNBLOCK on a non-CMU mainframe.
- 5) A new common deck COMSDDF was added which defines equivalences used by DDF and DDFILE.
- 6) New and significantly improved versions of DDF and DDFILE were installed.
- 7) New versions of BLOCK/UNBLOCK, COMCCCS and COMCEMP were added with minor changes.

Brad Blasing installed the following changes:

- 1) The USTL word in low core was eliminated. This word was used for bits indicating LOWRATE, NOFRILLS, AUTOUNLOAD, etc. All bits which use these were moved to SSTL, the official CDC locations for such things. The reason for the change is a problem in deadstart recovery. Whenever a L3 deadstart is performed, a limited amount of CM is recovered. Unfortunately, USTL was not included, hence after a L3 recovery we would sometimes get random data in USTL causing problems for operations because LOWRATE or NOFRILLS would get turned on or off (maybe). Since deadstart is too hard to fix we decided to move the bits.
- 2) With the elimination of the IQFT for LOWRATE and NOFRILLS jobs we no longer had a need for the ENABLE/DISABLE DELAYQ facility. This was removed.
- 3) Brad executed Bob Williams proposed change to remove CTPC (see DSN 6, 13 p. 119).
- 4) A small error in LRI processing of disk errors was repaired.

Andy Hastings changed CONTROL to allow OOB characters in KCL-IF statements. Andy also repaired a problem in the recently added SEND feature which allows a PN parameter on SEND. If TRANSIT sends a dayfile immediately after sending a permanent file to another machine the dayfile could end up on the same device as the permanent file. Additionally, Andy made a major change to TRANSIT which now calls DSP rather than QFM in order to enter a job into the input queue.

PROPOSED CHANGES TO THE SYSTEM

Timesharing COBOL - by Steve Reisman

Several years ago our COBOL subsystem was lost in the KRONOS to NOS system upgrade. We decided to drop it at that time because it was based on an old COBOL 3 compiler and statistics showed that usage was very low. We have thus been without a COBOL subsystem for over two years.

I would like to reinstall the COBOL subsystem (based on COBOL 5) into TELEX. The compiler will accept sequenced source, create a listing formatted for a 72 column output device, and automatically do a load-and-go if there are no fatal compilation errors. By default only errors will be listed and keeping in line with a recent MNF proposal, all debugging options will be turned on. The new subsystem name will be COBOLTS.

//////////

UNPURGE - by D. W. Mears.

I propose installing UNPURGE as described in the following writeup. This involves changes to PFILES and PFM. PFILES must make a call to the new UNPURGE PFM function. The UNPURGE function must be added to PFM. The PURGE function in PFM must be changed to save the original user index in the hole catalog entry.

UNPURGE - Restore purged indirect access files.

The UNPURGE control card has the following format:

UNPURGE,f1,f2, ... fn/PN=ppp,PW=www,na.

f1 to fn are the files to be UNPURGED. At least one file must be specified.

The following parameters are optional and may be selected in any order.

PN=ppp selects pack name ppp as the pack to perform the UNPURGE on.

PW=www causes www to be used as the password for the file being UNPURGED. The parameter is only required for files which were saved with file protection turned on (FP=ON).

NA selects *NO ABORT*. If NA is selected, UNPURGE will not abort on errors.

The UNPURGE command gives the user a chance to restore an indirect access permanent file which was purged or replaced. Explaining the usefulness and limitations of this command requires some description of how the permanent file system handles purged files.

When an indirect access file is purged, the file space is marked as being available for new permanent files. This unused space is referred to as a permanent file hole. The hole will remain intact until some permanent file request such as SAVE, REPLACE, or APPEND needs to create a new permanent file and that new permanent file fits into this hole. Only then is the data from the original file destroyed. This means that when a file is purged, its data remains in a recoverable state for a random amount of time ranging from zero seconds to several days.

The UNPURGE command searches for a permanent file hole with a name which matches the one specified on the UNPURGE command and changes that file into a real indirect access file. If more than one hole with a matching name is found, the hole with the most recent modification date is UNPURGED.

Here are some hints and cautions on using UNPURGE.

