TECHNICAL MEMORANDUM

106

Current and planned meteorological applications systems at ECMWF

J.K. Gibson

Research Department

October 1985

European Centre for Medium-Range Weather Forecasts
Europäisches Zentrum für mittelfristige Wettervorhersage
Centre européen pour les prévisions météorologiques à moyen

1. INTRODUCTION

The primary function of ECMWF is to produce "Medium Range Weather Forecasts" - numerical products indicating the expected values of meteorological parameters for forecast periods of up to 10 days ahead (see Fig. 1.1). This task poses three problems:

- how to obtain initial data of suitable quality
- how to obtain the forecast values
- how to form and distribute the results of the forecast.

Even if there existed a perfect numerical prediction scheme capable of computing an accurate 10 day forecast, given wrong initial data it would be unlikely to compute correct results. Again, if a perfect prediction model were given correct initial data and computed an accurate forecast, such results would only be useful if presented to the user in a suitable form. Such a suitable form must be correctly derived, or its derivation will introduce error.

Meteorological Applications Systems at ECMWF are the means by which the solution to the above problems is attempted. The ECMWF solution is not perfect. In consequence improvements are constantly being planned, tested, and if satisfactory, added. Nevertheless, the ECMWF systems have been used in daily operations since July 1979. During this period of operational activity forecast and analysis products have been produced each day without exception, and delivered to the ECMWF Member States. This unbroken record confirms the soundness of the overall systems design.

In the following, the current and planned Applications Systems will be introduced.

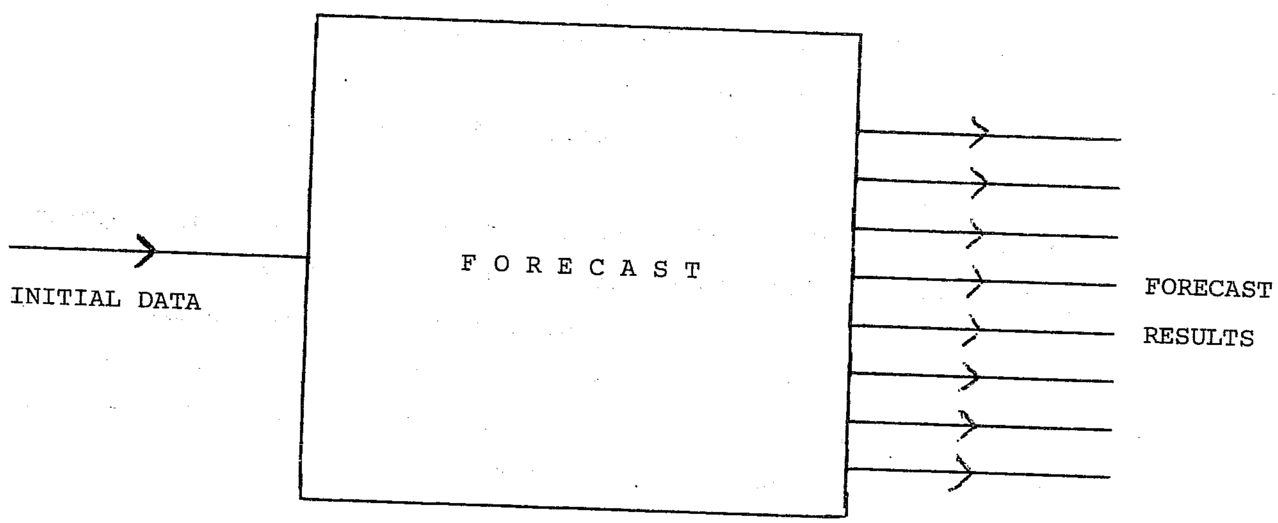

Fig. 1.1 A FORECAST SYSTEM - GROSSLY SIMPLIFIED!

2. A METEOROLOGICAL OPERATIONAL SYSTEM

The Meteorological Applications Systems are a set of systems which together comprise a Meteorological Operational System. The ECMWF Meteorological Operational System is known by its acronym "EMOS". Fig. 2.1 illustrates the major purpose of EMOS - to transform raw observational data into the forecast products required by the user, and to make an archive of data which can be used for later investigations.

To introduce the Applications Systems which are the components of EMOS, Fig. 2.2 illustrates the total Operational System with respect to the flow of data from sub-system to sub-system. This is a useful approach because

- it defines the necessary Applications Systems
- it indicates the required interfaces between systems
- it illustrates the need for several systems to manipulate common sets of data, indicating that data bases can improve efficiency.

Fig. 2.2 defines the following systems:

- Data Acquisition - to receive and protect the raw data
- Pre-processing - to decode and check the raw data, creating the Reports Data Base (RDB)
- Analysis - to analyse the reports and compute initial conditions for the forecast
- Forecast - to integrate the initial conditions with respect to time
- Post-processing - to manipulate the raw data produced by the forecast, and to create the Fields Data Base (FDB)
- Dissemination - to extract required products from the FDB and send them to the users in the required format
- Archive/Retrieval - to archive reports, fields, and other appropriate data, and to provide the means by which such data may be retrieved
- Operational Watch - to enable raw data to be corrected; to inspect results; to monitor data and processes, and to assess the total system performance.

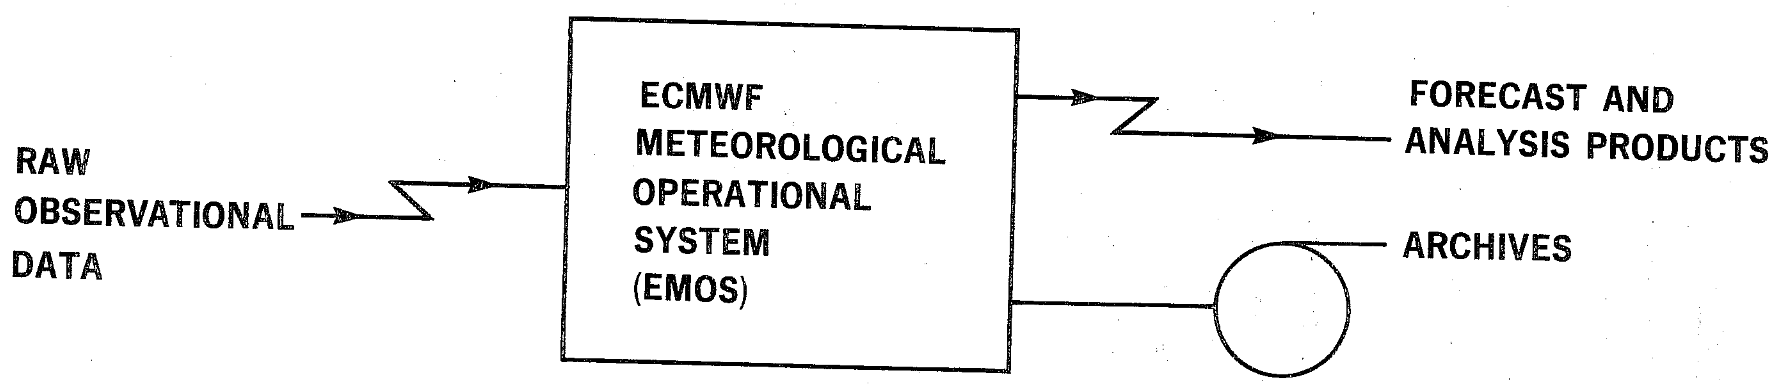

Fig. 2.1 EMOS Overview

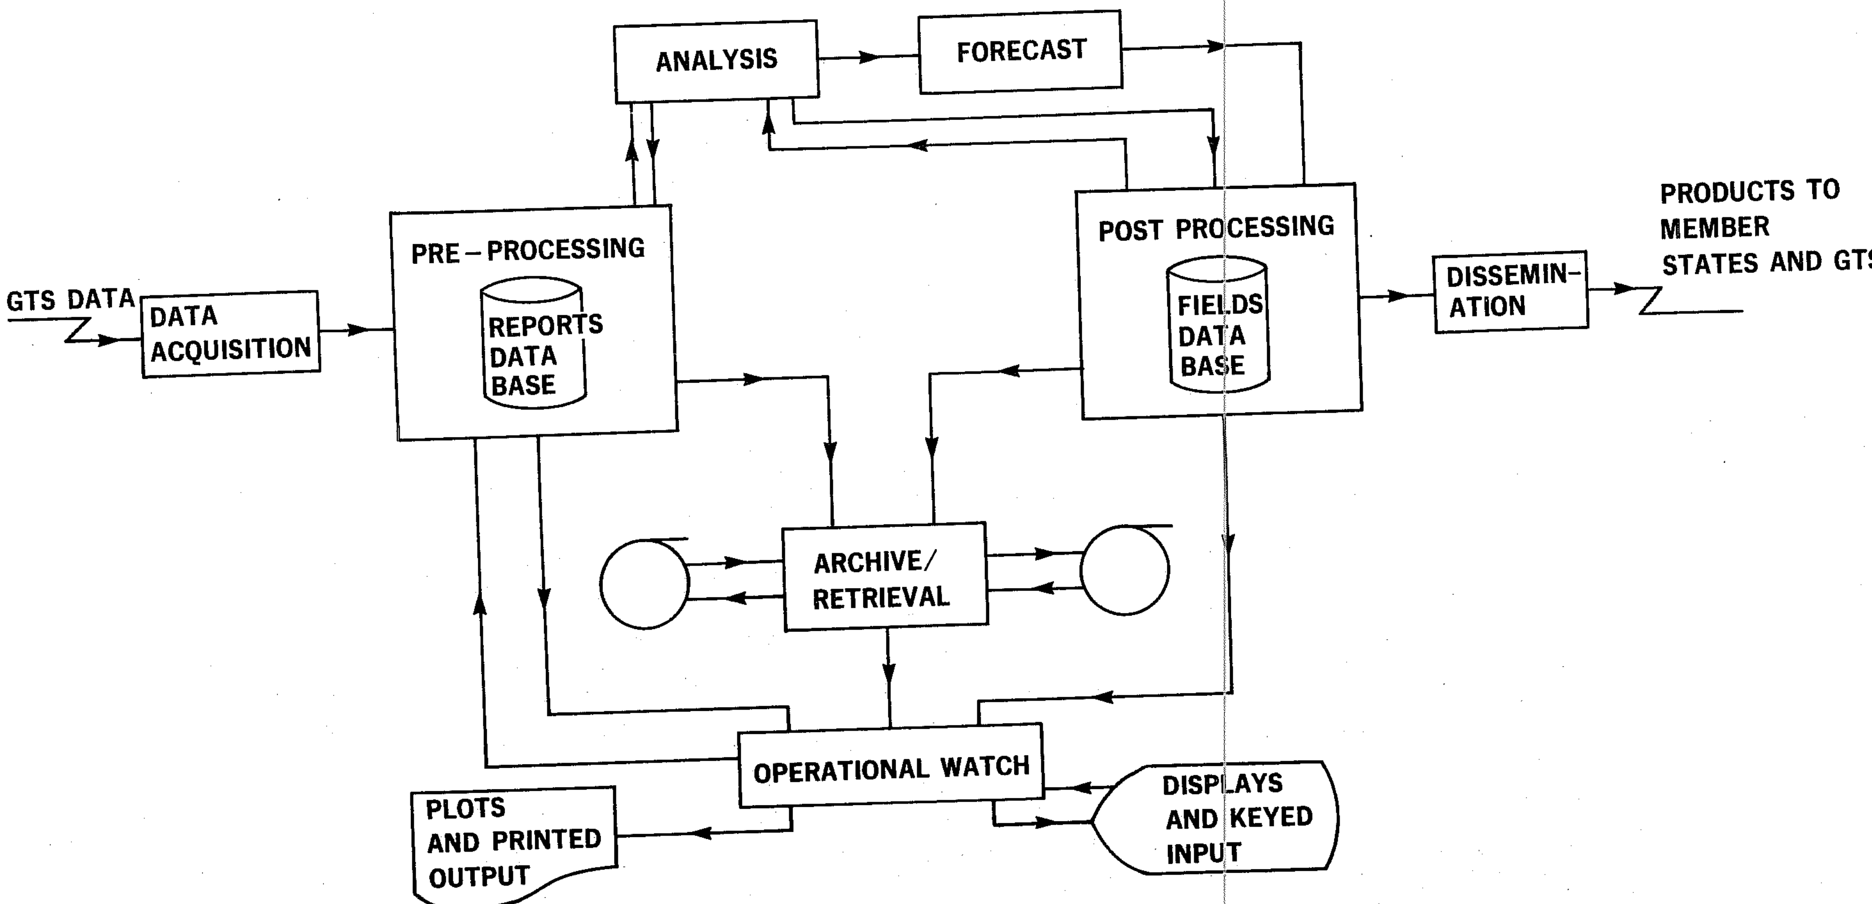

Fig. 2.2 ECMWF Meteorological Data Flow

A further system, not illustrated, concerns the need to control the processes involved. This is called the "Supervisor-Monitor-Scheduler" (SMS), because it provides an automated means of supervising the running of EMOS, of monitoring the progress of the individual components, and of scheduling the processing of the computer jobs within each system.

