technical memorandum Daresbury Laboratory

DL/SCI/TM16A

USER GUIDE TO THE SRS DATA LOGGING FACILITY.
I. STORING AND ARCHIVING DATA

by

B. E. TYSON, Daresbury Laboratory

FEBRUARY 1979

Science Research Council

Daresbury Laboratory

Daresbury, Warrington WA4 4AD

The state of the SRS is recorded every two minutes, thus providing a detailed History of its parameters.

Recording of History is done via the SRS Computer Network. This consists of a Master Computer, an Interdata 7/32, and three Minicomputers, Interdata 7/16s. Each of the Minicomputers controls one of the accelerators, Linac, Booster and Storage Ring. The Master Computer is connected to the Central Computer, an IBM 370/165, for jobs where greater computing power and storage are required.

The Master Computer has a total of 20 Megabytes of fixed and movable disc space but only about 5 Megabytes are available for History storage. The Minicomputers have no storage facilities.

2. HISTORY FILING SYSTEM

Fig. 1 shows the files used for the storage of SRS History records and their distribution throughout the Computer Network.

There are two sets of disc files, one on the Master Computer, and one on the Central Computer. The Central Computer also supports tape files. The disc files are used for temporary storage of the most recent data. This is the data the user is most likely to want in a hurry. The tape files are used for permanent storage and retrieving data from these files may take a little time while the tapes are being mounted.

The files chosen are sequential access files to allow the records to be formatted into large blocks for the most efficient storage. The Master Computer files contain one hour and the Central Computer files contain four hours of History. Thus no file contains more than four hours of History, so that when a search for History is initiated, the file on which the required History resides can be rapidly identified.

The dataset names for the tape files have been chosen to contain date and shift information. Thus, knowing the date and the time of the record required, the user can readily determine the dataset name of the

2.1 Master Computer files

There are nine History files in the Master Computer, named DATA:HISTORY.DAT and DATA:HISTORY2.DAT to DATA:HISTORY9.DAT. Each file can contain a maximum of just over one hour of History when all the SRS parameters are being recorded. The data is logged on file DATA:HISTORY.DAT. Once every hour DATA:HISTORY9.DAT, containing the oldest data, is deleted. The remaining files from DATA:HISTORY2.DAT upwards are renamed to the next highest number, e.g. DATA:HISTORY8.DAT becomes DATA:HISTORY9.DAT etc., and DATA:HISTORY.DAT becomes DATA:HISTORY2.DAT. A new DATA:HISTORY.DAT file is created. Thus a minimum of eight hours of History is normally available on the Master Computer discs.

2.2 Central Computer files

Every hour, the latest History is copied, via the network, to the file SRH.HISTORY in the Central Computer. This is a temporary file which can contain a large period of History. There are also nine permanent files on the Central Computer, SRH.HISTORY1 to SRH.HISTORY9. Data from SRH.HISTORY is then edited and copied to SRH.HISTORY1. When this file is filled with four hours of History, it is copied to Backup and Archive tape files. Any remaining History records in SRH.HISTORY are copied to SRH.TEMPHIST.

The nine History files on the Central Computer are renamed every four hours in a manner similar to that described for those in the Master Computer. SRH.HISTORY is deleted, and SRH.TEMPHIST is renamed SRH.HISTORY. This in effect removes the History records that have already been edited and copied from the front of SRH.HISTORY. A new SRH.TEMPHIST is created. The Central Computer filing system is then ready to receive the next four hours of History records from the Master Computer.

3. HISTORY STORAGE ON THE MASTER COMPUTER

The Interdata 7/32 has been named the Master Computer because it contains the master scheduling task SCHEDULE as shown in fig. 2. This task

controls all the data logging and transfer operations. To do this, SCHEDULE 'runs' two tasks in the Master Computer at selected time intervals

(a) MONITOR - Data Logger
(b) HISTMOVE - History File Rename Program

3.1 MONITOR

MONITOR is a task in the Master Computer scheduled every two minutes. This task polls each Minicomputer in turn to gather the SRS data.

Data is gathered in blocks of 32 parameters at a time. The database, DATA:DATABASE.DAT which contains normalisation factors is accessed by PROCESIO and a parameter table is made up for each of the 32 parameters. This parameter table contains the calibration data for the parameters. The analogue readings are calibrated. These calibrated values are combined with the status readings from the Minicomputers, and the parameter names and units from the parameter table, to form a record. This record is stored on the disc file DATA:HISTORY.DAT.

MONITOR also records the data and time at which the parameter values were read. This information, together with file pointers is formed into header records. These records are stored on DATA:HISTORY.DAT in association with the parameter records collected at the same time.

If the network to any of the Minicomputers fails, or if DATA:DATABASE.DAT cannot be accessed then the SRS operators are informed via the operator consoles. The loss of one or two scans of data per day is acceptable.

If contact between any of the Minicomputers and the Master Computer is lost, then History assumes that the parameters affected remain static at the last transferred values.

3.2 HISTMOVE

Every hour HISTMOVE renames all the History files in the Master Computer as indicated in section 2.1 above and then increments a file

name, which is the content of DATA:TRANSFER.DAT, by one to a maximum of nine.

The data is now in a state ready to be transferred to the Central Computer. HISTMOVE then starts a batch job, SRHHSTMV, in the Central Computer via the Network to perform the History transfer operation.

4. TRANSFER OF THE HISTORY TO THE CENTRAL COMPUTER

Efficient transfer is maintained between the Master Computer and the Central Computer even if there is a break in transfer of one file or even if contact is lost for up to 8 hours.

As shown in fig. 3, normally only DATA:HISTORY2.DAT in the Master Computer is transferred to the Central Computer. If the transfer fails, then the next time HISTMOVE is run in the Master Computer, it should tell the Central Computer that there are two files to transfer, i.e. DATA:HISTORY3.DAT and DATA:HISTORY2.DAT. Each time the transfer fails the number of files to be transferred is incremented by one. This is the reason why HISTMOVE increments the DATA:TRANSFER.DAT file as described above (see section 3.2). Thus, although HISTMOVE tries to start SRHHSTMV every hour, only one successful attempt in eight is necessary to ensure that all the History is transferred to the Central Computer.