1. If the file to be UNPURGED was a protected file (FP=ON), the password must be specified on the UNPURGE command.
2. The message "name ALREADY PERMANENT." is issued if an attempt is made to UNPURGE a file when a file of the same name already exists in the permanent file catalog. The CHANGE command can be used to change the name of the existing permanent file in order to avoid the name conflict problem.
3. When a indirect access file is REPLACed, the previous version of the file is not changed into a hole if the new file is the same size as the old file. REPLACE will always change the old file into a hole if the new file is bigger than the old file.
4. To get back to an earlier version of a file which has been REPLACed, use the CHANGE command to change the name of the existing permanent file and then use UNPURGE to restore the most recently modified matching hole.
5. UNPURGE updates the creation date and time, but leaves the last access and last modification dates and times unchanged. This allows the user to do a "CATLIST,LO=F,FN=name." in order to determine how up-to-date the UNPURGED file is.
6. UNPURGE causes all the permit information about the file to be lost. The permit information is stuff listed out on a "CATLIST,LO=FP,FN=name." command. That is the user numbers permitted to the file through the PERMIT command, and (if the file was semiprivate (CT=S)) the user numbers which accessed the file.

//////////

WRITEUP and CALLPRG as Two Different Programs - by M. Riviere

Originally, we had CALLPRG, an extension of the System's directory searcher that was retrieving packages not included in the system directory. Later, CALLPRG acquired the capability of obtaining machine retrievable documentation and listing it. This extension of CALLPRG was very natural since the major part of the searching procedures for documentation were similar to the ones used to retrieve any other kind of package. Eventually, new features were added to CALLPRG and existing features were enhanced (e.g. retrieving tape packages, setting up specific loading conditions, handling machine oriented packages, setting up user libraries, etc.). Although CALLPRG and its WRITEUP addition perform some equivalent tasks such as searching an index and such as following index directives in order to locate a given product, a large part of their other tasks have become less and less common as new capabilities have been implemented in both sections. The extension of both processors made many areas of code quite confusing and both processors are now in a stage where modifications to one can largely affect the performance of the other.

In one hand, we have CALLPRG, which is static and dependable but any modification to CALLPRG should be done very carefully and be thoroughly tested.

On the other hand, we have WRITEUP, which could be much more flexible and rich in options but we do not expand WRITEUP as much as we would like because of the difficulty of modifying CALLPRG.

I think that both processors should be now split into two different programs. The splitting will allow us to easily make modifications (and correct some long standing errors) to either one of them. Yvonne Murray and myself have been looking into the splitting process and we consider that as a result of the conversion, CALLPRG and WRITEUP should end up being as follows. CALLPRG should remain with all its existing capabilities, with the exception of course, of the WRITEUP processor. The removal of the writeup section from CALLPRG will allow us to clean up quite a lot of code from CALLPRG that was introduced in order to enhance WRITEUP. This removal will make CALLPRG much easier to understand and to modify, when needed. WRITEUP would be restricted to handle only specific WRITEUP related index parameters. Now, with the current CALLPRG/WRITEUP processor we could extend any CALLPRG index parameter to WRITEUP. All that may be needed is to include the parameter in a WRITEUP index entry. But most of the CALLPRG parameters do not make any sense for WRITEUP and the code to process them is completely useless in WRITEUP. Besides, although the CALLPRG index parameters seem to be available for WRITEUP, that does not necessarily mean that they all will work properly. It is much easier to add any parameter that we may remove now, later on, if we consider that we need it, then to keep all of them. The common CALLPRG/WRITEUP index parameters that should be maintained for the WRITEUP index entries are the ones used to identify the file and its location (DA/IA, PN, UN) and the tape parameters although these will be processed mainly by CALLPRG. The specific WRITEUP parameters, (DT, IX, LC, TR) will, of course, remain there.