Three concentrations of data are indicated in Fig. 2.2:

- the RDB
- the FDB
- the archive.

Each of these represents a set of data required to be accessed frequently and efficiently. EMOS uses simple data base and data access methods designed to achieve the efficiency required while taking into account the characteristics of the data involved.

3. EMOS DATA BASES - CURRENT STATUS

3.1 Introduction

The EMOS system was originally designed for a computer configuration, illustrated in Fig. 3.1, comprising

- a CRAY-1A for vector processing
- a CYBER 175 for general purpose computing
- a RC 8000 telecommunications system, later upgraded to a twin system.

When the CRAY-1A was first installed in 1978, the CRAY Operating System (COS) had not become sufficiently mature to be entrusted with the medium term retention of files within a data base. Additionally, all plotting, diagnostics and verification were performed on the CYBER 175. Thus the two on-line data bases (RDB and FDB) and the archive were developed for access using the CYBER front-end.

3.2 The Reports Data Base

The RDB is used to retain on-line observations which have been decoded into a defined data representation covering a period of 3 days. Fig. 3.2 shows the relationship between the RDB and the processes which manipulate reports.

The design of the data base took account of the data required to be stored, the predicted access pattern, the requirements of the systems using the data base, and processing efficiency. As a result of these considerations it was decided to use

- a packed data format
- 7 major observation types
- one file per observation type per day
- index sequential file organisation
- record keys enabling access at the single observation level
- supplementary files to record reception details, and to contain WMO Volume A station details.

In practice the data base has been extremely successful, and has fulfilled all of the original design aims. It is fed by the pre-processing sub-system, provides data for the Analysis, Operational Watch, Archives, and user jobs of a general nature. The data base statistics enable data reception to be monitored in real time. This enables decisions to be made concerning the early or late running of analyses. Quality control information is fed into the data base by the pre-processing and the analysis, in the form of flags and substituted values. Purpose written data base software enables users to access data in the original, packed form, or in fully expanded form (Meteorological Data Base Access, or MEDABA).