The JCL for SRHHSTMV is located in file SRH.CNTL(HISTMOVE) and has the following jobsteps:

(a) IEBCOMPR - O.S. Utility
(b) RECOVERY - History File Recovery Program
(c) HISTPASS - History File Transfer Program
(d) HISTEDIT - History File Edit Program
(e) ARKEDIT - History File Archive Initiate Program

We shall ignore the first two steps at present as these are not directly involved in the transfer of the History files (but see section 6).

4.1 HISTPASS

HISTPASS sets up a call to the Master Computer to read the

DATA:TRANSFER.DAT file. This file contains the name of the first History file that needs to be transferred to the Central Computer. HISTPASS then sets up a call to the Master Computer to read this file. The file is read and the records are stored in the temporary file SRH.HISTORY. When this is finished, the file number is decremented by one, and the program opens a call to the Master Computer to overwrite DATA:TRANSFER.DAT with this information. The program repeats the History file transfers until the file number is less than 2, or the transfer process fails.

In the event of a failure, the Master Computer knows how many file transfers have been successful. If only part of a file is transferred, i.e. the DATA:TRANSFER.DAT file was not updated in the Master Computer, then the whole file will be transferred again in one hour's time, or later. This may result in the data appearing twice on the Central Computer SRH.HISTORY file.

4.2 HISTEDIT

To eliminate redundant data, the History stored on the buffer file SRH.HISTORY is edited before being passed to more permanent user files. This editing is done by the jobstep HISTEDIT. This cuts the number of History records to between one half and one third of the total number recorded.

HISTEDIT is executed after the History file transfers from the Master Computer have been successfully completed by HISTPASS. It scans through SRH.HISTORY reading in the header records for each time period, and all the parameter records associated with that time. If some of these records have been lost in transit, this can be detected. The program then compares the time for these records with the time for the previous set of records.

If these records follow the previous set in sequence, then they are adjudged to be good. Each record in turn is compared with the previous record for that parameter, if any. If a change has occurred in either analogue or status value, then this record is marked for copying. When all the records have been checked, the file pointers in the header records are updated to account for any changes. The header and parameter

records are then copied to the permanent file SRH.HISTORY1.

If, however, the time sequence is found to be in error, then the whole of this set of records is discarded, and the program proceeds to read and check the next set of records.

HISTEDIT only edits the History for a selected four hour period. When the records in SRH.HISTORY go beyond this period, the remainder are copied directly to the file SRH.TEMPHIST. HISTEDIT then sets a flag to say that SRH.HISTORY1 is full. This is done by returning a completion code of 500.

4.3 ARCHEDIT

ARKEDIT modifies the JCL, SRH.ARCH, used to run the job, SRHARKIV, and then starts this job. ARKEDIT only runs if the return code from the HISTEDIT step was 500.

The program reads the first record of SRH.HISTORY1 to extract the date and time. The date and time are then processed to make up a tape file dataset name.

Records from SRH.ARCH are read, and the JCL is altered to contain the new dataset names. Tape labels are incremented, and the records in SRH.ARCH are overwritten with the new information. The program then issues a 'jobstart' to SRH.ARCH to start the job SRHARKIV.

When the label number reaches the maximum number of files allowed on the tape (see section 7.1 below), then ARKEDIT chooses the next free tape in the Master Catalogue and sets up the JCL in SRH.ARCH to use this tape.

5. TRANSFERRING HISTORY TO TAPES

With only nine four hour disc files on the Central Computer one can only store a maximum of 36 hours of History. Thus one cannot recover data that is more than a day and a half old. This presents us with a serious limitation. These files may possibly be doubled to give about two and a half days History. This does not, however, allow the user to recall a set of SRS parameters that were recorded, say, a week previously.

The problem has been overcome by saving permanent History records on tapes. Each tape will be able to contain about one month of edited SRS History. For greater integrity it has been decided to keep two copies of SRS History on separate tapes.

As shown in fig. 4, SRHARKIV performs the transfer of SRS History, held on the Central Computer disc files, to tape files. The JCL for this job is located in the dataset SRH.ARCH. SRHARKIV and has the following jobsteps:

(a) IEBCOMPR - O.S. Utility
(b) IEBGENER - O.S. Utility
(c) IEBGENER - O.S. Utility
(d) CATLGING - Tape File Cataloguing Program
(e) IEBGENER - O.S. Utility
(f) IEBGENER - O.S. Utility
(g) IEHPROGM - O.S. Utility
(h) SOURCE - SRS Utility
(i) IEFBR14 - O.S. Utility

We shall ignore the first step at present as it is not directly involved in the transfer of the History files (but see section 6).

5.1 IEBGENER

There are two jobsteps, (b) and (c), that use the O.S. Utility Program IEBGENER. These two jobsteps copy the disc file SRH.HISTORY1 to the Archive and Backup tapes respectively. The names of these tapes are not important. It is intended that one of these tapes should be stored away from the Central Computer in case of accidents.

5.2 CATLGING

If the two previous IEBGENER jobsteps have been successful, then the JCL in dataset SRH.ARCH.SPHARKIV containing the tape file dataset names and the tape labels, is copied into two tape file catalogues, SRH.CATLGA, and SRH.CATLGB for Archive and for Backup respectively. The program also updates the Master tape catalogue SRH.TAPES.DATA.

A more detailed description of the tape catalogue system will be described below (see section 7).

5.3 IEBGENER

The jobsteps, (e) and (f), use the O.S. Utility program IEBGENER to copy the Working Catalogues, SRH.CATLGA and SRH.CATLGB (see section 7.3), to the SRS disc pack. These are the permanent tape catalogues. This copy also helps as a backup facility.

5.4 IEHPROGM

As with the files in the Master Computer, the disc files on the Central Computer must be rotated, with the oldest data being deleted. This is performed by the RENAME facility of the O.S. Utility IEHPROGM.

SRH.HISTORY9 is deleted, and files SRH.HISTORY1 to SRH.HISTORY8 are renamed to increment the number by one to give the files SRH.HISTORY2 to SRH.HISTORY9. A new SRH.HISTORY1 is created.