WRITEUP should also be restricted from handling the retrieval of tape resident documents. The tape feature should be partially removed from WRITEUP. Currently the CALLPRG/WRITEUP processor submits a System origin job that handles the tape mounting and copying requirements independently of the user's validation. WRITEUP should not need to have System origin privileges any longer in order to access tape resident packages. WRITEUP could submit a job that can be processed by CALLPRG and then CALLPRG will be the one that will carry on all the tape handling processes using its own system origin privileges. WRITEUP should submit then only jobs that will run within the user's validation. That is, users without submit validation will not be able to request a tape resident writeup. (I will expand on this fact later.) The information needed to locate a tape resident writeup will be, however, in the WRITEUP index.

Before submitting a job to request CALLPRG to locate a tape resident document WRITEUP will explain to the user how the retrieval processes will be (see message "A"). If the user agrees to issue the job for the request, WRITEUP will ask the user for his or her account number and password. Then, upon issuing the job, WRITEUP will ask the user if s/he wants to wait for the document to be retrieved (see message "B"). Should the user not be validated to submit the simple WRITEUP tape job, (that will be a CALLPRG Fetch type of request), WRITEUP will issue another message about this fact (see message "C"). A user without validation could, in this case, ask a consultant to request the writeup retrieval and access it, later on, when available. (Tape resident writeups and CALLPRG packages remain available as disk pack files for a while once that they are retrieved from the tape.)

I know that this validation conflict may seem a little messy and confusing, but we should consider the facts that there are very few writeups on tape (only one for the time being) and also that most of our users are validated to submit jobs. Besides, the chances that a user not validated for submitting jobs will be requesting a tape resident writeup are quite small since tape resident writeups are usually long ones and they are referenced only by users with a certain degree of sophistication which is associated in most cases with submit validation.

Following is the text of my suggested messages:

Message "A":

The writeup that you are asking for resides on a tape. It is necessary, in order to retrieve the writeup, to issue a job. Do you want the job to be issued? Please type yes or no.

Message "B":

It will take between 5 and 15 minutes to have the writeup available. You have the option of not waiting for it and to request the writeup later on, when available. Do you prefer to wait? Please type yes or no.

Message "C":

You are not validated for submitting jobs or the number of submitted jobs, at this moment, is at its limit. If you are not validated see a consultant in order to obtain the writeup. If you are validated, request the writeup again once the number of submitted jobs becomes smaller.

SYSTEM MAINTENANCE: People and Procedures

Last Week's Systems Group Meeting - by T. W. Lanzatella

Preceding our customary discussion of current proposals, we had a brief chat about the content and form of proposals appearing in the DSN. Proposals should include:

- 1) History or context of the problem;
- 2) Statement of the problem;
- 3) Proposed solutions;
- 4) Consequences of this solution.

In the future, TWL, as editor of the DSN, will expect all proposals to contain this minimal information. For those individuals not inclined or unable to formulate a complete proposal, an avenue will be provided so that their ideas can be heard by the Systems Group. A new section will be added to the DSN called Suggested Changes To The System. Topics appearing in this section will be taken up at the System's Group meetings as discussion topics only. People making the suggestions will have to decide, based on the outcome of the discussion, whether to proceed with a proposal. People who make suggestions will also have to find someone in the Systems Group willing to write a detailed proposal and perform the work of installation.

Next, we took up Jeff Drummond's Dayfile Madness proposal left over from the last meeting. After a brief recap of the facts we readily approved the first suggestion that all dayfiles should be truncated to some length which is less than the size of the permanent file DAYFILE. We then laboriously decided to adopt the most radical (but neatest) scheme for adding dayfiles to DAYFILE. We decided to add all dayfiles to the beginning of DAYFILE thus the oldest dayfiles will be scrolled off of the end of the file (see DSN 6, 13 p. 122). This scheme requires quite a bit of user notification. Writeup SUBMIT will be updated to describe the new procedure.

Larry Liddiard's proposal dealing with priority scheduling and a screen editor was taken up as a discussion topic (see DSN 6, 14 p. 134). Regarding priority scheduling, we concluded that implementation would be easy using a job class and some small changes to the job scheduler. Accounting would also be easy. Our biggest concern was with the potential conflict between commercial users with plenty of money hampering service to University users. We also wondered what the charging rate might be. We all thought that a rate at least double the usual rate would be appropriate. Regarding screen editors we wondered about need. Mike Frisch indicated that any decent word processing capability on the Cybers would require a screen editor. The dominant opinion was that there is no general need for a screen editor on the Cybers considering the burgeoning use of small computers like Terak.