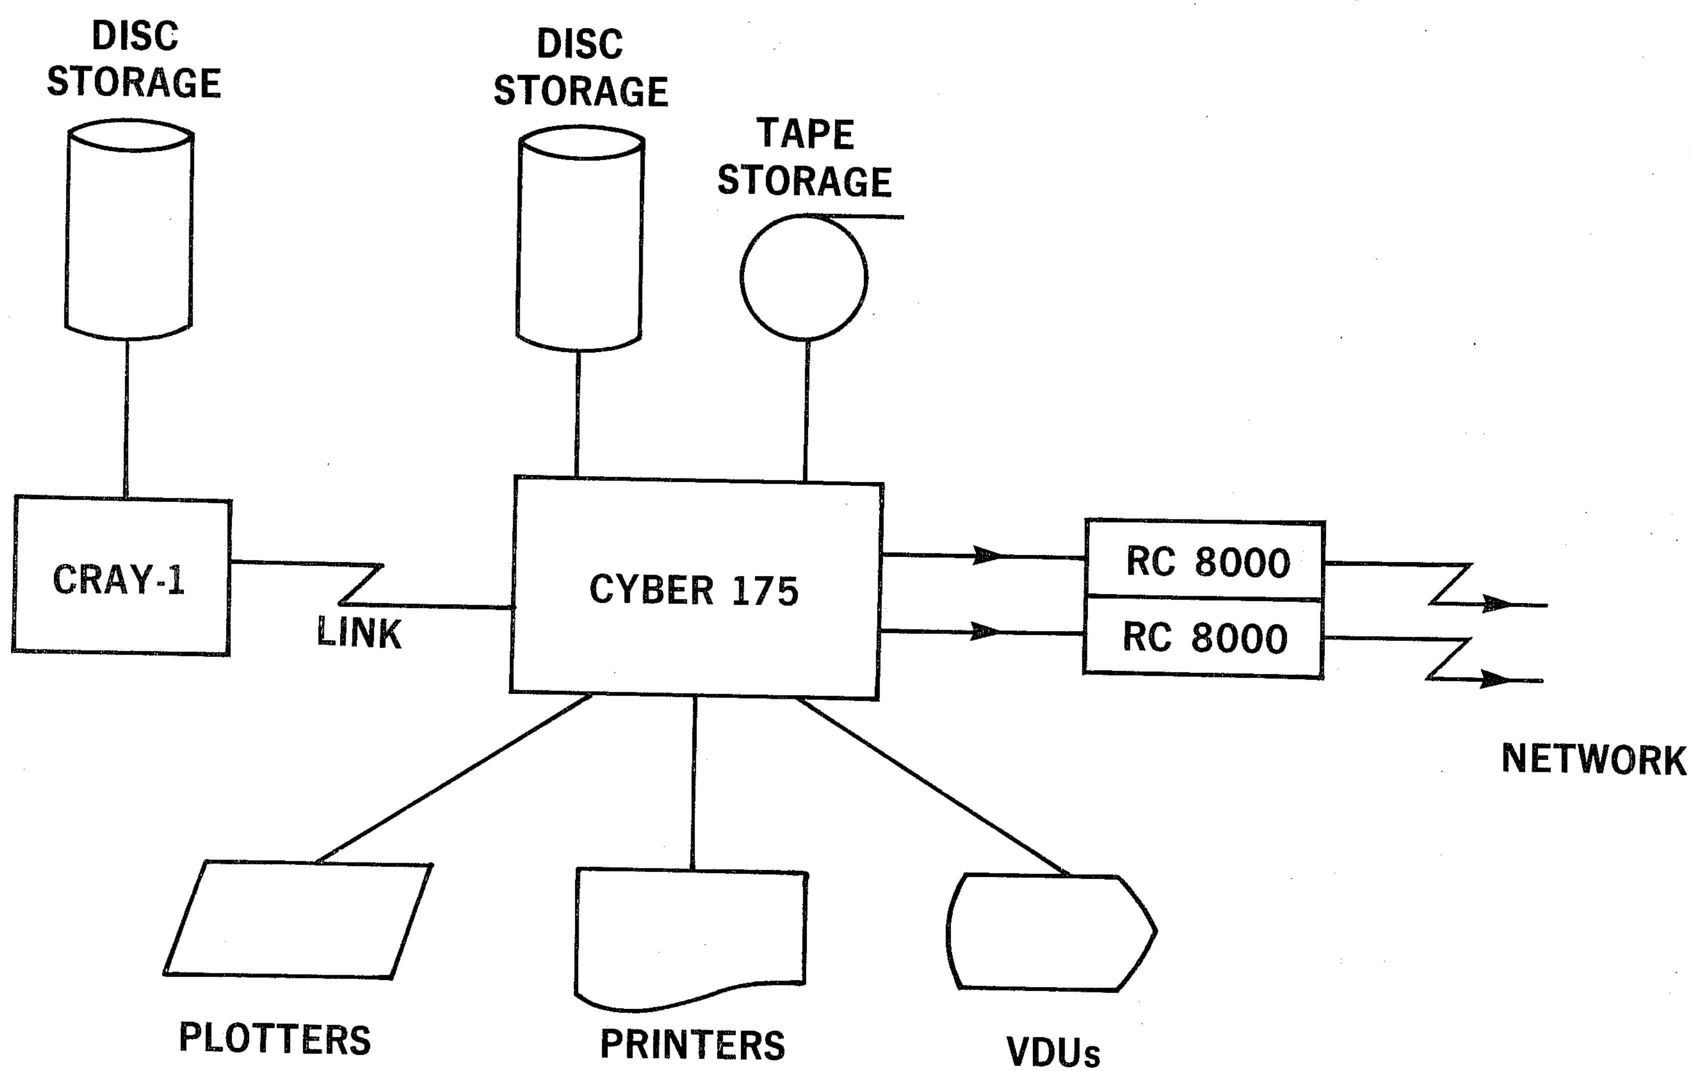

Fig. 3.1 First ECMWF Computer System

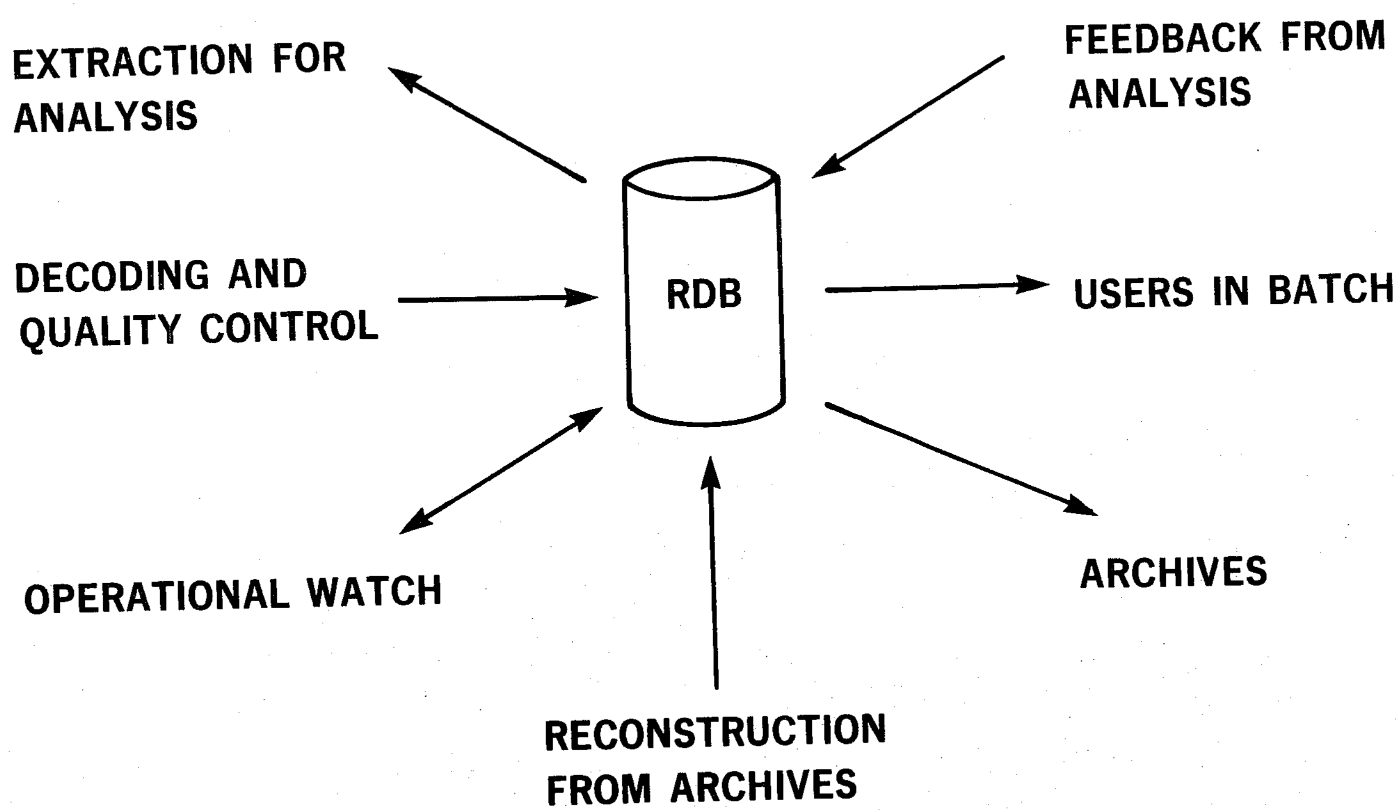

Fig. 3.2 Processes Using RDB

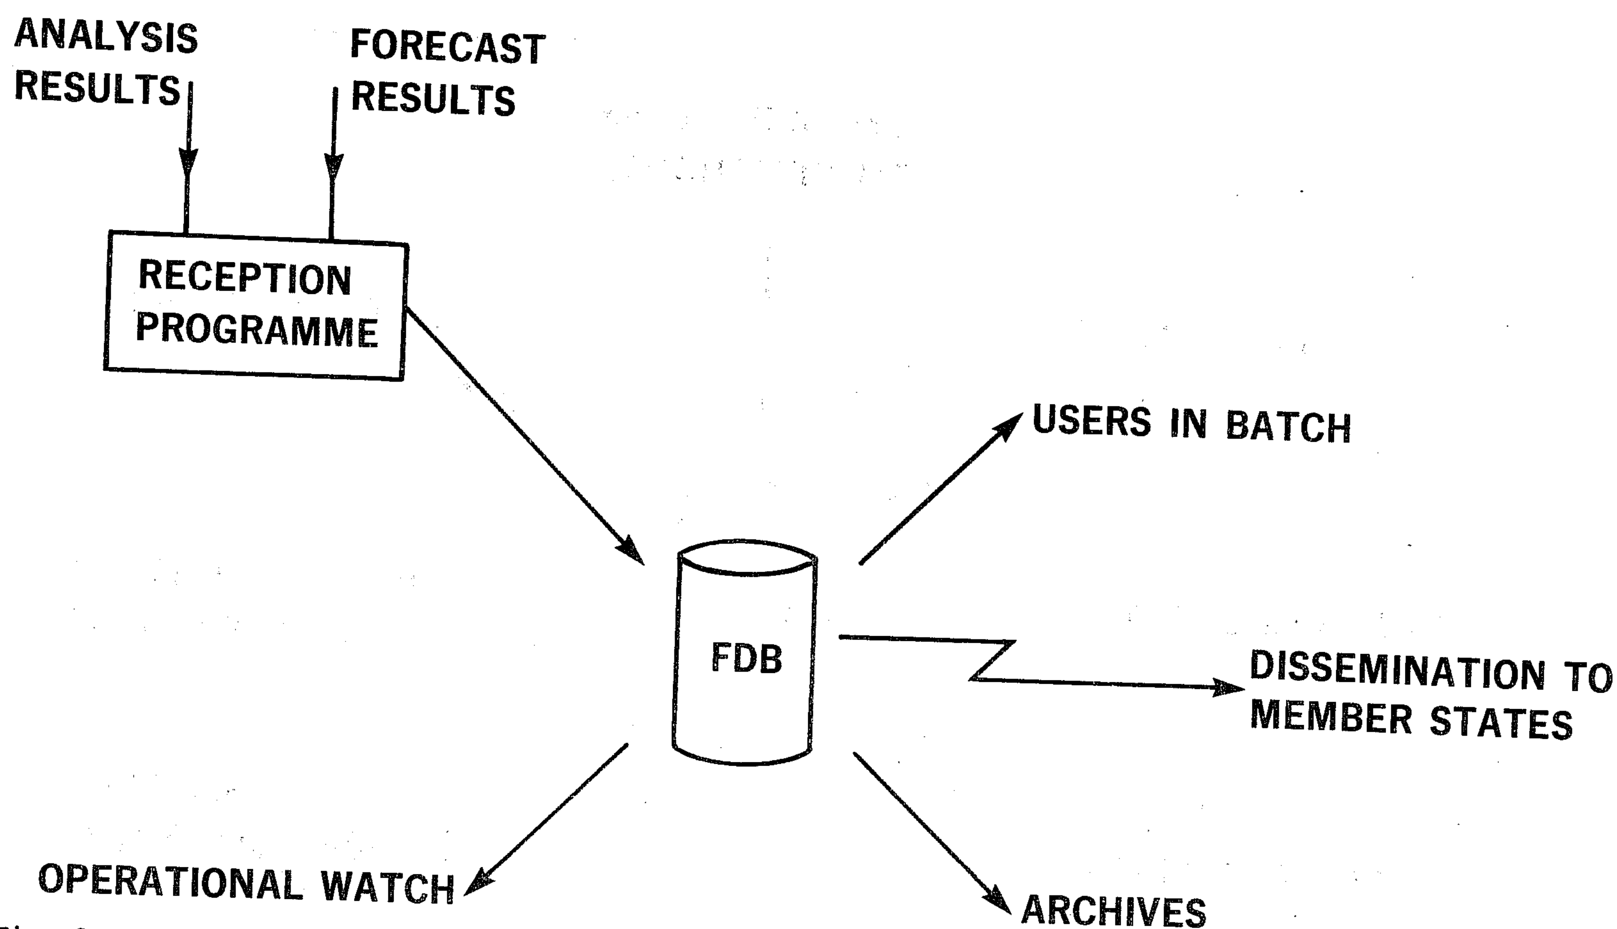

Fig. 3.3 Processes Using FDB

3.3. The Fields Data Base

The FDB is used to store forecast and analysis results in "field" form on-line. A "field" is defined as a representation of a single meteorological parameter at a single level over an area of interest. For ECMWF data two types of field are included - global fields, and fields of dissemination products. Fig. 3.3 shows the relationship between the FDB and processes within the EMOS system.

Since fields of data are somewhat larger than observations a different set of data base techniques have been used. The data are stored in the order in which they are produced in "word addressable" files. A standard length index is maintained as part of each file. The index records the word address and length of each field.

The data are stored in a packed form, using 15 or 20 bits for each data item. Accuracy is maintained by storing the minimum value of all of the data items of a field at full precision, and packing only the differences between this minimum value and appropriate data items. Data items represent grid point values for global surface fields and for dissemination products, and spherical harmonic coefficients for upper air values.

The index based access method is particularly suited to the random reception and random access required to be supported. It enables current fields to be examined and plotted at will, and simplifies the sorting required for ordered, sequential archiving. A comprehensive data base software package enables users to access the data in packed or expanded form (Fields Data Base Access, or FIDABA).

3.4 The Archives

The original ECMWF archives reflected closely the data of the on-line RDB and FDB. Data is archived from the data bases directly to tape. The ordering of the archive data is pre-determined; the data base access achieves the function of sorting into order. Archive/retrieval software is provided to facilitate access to the data (GETDATA - see Fig. 3.4).

These original ECMWF archives have performed well, but some deficiencies have become apparent. Access to observational data requires reconstruction of a considerable portion of the on-line RDB because of the mode of data

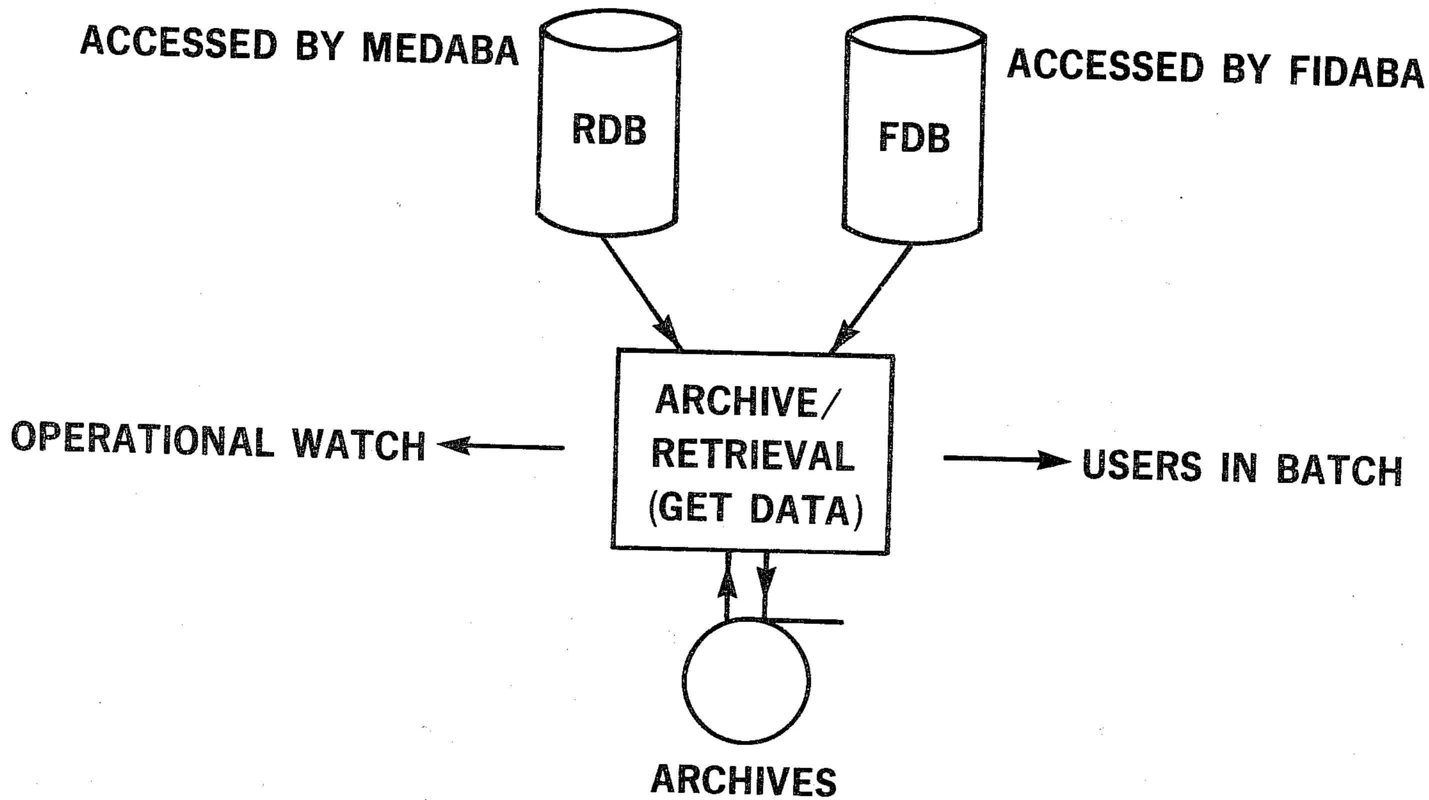

Fig. 3.4 The "GETDATA" Archive/Retrieval

representation. Access to fields data is especially expensive in terms of tape mounts and computer resources if time series data are required. Both observations and fields are stored in a representation unique to ECMWF, requiring transformation before being supplied to most external users.

To begin to correct these deficiencies, work is now well advanced on a Meteorological Archive and Retrieval Service (MARS). This will be described more fully in section 4.4 below. Currently fields data from 1 July 1985 have been added to both archives (GETDATA and MARS).

4. EMOS DATA BASES - FUTURE PLANS

4.1 Introduction

The efficiency of data handling is crucial to the total efficiency of EMOS from a technical standpoint. Data handling must

- provide the data required in a suitable form
- provide the data required to a suitable accuracy.

Efficient data handling requires that

- data transformations be minimised
- data transmissions be minimised
- data volumes be minimised.

Since interpolation of data introduces error, minimising data transformations is consistent with the requirement to provide data of suitable accuracy. Accuracy is also the limiting constraint on reducing data volumes. Data transmission may be reduced by

- storing on-line data, where practical, on the computer that will later use the data
- creating the data in a compact, machine-independent form.

Thus, the "data form", or "data representation" is the key both to the usefulness and the efficiency of use of the data. From this standpoint the <u>most desirable solution</u> would be to use a data representation for observations that would be

- efficient (binary)
- internationally agreed
- used for acquisition (i.e. by the GTS)
- used on-line for all purposes
- used for archiving

together with a data representation for fields that would be

- efficient (binary)
- internationally agreed
- used on-line for all purposes
- used for archiving
- used for dissemination to all users.

It is the intention of ECMWF to seek to follow this philosophy of adopting international standards for data representation as speedily as is practical. To this end ECMWF have participated in the activities of the WMO Commission for Basic Systems working groups, especially those of the Global Data Processing System (GDPS), and welcome the advances made in establishing the FM 92 GRIB to represent fields data, and those currently being made towards establishing a binary representation for observational data (binary observations, or BIOBs (binary universal form for reports, or BUFR)).

The second major influence on future plans for EMOS is the continued development of the ECMWF computer system. Fig. 4.1 indicates the projected configuration for 1986.