History that has been copied from SRH.HISTORY to SRH.HISTORY1 must also be deleted, otherwise SRH.HISTORY will keep on growing to fill all its extents. SRH.HISTORY is deleted, and SRH.TEMPHIST is renamed SRH.HISTORY. A new SRH.TEMPHIST is created.

5.5 SOURCE

SOURCE is used to copy the updated Master Catalogue file, SRH.TAPES.DATA, to the Master Computer file DATA:TAPELIST.DAT. This copy is needed when recovering History from tapes via the Master Computer.

5.6 IEFBR14

This utility creates the tape catalogue names for the next set of tapes required by this service. This has been done in separate jobstep as this only needs to be created once. After the name has been successfully catalogued this jobstep will fail. As it is the last step in SRHARKIV job it will not affect the preceeding jobsteps.

6. JOB INTEGRITY

In sections 4 and 5 certain jobstep descriptions were ignored, i.e. those using the O.S. Utility IEBCOMPR. The reason for the inclusion of these jobsteps is described in this section.

When using tape storage on the Central Computer, some delays are inevitable in running a job. The SRHARKIV job may still be on the input queue when the next SRHHSTMV job is started by HISTMOVE in the Master Computer. This could lead to file conflict, especially as the disc files have not been renamed. To overcome this problem a set of message files has been created to allow the jobs to communicate with each other.

On entering the job SRHHSTMV the first jobstep checks the messages to see if SRHARKIV has been completed successfully. This allows jobsteps HISTEDIT and ARKEDIT to be executed. An error message, or no message from SRHARKIV forces the RECOVERY step to operate. This should restore the system to a useable state. The jobstep HISTPASS is run regardless of the state of the messages from SRHARKIV. It may, however, fail because there is no SRH.HISTORY file present. This is not serious, as we have seen this has been taken care of by the way we handle the History transfer.

Similarly SRHARKIV checks the messages from SRHHSTMV to see if it should continue. There may be instances where the last SRHARKIV has not run before the next SRHHSTMV is started. In this case SRHHSTMV will execute a RECOVERY action and jobstart another SRHARKIV. The first job on the queue will run satisfactorily, but the second will fail because it will now receive error messages to the effect that it has already performed the required service. Thus, two calls to SRHARKIV with the same JCL will not cause the renaming of the files to be repeated.

7. THE SRS TAPE CATALOGUE SYSTEM

The reader may be wondering why one should wish to develop a cataloguing system when there is already one on the Central Computer. The reason is that there is not enough space in the present catalogue to

accommodate the large number of dataset names that will be generated. The system is creating two new files every four hours. This represents a staggering total of 4400 new files per year for a period of about 20 years. This number would swamp the present cataloguing system.

A much simplified cataloguing system has been devised which uses only one disc track per month. The whole catalogue of files for the life of the SRS can be contained on about 250 disc tracks.

The catalogue system consists of a Master Catalogue, and an Individual Catalogue for each tape that has been used. The Master Catalogue contains pointers to all the Catalogues.

7.1 The Master Catalogue

The Master Catalogue is kept in the file SRH.TAPES.DATA. This file contains a complete record of the state of all the tapes that are available to the system.

The record length for this file is 80 bytes. Each record contains all the information pertinent to a single tape. The record layout is shown below.

| Tapeno | Labno | Max | DSN1 | DSNL | Catno | Catname |

where

Tapeno - DNPL Tape Number
Labno - Tape Label Number of Last File
Max - Maximum Number of Files
DSN1 - Dataset Name of First File
DSNL - Dataset Name of Last File
Catno - Catalogue Number
Catname - Catalogue Name

Unused tapes contain only the first three parameters, with Labno set to zero.

Using the dataset names, which relate to the time and date of the first record on the file, a search of this catalogue will give the catalogue name of the tape on which the file can be found. The catalogue number is a number that is part of the catalogue name when the tape has been filled and is no longer being updated. The current tapes hold the Working Catalogue name in the Catname location. The Working Catalogue will be described more fully below.

The Master Catalogue is continually being updated every four hours. Because of this it is kept on a TSO disc pack, where it is backed up regularly. Hence recovery of this file is straightforward. In the event of any mishap the data on the backup file should be up to date.

7.2 Individual Catalogues

Each tape has its Individual Catalogue held on a Central Computer disc pack. The catalogue records are 80 bytes long. Each record represents a four hour period of History.

The records contain a JCL statement that includes the dataset name, tape number, and label number of the file for that period. If no data has been recorded for a given period, then the record contains a statement

NO DATA AVAILABLE FOR THIS SHIFT

Reading the first record of this catalogue gives the date and time of the first file. Knowing the date and the time of the file sought, one can subtract the two figures to give the number of the record containing the JCL of the file from which data is required. Thus the data can be found with the minimum of delay in the searches.

Individual Catalogues are removed to a private disc pack once the tapes are full. The private disc pack is only backed up once a week. This presents no problems as the data on the catalogues is static.

7.3 Working Catalogue

The Working Catalogue is a tape catalogue kept on the TSO disc pack. There are two of these catalogues SRH.CATLGA and SRH.CATLGB for the current Archive and Backup tapes respectively. These are the catalogues referred to in section 5.2.

The Working Catalogues are kept on the TSO disc pack for the same reason that the Master Catalogue is kept on the TSO disc pack, namely to have them regularly backed up. This is necessary as they are being updated every four hours.