Mike Frisch's proposal to install a facility to automatically reload archived permanent files was rejected without discussion as time was growing short. After the meeting two serious drawbacks to Mike's suggestion were pointed out.

- 1) Under Mike's scheme, a user would have no way of purging a permanent file which had been archived.
- 2) The proposal did not address or estimate the probable increase in tape mounting.

Larry Liddiard concluded the meeting with a brief discussion of our work on networking and suggested that we may not want to use a VAX for a communications device as PDP-11 appears to be a better choice. Larry also announced that we have bought TIELINE, a package which allows Cybers to talk to IBM machines. We will be using the SUPIO PDP-11 as the communications device for TIELINE.

//////////

Callprg and Library Tape News - by M. Riviere

The only modification scheduled to take place among Library Tape and Callprg index products is the inclusion of the LINPACK library on the Cyber 720 Library Tape and the needed entries in the Callprg index to define LINPACK as a user library. This modification was requested by M. Frisch.

The next set of Callprg and Library Tape modifications will be implemented on September 2. An additional, unscheduled change will be taking place, however, on August 31.

The August 31 change concerns the replacement of the MNF, TSF, M77, and FTN compilers with versions where the default specifications for the trace parameter on MNF, TSF, and M77, and for the optimization level on FTN have been changed. A description of this change was included in the July issue of the UCC Newsletter. I have chosen the date for the change to be August 31 in order to make the change take place at the end of the summer quarter but still during the month of August, as mentioned in the Newsletter article.

Modifications for August 31 as well as for September 2, should be submitted no later than August 22, at noon.

After the September 2 changes, the next dates for Library Tape and Callprg index modifications will be September 16 and September 30. The deadline for submitting modifications for those dates are September 4, and September 18, respectively. Please, plan your end of the summer quarter modifications early enough in order to make them fit into these dates.

//////////

16 Word PFC System - by K. C. Matthews

On 1 September, UCC is going to implement mods to use a 16 word catalog entry for permanent files; currently our catalog entries are only 8 words long. The catalog entry is a data block on the disk which identifies and locates a permanent file. A 16 word entry is returned for each permanent file in response to the CATLIST request in PFM. Any routines which use this function internally have to be modified. The following routines have been modified and tested here

already:

ARCHIVE (past and current), CATLIST, CATLSYS, COPYCAT, DDF, DUMPPF, RELOAD, UNBUSY, PFILES (which contains the ATTACH, DEFINE, ... entry points), and all of the CDC permanent file dumping and loading routines.

The following fact should be noted by System Programmers. The 16 word PFC mods add a system sector to each indirect access permanent file. So now, for each file in the indirect access data chain, we have a system sector, hopefully some data sectors, and an EOI sector. This has required some changes in common decks COMPRSS and COMPWSS.

CDC supplies a modset called KRAPFC which changes the CDC routines for the 16 word PFC. Modsets KRAPFCA, KRAPFCB, KRAPFCC, and PFUI are also supplied which make corrections to bugs in KRAPFC. These modsets will be applied in their stock CDC form. They will not have HISTORY entries in the decks they modify. Our local mods can modify these mods as if they were on the stock CDC OPL. We decided to do this because it leaves the sequence numbers and modnames the same as those published by CDC.

All the permanent files in the system have to be dumped under our current permanent file system and then reloaded under the 16 word PFC system. This will be done on Sunday, 31 August. This implied that the BL level deadstart tape cannot be installed on Thursday, August 28 - its normally scheduled date. It will be installed for production use on Sunday, 1 September. The BL tape cannot be used on system time before 1 September, because of the difference in permanent file catalog sizes.

The CDC permanent file utilities correctly handle the differences between 8 word PFC dumps and 16 word PFC dumps. The 8 word PFC tapes that we currently have saved for archive purposes can always be loaded by the 16 word PFC system. We have also made sure that our current 8 word system can handle both full and incremental dumps from the 16 word system. So if this 16 word system turns out to be a fiasco, we can give up and go back to the 8 word system. (We would have to do a lot of dumping and reloading again, however.) It is unlikely that this will happen.