In a changing computer environment it is not only the support hardware that changes. The CRAY Operating System (COS) has developed from a crude, raw, unreliable system to a sophisticated, reliable system. Also, the power of the "front-end" CYBER system has decreased in proportion from about 15% of the total available power to less than 5%.

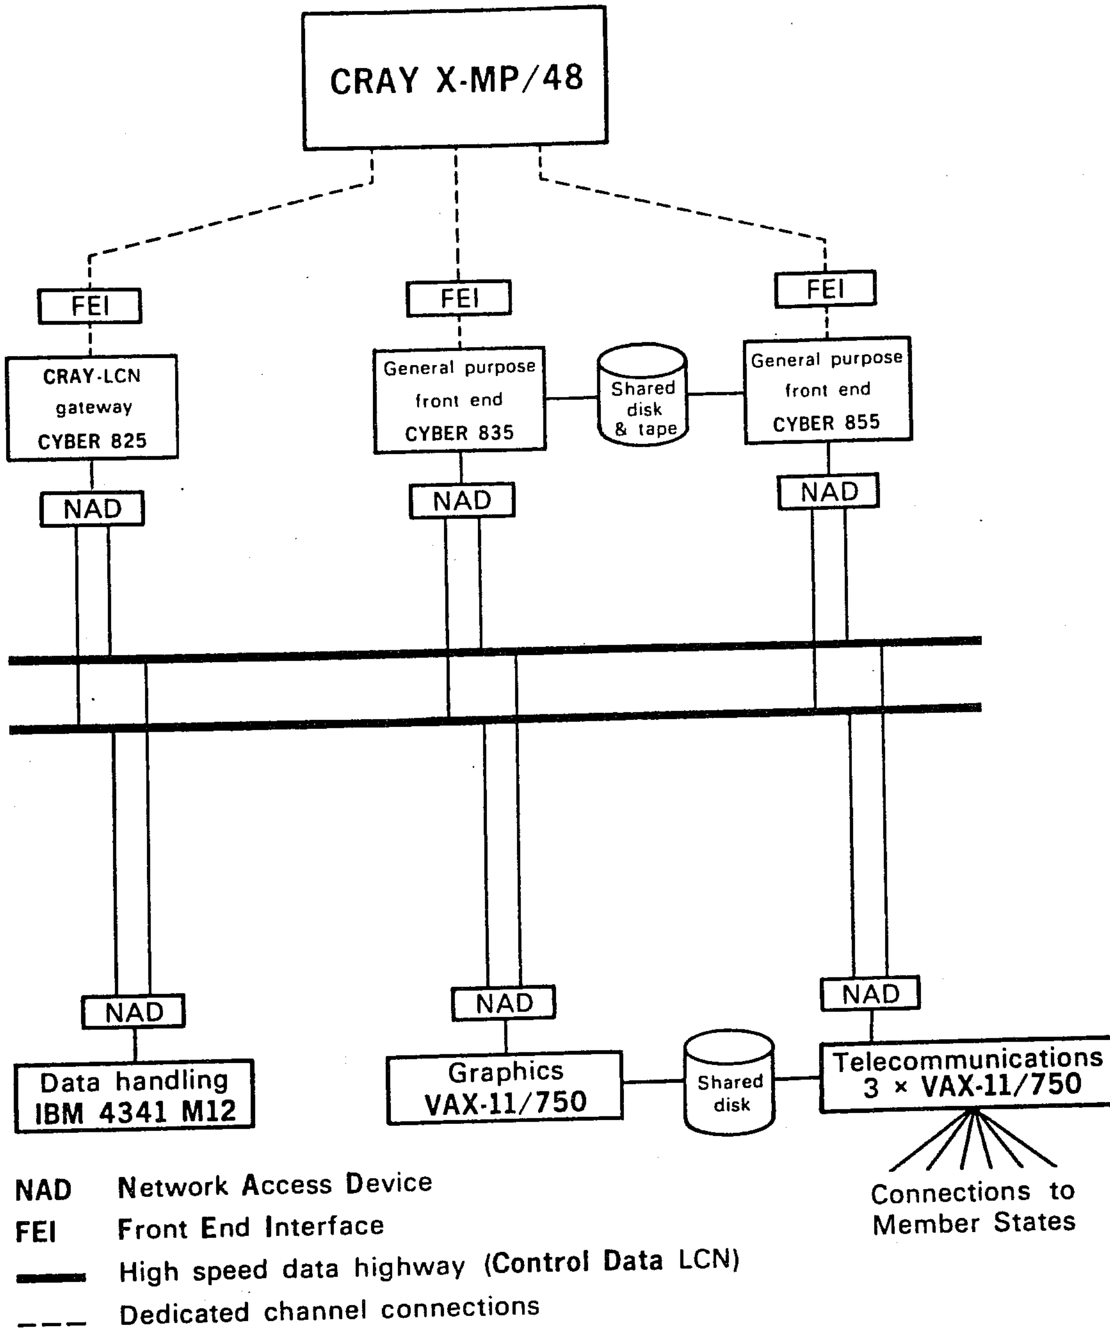

Fig. 4.1 ECMWF Computer System in 1986

The obvious conclusions to be drawn are:

- EMOS should not be designed to optimise the use of a particular computer configuration
- EMOS sub-systems should be transportable between machines where possible, to balance the operational work load
- EMOS interface data must thus be similarly transportable
- each sub-system should be moved as necessary to the computer best able to execute that system
- the on-line data bases should be located on the computers from which they will be most frequently accessed.

4.2. The Reports Data Base

Although the exact format for representing binary observations is not yet known, it is likely that

- sets of similar observations will be closely associated to facilitate data packing
- such sets could be regarded as "pseudo fields" in terms of storage characteristics
- either direct indexing, as currently used for fields data, or indexed sequential (with key) type access methods could be used
- direct indexing (with possibly two levels of index) is the probable basis for ECMWF's next RDB
- ECMWF may move the RDB to the CRAY.

In addition to reported information, the next RDB will provide increased facilities for

- quality criteria
- substituted values
- analysis feed-back
- bias values
- statistics.

4.3 The Fields Data Base

Data is already being created on the CRAY in GRIB representation for the new MARS archives. These data will be organised into a formal data base structure, using similar principles to those of the current FDB. This new FDB will be CRAY resident. A sub-set of the FDB will be transferred to the New Telecommunications Complex (NTC) to feed a demand-driven dissemination scheme.

The access software for the new FDB will include facilities to enable users of the MARS archives to access on-line or archived data through an identical interface. Data will be retrieved in the GRIB representation used for storage, or fully expanded, according to the user's request.

A "Meteorological Applications Graphics In Colour System" (MAGICS) will gradually replace the current graphical aspects of the Operational Watch. This will initially be implemented on the CRAY X-MP, will produce output in a metafile form, and the metafiles produced will be directed to appropriate machines and devices for post-processing and plotting. As GKS standards, including standards for metafiles, become internationally accepted, those standards will be incorporated into the metafile form used. Thus it is intended that metafiles be added to the fields data base, enabling them eventually to be included in the dissemination process.

4.4. The Archives

The Meteorological Archive and Retrieval Service (MARS) was introduced in 3.4 above. The aims of MARS are to

- use standard representations to store data
- enable simplified access to data
- continue to provide daily fields data for case studies
- enable more efficient access to time series of fields data
- minimise tape mounts by using a hierarchy of on-line disks, on-line cartridge, and tape storage efficiently
- centralise data handling on a specially procured data handling machine capable of supporting state of the art storage devices.

MARS will achieve these aims by

- using GRIB for fields data
- using "BIOB" for observations
- storing time series and daily data
- using a hierarchical file migration and management package called "Common File System" (CFS)
- interfacing to CFS on a dedicated IBM computer equipped with appropriate storage devices.

The MARS software package enables users to specify required data in meteorological terms, without reference to the physical location of the data. The MARS system will be capable of retrieving the data from MARS archives or from on-line data bases, as appropriate. Fig. 4.2 illustrates the principles of the MARS archive.

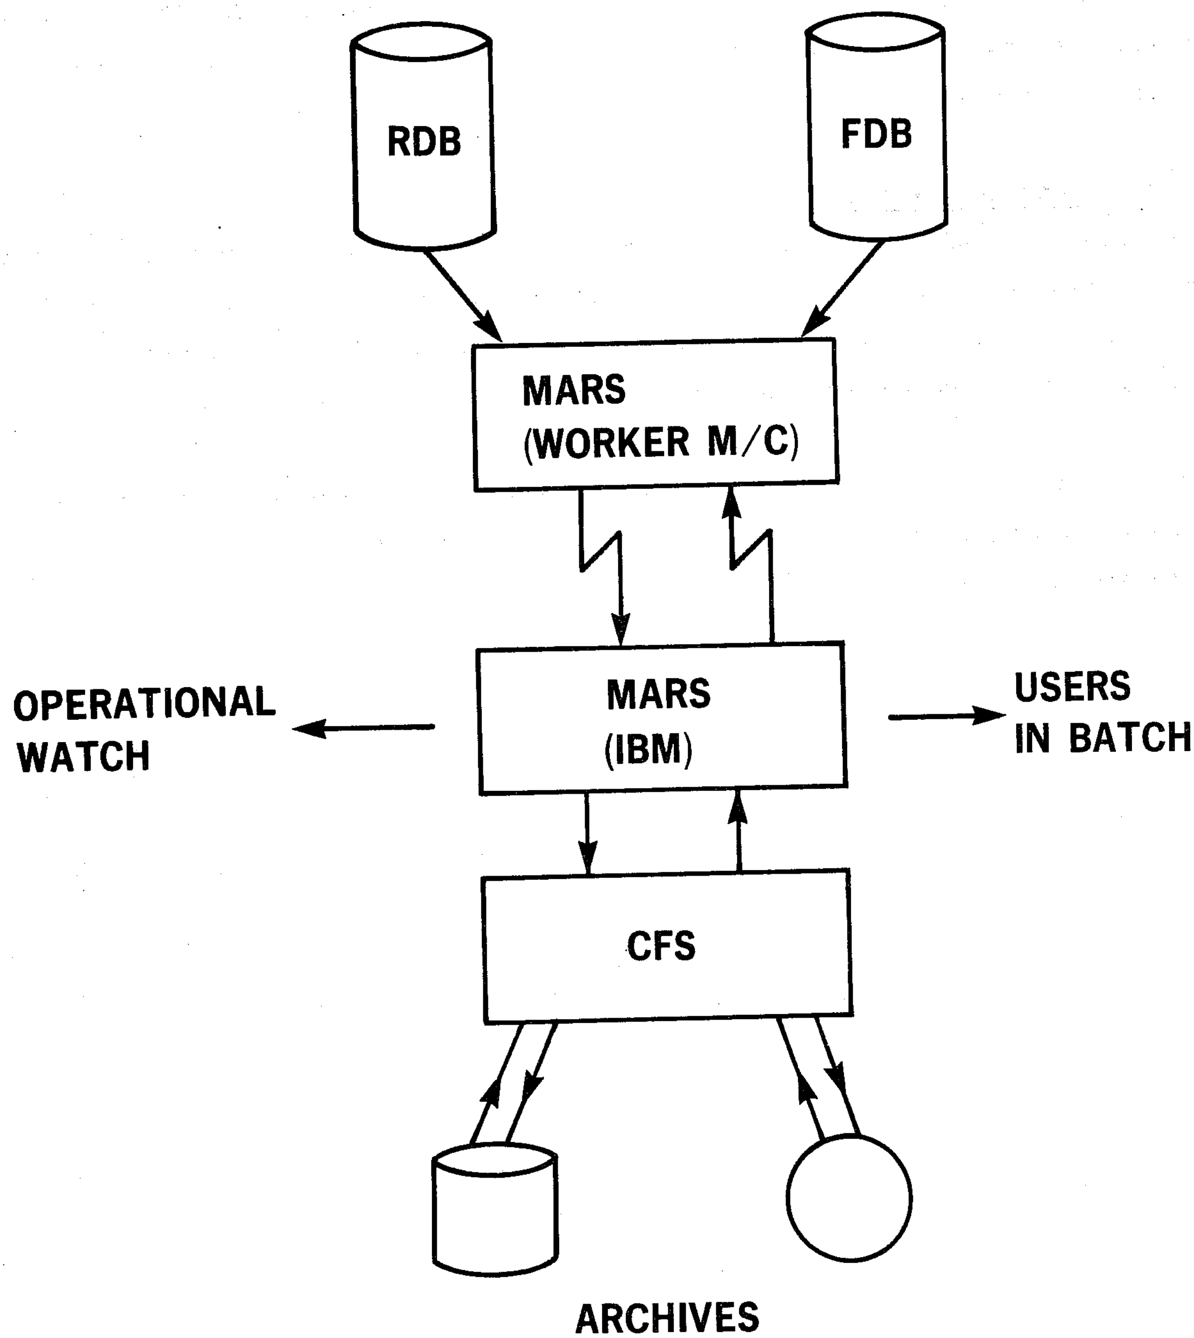

Fig. 4.2 The MARS Archive/Retrieval

5. EMOS APPLICATIONS SYSTEMS

5.1 Introduction

Section 2 above introduced the Applications Systems within EMOS. It is beyond the scope of this paper to describe each system in detail. In the following, an outline of each system will be presented with some indications of plans for future changes.