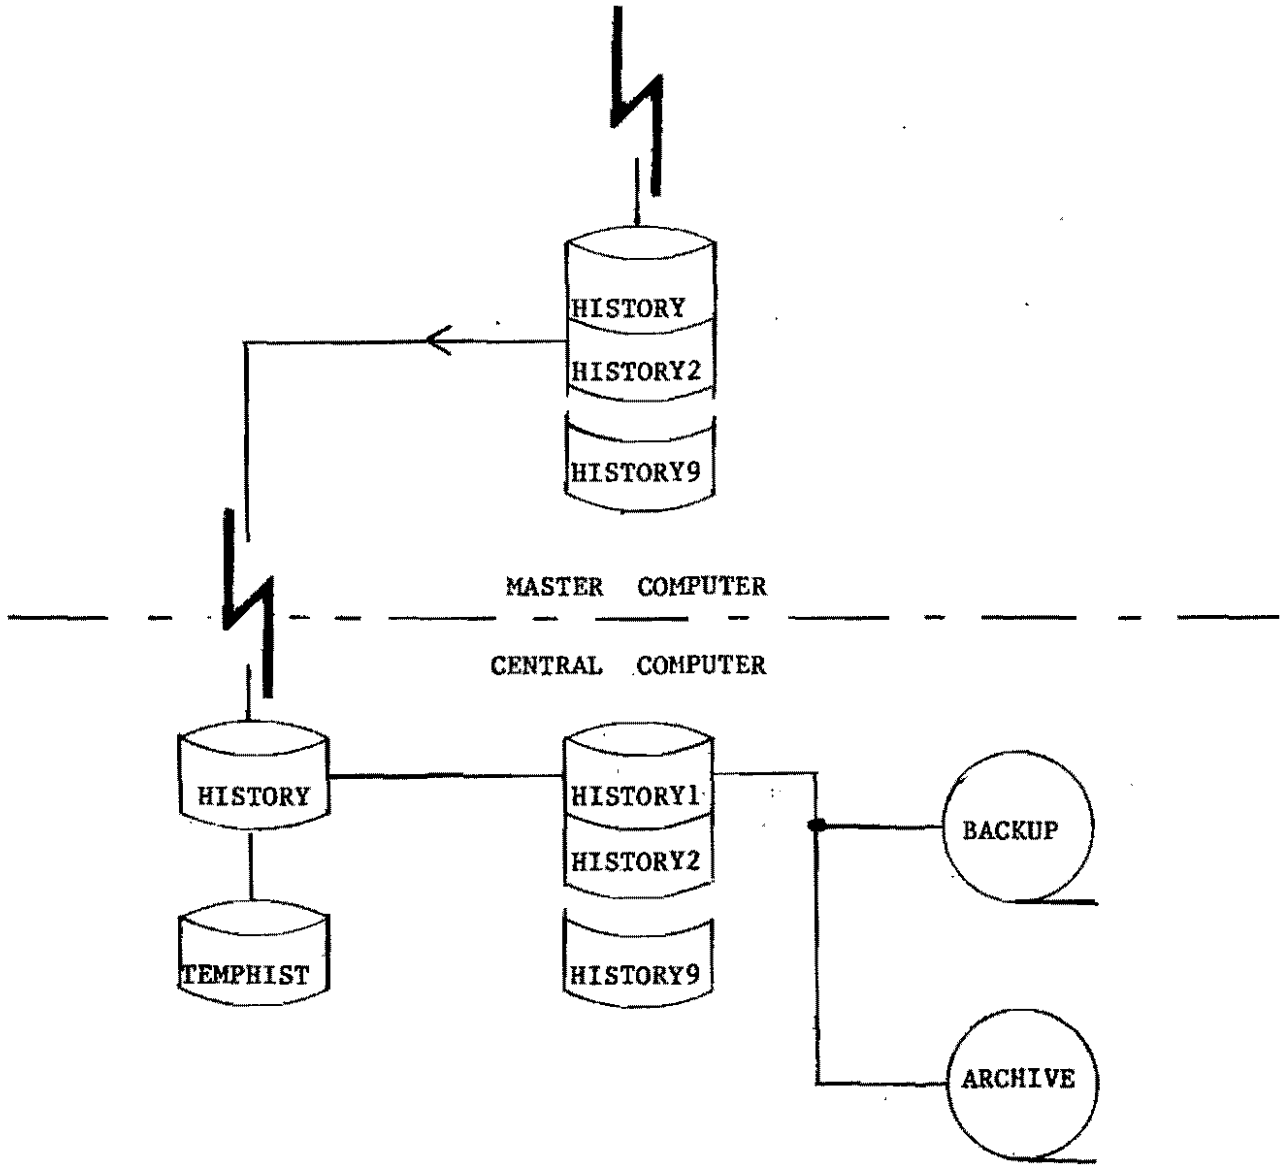
HISTORY
HISTORY2
HISTORY9
MASTER COMPUTER
CENTRAL COMPUTER
HISTORY
TEMPHIST
HISTORY1
HISTORY2
HISTORY9
BACKUP
ARCHIVE