Remember that all DI removable packs will be turned into removable DY packs on 1 September. This gives us 5% more space on each of those packs, and saves a little memory. This will affect removable pack users who need a RESOURC statement. Those who have thing like:

```
RESOURC(DI=1,...)
```

will have to change the statement to

```
RESOURC(DY=1,... )
```

I have not received a single response to my July newsletter article where I mentioned this, so we'll have to assume it is an acceptable change to removable pack users. Helpline staff should expect a few calls from people who either forgot or never heard about the change.

Finally, if you have some CALLPRG package which needs to be changed on 1 September, I will make sure it is changed on that day. If you leave the new version of your program somewhere, and leave me directions on what to do, I'll try to follow those directions on 1 September.

//////////

Cyber 74/172 Deadstart Dump Analysis from Friday, 18 July through Thursday, 7 August - by K. C. Matthews

Friday, 18 July

02:48 (DD2002) Cyber 74
DDF hung the system while trying to read an illegal equipment. We have been unable to duplicate the problem. A level 3 recovery deadstart was required.

Tuesday, 22 July

09:50 Both Machines
The system hung on a shared queue device shortly after some cables had been moved. This should never have caused a software hang of the device, but it did.

Wednesday, 23 July

11:08 (DD-2021) Cyber 74
PP program QAP hung at the BATCHIO control point. The dump appears to have nothing to do with the problem. No information available to analyze.

Thursday, 24 July

08:29 Cyber 74
A level 0 deadstart was required when it was noticed that one of the production disk packs was not used at deadstart time - a system time pack was left up in its place. Brad Blasing has added a check for this in the deadstart procedure. The problem seemed to occur every few months.

20:53 (DD2023) Cyber 74
The system went into STEP mode because it thought a power problem was imminent. When it was unstepped, LHS (the EXPORT PP program) was hung and could not be dropped.

Monday, 28 July

11:35 Both Machines
A long power failure occurred. We were down for 3 1/2 hours.

Wednesday, 30 July

08:15 Cyber 74
The FMD disk unit for DN 12 suddenly hung up. A deadstart cleared up the problem. This problem has occurred before on this unit.

Tuesday, 5 August

00:25

Cyber 74

The FMD disk for DN30 failed. The failure was solid. Since FMD packs cannot be moved, the device had to be reloaded to our spare FMD. Since an incremental existed from 22:00, only 2 1/2 hours of users permanent file work were lost.

Wednesday, 6 August

02:35

Both Machines

A short power failure brought down both machines for a few minutes.

//////////

Cyber 170-720 Deadstart Dump Analysis (7/21-8/10) - by R. A. Williams

<u>Date</u>	<u>Description</u>	<u>Tape</u>
800724	A power failure caused a 6676 to go down, resulting in a TELEX abort.	Fixed
800726	844 disk unit 1 (just received from the Cyber 172) hung the system with a negative voltage fault.	N.A.
800728	A power failure caused the system to go down.	Fixed
800730	844 disk unit 1 again hung with a negative voltage fault.	N.A.
800731	The system went down twice to allow the CDC Engineer to work on a problem with A 6676 multiplexer. The problem also was responsible for making 1/4 of the system unavailable for half the day.	Fixed
800807	The scopes went blank for unknown reasons. After deadstart, the SCR showed many PPU parity errors and a CM SECEDED error. It is suspected that the CDC Engineer who was pulling cables from under the floor, caused the problem.	DDT-13

//////////

TELEX and TELEX PDP11 Crash Analysis (6/16-8/7) - by D. W. Mears

6/11	15:50	The link on the 172 failed to send any data after TELEX requested a mode 1 read.
6/19	0800-0945	After installing an FCO in our Rianda mux, the link on the 172 seemed to stop working. The actual problem was an error in the Rianda FCO which caused the mux to prevent the link from getting access to the UNIBUS. We had to remove the new Rianda from the system and consequently run with only 16 of our 24 1200 baud ports for the rest of the day.
6/23	12:45	Same as 6/11 15:50.