5.2 Data Acquisition

As observational data is received from the GTS it must be stored as raw material for the pre-processing system. Care must be taken to guard against outages, and status information must be updated to take account of the data received. Fig. 5.1 indicates the main processes involved.

Some changes to the data acquisition system will be necessary to accommodate the new telecommunications complex (NTC); however, the strategy involved will remain broadly unchanged.

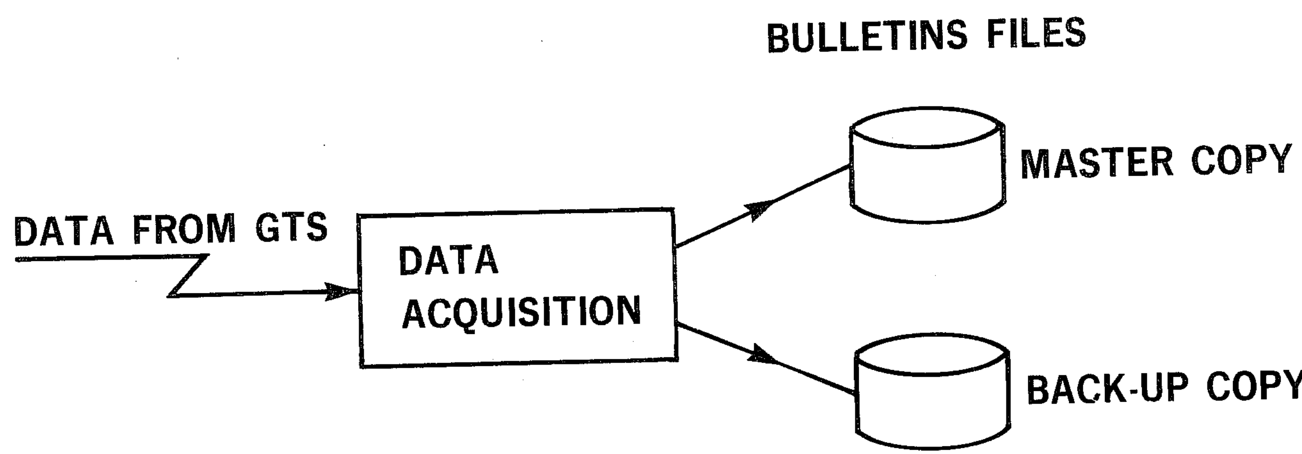

Fig. 5.1 Data Acquisition

5.3 Pre-Processing

The bulletins accumulated by the data acquisition are first checked and decoded. Each bulletin is checked for coding errors, rejected bulletins being copied to an error file. The many WMO FM character codes are classified into 7 major observation types. Each observation type is decoded onto a separate interface file for input to the quality control programme. The error file is passed to the operational watch system for manual correction. Corrected bulletins are fed back into the decoding programme.

The decoded reports are next checked for internal consistency, and, in the case of ship reports only, for temporal consistency. Where discrepancies are noted these are flagged, and, if possible, corrected values are computed and added as substitute values. Finally, the fully checked reports are added to the RDB. Fig. 5.2 illustrates these processes.

During the last 3 years much of the decoding has been re-written in FORTRAN 77. However, due to the inability of CDC FORTRAN to manipulate ASCII characters, the CCITT Alphabet No. 5 format of the bulletins is converted, character by character, into an integer representation.

Future plans are to provide alternative character handling code, using the character processing features of FORTRAN 77. The code will be made as machine independent as possible to enhance its transportability. The RDB will be re-designed to conform with new trends and standards for binary representation of observational data. Decoding will be enhanced to accept data in the new forms. Consideration will be given to moving the pre-processing system either to the CRAY or to the NTC. Temporal checks may be extended to additional observation types.

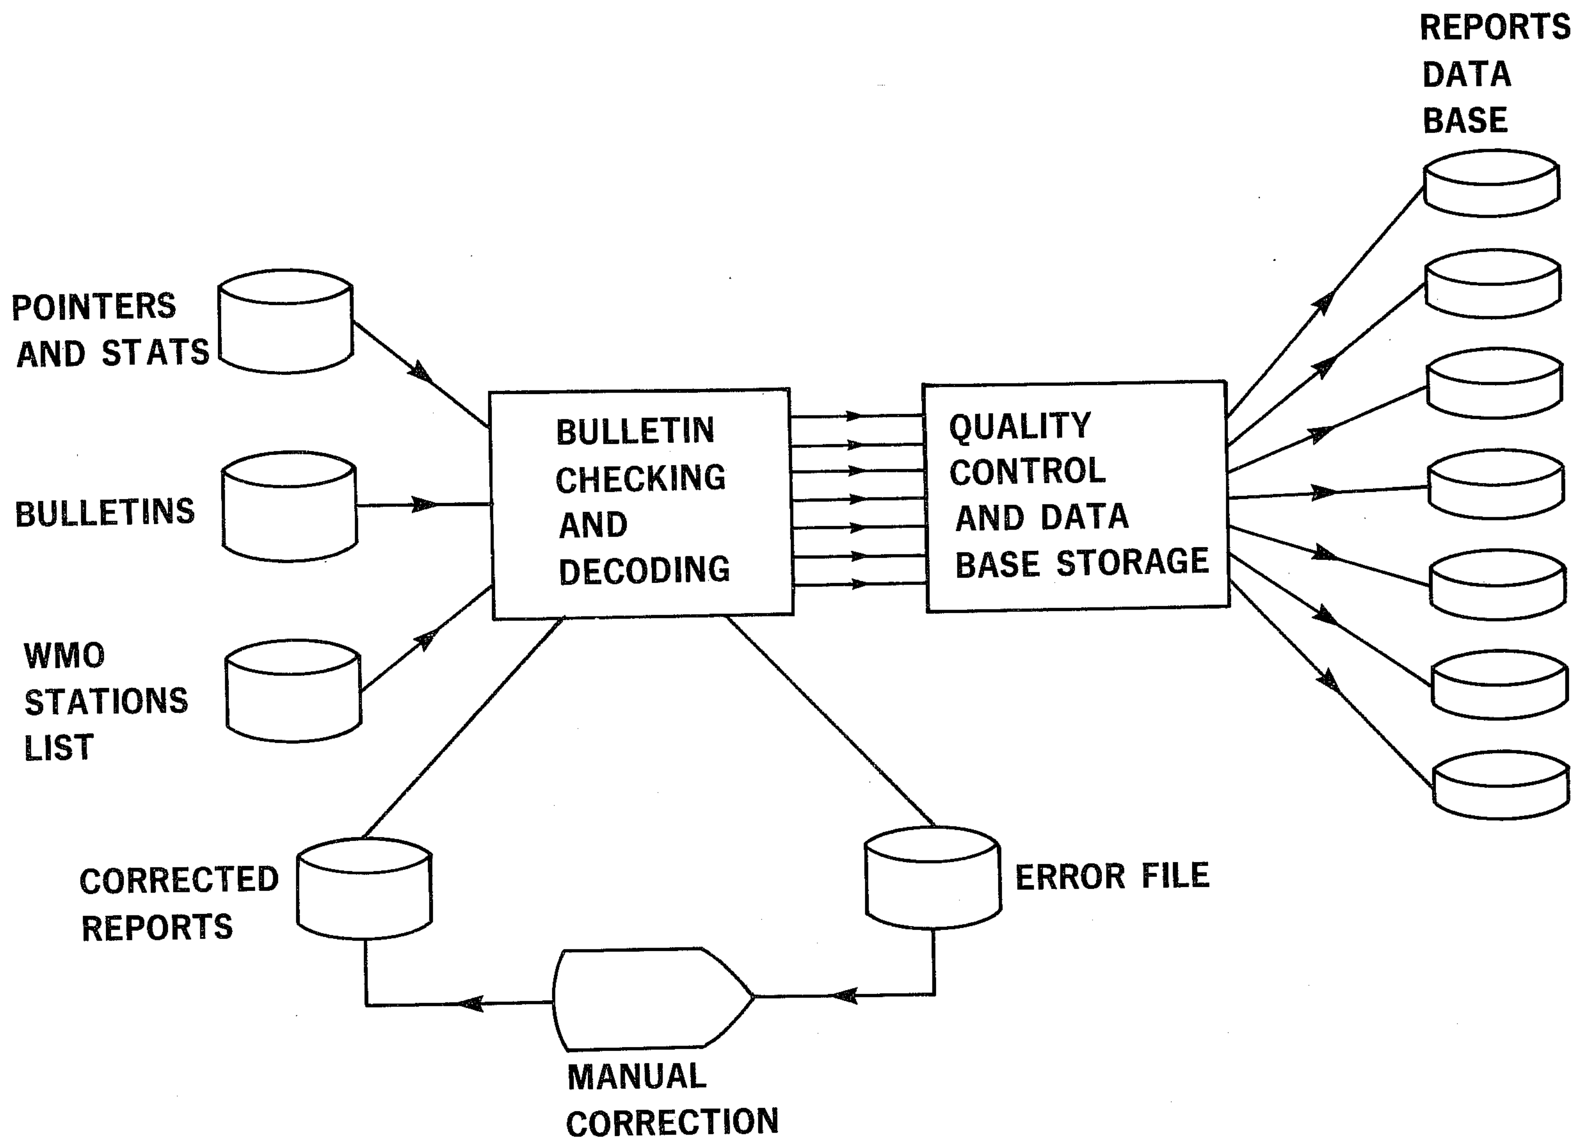

Fig. 5.2 Pre-Processing

5.4 The Analysis

An extraction programme obtains and sorts the appropriate reports required by the analysis from the RDB.

These reports are then analysed, using a sophisticated optimum interpolation data assimilation scheme. This involves first checking the reports for mutual consistency, then computing increments to be applied to correct a first-guess field. The scheme takes account of reported values, first-guess information, persistence, and climatology. As first-guess, a 6 hour forecast based on the previous analysis is used. In this way, reported data from data rich areas is "assimilated", or carried forward by the forecast model processes to data sparse regions.

In addition to the analysis results, the analysis produces information on the quality of reports for feed-back to the RDB, and departure values comparing reported values with those of the first-guess, analysis, and initialisation. Fig. 5.3 illustrates the analysis process.

Currently the ECMWF analysis is being completely re-written, to enable a more flexible, higher resolution scheme capable of exploiting the multi-tasking features of the CRAY X-MP to be used. Projected changes to the RDB will also require corresponding changes to the reports interface to the analysis. The information produced by the analysis has expanded considerably over the years, especially with respect to observational quality criteria. Attempts will be made to enhance the feed-back to the new RDB, and to accommodate the observational departures. There is a good case for a direct interface between the RDB and the analysis.