Fig. 1

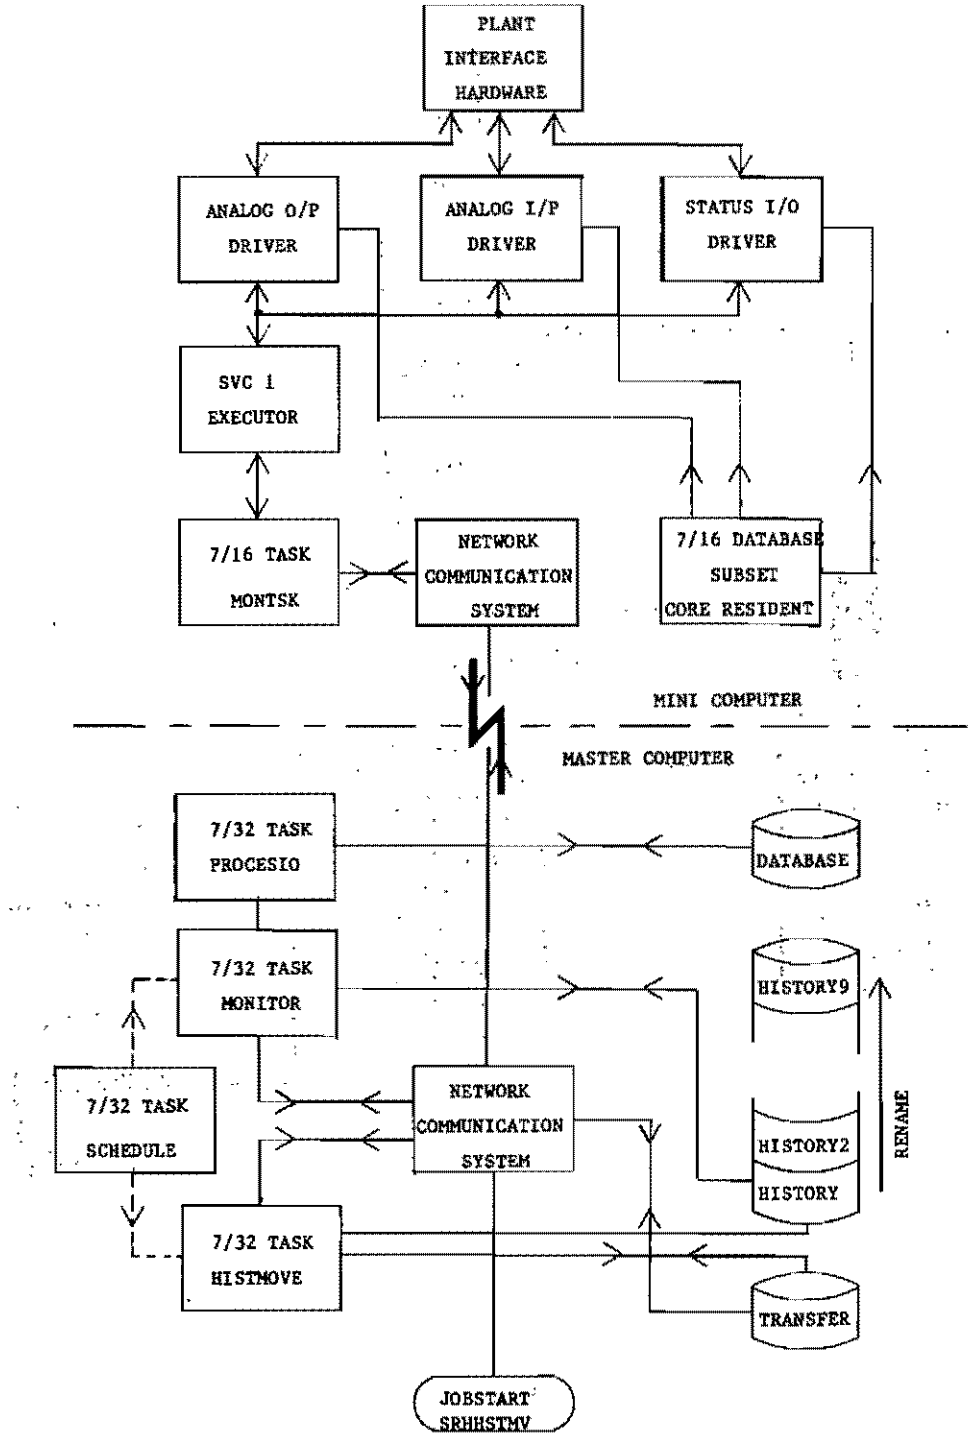
PLANT
INTERFACE
HARDWARE
ANALOG O/P
DRIVER
ANALOG I/P
DRIVER
STATUS I/O
DRIVER
SVC 1
EXECUTOR
7/16 TASK
MONTSK
NETWORK
COMMUNICATION
SYSTEM
7/16 DATABASE
SUBSET
CORE RESIDENT
MINI COMPUTER
MASTER COMPUTER
7/32 TASK
PROCESIO
DATABASE
7/32 TASK
MONITOR
HISTORY9
7/32 TASK
SCHEDULE
NETWORK
COMMUNICATION
SYSTEM
HISTORY2
HISTORY
RENAME
7/32 TASK
HISTMOVE
TRANSFER
JOBSTART
SRHHSTMV

Fig. 2

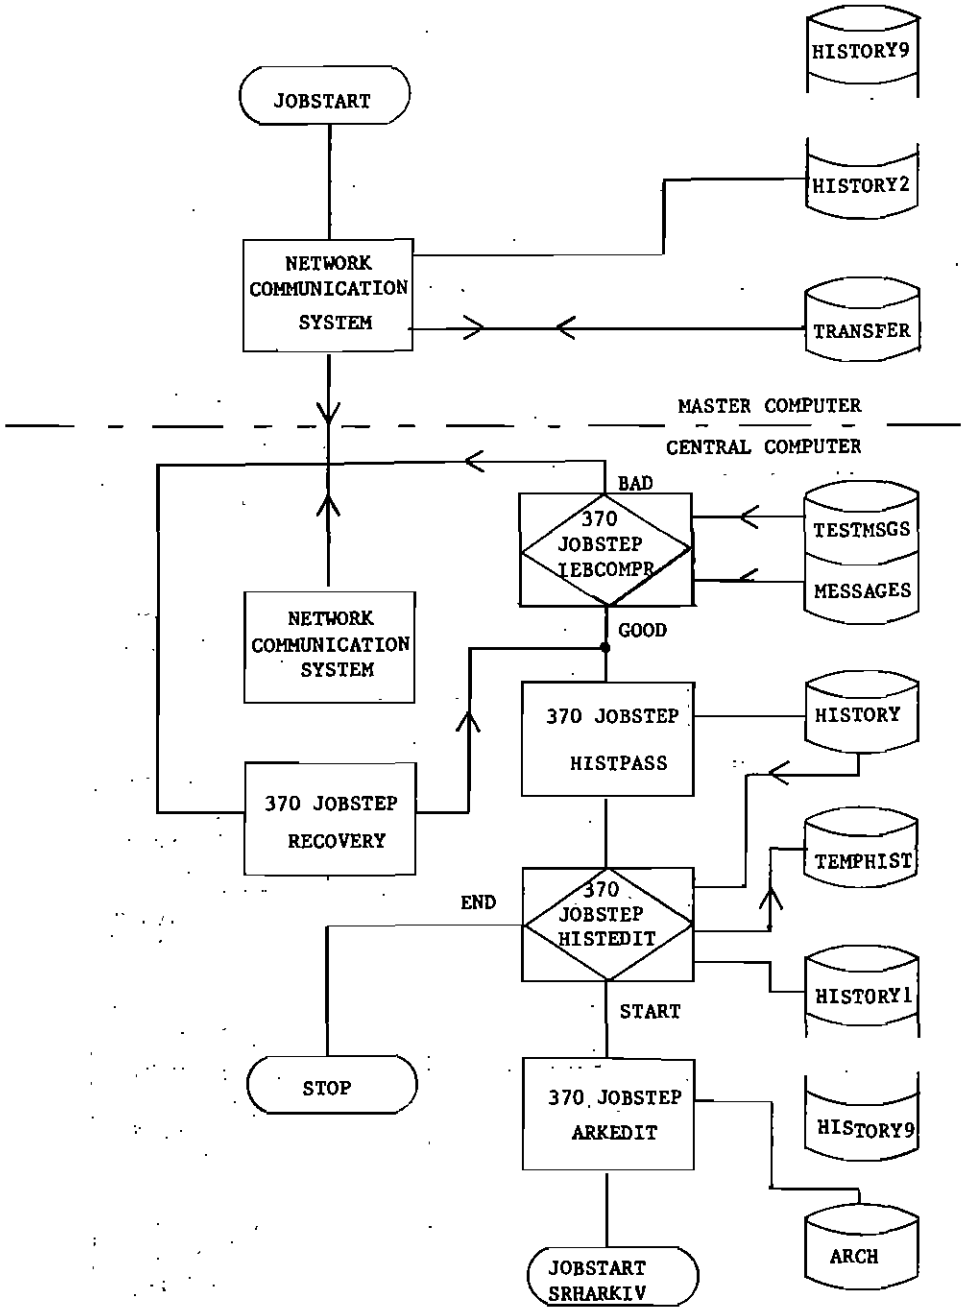
HISTORY9
JOBSTART
HISTORY2
NETWORK COMMUNICATION SYSTEM
TRANSFER
MASTER COMPUTER
CENTRAL COMPUTER
BAD
370 JOBSTEP IEBCOMPR
TESTMSGS
MESSAGES
NETWORK COMMUNICATION SYSTEM
GOOD
370 JOBSTEP HISTPASS
HISTORY
370 JOBSTEP RECOVERY
TEMPHIST
END
370 JOBSTEP HISTEDIT
HISTORY1
START
STOP
370 JOBSTEP ARKEDIT
HISTORY9
ARCH
JOBSTART SRHARKIV

Fig. 3

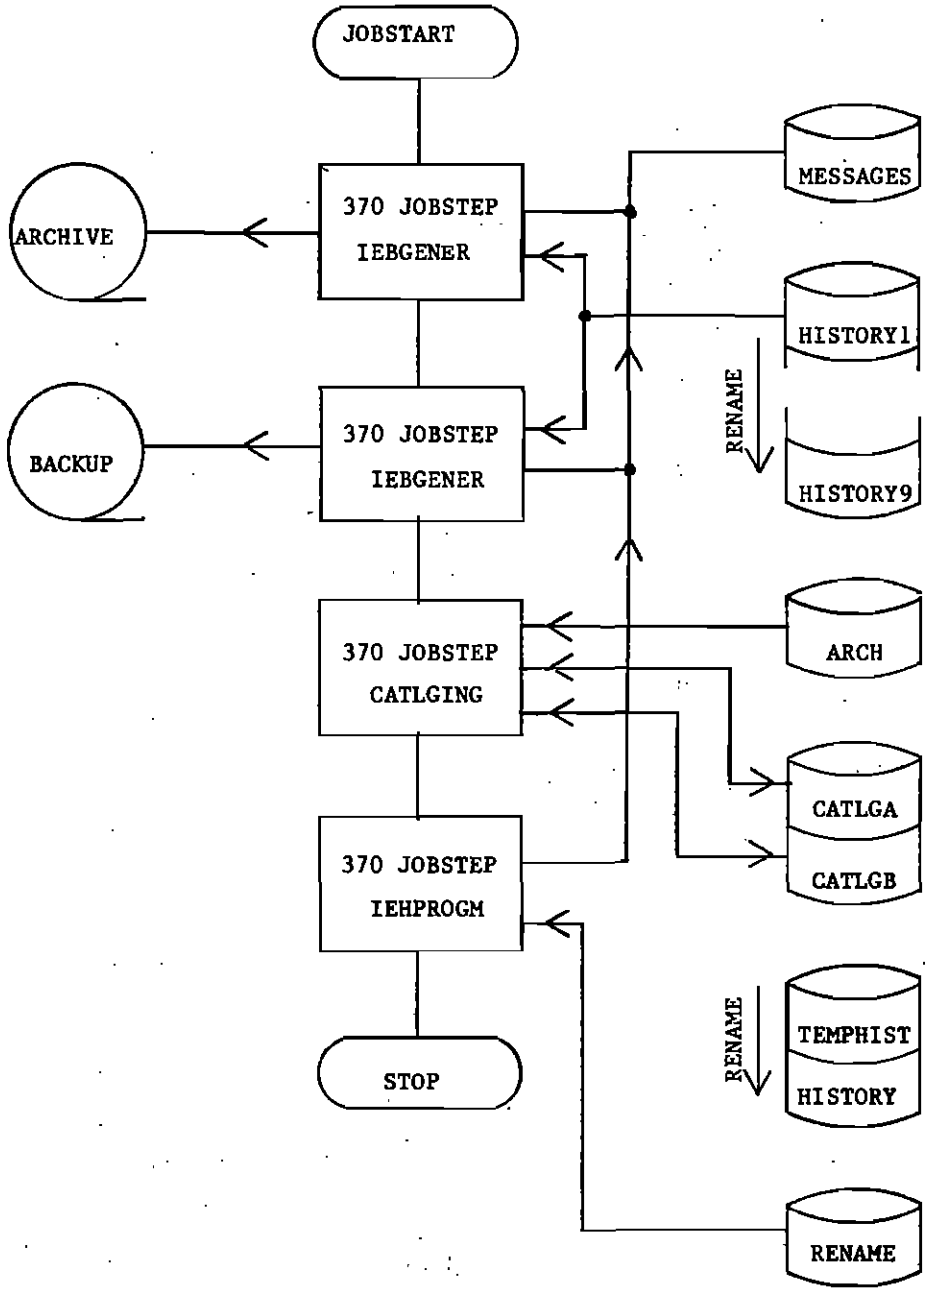
JOBSTART
370 JOBSTEP
IEBGENER
ARCHIVE
MESSAGES
HISTORY1
RENAME
HISTORY9
370 JOBSTEP
IEBGENER
BACKUP
370 JOBSTEP
CATLGING
ARCH
CATLGA
CATLGB
370 JOBSTEP
IEHPROGM
RENAME
TEMPHIST
HISTORY
STOP
RENAME