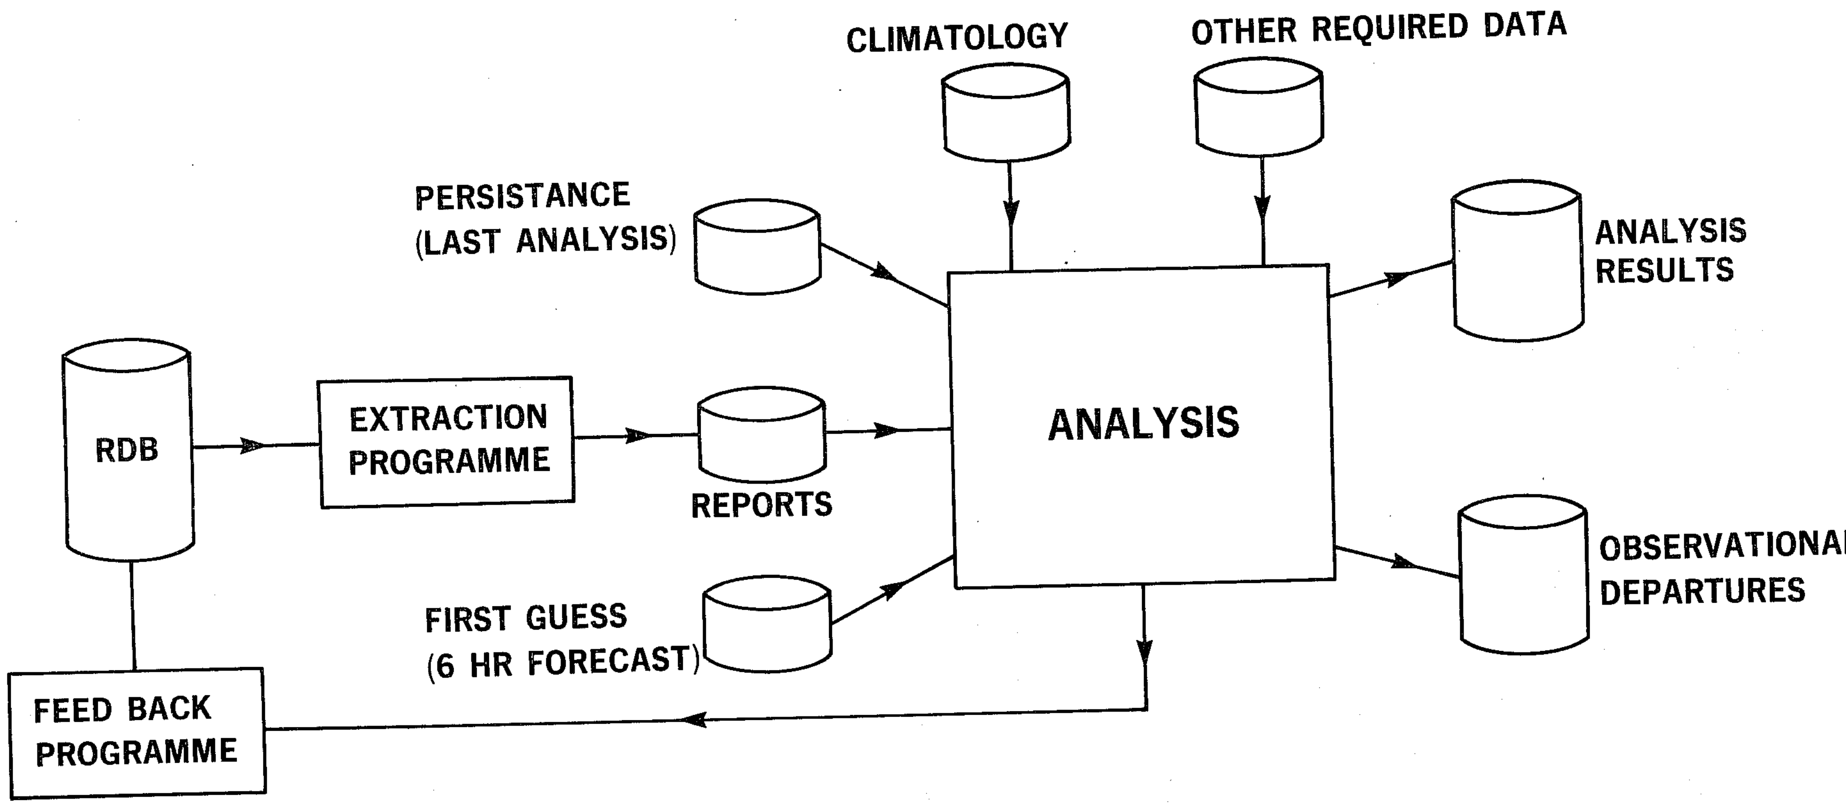

Fig. 5.3 Analysis

5.5 The Forecast

There have been 2 major changes to the forecast during the history of ECMWF operations. In each case modifications to the interfaces with the analysis and post-processing systems enabled the changes to be introduced with a minimum of disturbance to the overall EMOS system.

Currently a T106 spectral model is used (see Fig. 5.4). Analysed data is first subjected to a process of normal mode initialisation. The forecast is then run, using the full dual-processing facilities of the CRAY X-MP/22.

Results of the 6 hour forecast are passed to the next analysis, to be used as a first-guess. These and other results form the input data for the post-processing system.

Future plans for the forecast include the use of all four processors of the CRAY X-MP/48 when this becomes available. Since the T106 forecast model was introduced into operations as recently as 1 May 1985, further radical changes are not envisaged in the immediate future. It is likely that some enhancements will be made by mid 1987.

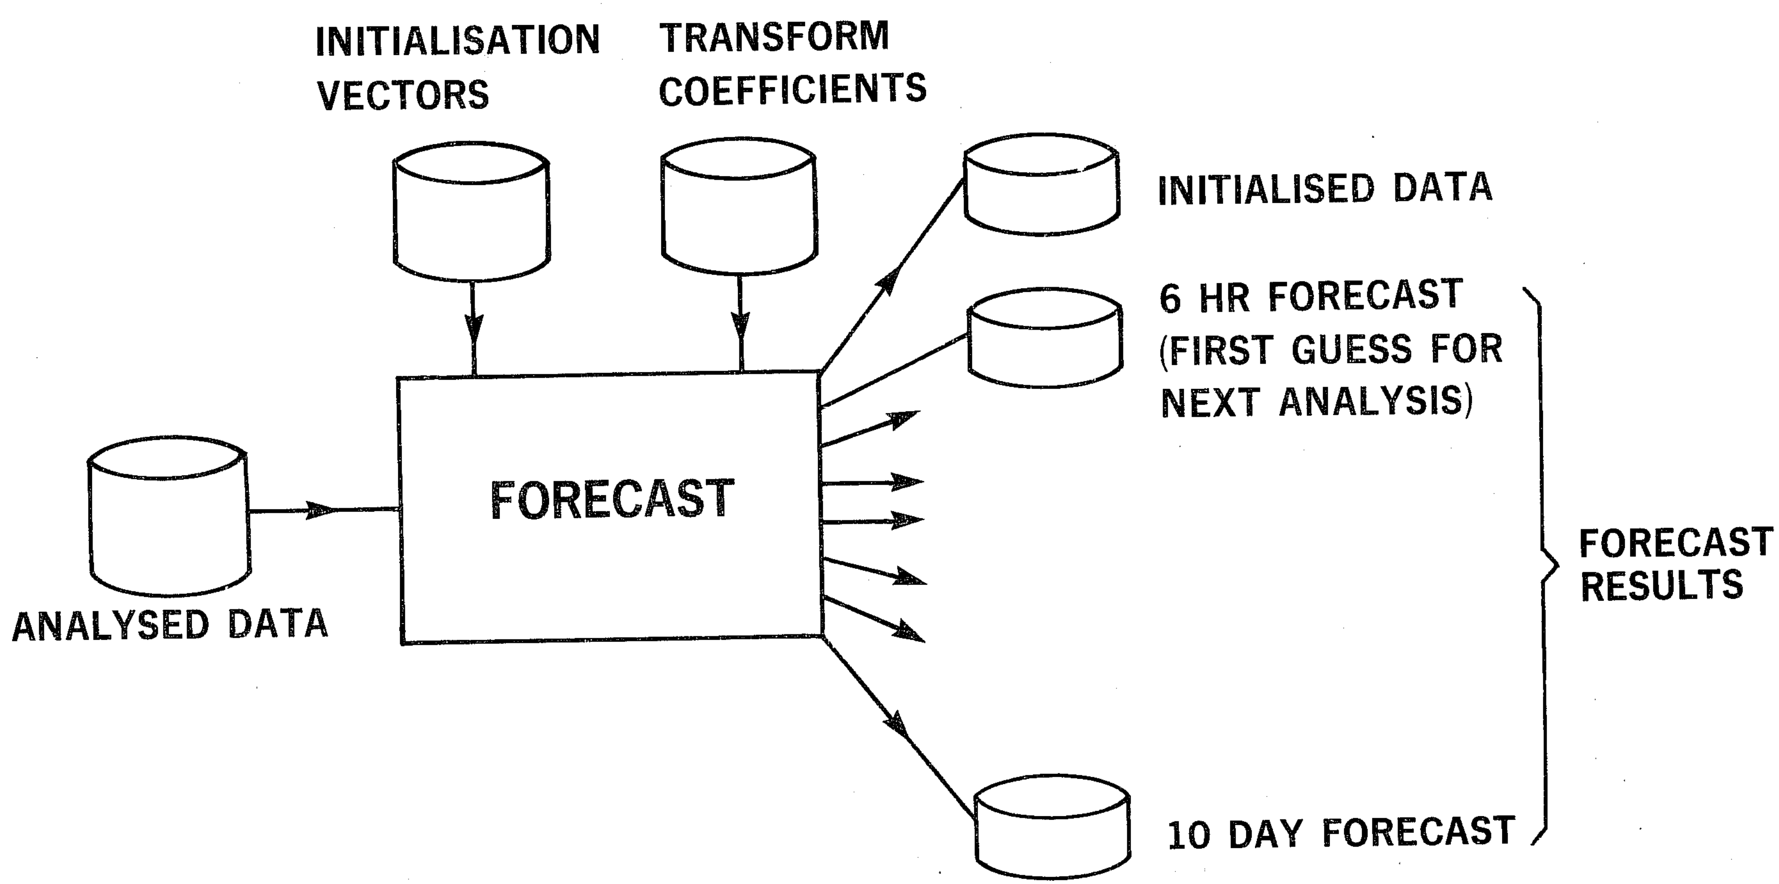

Fig. 5.4 Forecast

5.6 Post-Processing

The current post-processing scheme is illustrated in Fig. 5.5.

The forecast and analysis results are produced in a form suitable for re-starting the forecast model; this form is most unsuitable for the majority of users. In consequence, the data is transformed, and, if necessary, interpolated to produce global fields of grid point values and spherical harmonic coefficients. Two representations are derived:

- fields in GRIB form for input to MARS archives
- fields in GETDATA archive form for input to the CYBER reception and transformation programme.

The reception and transformation programme writes global fields to the FDB, and generates dissemination fields according to the contents of a post-processing requirements file. As part of the modifications made on the introduction of new forecast models some transformation and interpolation is also required, especially the conversion of Gaussian grids to standard latitude/longitude grids. Where polar stereographic products are required this conversion is also performed.

Future plans are indicated in Fig. 5.6. Using GRIB as the data representation within the FDB would enable data base software to be used directly within the CRAY interpolation and transformation programme. A sub-set of the FDB on CRAY would be transmitted to the NTC, where it would be used as a basis for the extraction of dissemination products. This would result in the bulk of the post-processing being moved from the CYBER computers to the CRAY X-MP. Since these processes can be highly vectorised, and since the results are required on CRAY for plotting and MARS archiving, this is not unreasonable. Current calculations indicate that the programmes involved will use small amounts of memory, and small amounts of CPU. It can thus be arranged that they share memory with the forecast, using the CPU time available surplus to the forecast's requirements. Thus an extremely efficient job mix would result.