Fig. 4

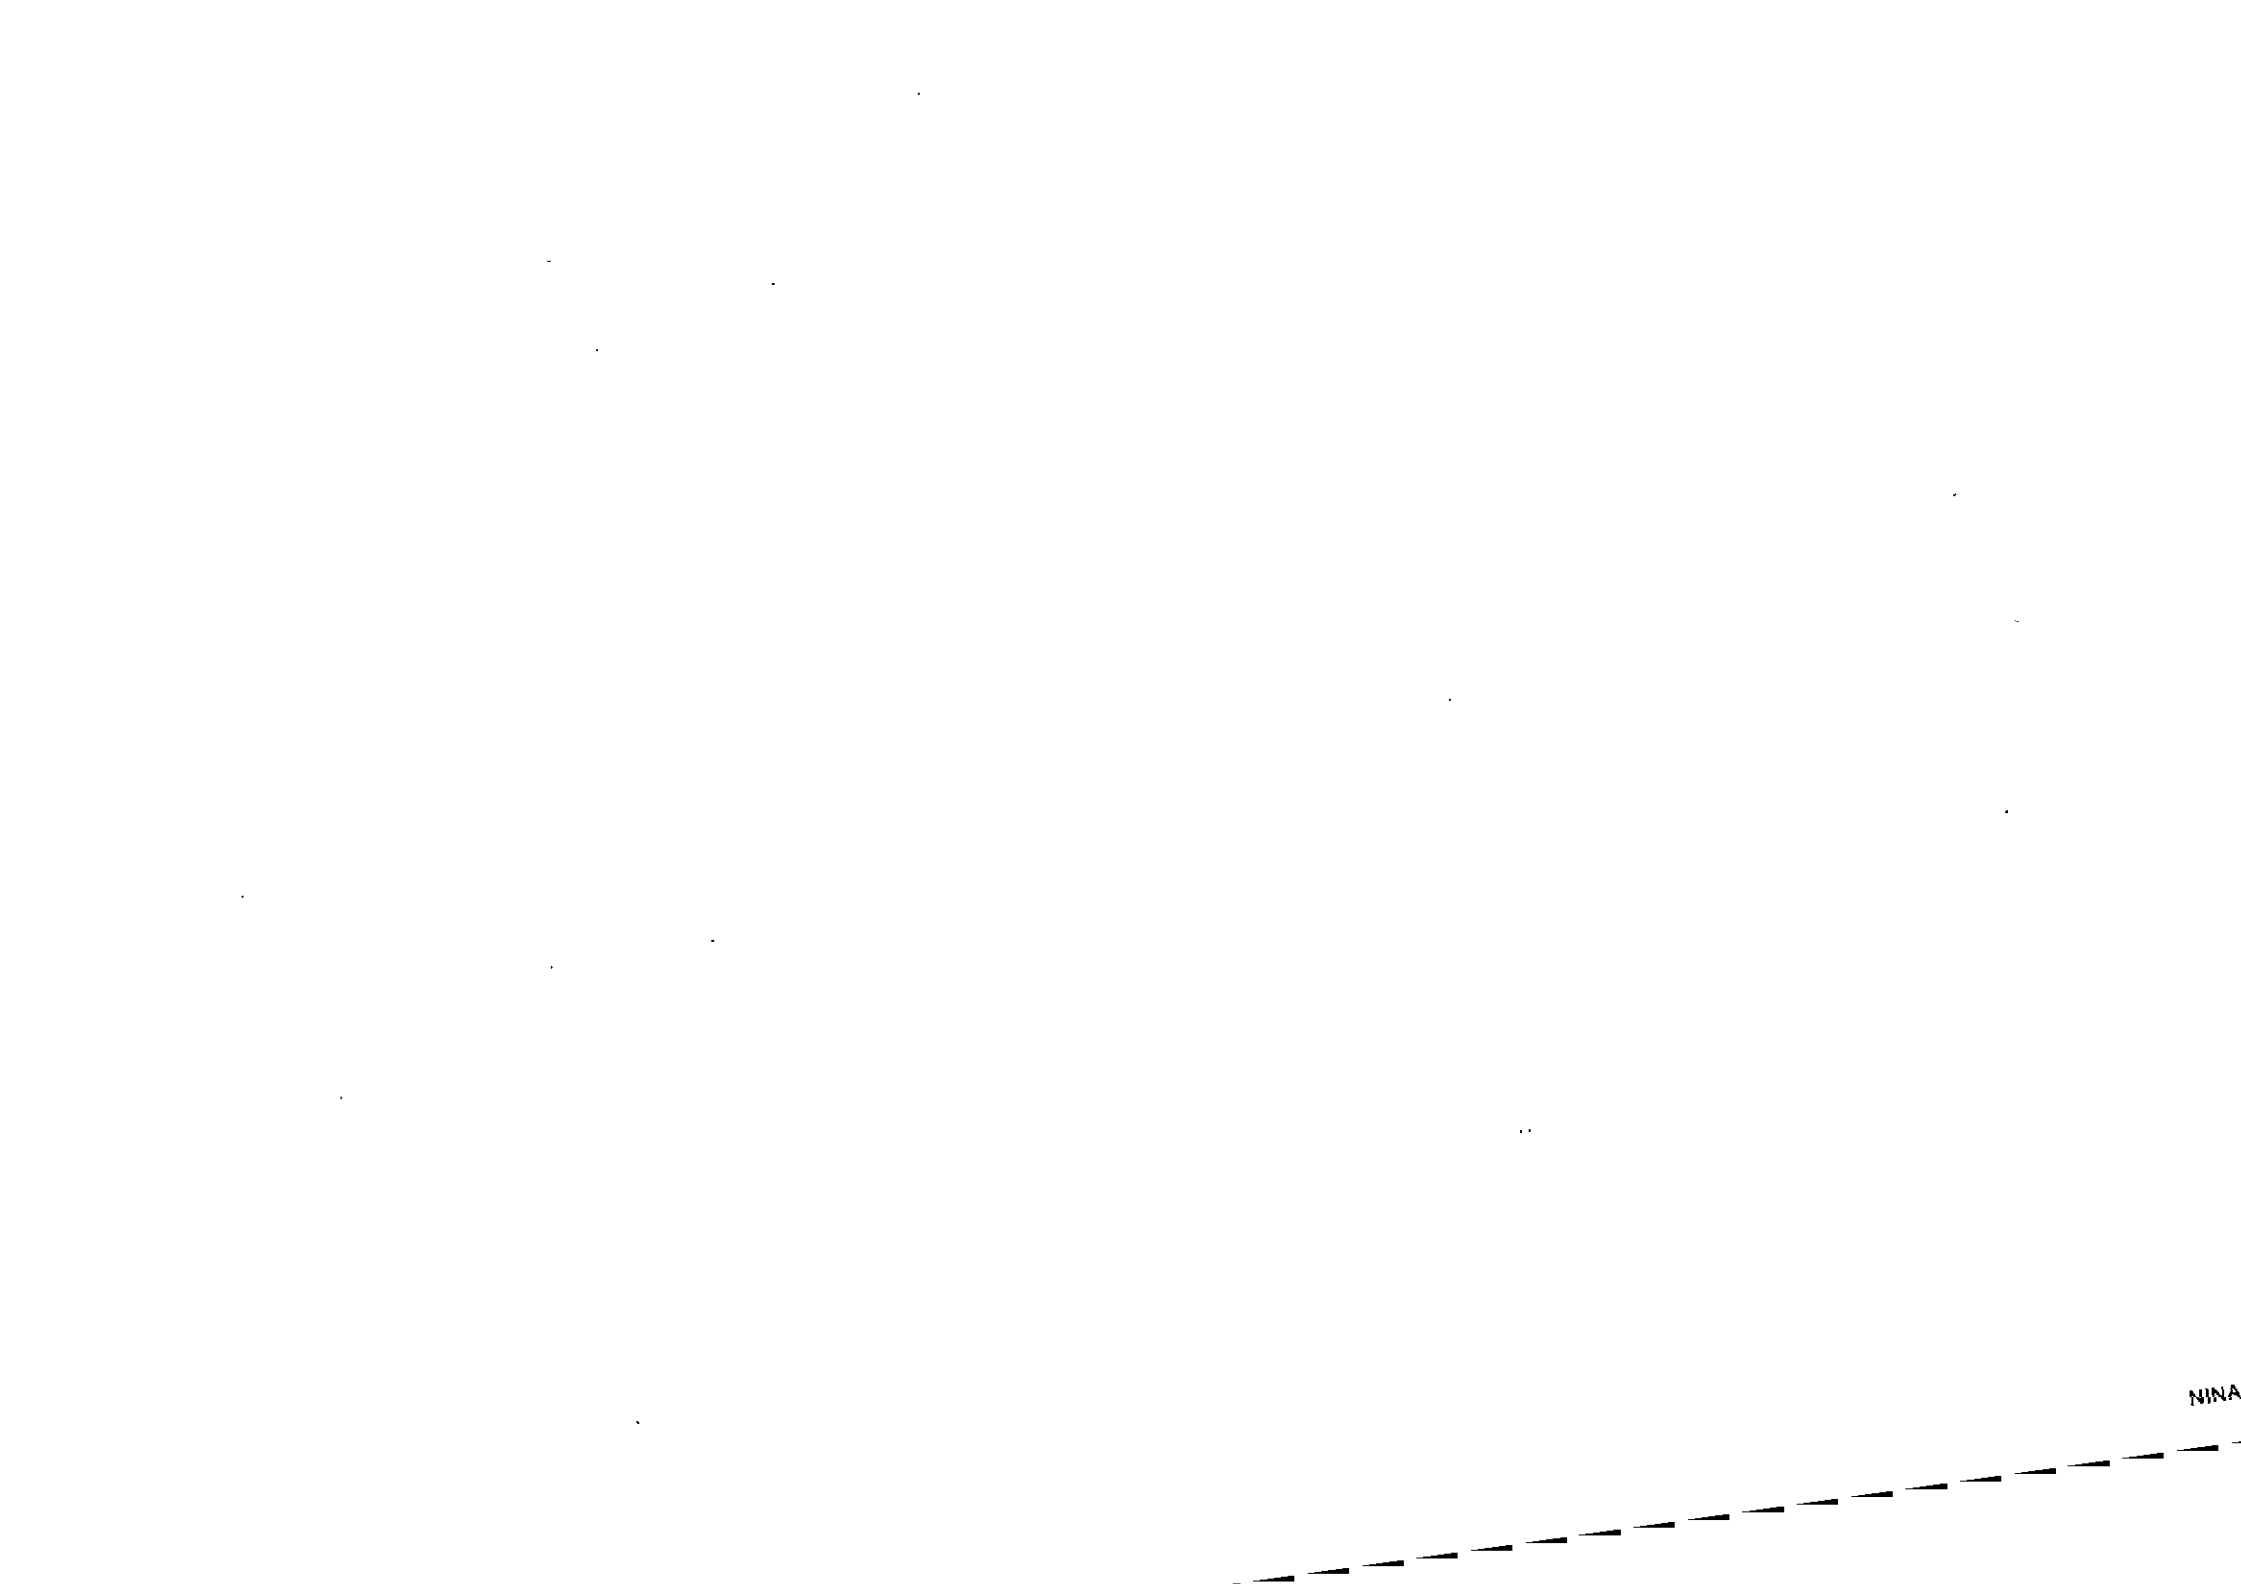
NINA