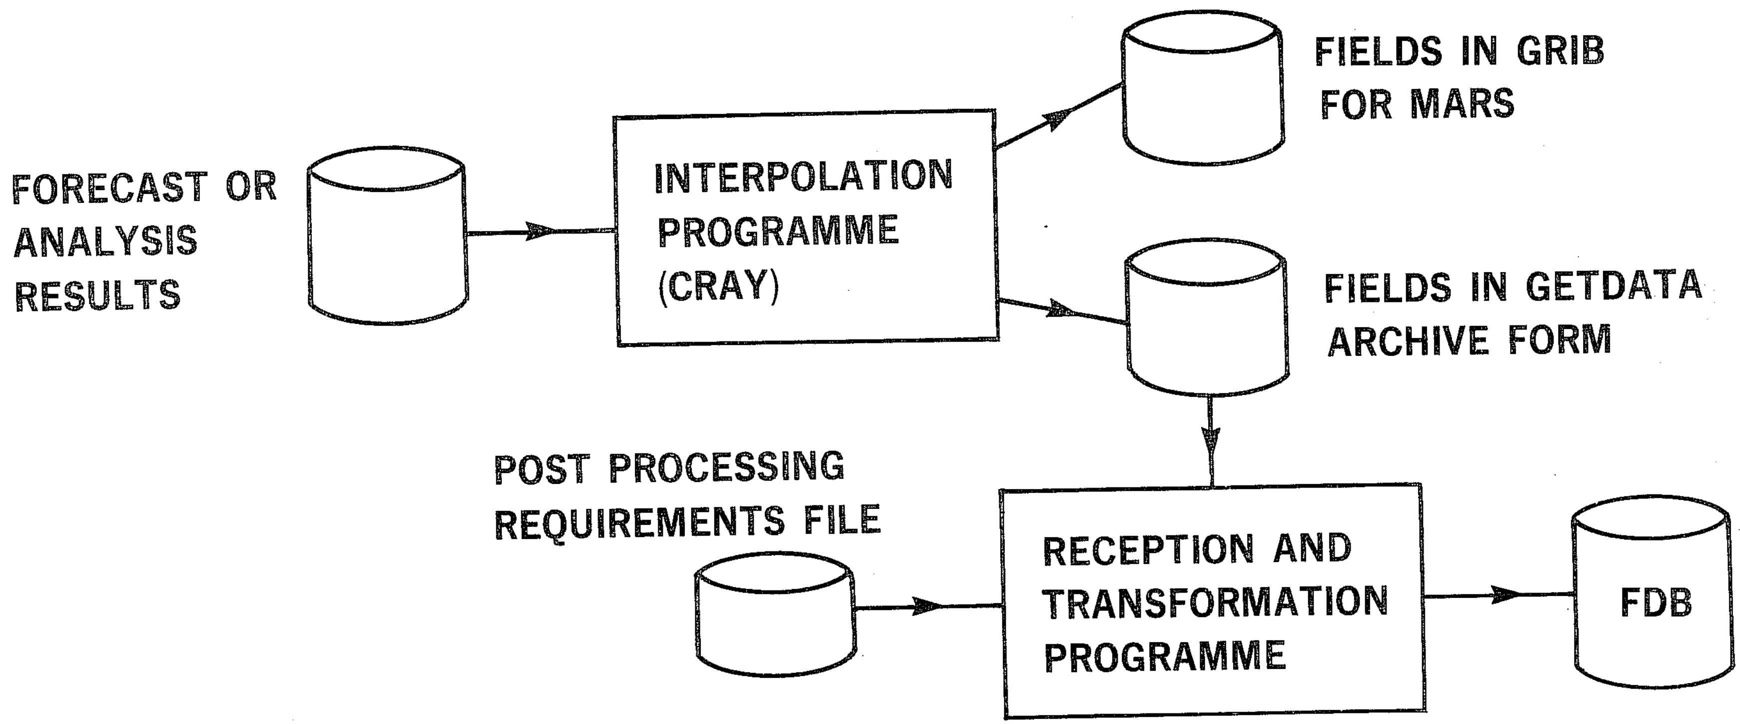

Fig. 5.5 Post-processing - Current Status

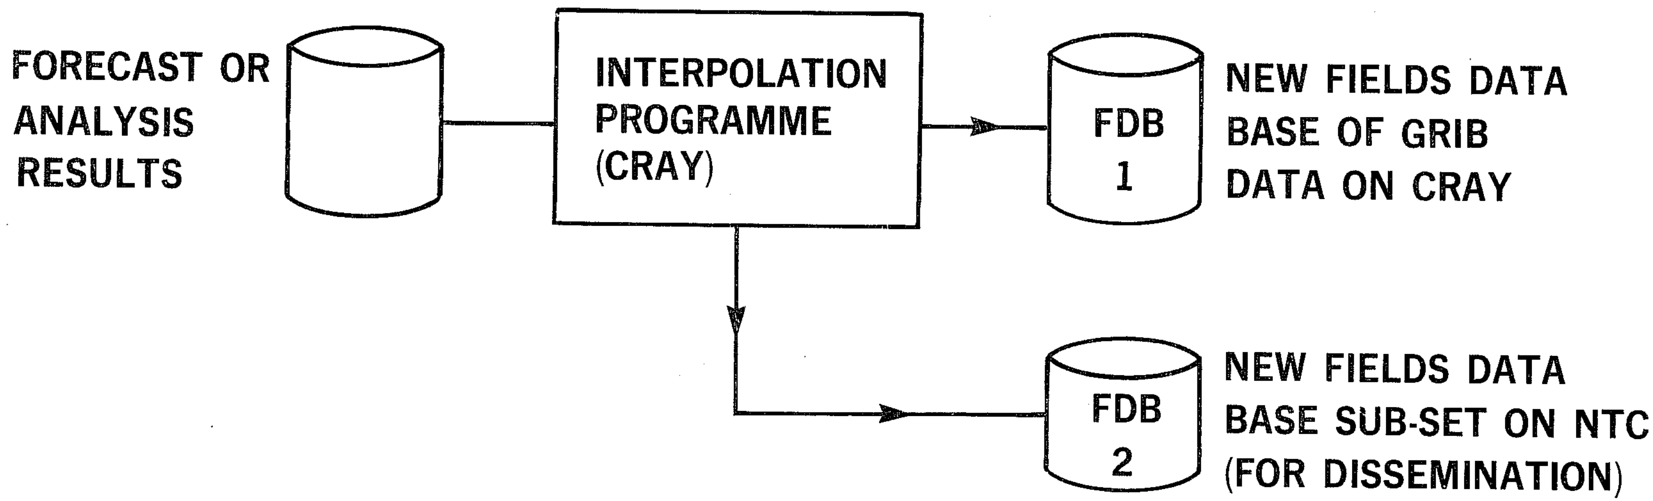

Fig. 5.6 Post-processing - Planned System

5.7 The Dissemination System

The current dissemination system, illustrated in Fig. 5.7, is straight forward and comparatively simple. The dissemination products, formed during the post-processing, already exist in the FDB. The dissemination programme extracts appropriate products according to a catalogue, appends headings, and routes the resulting files to the recipient Member States.

This scheme has the advantage that each user receives a "tailor made" service. The disadvantages are:

- only products from a specific (but comprehensive) product set are available
- the product catalogue is maintained by ECMWF - changes must be notified, and implementation performed by ECMWF staff
- bit format products do not yet conform to WMO standards
- products cannot be obtained on demand.

Nevertheless, the scheme has worked well, and in excess of 12,100 products are disseminated daily.

The aim of future dissemination strategy will be

- to provide a comprehensive, flexible service
- to support an element of demand-driven dissemination
- to use WMO approved standard data representations
- to enable users to request dissemination and archive products through an identical interface
- to provide dissemination by satellite when support technology so permits.

It should be noted that the NTC is expected to be the most reliable item of computer hardware from the viewpoint of the Member States. Thus it is intended to move appropriate FDB data to the NTC as it is produced, as raw material for the dissemination system. From this information standard, predefined products and products requested specifically on demand will be created and transmitted by the dissemination programme. Products for satellite transmission will also be created and transmitted to the appropriate ground station. Products will be expanded to include suitable products for boundary conditions for limited area models, and graphical products. Fig. 5.8 indicates some of the future plans for dissemination.

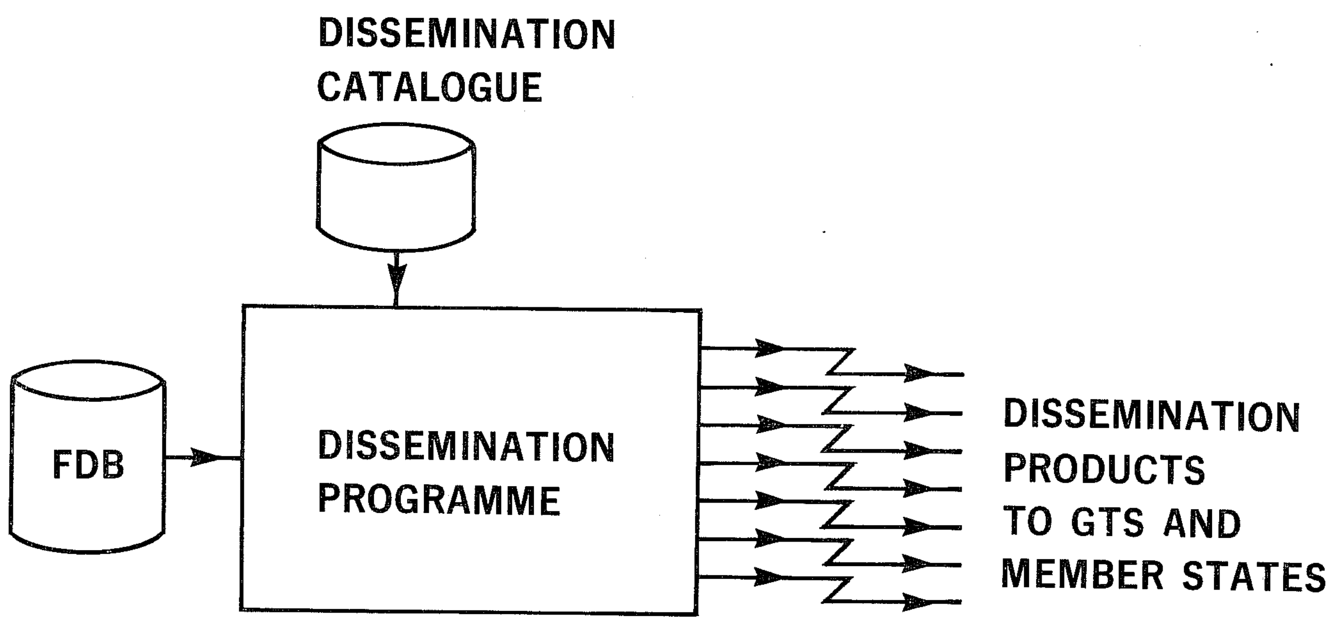

Fig. 5.7 Dissemination - Current Status

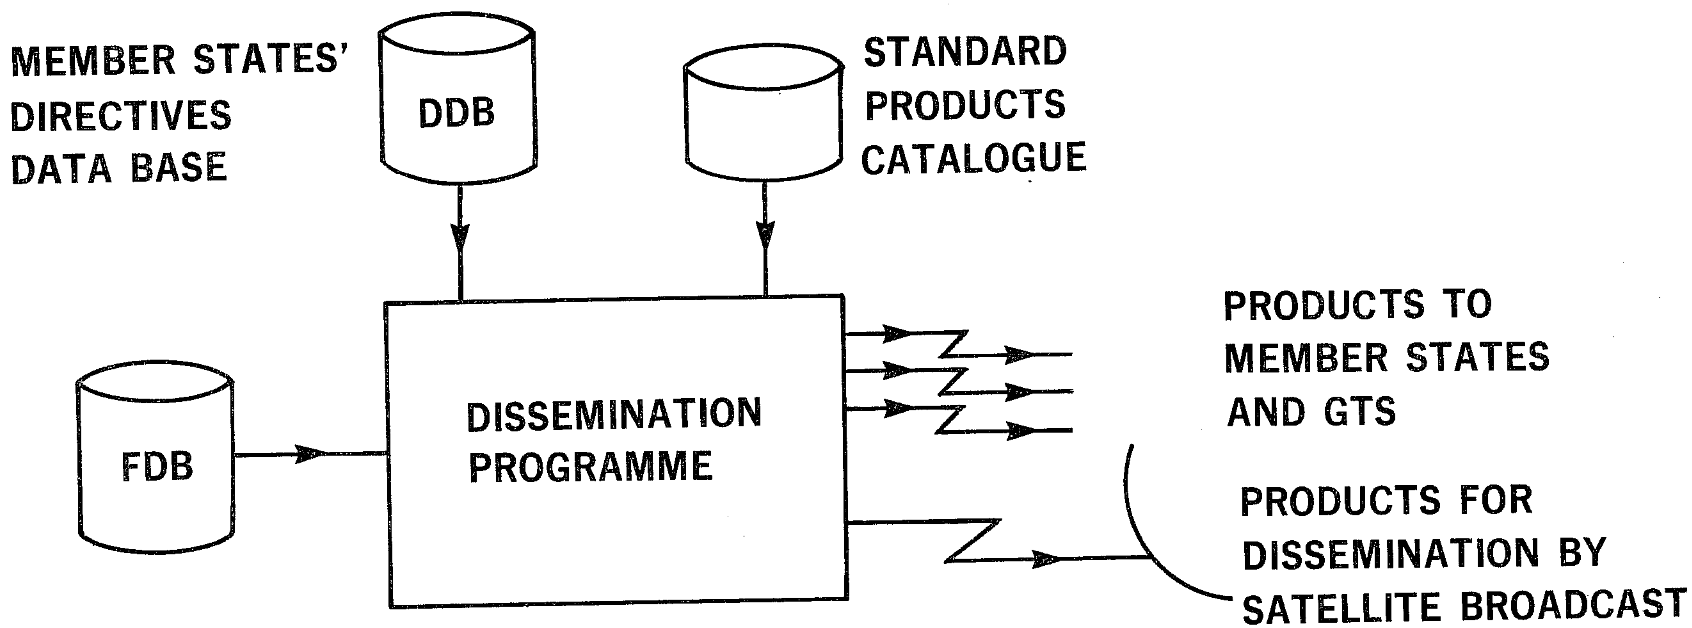

Fig. 5.8 Dissemination - Future Plans

5.8 Archive/Retrieval

The current and planned archival/retrieval services have been covered in sections 3.4 and 4.4 above.

5.9 The Operational Watch

Whereas the systems described above follow a set pattern of processing in an automated way, many features of the Operational Watch require human invocation. Broadly speaking, OPWATCH provides a set of software tools for monitoring, interrogation, plotting, correcting, and extracting operational information with respect both to the progress of EMOS and the data in use (see Fig. 5.9).

To this end OPWATCH uses the access software to the on-line data bases and the archives; it requires considerable support from graphics software; it needs to be used by meteorological rather than computer staff, and thus must interface to the user in meteorological rather than computer based terms.

Currently, new graphics packages are being developed, and will be extended to interface to the planned new data representations of the RDB and FDB. Colour will be used increasing to enhance the usefulness of the graphical products. On-line, intelligent graphical devices are being integrated into the system to enable foreground manipulation of graphical displays. The resulting system will use the GKS standards currently being agreed for graphics, maximising the supportability and adaptability of both the system and its products.

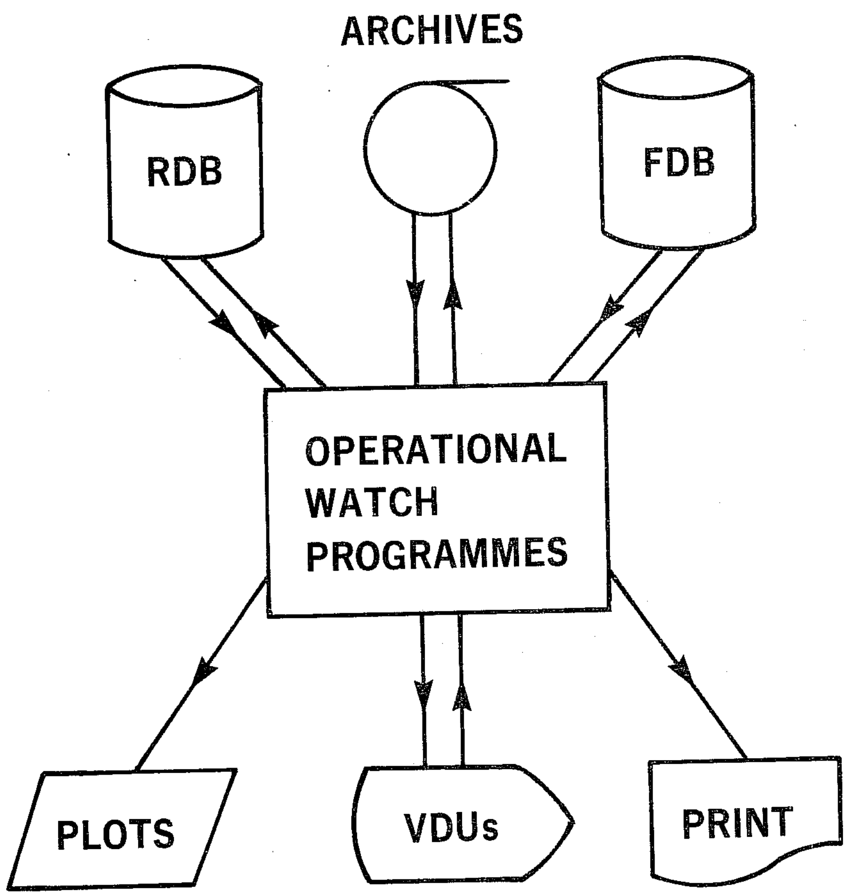

Fig. 5.9 The Operational Watch System

5.10 The Supervisor-Monitor-Scheduler (SMS)

The computer jobs which comprise the individual processes or "tasks" within EMOS need to be submitted for execution at appropriate points within the daily operational schedule. Since there are in excess of 400 different jobs involved, some of which needing to be run as many as 40 times during the operational cycle, manual submission would be extremely inefficient, and manual scheduling would be a tiresome task - so complicated are many of the conditions for task submission.

To overcome the problem of automating the supervision of EMOS, the SMS system was implemented. A control programme provides facilities for the maintenance and updating of a task table (see Fig. 5.10). This table contains the conditions required to be met prior to the submission of each task. Tasks are grouped into families, to simplify the conditional structure, and to aid task identification. Thus a single task may appear within many families - unique identification requiring the family/task combination. Events are set by tasks, and the setting of events recorded by SMS. Such events contribute to the conditions that may be set for the submission of subsequent tasks. Other factors taken into account may be

- time of day
- day of week
- date of month.

Tasks which fail are normally intercepted and re-submitted automatically once only; this automatic re-submission can be suppressed. By default, only failed job output is printed; again, this can be annulled if desired. SMS enables up to 9 operational suites (i.e. EMOS systems) to be controlled simultaneously.

The current SMS programme operates at a CYBER control point, and is coded in a highly machine dependent manner. However, it is well defined and documented, and could be implemented in a more general, machine independent way.

SMS has been the corner-stone on which the reliability of EMOS has been built. It has not only succeeded in providing automated process control it has been invaluable in the monitoring of the EMOS system and the reporting and subsequent correction of problems.

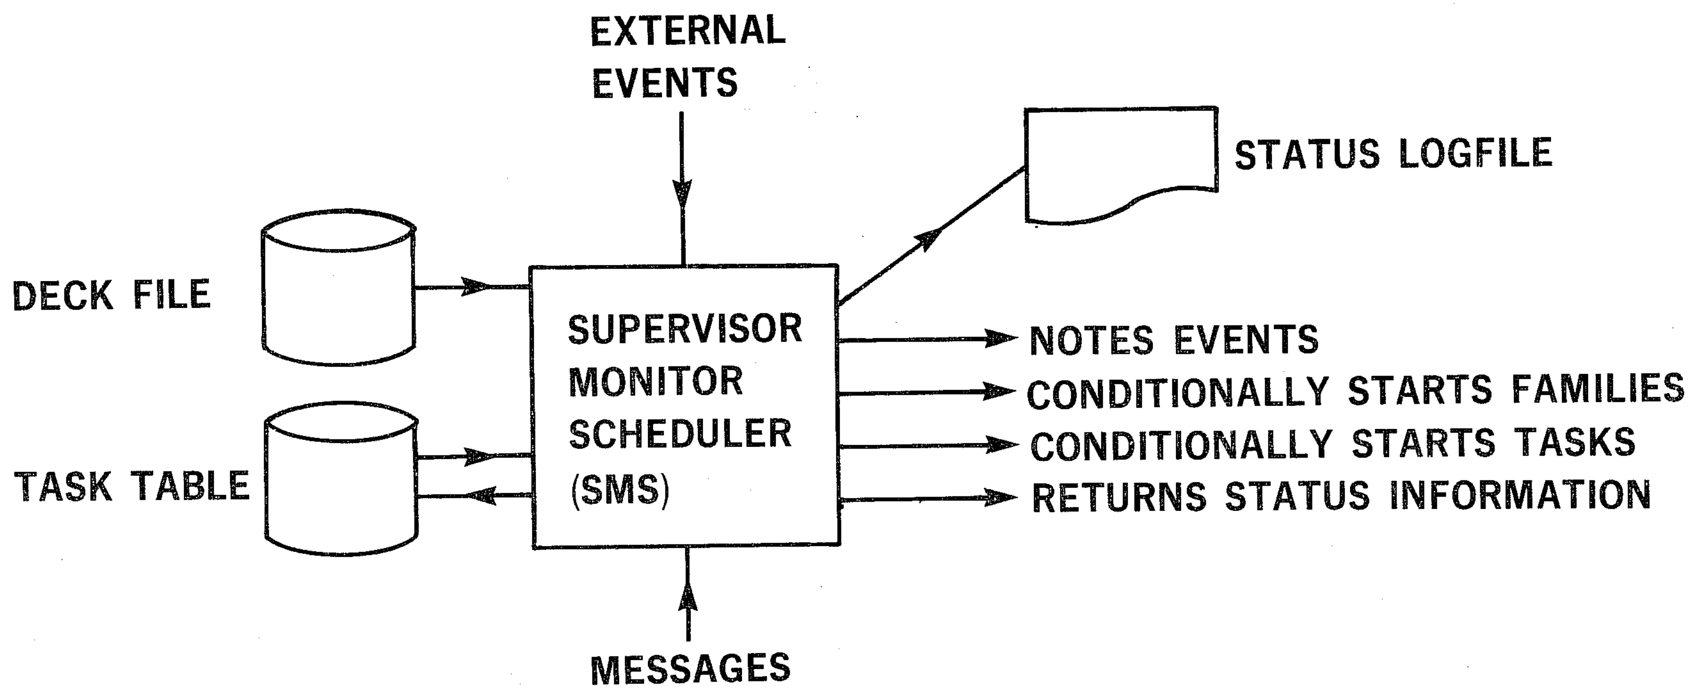

Fig. 5.10 The Supervisor-Monitor-Scheduler

Future plans include the addition of the ability to update SMS task tables without stopping a suite, the incorporation of facilities to delegate control of some families to remote computers, and the ability to control more suites for research purposes.

6. CONCLUDING REMARKS

6.1 The Importance of Design

The success of EMOS has been in no small measure due to the soundness of the original design. The original system was designed as a whole, then in detail with respect to the component systems and data entities. Subsequently, modifications were inevitably required from time to time. While no modification is consciously made in violation of the design principles, the sum total of such modifications eventually leads to a corruption of the parity of the overall system. Future plans require the total EMOS design to be re-worked in the light of changed and changing conditions. Only then can individual component systems be re-designed.

6.2 The Importance of Interfaces

The component systems are dependent on other systems only where they share a common interface. This requires that careful planning be given to the design and modification of all such interfaces.

6.3 The Importance of Standards

Standards for data representation, where applied, not only enhance the usefulness of the data but lead to well defined, stable interfaces between systems. Usually data representations accepted as international standards will be well designed and efficient. Using them wherever practical can often lead to a rationalisation of data formats and a reduction in costly transformations.

6.4 The Importance of Process Control

Efficient process control is at the heart of every good operational system. The concept of SMS has proved invaluable in this respect.

6.5 The Importance of Documentation

Documentation is essential to the maintenance and understanding of any complex system. EMOS documentation, though not of the highest standard, has proved sufficient. In particular, three classes of documentation have proved essential:

- design papers
- source code documentation
- operational manual.

6.6 The Importance of Integration

There can be no weaknesses in an operational suite. The best forecast in the world is useless without initial data of appropriate quality. It is also useless if its results cannot be disseminated. Each component system plays its part, and that part is vital to the whole. An inefficiency in one system makes the total system inefficient. Thus the component systems must be carefully integrated. Their requirements must be balanced, and they must be allocated the appropriate resources to their needs. Critical path analysis over the total system is the only means of identifying bottlenecks, and of optimising performance. Modifications must not be introduced into one area without proper consideration of the effects they may have on other systems. Only a properly integrated system can produce the required results, make them available where they are required and in an appropriate form, and complete the processes efficiently by the required time.